

Editors:

Elena Irimia • Verginica Barbu Mititelu • Maria Mitrofan
Daniela Gîfu • Diana Trandabăţ • Dan Tufiş • Dan Cristea

**LINGUISTIC RESOURCES AND TOOLS
FOR NATURAL LANGUAGE
PROCESSING**

Proceedings of the 20th International Conference,
Bucharest, 8-10 October 2025



Editura Universităţii “Alexandru Ioan Cuza” din Iaşi
2025

Organisers

“Mihai Drăgănescu” Research Institute for Artificial Intelligence,
Bucharest

Institute for Computer Science, Romanian Academy, Iași branch
Faculty of Computer Science, “Alexandru Ioan Cuza” University of Iași
“Alexandru Philippide” Institute of Romanian Philology, Romanian
Academy, Iași branch

Romanian Association of Computational Linguistics
Academy of Technical Sciences

Under the auspices of the Romanian Academy

Endorsed by RACAI4RO Clarin-K Centre

This volume was published with the support of the Ministry of Education
and Research.

ISSN: 3045-2074
ISSN-L: 3045-2074
ISSN (online): 3045-2538

PROGRAM COMMITTEE

Adriana Stan – The Technical University of Cluj-Napoca, Romania

Anca Dinu – Faculty of Foreign Languages, University of Bucharest, Romania

Anca Vasilescu – Transylvania University of Braşov, Department of Mathematics and Computer Science, Romania

Constantin Orasan – University of Surrey

Corneliu Burileanu – University POLITEHNICA of Bucharest, Romania

Dan Cristea – Institute for Computer Science, Romanian Academy, Iaşi

Dan Tufiş – Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

Daniela Gîfu – Institute for Computer Science, Romanian Academy, Iaşi & UAIC-FII

Diana Trandabăţ – Faculty of Computer Science, “Alexandru Ioan Cuza” University

Diana Zaiu Inkpen – School of Electrical Engineering and Computer Science, University of Ottawa, Canada

Elena Irimia – Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

Elena Isabelle Tamba – ARFI, Philippide Institute of Romanian Philology

Gabriela Haja – ARFI, Philippide Institute of Romanian Philology

Horia Cucu – University POLITEHNICA of Bucharest, Romania

Horia-Nicolai Teodorescu – Institute for Computer Science, Romanian Academy, Iaşi & Faculty of Electronics, Telecommunications and Information Technology, “Gheorghe Asachi” Technical University of Iaşi

Ivelina Stoyanova – Department of Computational Linguistics, Institute for Bulgarian Language, Bulgarian Academy of Sciences, Bulgaria

Liviu Dinu – Department of Computer Science, Faculty of Mathematics and Computer Science, University of Bucharest, Romania

Mădălina Chitez – Department of Modern Languages and Literatures, CODHUS – Centre for Corpus Related Digital Humanities, West University of Timisoara, Romania

Maria Mitrofan – Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

Marius Popescu – Department of Computer Science, Faculty of Mathematics and Computer Science, University of Bucharest, Romania

Mihaela Colhon – Faculty of Mathematics and Natural Sciences, University of Craiova, Romania

Mircea Giurgiu – The Technical University of Cluj-Napoca, Romania

Mircea Petic – Vladimir Andrunachievici Institute of Mathematics and Computer Science Moldova State University & ACETI

Radu Ion – Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

Ruxandra Cosma – Faculty of Foreign Languages, University of Bucharest, Romania

Ștefan Trăușan-Matu – Faculty of Automation, Control and Computer Engineering, Politehnica University of Bucharest & Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

Svetlana Cojocar – Academy of Sciences of Moldova, Institute of Mathematics and Computer Sciences

Svetlozara Leseva – Department of Computational Linguistics, Institute for Bulgarian Language, Bulgarian Academy of Sciences, Bulgaria

Tiberiu Boroș – Adobe, Romania

Vasile Florian Păiș – Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

Verginica Barbu Mititelu – Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

Victoria Bobicev – Technical University of Moldova (UTM)

ORGANISING COMMITTEE

Adrian Iftene – Faculty of Computer Science, “Alexandru Ioan Cuza” University

Andrei Scutelnicu – Institute for Computer Science, Romanian Academy, Iași & Faculty of Computer Science, “Alexandru Ioan Cuza” University

Cecilia Bolea – Institute for Computer Science, Romanian Academy, Iași

Corina Forăscu – Faculty of Computer Science, “Alexandru Ioan Cuza” University

Dan Cristea – Institute for Computer Science, Romanian Academy, Iași

Dan Tufiş – Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

Daniela Gîfu – Institute for Computer Science, Romanian Academy, Iași & Faculty of Computer Science, “Alexandru Ioan Cuza” University

Diana Trandabăţ – Faculty of Computer Science, “Alexandru Ioan Cuza” University

Elena Irimia – Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

Lucian Gâdioi – Faculty of Computer Science, “Alexandru Ioan Cuza” University

Maria Mitrofan – Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

Maria Ungureanu – Doctoral School, State University of Moldova and Theoretical College ‘Miron Costin’ Paşcani

Vasile Apopei – Institute for Computer Science, Romanian Academy, Iași

Verginica Barbu Mititelu – Research Institute for Artificial Intelligence, Romanian Academy (RACAI)

TABLE OF CONTENTS

FOREWORD	XI
KEYNOTE SPEECHES	XV
EDUCATIONAL PARTNERSHIPS FOR THE FUTURE OF AI: CONSILR'S EXPERIENCE IN TRAINING YOUNG RESEARCHERS MARIA UNGUREANU	1

CHAPTER 1. SPEECH PROCESSING

BRIDGING NEUROSCIENCE AND AI: ADAPTIVE, CULTURALLY SENSITIVE TECHNOLOGIES TRANSFORMING APHASIA REHABILITATION ANDREEA I. NICULESCU, JOCHEN EHNE, MINGHUI DONG	13
WHO DROPS /ʁ/ AND WHERE? INSIGHTS FROM SPEAKER VARIATION IN FRENCH RHOTIC DELETION ANISIA POPESCU, YARU WU, LORI LAMEL, IOANA VASILESCU.....	31
F5-TTS-RO: EXTENDING F5-TTS TO ROMANIAN TTS VIA LIGHTWEIGHT INPUT ADAPTATION RADU GABRIEL CHIVEREANU, TIBERIU BOROS.....	40

CHAPTER 2. LANGUAGE VARIETIES

A REVIEW OF ADVANCES IN COMPUTATIONAL ANALYSIS OF KAZAKH POLITICAL DISCOURSE LAURA ORYNBAY, GULMIRA BEKMANOVA, ALTANBEK ZULKHAZHAY, AYAUPLYM SAIRANBEKOVA	55
GEOGRAPHIC DISTRIBUTION OF DACO-ROMANIAN SUBDIALECTS THROUGH COMPUTATIONAL METHODS SILVIU-IOAN BEJINARIU, FLORIN-TEODOR OLARIU, RAMONA LUCA	67

CHAPTER 3. DICTIONARY-BASED METHODS AND LEXICOGRAPHIC RESOURCES

FROM TRANSITIONAL SCRIPT TO CRITICAL DIGITAL EDITION: DIGITIZING <i>DISIONĂRAȘ ROMÂNESC DE CUVINTE TEHNICE ȘI ALTELE GREU DE ÎNȚELES</i> (1851) TUDOR BUMBU, ELENA ISABELLE TAMBA, ANA VULPE, MARIUS-RADU CLIM, MARIANA NASTASIA, LIVIA CĂRUNTU-CARAMAN, LILIANA BOTNARU, CRISTINA CATEREV, CLAUDIUS TEODORESCU.....	87
EDTLR – A NEW BEGINNING PETRU REBEJA, DAN CRISTEA.....	98

ENHANCING THE PARSEME-RO CORPUS WITH NOMINAL, MODIFIER AND FUNCTIONAL MULTIWORD EXPRESSIONS VERGINICA BARBU MITITELU, IOANA-MARIA BIOLAN, OANA BUHNILĂ, MIHAELA CRISTESCU, ELENA IRIMIA, CĂTĂLIN MIHĂILĂ, ISABELLA ȘINCA, AMALIA TODIRAȘCU, CARMEN VASILE	113
THE AUTOMATIC DIGITALIZATION OF THREE ROMANIAN DICTIONARIES: DELR, DOOM AND DLR CĂTĂLIN MITITELU, HORIA CUCU	126

CHAPTER 4. LLM-BASED EXPERIMENTS & EVALUATION

AI REVIEWER'S CONFIDENCE DANIELA GÎFU, SPERANȚA CECILIA BOLEA, PAULA ALEXANDRA CRUCIANU, SILVIU-IOAN BEJINARIU	145
CROSS-PLATFORM GENAI EVALUATION FOR ESP LEARNING SCENARIOS: A COMPUTATIONAL ANALYSIS OF LINGUISTIC AND CULTURAL APPROPRIATENESS ANDREEA NECHIFOR, ALINA TURCULET, ANCA VASILESCU	162
A CROSS-LINGUAL ANALYSIS OF BIAS IN LARGE LANGUAGE MODELS USING ROMANIAN HISTORY MATEI-IULIAN COCU, RĂZVAN-COSMIN CRISTIA, ADRIAN MARIUS DUMITRAN.....	185
PRESERVING EMOTIONS ACROSS LANGUAGES: EVALUATING LITERARY TRANSLATIONS WITH EMOTION ANALYSIS ALEXANDRA ILIE, MARIA-CORINA DIMITRIU.....	199
DO LLMS UNDERSTAND ROMANIAN DRIVING LAWS? A STUDY ON MULTIMODAL AND FINE-TUNED QUESTION ANSWERING EDUARD BARBU, ADRIAN MARIUS DUMITRAN	222

CHAPTER 5. LARGE LANGUAGE MODELS

TOWARDS ROMANIAN LARGE LANGUAGE MODELS USING THE COROLA CORPUS VASILE PĂIȘ, DAN TUFÎȘ, VERGINICA BARBU MITITELU, RADU ION, MARIA MITROFAN, ELENA IRIMIA, ERIC CUREA, ANDREI PĂUN	235

AUTOMATIC DETECTION OF AI-GENERATED TEXTS FROM MOLDOVA AND ROMANIA (ADAMO). A PROJECT PRESENTATION	
VERGINICA BARBU MITITELU, VICTORIA BOBICEV, VICTORIA ALEXEI, RODICA BRANIȘTE, OLESEA CAFTANATOV, MARIA NITROFAN, RADU ION, ELENA IRIMIA, DANIELA ISTRATI, LUDMILA MALAHOV, SERGIU NISIOI, ALEXANDR PARAHONCO, VASILE PĂIȘ	249
PERSONALIZED VIRTUAL TUTORS THROUGH REASONING LLMS	
ALEXANDR PARAHONCO, OLESEA CAFTANATOV, TUDOR BUMBU, MARIUS CERESCU, NICHTA DEGTEARIOV, DAN BUMBU, ION BOSTAN, IURIE CIALA.....	265
 CHAPTER 6. CULTURAL HERITAGE	
QUANTUM-ASSISED PARSING OF HISTORICAL ROMANIAN TEXTS	
VASILE VANCEA	281
THE WISDOM EXTRACTOR: MINING CROSS-CULTURAL PROVERBS TO ELICIT TIME-TESTED HEURISTICS	
VLAD BELCIUG, ELENA PELICAN	281
EXPLORING THE BESSARABIAN IDIOM UNDERSTANDING BY LLMS AND AI TOOLS	
INGA TITCHIEV, OLESEA CAFTANATOV	319
FOLKAI: A DIACHRONIC CORPUS OF BESSARABIAN FOLKLORE TEXTS	
OLESEA CAFTANATOV, SVETLANA COJOCARU, LUDMILA MALAHOV, MIRCEA PETIC	332
THE MAPPING GREETING FORMULAE IN THE HISTORICAL MOLDOVA PROJECT: AN OVERVIEW	
ANCA-DIANA BIBIRI, MIHAELA MOCANU, LIVIU-ANDREI SCUTELNICU, LIVIU PETCU, VIORICA POPA, ELENA LACUSTA, ANGELA BEJAN.....	346

FOREWORD

Today's communication is more and more mediated by information technology and, by analogy with the already established term *international communication language*, one could speak about *e-communication languages* with direct reference to the languages used in cyberspace. This concept, beyond its direct significance, has profound cultural, social, economic and ethical implications, stating the right of any citizen to have access in his/her own language to the knowledge and services of cyberspace. Thus, it is not surprising the huge volume of research in the area of human language engineering, consistently supported through the international or national research programs which produced, in the last 20 years or so, practical results which paved the way to multilingual industry. The unprecedented increase of computers performances and of the data storage capacities, the success of neural networks and deep-learning approaches generated in the last decade not only a research paradigm change, but also more consistent and convincing results than ever.

The ConsILR 2025 conference was an anniversary one, the 20th edition in a row of itinerant successful events. This year the Conference was organised in Bucharest at the Romanian Academy House, an impressive building in the heart of the capital, next to the Parliament and the National Cathedral.

The scientific events were organized by two institutes of the Romanian Academy, the Institute of Artificial Intelligence "Mihai Drăgănescu" in Bucharest and the Institute of Computer Science in Iași, together with the Faculty of Computer Science of the University "Alexandru Ioan Cuza" Iași and the Romanian Association of Computational Linguistics.

From its inception, in 2001, the ConsILR Conference (traditional acronym for *Consortiul de Informatizarea pentru Limba Română*, an initiative born in the Section for Information Science and Technology of the Romanian Academy) was meant as a meeting place for linguists and computational linguists, but also for researchers of the humanities, PhD students and master students in Computational Linguistics, all with a major interest in the study of the Romanian language from a computational perspective. The series of events have run, with few exceptions, once every year, first in the format of a workshop, and since 2010 – as a conference. In order to reach wider visibility, the organisers decided to make the Conference itinerant and to publish its Proceedings in English. Thus, ConsILR is not strictly addressed to researchers working on Romanian language but also to other scientists, from any part of the world, which could find sources of inspiration in the models and techniques developed for our language and apply them for their own languages. Opening the gate for researchers working on languages other than Romanian to participate in the Conference and publish their work in these Proceedings, a reverse influence is also facilitated, namely that their work inspires scientists working on the Romanian language.

FOREWORD

The conference program, mirrored in this volume, was dense, with 22 presentations of the latest results of researchers in this field. The contributed articles were organized in six chapters:

- 1. Speech technology, phonology, and applied linguistics for speech therapy,**
- 2. Language varieties,**
- 3. Dictionary-based methods and lexicographic resources,**
- 4. LLM-based experiments and LLM evaluation,**
- 5. LLMs building and their applications,**
- 6. Cultural Heritage.**

Most of the results presented at the conference are public, open to anyone interested in taking them over and using them.

The organisers invited four researchers in our field, who accepted to deliver keynote speeches:

Agata Savary is a professor in computer science at the Paris-Saclay University in France. Her domains of interest include natural language processing (NLP), universalist and NLP-applicable modelling of idiomaticity in language, as well as construction of language resources and tools for multiword expression identification, named entity recognition, and coreference resolution. She currently chairs the CA21167 COST action UniDive on universality, diversity and idiosyncrasy in language technology (2022-2026). In the past, she also served as the chair of the IC1207 COST action PARSEME, on PARSing and Multiword Expressions (2013-2017), where steps were taken towards evaluating diversity next to performance in NLP tools. She authored 21 refereed journal papers, 40 conference papers, 20 workshop papers, and 10 book chapters.

Amalia Todiraşcu is Professor at University of Strasbourg, and obtained a bachelor in Computer Science from the “Alexandru Ioan Cuza” University (Iaşi, Romania), in 2001, and PhD degree in Natural Language Processing (NLP) and Artificial Intelligence from the University Louis Pasteur (Strasbourg, France). Her thesis was about extracting ontologies from texts by using NLP tools. She works in several research projects in the domains of Information Extraction and Information Retrieval. Since 2004, she has worked at the University of Strasbourg, teaching courses on Corpus Linguistics, NLP, XML, and annotation tools. Her research interests focus on creating automatic annotation tools: automatic coreference detection tool (CoFr, developed jointly in the context of DEMOCRAT and ALECTOR project), multiword expression annotation. In linguistics, she works on the domains of textual cohesion (coreference models) and lexical semantics (classes of humans).

After graduating as a medical doctor at the University of Medicine and Pharmacy (Iaşi, Romania) and following a master's degree in Theater, **Marius Ursache** devoted himself to creating value as an entrepreneur. For the past 25 years, he has worked in media, design, branding, healthcare, marketing, fintech,

FOREWORD

AI, innovation consulting, investment advisory, and more. He founded four companies, an NGO, and a music festival. In the past he was, as a designer & art director, with Sally Johns Design LLC (Raleigh, North Carolina, United States), one of the top design agencies in the state of North Carolina, has founded Grapefruit SRL (Iași, Romania), one of the most successful Romanian digital marketing and software development agencies, had the idea and founded Eternime, Inc. (San Francisco, California, United States), an artificial intelligence application that creates a digital avatar for people, based on content generated by them in digital environments, has performed training, mentoring and curriculum development for MIT Bootcamps (Cambridge, MA, United States), an educational initiative for innovators and entrepreneurs, having taught over 1,300 individuals in 100+ countries, etc. His latest venture, Performikon, is an innovation consultancy helping companies to take their ideas from R&D lab to market, using a combination of AI agents and the Disciplined Entrepreneurship from MIT.

Paula Gradu is a fourth year PhD student at UC Berkeley, where she is advised by Michael I. Jordan and Ben Recht as part of the Berkeley AI Research lab (BAIR). Her research focuses on designing adaptive methods for healthcare applications with provable guarantees characterizing their expected behaviour. Some concrete examples include designing individualized treatment protocols using online control, developing methods for accelerating discovery through online learning, and proposing methodology for identifying disparate harm in incident databases via sequential hypothesis testing. When an individual reports a negative interaction with some system, how can their personal experience be contextualized within broader patterns of system behaviour? She is studying the incident database problem, where individual reports of adverse events arrive sequentially, and are aggregated over time. In this work, her goal is to identify whether there are subgroups – defined by any combination of relevant features – that are disproportionately likely to experience harmful interactions with the system. She formalises this problem as a sequential hypothesis test, and identifies conditions on reporting behaviour that are sufficient for making inferences about disparities in true rates of harm across subgroups.

As in the previous edition of ConsILR, the organisers had the initiative to dedicate a special session to high school young students, in order to raise their interest for this fascinating field of science, so much linked to both sides of school didactic activities, humanities and exact sciences. As such, the first paper in the volume resumes the interventions in this special session.

The current selection of papers is inherently incomplete and its coverage is far from being ideal. There are many more significant results in the NLP area dedicated to Romanian language, but some of them were already published elsewhere. There are several other reputed computational linguists, who worked

FOREWORD

on Romanian or with Romanians, who, unfortunately, are not present in this volume, but we acknowledge their work and bring here our homage to them.

As it happened with more volumes of previous editions, this volume will also be sent for indexation to Clarivate Analytics in order to obtain the prestigious recognition of this agency.

To conclude, this volume offers a valuable collection of contributions presenting research and development results. Several issues discussed throughout the book are present-day concerns of the scientific community, irrespective of the investigated language. Therefore, we believe that this volume will be a very useful instrument for NLP researchers, for students preparing their MSc dissertations or PhD theses in the area of computational studies of natural languages.

We would like to thank the Organising Committee, the program chairs, our sponsors, our reviewers, and the representatives of local host institution for their dedication in putting this event together. We also express gratitude to all the keynote speakers, session chairs, and volunteers whose efforts make this conference possible. Thanks to all the participants who have travelled to be with us or have participated on-line. Your presence, in person or online, enriched the content of this anniversary conference. The many feedback that we received after the end of the Conference made us believe that everybody has left with new ideas, new interlocutors, and renewed research energy. It is also our hope that you, readers of this Proceedings who could not be present in the event, will find here a valuable source of inspiration for your own projects.

December 2025
The editors

KEYNOTE SPEECHES

AGATA SAVARY

Short bio: Agata Savary is a professor in computer science at the Paris-Saclay University in France. Her domains of interest include natural language processing (NLP), universalist and NLP-applicable modelling of idiomaticity in language, as well as construction of language resources and tools for multiword expression identification, named entity recognition, and coreference resolution. She currently chairs the CA21167 COST action UniDive on universality, diversity and idiosyncrasy in language technology (2022-2026). In the past she also served as the chair of the IC1207 COST action PARSEME, on PARSing and Multiword Expressions (2013-2017), where steps were taken towards evaluating diversity next to performance in NLP tools. She authored 21 refereed journal papers, 40 conference papers, 20 workshop papers, and 10 book chapters.

Contact: <https://perso.limsi.fr/savary/>, agata.savary@universite-paris-saclay.fr

Title: Quantification of Diversity in Natural Language Processing

Abstract: The concept of diversity has received increased consideration in Natural Language Processing (NLP) in recent years. This is due to various motivations like promoting equity and inclusion, approximating human linguistic behaviour, and increasing systems' performance. Diversity has however often been addressed in an ad hoc manner in NLP, and with few explicit links to other domains where this notion is better theorized. We survey articles in the ACL Anthology from the past 6 years, with "diversity" or "diverse" in their title. We find a wide range of settings in which diversity is quantified, often highly specialized and using inconsistent terminology. We put forward a unified taxonomy of why, what on, where, and how diversity is measured in NLP. Diversity measures are cast upon a unified framework from ecology and economy (Stirling, 2007), with 3 dimensions of diversity: variety, balance and disparity. We discuss the trends which emerge due to this systematized approach. We believe that this study paves the way towards a better formalization of diversity in NLP, which should bring a better understanding of this notion and a better comparability between various approaches.

This is a joint work with Louis Estève (Université Paris-Saclay, France), Marie-Catherine de Marneffe (Université Catholique de Louvain, Belgium), Nurit Melnik (The Open University of Israel) and Olha Kanishcheva (Jena University in Germany and the STEP University in Ukraine), within the framework of the UniDive COST Action on Universality, Diversity and Idiosyncrasy in Language Technology.

AMALIA TODIRASCU

Short bio: Amalia Todirascu obtained a PhD degree in Natural Language Processing (NLP) and Artificial Intelligence at University Louis Pasteur (Strasbourg, France) and University Gh. Asachi (Iasi, Romania), in 2001. Her PhD thesis was about extracting ontologies from texts by using NLP tools. She worked in several research projects in the domains of Information Extraction and Information Retrieval. In 2004, Amalia Todirascu obtained an Assistant Professor position at the University of Strasbourg, and a Full Professor (in 2016). She teaches courses on Corpus Linguistics, NLP, XML, Annotation tools. Her research interests in NLP focus on creating automatic annotation tools: automatic coreference detection tool (CoFr, developed jointly in the context of DEMOCRAT and ALECTOR project), and multiword expression annotation. Moreover, she developed several resources and tools for a more inclusive society: automatic text simplification systems (for people with dyslexia) (syntax and discourse level), gender bias in NLP and Computer-Aided Language Learning systems (SimpleApprenant project and the STAR-FLE, funded by ANR). Moreover, she participates in an interdisciplinary project (GREMO-Ling) aiming to build tools to evaluate the impact of a therapy by analysing narratives of patients with brain injuries. In linguistics, she provides contributions in the domain of textual cohesion (coreference models) and lexical semantics (classes of humans). **Contact:**<https://lilpa.unistra.fr/theme-1-lexiques-discours-et-transpositions/membres/enseignants-chercheurs/todirascu-amalia/>, todiras@unistra.fr

Title: Automatic Text Adaptation for Target Audiences

Abstract: The talk aims to present the main notions required by automatic text adaptation systems for various audiences. The task of automatic text adaptation requires tools and resources for various target audiences which have various expectations in terms of text adaptation. Thus, people with cognitive disabilities, such as people with dyslexia, will expect simplified texts at syntactic, discourse and lexical level to help them to improve their comprehension while reading. Specialized texts such as medical ones, are difficult to understand for patients and their families, so texts should be transformed in order to replace complex medical terms with their paraphrases. Finally, language learners expect pedagogical materials adapted to their CECRL level. Resources such as lexical databases and parallel corpus (original and simplified texts) should be designed for text adaptation. I will then present several samples of automatic text adaptation systems and resources for various target audience: HECTOR for improving reading skills for children with dyslexia (Alector project), Adaptmed designed for producing medical terms paraphrases for helping lay persons to understand their treatment, and SimpleApprenant and STAR-FLE resources for language learners.

MARIUS URSACHE

Title: How I 10X'ed Myself with AI and What I Learned in the Process: A very non-academic but hopefully interesting approach to using AI tools to make my work better and my life easier.

Abstract: Marius Ursache, a serial tech entrepreneur, and MIT teaching fellow, showcases how using AI and agents has exponentially improved his professional and personal endeavors. The presentation explores practical strategies for dramatically amplifying productivity using AI across three areas: advanced prompting, workflow automation, and AI-assisted coding.

Moving beyond basic ChatGPT usage, structured prompts with clear roles, goals, context, workflows, constraints, and validation steps transform AI from a search engine into a capable assistant. Custom instructions and personalization settings eliminate fluff and demand direct, critical responses.

Real automation workflows built with n8n, such as YouTube summarization bots, meeting transcript processors, LinkedIn content generators, and invoice handling systems combining multiple AI models with traditional tools, radically shorten task times and automate otherwise tedious, repetitive work.

AI-powered coding through platforms like Cursor, Lovable, and Windsurf enables building complex applications – from research assistants to multi-agent orchestrators – using conversational interfaces and agentic coding tools.

These are not highly advanced projects, but very practical ways to use AI in multiple personal and professional areas to explore, learn, and improve, that are accessible to everyone today. Curiosity (or rather lack of it) is the only limitation.

PAULA GRADU

Short bio: Paula Gradu is a fourth-year PhD student at UC Berkeley, where she is advised by Michael I. Jordan and Ben Recht and part of the Berkeley AI Research lab (BAIR). Her research focuses on designing adaptive methods for healthcare applications with provable guarantees characterizing their expected behaviour. Some concrete examples include designing individualized treatment protocols using online control, developing methods for accelerating discovery through online learning, and proposing methodology for identifying disparate harm in incident databases via sequential hypothesis testing.

Contact: <https://www.linkedin.com/in/paula-gradu-7505591b0>,
pgradu@berkeley.edu

Title: From Individual Experience to Collective Evidence: A Reporting-based Framework for Identifying Systemic Harms

Abstract: When an individual reports a negative interaction with some system, how can their personal experience be contextualized within broader patterns of system behaviour? We study the incident database problem, where individual reports of adverse events arrive sequentially, and are aggregated over time. In this

KEYNOTE SPEECHES

work, our goal is to identify whether there are subgroups--defined by any combination of relevant features--that are disproportionately likely to experience harmful interactions with the system. We formalize this problem as a sequential hypothesis test, and identify conditions on reporting behaviour that are sufficient for making inferences about disparities in true rates of harm across subgroups. We show that algorithms for sequential hypothesis tests can be applied to this problem with a standard multiple testing correction. We then demonstrate our method on real-world datasets, including mortgage decisions and vaccine side effects; on each, our method (re-)identifies subgroups known to experience disproportionate harm using only a fraction of the data that was initially used to discover them.

EDUCATIONAL PARTNERSHIPS FOR THE FUTURE OF AI: CONSILR'S EXPERIENCE IN TRAINING YOUNG RESEARCHERS

MARIA UNGUREANU^{1,2}

¹ *Doctoral School, State University of Moldova*

² *Theoretical College "Miron Costin" Pașcani*
umariamarrynna@gmail.com

Abstract

This article describes an initiative of ConsILR, currently unique in Romania, materialized in an educational partnership with several pre-university education institutions in the country. The approach aims to familiarize high school students with the specifics and applicability of the field of computational linguistics. Now in its second edition, the project confirms its relevance through the significant contributions made to the school and professional orientation process of students, facilitating their access to specialized information and real perspectives on academic training and career opportunities. Providing school counselling at the level of pre-university education is an essential element for the informed development of students' educational path. Through it, young people benefit from support in identifying personal interests and skills, in understanding the particularities of various professional fields and in making decisions about continuing their studies. At the same time, the collaboration between university and pre-university institutions becomes a catalyst for educational openness, facilitating the transfer of expertise, creating authentic learning contexts and stimulating students' interest in emerging areas of knowledge. Such partnerships contribute to strengthening a coherent educational ecosystem, in which school guidance, academic training and professional integration are approached in an integrated and sustainable way.

Keywords: educational partnership, ConsILR 2025, school and career guidance, computational linguistics

1 Introduction

In the context of the acceleration of the digitization process and the transformations generated by new technologies, school and professional guidance take on new valences, overcoming the traditional delimitations of disciplines. In this framework, it becomes necessary to direct students' attention to interdisciplinary fields, such as computational linguistics, which integrate the

rigor of linguistic analysis with the tools and methods of computer science. Computational linguistics deals with the formal modelling of natural language, the development of technologies for automatic text and speech processing, as well as the development of applications capable of enabling advanced interactions between users and intelligent systems. Specialists working in this field combine solid language skills with technical, programming and data analysis skills, complemented by a pronounced capacity for abstraction, algorithmic thinking and adaptability to the rapid evolution of technology. The early interaction of high school students with information, concepts and application examples from the field of computational linguistics contributes to the development of a realistic understanding of the related educational and professional opportunities, thus facilitating the formation of an informed academic path in line with the dynamics of the labour market. In addition, early exposure to emerging fields stimulates interest in research, interdisciplinary spirit and the development of advanced digital skills, which are essential in today's society.

In this context, the ConsILR initiative to promote computational linguistics among students from high school level can be considered not only opportune, but even necessary to align school orientation with the evolution of interdisciplinary fields. Starting with 2024, the annual conference of ConsILR integrates a special section dedicated to high school students, designed as a space for direct interaction and dialogue between young people and senior scientists. This structure facilitates students' access to authentic expertise and up-to-date perspectives, allowing them to formulate questions, explore relevant themes and outline a clearer picture of how technology and linguistics intersect. The discussions are explicitly oriented to the knowledge needs and interests of the students, ensuring an educational framework adapted to their level of understanding and their potential options for academic and professional development.

The partnership developed within this project takes the form of ConsILR's collaboration with the "Miron Costin" Theoretical High School in Pașcani, Iași County, an institution that had, in both editions of the project, the role of identifying participating schools from any county of the country and coordinating the organization of the section dedicated to students. In this year's edition, approximately 350 high school students from Iasi and Alba counties, as well as from Bucharest, participated in the project activities, either on-site directly or online, demonstrating the national interest in such initiatives. The conference was hosted by the Research Institute for Artificial Intelligence „Mihai Drăgănescu” (RACAI) in Bucharest. Thus, the participants had the opportunity to get to know an institutional framework of high professional and scientific standing. The participants had the opportunity to interact with prestigious researcher in the field, including Prof. Dan Cristea, PhD, Corresponding Member of the Romanian Academy; Dr. Andrei Păun, Director at RACAI; Dr. Verginica Barbu Mititelu, Deputy Director at RACAI; Dr. Diana Trandabăț from the Faculty of Informatics

of the "Alexandru Ioan Cuza" University of Iași; Dr. Daniela Gîfu from the Institute of Theoretical Informatics of the Romanian Academy, Iași Branch, Dr. Maria Carp, Senior Researcher at RACAI; and PhD candidate Paula Gradu from the University of California, Berkeley.

2 The role of the Research Institute for Artificial Intelligence (RACAI) in training the new generation of researchers

The session was opened by Dr. Verginica Barbu Mititelu, who offered the students an extensive presentation of RACAI, an institution located at the centre of Romanian fundamental research in the field of artificial intelligence. Within this exposition, the role of the Romanian Academy – founded in 1866 and responsible, since its beginnings, in the standardization of the Romanian language and for the realization of defining works for national history and culture – as the main forum of scientific recognition and privileged space for interdisciplinary research was also underlined.

The main research directions of the institute were presented, including human-AI interaction, AI applications in medicine, language technologies, neuromorphic structures and brain activity modelling developed within the National Centre for Brain Research a branch of RACAI. These directions are at the intersection of artificial intelligence, computer science, neuroscience, linguistics and cognitive sciences, illustrating both the complexity of the field and its transformative potential. The main message of the presentation highlighted the idea that the training of new generations of artificial intelligence specialists involves early access to advanced research, emerging technologies and institutional frameworks dedicated to innovation.

The relevance of these research directions is naturally articulated with the findings of the specialized literature, which underlines the deeply interdisciplinary nature of artificial intelligence and the need for early training of the skills associated with this field. Significant studies in the area of artificial intelligence and cognitive sciences (Russell and Norvig, 2003; Manning, 2008) highlights that technological progress depends on the collaboration between linguistics, computer science, neuroscience and computational modelling, as well as on the ability of young researchers to integrate multiple perspectives in approaching problems. At the same time, recent research on the development of digital skills and preparation for emerging professions (Williams *et al.*, 2018; Williams *et al.*, 2019; Borgonovi, 2023) emphasizes the importance of early exposure to AI concepts and practices, considered a determining factor in stimulating interest in research and shaping a solid academic path. In this regard, the presentation offered to the students by the representatives of the Research Institute for Artificial Intelligence is part of a modern educational paradigm, based on access to advanced knowledge, interdisciplinary dialogue and the development of the

necessary skills to build a new generation of specialists capable of contributing to the evolution of the field.

2.1 A researcher trained at RACAI, Bucharest

In the session dedicated to students, discussions focused on how young people can access and capitalize on training and research opportunities. Paula Gradu, a former intern of the institute, spoke about the beneficial impact of participating in research activities since high school, emphasizing as benefits the development of critical thinking and problem-solving skills. She highlighted the contribution of the research projects in which she participated, to strengthening her academic and professional training.

Speaking about the importance of mentors, both Paula Gradu and Diana Trandabăț highlighted the contribution of teachers and researchers in guiding young people and facilitating their intellectual progress. In addition, volunteer projects within the institute were presented – including the creation of an oral corpus of voices and the training of voice processing systems – which also offer certified participants the possibility of further involvement in writing scientific articles or presenting at conferences.

3 Familiarizing students with the field of computational linguistics and natural language processing processes

Dr. Diana Trandabăț (professor at the Faculty of Computer Science of the "Alexandru Ioan Cuza" University, Iași) and Daniela Gîfu (principal researcher at the Institute of Theoretical Informatics in Iași), initiated a consistent dialogue with the students present, about computational linguistics, aiming to familiarize them with the internal mechanisms of natural language processing and with the essential role of this field in contemporary society.

The authors' discourse started from an essential question: why computational linguistics? In a world where algorithms intervene in almost every area of life – from education and communication to science and ethics – the understanding of natural language by machines has become a central focus of modern research. Language, which for millennia has been the exclusive prerogative of the human being, is now becoming a technological frontier. The fundamental idea that the authors convey to students is that *learning how machines communicate actually means understanding more deeply how humans communicate*.

To introduce the notions of *meaning* and *linguistic representation*, the authors resort to a classic metaphor: the Shakespearean quote "What is a name? What we call a rose by any other name would smell as sweet?" The choice is not accidental. This reflection on the relationship between word and reality serves as a starting point for understanding the fundamental problem of meaning in computational linguistics. Essentially, a computer program must learn that the "rose" is not just a succession of letters, but an entity with smell, shape, context, and cultural

connotation. This is where a first great challenge of the field emerges: how can the richness of human meaning be translated into numerical representations that can be interpreted by machines? The question was able to provoke the interest of the students who reacted immediately, advancing not only questions, but also interesting opinions.

The discussion then expanded to how natural language structures everyday life in the digital age. In a connected world, linguistic algorithms are already present in trivial actions: Google Maps calculates routes by interpreting data generated by millions of users, commercial platforms anticipate customer preferences by analysing descriptions and reviews, and automatic recognition systems identify license plates, locations or facial expressions in real time. All these examples demonstrate a complex reality: *contemporary man is surrounded by technologies that "understand" language.*

At the students' proposal, the discussion shifted its focus from the theoretical sphere to ethical and social applications, addressing a topical topic: aggressive behaviour in the online environment. Because, currently, the digital space is becoming the terrain of manifestation of hatred, manipulation and disinformation, NLP tools acquire a moral dimension. The examples provided – from offensive language and direct threats, to illegal forums or online recruitment for criminal activities – illustrate how urgent the need for technological solutions to protect virtual environments is. The analysis of these phenomena involves advanced linguistic processing techniques: identification of feelings, recognition of named entities (people, organizations, places), analysis of emotional polarity or detection of hidden intentions. These methods, developed in projects such as Perspective API (Google Jigsaw) or HateBERT (MIT CSAIL), aim to recognize patterns of hatred and automatically isolate them, while maintaining the balance between freedom of expression and public safety. Studies by Vidgen *et al.* (2022) show that aggressive speech detection models need to be trained on multilingual datasets that are sensitive to cultural variations. In the context of the Romanian language, this task is particularly complex, as large bodies of labeled texts and standardized tools are missing. Therefore, Romanian research is at a crucial point: the development of local linguistic resources becomes a condition for real participation in the global community AI. Pe European level, initiatives such as CLARIN-ERIC¹, in which Romania is an active part, build common infrastructures for the analysis of toxic discourse, offering annotated corpora and research tools accessible to the academic public. In this way, computational linguistics becomes not only a language science, but also *a form of civic responsibility*, helping to maintain the health of online communication.

¹ Pagina principală CLARIN | CLARIN ERIC – Resurse lingvistice comune și infrastructură tehnologică.

Among the realizations that the students had was the idea that language is more than a code – it is an expression of thought, culture and ethics. Therefore, processing it with the help of technology means understanding its internal mechanisms and respecting its power. In an era where language models are becoming more and more performant, the challenge is no longer just to create machines that talk, but *to cultivate people who think critically about how these machines learn to speak*. In the end, she enjoyed a special appreciation from the students, the fact that the authors managed to transform a scientific presentation into a lesson about humanity. Without resorting to excessive jargon, they translated complex concepts – from syntactic analysis to entity recognition – into an accessible but rigorous discourse. At the same time, they introduced philosophical and ethical dimensions into the debate, inviting students to perceive technology not as a cold tool, but as an extension of the human capacity to understand and create meaning.

In light of recent research (OECD, 2024; UNESCO, 2024), responsible digital education involves the formation of "AI literacy" – the ability to understand how language models work and what implications they have on thinking and communication. The discussion in the special session for students of the ConsILR conference is perfectly in line with this: it prepares a new generation of young people who not only use technology, but understand it in depth.

4 Artificial Intelligence Literacy Assessment: AILQ Quiz Results

During the special session dedicated to students, held during the ConsILR Conference, the results of a research conducted at RACAI were presented, a research aimed at assessing the level of literacy in artificial intelligence among middle and high school students in Romania. The study was part of a broader initiative of the institute, aimed at introducing AI education in schools (Carp *et al.*, 2025), in a context in which intelligent systems are becoming an integral part of everyday life, and understanding them becomes essential for the training of future generations. The purpose of the session presentation was to highlight how the research measured students' real level of understanding, attitudes, behaviours and ethical competences in relation to artificial intelligence technologies. The research team used **the AI Literacy Questionnaire (AILQ)**, an internationally validated tool built on four fundamental dimensions — affective, behavioural, cognitive, and ethical — that provide a detailed picture of how students perceive and relate to AI.

The questionnaire was completed by 100 high school students, enrolled in grades IX–XII, coming from educational institutions in Bucharest, Braşov and Paşcani. The choice of a diversified geographical sample aimed to capture the differences between various educational contexts and to obtain a more accurate representation of the level of AI literacy at national level. The results presented during the session highlighted several major directions. Students demonstrated a

high level of cognitive understanding, with most reporting that they understand the basic concepts of AI and can identify examples of concrete applications. As for the ethical dimension, it recorded the highest scores: over 53% of students showed interest and awareness of the moral implications of the use of AI, such as data protection, developer responsibility or the social impact of algorithms. On the other hand, the behavioural dimension, which aims at the effective application of knowledge, involvement in practical activities and collaboration, presented the lowest values. This discrepancy indicated that, although students understand AI and perceive it positively, they do not yet have sufficient educational contexts to transform theoretical knowledge into applied skills — through projects, practical exercises or extracurricular activities. The conclusions presented during the session highlighted the need for early and structured introduction of AI literacy in the pre-university curriculum. The research showed that students have potential, interest and motivation, but they need a coherent learning framework, resources in Romanian and teachers trained in the field. The study thus provided a solid empirical basis for the development of future educational programs and for arguing the need for a national AI curriculum. The contribution of the research was directly reflected in the development of the field of AI and education, by providing a current x-ray of the level of preparation of students, as well as by highlighting the real needs of pre-university education in relation to emerging technologies. The results strengthened the strategic directions of the institute's subsequent initiatives, including educational projects at national level and extensive collaborations with academic and institutional actors.

The presentation given during the student session of ConsILR, illustrated the importance of associating scientific research with educational practice and demonstrated the potential that AI literacy has to form a generation capable of critically, responsibly and creatively using the technologies of the future.

5 Dialogue with students: mentoring, volunteering and reflections on the role of AI in education

Students showed a constant interest in integrating artificial intelligence into the educational process, expressing a desire to identify effective tools that can support learning and to critically analyse the limits of generative technologies. The dialogue carried out at this stage highlighted the need to train critical thinking skills and information verification practices, elements considered essential in an educational context marked by unfiltered access to automatically generated content. Equally, the participants stressed the importance of cultivating logical and algorithmic skills, which are perceived as fundamental for understanding how intelligent systems work and for developing intellectual autonomy in relation to emerging technologies.

In a later stage of the discussions, the students reflected on the evolution of programming in a digital ecosystem where AI tools are becoming more

sophisticated and accessible. Their interventions covered the full spectrum of possible implications: from the risks of progressively diminishing coding skills among future specialists, to the innovation opportunities opened up by human-machine collaboration and the automation of repetitive tasks. The ethical implications of the use of AI were also discussed, including developer responsibility, algorithmic transparency and personal data protection, along with the ecological consequences of generative technologies, such as the high energy consumption of large-scale models. Part of the discussion also focused on the differences between the European and American regulatory frameworks, with the students noticing the distinct way in which the two legal spaces address the balance between innovation, user protection and the social responsibility of technologies.

6 Conclusions

The project "Today's student, tomorrow's researcher", carried out under the aegis of ConsILR, stood out through a significant contribution to the school and professional orientation of high school students, giving them direct access to the experience of specialists in the field, research institutions and emerging fields of knowledge. Through structured exposure to valid information about computational linguistics and artificial intelligence, students benefited from a rare opportunity in the Romanian pre-university system: to explore, in an authentic academic setting, an interdisciplinary field in full expansion. Interaction with researchers and academics facilitated not only the understanding of scientific concepts, but also the familiarity with real academic backgrounds, with the way of organizing the research activity and with the professional responsibilities specific to this sector.

For pre-university education, the project demonstrated the value of early integration of advanced learning experiences, contributing to the development of students' digital skills, critical thinking and ethical sensitivity. At the same time, the activities carried out strengthened the role of school counselling as an essential tool in supporting informed educational choices, offering students the opportunity to identify their academic interests in a modern context, connected to the current scientific evolution.

For the university environment, the project created the premises for the formation of a coherent educational flow between pre-university and higher education, generating a collaborative framework through which universities and research institutes can adapt their training strategies to the needs and profiles of the new generations of students. Through direct interaction with students, researchers benefited from a better understanding of the expectations and gaps in pre-university training, useful information for the development of more flexible and relevant university programs.

For the participating students, the benefits of the project were multiple: access to current scientific resources, clarification of some study directions, stimulating interest in STEM fields and, in many cases, developing an authentic motivation for involvement in research activities. Direct contact with experts and institutions such as the Artificial Intelligence Research Institute has had a strong inspirational effect, confirming the potential of experiential education to transform young people's perceptions and professional aspirations.

For the researchers, the dialogue with the students offered a double gain: on the one hand, it generated a space for reflection on how complex concepts can be communicated to young audiences; on the other hand, it highlighted adolescents' genuine interest in emerging technologies, providing valuable insight into the educational needs of future cohorts of students.

Overall, the project "Today's Student, Tomorrow's Researcher" managed to demonstrate, through the nature of its activities, the importance of dialogue between the pre-university and university systems, as well as the essential role of interdisciplinary education in preparing young people for the professions of the future. The ConsILR initiative thus represents a model of good practice that deserves to be expanded and strengthened, offering students early access to science, technology, research and ethical responsibility – essential elements in forming a generation capable of innovating and actively contributing to the development of the digital society.

AI Use statement

Generative-AI tools were used solely for minor stylistic refinements and copyediting. All content, analysis, and conclusions in this paper are the author's own work.

References

- Andreea Nechifor and Cristina Dimulescu. 2024. External Factors Influencing First-Year Students' Usage of Digital Learning; Tools at 'Transilvania' University of Braşov. In Blažević I., M. Diković, I. Kovačić, L. Lazarić, and J. Gugić
- Andreea Nechifor and Maria Ungureanu. 2024. The discourse of professional communication with young learners: career orientation practices in the field of computational linguistics, Proceedings of the 19th International Conference, Alba Iulia, 14-16 November 2024, pages 221-237.
- Bertie Vidgen, Scott A. Hale, Ella Guest and Leon Derczynski., Challenges in Building Online Hate Detection Models. 2022
- Bertie Vidgen and Leon Derczynski. 2023. HateCheck 2.0: Benchmarking Hate Speech Detection Systems Multilingually. arXiv:2305.04112. [<https://arxiv.org/abs/2305.04112>] (<https://arxiv.org/abs/2305.04112>)
- Christopher D. Manning, 2008. Introduction to information retrieval. New York

- Francesca Borgonovi *et al.* 2023. Emerging trends in AI skill demand across 14 OECD countries”, OECD Artificial Intelligence Papers, No. 2, OECD Publishing, Paris. (Emerging trends in the demand for AI skills in 14 OECD countries | OECD)
- Maria Carp, Vasile Păiș and Verginica Barbu Mititelu, 2025. Bringing AI to schools: an initiative to promote AI Literacy and emerging technologies in Romanian education, 17th International Conference on Education and New Learning Technologies, Palma, Spania
- Patterns, 3(8), 100633. [<https://doi.org/10.1016/j.patter.2022.100633>] (<https://doi.org/10.1016/j.patter.2022.100633>)
- Randi Williams, Hae Won Park and Cynthia Breazeal. 2018. How smart are smart toys?: the interaction of children's and parents' agents and the attribution of intelligence. In the Proceedings of the 17th ACM Conference on Interaction Design and Children. ACM, 231-240.
- Randi Williams, Hae Won Park and Cynthia Breazeal. 2019. A is for Artificial Intelligence: The impact of artificial intelligence activities on young children's perception of robots. Proceedings of the 19th CHI Conference on Human Factors in Computing Systems, Glasgow, Scotland, UK. (A stands for Artificial Intelligence)
- Stuart Russell and Peter Norvig. 2003. Artificial Intelligence- A Modern Approach, Second Edition, Prentice Hall, pages 78-79.

CHAPTER 1

SPEECH PROCESSING

BRIDGING NEUROSCIENCE AND AI: ADAPTIVE, CULTURALLY SENSITIVE TECHNOLOGIES TRANSFORMING APHASIA REHABILITATION

ANDREEA I. NICULESCU¹, JOCHEN EHNES², MINGHUI DONG³

¹*A*STAR Institute for Infocomm Research (I2R)*
andreeaniculescu@sigchi.org

²*A*STAR Institute for Infocomm Research (I2R)*
jehnes@gmail.com

³*A*STAR Institute for Infocomm Research (I2R)*
dongminghui@gmail.com

Abstract

Aphasia, a language impairment primarily resulting from stroke or brain injury, profoundly disrupts communication and everyday functioning. Despite advances in speech therapy, barriers such as limited therapist availability and the scarcity of personalized, culturally relevant tools continue to hinder optimal rehabilitation outcomes. This paper reviews recent developments in neurocognitive research and language technologies that contribute to the diagnosis and therapy of aphasia. Drawing on findings from our ethnographic field study, we introduce two digital therapy prototypes designed to reflect local linguistic diversity and enhance patient engagement. We also show how insights from neuroscience and the local context guided the design of these tools to better meet patient and therapist needs. Our work highlights the potential of adaptive, AI-enhanced assistive technologies to complement conventional therapy and broaden access to therapy. We conclude by outlining future research directions for advancing personalized and scalable aphasia rehabilitation.

Keywords: applied linguistics, NLP integration, assistive technology, speech therapy, aphasia

1 Introduction

Aphasia is a complex language disorder, typically caused by stroke or brain injury, that significantly impairs an individual's ability to communicate. It affects core language functions – speaking, understanding, reading, and writing – and has profound implications for quality of life. Each year, approximately 15 million people suffer strokes, and up to 40% develop aphasia (World Health Organization, 2004).

This condition can seriously hinder communication, making it difficult for individuals to express themselves or understand others, often leading to social isolation and emotional distress (Hilari and Byng, 2009; Parr, 2007). Aphasia presents in various forms, each affecting language differently. Broca's aphasia (non-fluent) involves slow, effortful speech with preserved comprehension, while Wernicke's aphasia (fluent) results in fluent but often nonsensical speech and impaired understanding (Damasio, 1992). Another major subtype is *Primary Progressive Aphasia* (PPA), a neurodegenerative condition marked by gradual atrophy in language-related brain regions. Unlike stroke-induced aphasia, PPA stems from progressive diseases like frontotemporal dementia or Alzheimer's, and it worsens over time. It often begins with subtle word-finding issues and progresses to more severe impairments, while other cognitive functions remain initially intact. PPA includes three variants – nonfluent/agrammatic, semantic, and logopaedic – each with distinct linguistic and neural profiles (Mesulam, 2003).

Currently, no medical or surgical treatments exist for aphasia, making speech therapy the primary intervention. Research has shown that consistent speech therapy can lead to significant improvements in patients' conditions (Roberts *et al.*, 2022). Rehabilitation programs are tailored to address the specific challenges experienced by individuals with different types of aphasia. These highly personalized interventions target each patient's unique communication needs, often strengthened through a collaborative therapeutic alliance between the clinician and the individual. To guide therapy effectively, clinicians use assessments such as the **Boston Diagnostic Aphasia Examination (BDAE)** to evaluate language abilities and design customized treatment plans (Marie and AISwaiti, 2024; Cherney *et al.*, 2011). The success of aphasia therapy depends on factors such as aphasia type, severity, patient motivation, cognitive profile, and therapy intensity, with evidence indicating that early, intensive, and tailored interventions yield the best outcomes (Berthier, 2005; Biel *et al.*, 2022).

Despite its importance, many individuals receive insufficient therapy. Barriers such as therapist shortages, limited insurance coverage, mobility constraints, and caregiver availability, along with traditional sessions – often lasting up to two hours daily for optimal results – pose further challenges (Stahl *et al.*, 2018; Guo *et al.*, 2014). Tele-rehabilitation emerged as a promising way for patients to get therapy from home, helping to overcome distance or other practical problems (Towey, 2012). Computer programs can also help with rehabilitation, training skills like speaking, reading aloud, writing, understanding, memory, and finding words (Katz, 2010; Loverso *et al.*, 1992). However, their repetitive nature and lack of personalization often reduce patient motivation over time (Niculescu *et al.*, 2023).

This paper integrates insights from neuroscience, language technology, and clinical practice, presenting findings from our fieldwork and prototype development of culturally and linguistically adaptive tools for aphasia

rehabilitation. Specifically, we explore the integration of neurolinguistic principles and natural language processing (NLP) methods into assistive technology to support speech therapy. Our work focuses on developing a prototype tailored to the needs of aphasia patients in Singapore, aiming to bridge the gap between clinical intervention and home-based practice through adaptive, user-centred design.

2 Background and Related Work

Breakthroughs in neuroscience have significantly deepened our understanding of the neural foundations of language, while innovations in language technologies are opening new avenues for diagnosis and intervention, positioning aphasia research at the intersection of multiple disciplines. This section provides a two-part review: the first part surveys cognitive neuroscience research that underpins current therapeutic approaches; the second examines emerging developments in language technology – from automatic speech recognition and AI-driven assessments to digital therapy platforms – that are reshaping diagnostic and treatment practices.

2.1 Neuroscience insights

Progress in neuroscience and neurolinguistics reveals how the brain processes language, the ways aphasia disrupts these functions, and possible pathways for recovery. Neuroplasticity plays a key role in how the brain reorganizes after injury, supported by insights from neuroimaging and electrophysiological markers that track language deficits and recovery. Additionally, recognizing that language involves multifunctional brain networks highlights the importance of addressing cognitive impairments alongside linguistic rehabilitation.

Neuroplasticity and non-invasive brain stimulation (NIBS)

A main focus in neuroscience and language research is **neuroplasticity**, the brain's remarkable ability to reorganize and form new connections after injury. Brain imaging shows that the brain can compensate for language problems by recruiting alternative paths, often in the right hemisphere or areas near the injury. Understanding these compensatory processes helps researchers learn how the brain adapts after stroke or neurodegenerative diseases (Turkeltaub *et al.*, 2025). Based on this knowledge, researchers are developing personalized rehabilitation programs tailored to each person's type of aphasia and brain injury. Studies show that therapy matched to a patient's neural profile is more effective at promoting recovery (Kiran and Thompson, 2019).

Building on this, **non-invasive brain stimulation (NIBS)** – including transcranial direct current stimulation (tDCS) and repetitive transcranial magnetic stimulation (rTMS) – is being used alongside speech therapy. These techniques modulate brain activity and support the reorganization of language networks (Fridriksson *et al.*, 2018; Sebastian *et al.*, 2016). Clinical trials suggest that tDCS,

combined with therapy, could help patients with both post-stroke and primary progressive aphasia (PPA) improve skills such as naming, fluency, and communication (Meinzer *et al.*, 2014; Tsapkini *et al.*, 2018). Research on post-stroke aphasia also suggests tDCS may reduce mental fatigue and help patients maintain attention, supporting longer-term benefits (Fridriksson *et al.*, 2018; Kang *et al.*, 2009).

Neuroimaging and electrophysiological markers

Advanced neuroimaging techniques, including functional MRI (fMRI), diffusion tensor imaging (DTI), and magnetoencephalography (MEG), are crucial for understanding how the brain reorganizes after stroke, particularly in chronic aphasia (Stockert *et al.*, 2020; Li *et al.*, 2022). These methods are often combined with electrophysiological techniques such as event-related potentials (ERPs), which provide detailed timing information about language processing.

Two ERP components – N400 and P600 – are especially informative. The N400, a negative deflection around 400 milliseconds after a word is presented, reflects **semantic processing**, or the retrieval and integration of word meaning. It occurs relatively quickly because understanding meaning relies on direct access to memory and context. The P600, a positive deflection around 600 milliseconds, is linked to **syntactic processing**, or the analysis of grammatical structure, which is computationally more complex and therefore slower. In aphasia, the N400 is often delayed or reduced, indicating semantic impairments, while the P600 may be diminished or absent, signalling syntactic deficits. Importantly, both components can change with therapy: successful treatment can increase N400 amplitude and normalize its timing, while re-emergence of the P600 may indicate recovery of syntactic integration (Stalpaert *et al.*, 2022). These ERP patterns can therefore be used to diagnose language impairments and to track progress during rehabilitation, serving as sensitive neural markers of both deficits and neuroplasticity in aphasia recovery.

Neural multifunctionality and cognitive impairments

While these neurophysiological markers provide important insights into core language processes, research increasingly shows that aphasia often affects broader cognitive functions beyond language. Studies in neurolinguistics have found that deficits in attention, theory of mind, and executive function frequently occur alongside morphosyntactic, semantic, and pragmatic difficulties (Caplan, 2020). These cognitive impairments are associated with the severity of aphasia and with structural changes in the brain, particularly in frontotemporal regions (Strikwerda-Brown *et al.*, 2019; Fittipaldi *et al.*, 2019). For example, individuals with primary progressive aphasia (PPA) may show early problems with emotion recognition and empathy, even when language abilities remain relatively intact. Similarly, post-stroke aphasia is linked to reduced performance on attention and executive

function tasks, which closely relates to communication difficulties (Murray, 2012). These findings support the concept of neural multifunctionality – the idea that language operates within broader domain-general brain systems – and highlight the importance of interventions that target both linguistic and cognitive aspects of aphasia (Cahana-Amitay and Albert, 2014; Kytarova *et al.*, 2017).

From this broader perspective, pharmacological interventions are being investigated as complements to behavioural therapy, with the potential to improve not only language deficits but also co-occurring cognitive impairments (Stockbridge, 2022). Although still in early clinical stages, some compounds aim to enhance cerebral blood flow, modulate neurochemical balance, or support neural recovery, thereby promoting improvements in both language and cognitive function (Berthier, 2021; Mayo Clinic, 2025). As understanding of the neural basis of language and cognition grows, pharmacological strategies may become an increasingly important component of personalized, integrated aphasia rehabilitation programs.

2.2 Language technology developments

Language technologies are increasingly important for supporting aphasia diagnosis, therapy, and communication. Advances in artificial intelligence (AI) have made these tools more scalable, adaptive, and personalized, complementing traditional clinical care. They can be used for automated assessment and diagnosis, personalized or adaptive therapy, and augmentative or alternative communication (AAC) to help patients communicate when natural speech is severely limited.

Aphasia assessment and diagnosis

AI technologies are increasingly used for early detection, accurate diagnosis, and ongoing monitoring of aphasia. Machine learning models can now extract detailed linguistic and acoustic features from speech to identify language impairments with high precision. For example, (Wagner, 2023) developed an automated speech recognition (ASR) and natural language processing (NLP) pipeline that classified aphasia subtypes with up to 90% accuracy – comparable to expert clinicians. Similarly, (Herath *et al.*, 2022) used deep neural networks and Mel Frequency Cepstral Coefficients (MFCCs) to classify ten levels of aphasia severity with over 99% accuracy, showing promise for more precise severity grading and personalized therapy planning.

Beyond classification, AI models are increasingly being applied to fine-grained error analysis and discourse-level interpretation. (Salem *et al.*, 2023) fine-tuned a large language model (LLM) on picture elicited narratives to identify target words in paraphasic speech, achieving 50.70% accuracy. While further work is needed, this study demonstrates how AI can automatically detect and analyse paraphasic

expressions in connected speech, taking a step closer to automatic aphasic discourse analysis.

AI is also advancing longitudinal tracking in progressive language disorders such as PPA. Subtle changes in lexical diversity, syntactic complexity, or speech rate can signal cognitive decline. As noted by (Zhong, 2024), AI-integrated systems offer non-invasive, cost-effective tools for continuous monitoring, supporting earlier intervention and personalized therapy adjustments.

Augmentative & alternative communication

Language technologies are increasingly providing new ways for people with aphasia to communicate, through both AI-enhanced Augmentative and Alternative Communication (AAC) systems and emerging Brain–Computer Interfaces (BCIs). Modern AAC devices now use machine learning and natural language processing to generate contextually appropriate phrases and predictive text. For example, Transformer-based models can adapt vocabulary to a user’s communication history and conversational context, improving fluency and reducing cognitive effort (Caria, 2025). User-centred studies also show that features such as real-time grammar support and visual feedback make communication more accurate and boost user confidence (Mao *et al.*, 2025).

At the same time, non-invasive BCIs are showing promise in decoding imagined speech – translating neural activity into words. Recent EEG-based research highlights major progress in signal processing, feature extraction, and classification accuracy, especially using attention-based deep learning models (Su and Tian, 2025; Panachakel and Ramakrishnan, 2021). Some clinical trials have begun using BCI-driven neurofeedback in aphasia therapy. For instance, (Musso *et al.*, 2022) found that EEG-based feedback during word detection tasks improved language abilities across several domains in chronic poststroke aphasia, suggesting these systems may both restore communication and enhance neural plasticity. Together, these developments point toward a shared goal: bridging the gap between thought and expression for individuals with severe speech or motor impairments. Although most BCI applications are still experimental, the combination of adaptive AI and brain-based communication technologies holds great promise for improving expressive communication in aphasia related disorders.

Language technologies for aphasia therapy

Recent advances in language technologies – such as automatic speech recognition (ASR), neural machine translation (NMT), and natural language processing (NLP) – are increasingly being applied to aphasia therapy to strengthen pronunciation, grammar, vocabulary, and overall communication skills. For example, in a small study of an iPad-based speech therapy app, the ASR system correctly identified whether participants pronounced each word accurately in the

same way as human raters about 80% of the time, and participants showed improved word production after four weeks of ASR-guided practice (Ballard *et al.*, 2019). NMT models have also been used to automatically translate aphasic sentences into fluent ones (BLEU $\approx$ 38), demonstrating the feasibility of providing grammatical and lexical support to aid communication (Smaili *et al.*, 2022).

In addition, integrated ASR–NLP approaches can transcribe spontaneous aphasic speech with a word error rate of 24.5% and, when combined with feature-based classifiers, can detect aphasia with 86% accuracy (Barberis *et al.*, 2024).

Building on these capabilities, mobile and gamified apps have made therapy more accessible and engaging for long-term self-directed use. Virtual reality app, such as (Bu *et al.*, 2022), offer immersive language and cognitive training targeting oral expression, comprehension, and cognitive skills; they have been pilot-tested with patients, healthy volunteers, and clinicians. Commercial programs including (VAsTtx, 2017; Constant Therapy, 2023; Bungalow Software, 2021; AphasiaScripts, 2021; Sentence Sharper, 2023 and iTalkBetter, 2024) offer structured exercises targeting multiple language and cognitive skills (reading, writing, speaking, and problem-solving). Many now use adaptive algorithms that adjust difficulty based on performance (Zhong, 2024).

Last but not least, Melodic Intonation Therapy (MIT) remains an important complementary intervention for non-fluent aphasia. By using melody, rhythm, and pitch to stimulate speech, it has shown effectiveness in both subacute and chronic stages. A randomized controlled trial conducted across multiple therapy centres reported significant gains in repetition and communication after six weeks of MIT (van der Meulen *et al.*, 2014). A crossover trial also confirmed functional improvements using the Communicative Activity Log (CAL), a standardized tool measuring how well patients communicate in daily life rather than only in the clinic (Haro-Martinez *et al.*, 2018). Neuroimaging evidence shows that MIT promotes right-hemisphere reorganization and increases structural connectivity of the *arcuate fasciculus* in patients with non-fluent aphasia (Zhang *et al.*, 2023). New approaches are now combining MIT with ASR-based prosody feedback, though these innovations have yet to be fully evaluated.

3 An aphasia training tool for Singaporean patients

3.1 Motivation

Despite advances in language technologies and the growing availability of digital tools for home-based aphasia therapy, significant limitations remain – particularly in generalizing rehearsed speech to spontaneous, real-world communication (Simmons-Mackie *et al.*, 2014). Many existing systems still lack adaptive feedback, contextual relevance, cultural and linguistic personalization. Motivated by these challenges, our team – drawing on expertise in speech recognition under noisy and reverberant conditions (Chong *et al.*, 2019; Niculescu

et al., 2010; D’Haro *et al.*, 2017; Niculescu *et al.*, 2016a; Niculescu *et al.*, 2015; Niculescu *et al.*, 2016b), as well as in chatbot development across various domains (Niculescu *et al.*, 2020; Niculescu *et al.*, 2009) – engaged in a focused collaboration with Tan Tock Seng Hospital in Singapore to develop assistive tools to local patient needs.

3.2 Field study

To gain a nuanced understanding of the challenges faced by both patients and clinicians in aphasia therapy, we conducted an ethnographic field study. The study included structured observations of therapy practices, qualitative analysis of data samples from the Open AphasiaBank corpus (AphasiaBank, 2017), and interviews with speech-language pathologists from Tan Tock Seng Hospital (TTSH). The results of this study, published in (Niculescu *et al.*, 2023) identified several structural and interactional challenges in aphasia therapy. Patients struggled with sentence construction, word retrieval, and pronunciation. Therapists, in turn, emphasized the importance of individualized therapy plans incorporating storytelling, conversational practice, grammatical exercises, and music-based tasks – such as singing familiar phrases – to activate automatic speech sequences.

Motivating patients, adapting materials to local cultural contexts, and supporting long-term independent training were recurrent themes. Therapists frequently used mental connections between words and images to stimulate speech production and relied on family input to establish a core vocabulary drawn from the patient’s environment and everyday routines. However, technological limitations – such as ASR tools that fail with non-American accents or therapy apps that use culturally irrelevant imagery – continue to hinder effective therapy delivery.

Further interviews highlighted that local patients present wide variation in language proficiency, literacy levels, and dominant languages (e.g., English, Malay, Chinese dialects, Tamil, etc.). Therapy must often be adapted to suit these differences, especially when literacy impairments or apraxia of speech are present. Yet current therapy apps frequently rely on Western-centric content or accent-dependent recognition, limiting their usability in the local context.

As patients improve, they seek more advanced training – such as more complex sentence construction – but lack appropriate digital tools beyond therapist-guided sessions. The absence of structured feedback during home practice poses a risk of reinforcing errors, especially for individuals living alone. These insights underscore the need for culturally relevant, feedback-rich tools that can support both patients and therapists in tracking progress and promoting recovery.

3.3 Mapping neurolinguistic and contextual insights to therapy goals and prototype features

Drawing on ethnographic findings and neurocognitive research on aphasia, we mapped key neurolinguistic and contextual insights to corresponding NLP applications and speech therapy goals, highlighting potential solution strategies realized through prototype implementation. This mapping guided the design of two digital therapy prototypes, emphasizing targeted, adaptive, and linguistically grounded interventions. Table 1 summarizes these connections, with implemented features highlighted in **bold** and planned features in *italic*, showing how advances in speech recognition and NLP are applied to support personalized and engaging home-based rehabilitation that complements traditional clinical therapy.

Table 1: Mapping neurolinguistic and contextual insights to speech therapy goals and prototype features

Neurolinguistic insights	NLP applications	Speech therapy goals	Prototype implementation
1. Knowledge of syntactic deficits	Use parsers to detect malformed sentences	Sentence-level training	Prototype 2: 'Complete the Sentence' targets syntax via sentence construction (1:1 match)
2. Lexical-semantic mapping in the brain	Semantic vector models	Suggest related words during naming	<i>Planned:</i> Future versions of both prototypes may use LLMs to offer flexible suggestions
3. Multimodal association in lexical access	Vision-language models	Use image cues to trigger speech	Prototype 1: Current 1:1 image word pronunciation match; visual stimuli planned
4. Real-time language processing deficits	ERP-based dialogue pacing	Adapt turn-taking pace	<i>Planned:</i> Potential for future interaction design
5. Pronunciation errors	Classification models on aphasic speech	Error-specific feedback	Prototype 1: Custom scoring + DTW detect phonetic errors, give personalized feedback
6. Reduced speech fluency & word retrieval	Rhythmic/formulaic language models	Support automaticity; use fixed patterns	<i>Planned:</i> Singing-based or patterned repetition exercises may be added
7. Storytelling & conversation deficits	Use LLMs	Support long sentence formation	<i>Planned:</i> GPT model embeddings
Other contextual insights	Strategies	Speech therapy goals	Prototype implementation
1. Lack of personal & cultural adaptation	Enable caregiver involvement	Engage patients with relevant content	Prototype 2: Exercises tailored to personal, cultural context via caregiver input
2. Lack of remote progress monitoring	Enable therapist involvement	Track improvement outside the clinic	Prototype 1: pronunciation scoring Prototype 2: Session logs + therapist view

Neurolinguistic insights	NLP applications	Speech therapy goals	Prototype implementation
3. Fatigue & repetitive exercise routine	Gamification	Increase patient engagement	Prototype 2: Uses tree-growth metaphor to encourage consistent training
4. Local accent challenges in ASR	Accent-adapted ASR models	Broader accessibility for diverse speakers	Prototype 1: uses ASR adapted to Singaporean English

3.4 Prototype 1

The first prototype, as presented in (Niculescu *et al.*, 2023), aims to facilitate home-based speech training for individuals with mild aphasia by addressing the lack of immediate and personalized feedback in existing digital interventions. It is a web-based system that employs multimodal prompting – comprising text, images, audio, and synchronized lip videos – to accommodate patients with varying degrees of aphasia and to enhance comprehension of therapy tasks (see Figure 1).




Figure 1: Prototype 1 – exercises for word reading, object naming, and pronunciation with lip sync

Each word is paired one-to-one with an image cue. Patients are shown an image and invited to speak the name of the object. They can practice reading, naming, listing, and repetition. At the end of each task, patients receive a **pronunciation score** along with an **explanation of errors**, and can replay a lip-synchronization video to review the **correct pronunciation**¹.

At the core of the system is a speech recognition engine adapted specifically for Singaporean-accented English to better reflect the phonetic characteristics of the target user population. While British English is often preferred in applications for its perceived authority and prestige (Niculescu *et al.*, 2008), Singapore English – a local variety influenced by Chinese and Malay – is commonly used in everyday interaction. The initial speech-to-text conversion facilitates basic content matching, while a secondary evaluation layer assesses pronunciation accuracy using Dynamic Time Warping (DTW) (Senin, 2008). This method compares patient

¹ A video demonstration of the prototype is available here: <https://www.youtube.com/watch?v=JkL7lXQm7z8>

utterances against multiple prerecorded templates, enabling detailed and nuanced feedback that accommodates accent variability and low-resource language conditions.

To enhance semantic generalization – particularly in object-naming tasks – the system compares recognized speech against a set of acceptable alternative responses, minimizing false-negative evaluations. While the current semantic matching operates via rule-based logic, future iterations aim to incorporate large language models (LLMs) to support more natural and flexible user input.

The system is cross-platform, supports patient progress tracking, and includes **therapist involvement**.

3.5 Prototype 2

Building on the foundational Prototype 1, Prototype 2, **SpeakNow**, was designed and implemented by (Lim, 2024) as a web-based speech therapy platform featuring dual interfaces: one for patients engaging in therapy, and another for caregivers and therapists to create, assign, and monitor exercises.

A key feature is the customizability of therapy decks, which allows caregivers to design exercises tailored to the patient’s **personal history, cultural background, and language proficiency**. For instance, practice materials can reference familiar local dishes, hawker foods, or recognizable streets and neighbourhoods in Singapore, alongside meaningful life events. For example, as shown in Figure 2a, one exercise refers to Kopi C – a popular local coffee made with evaporated milk and sugar – reflecting everyday cultural references familiar to many Singaporeans. This personalization enhances emotional relevance and motivation by grounding therapy in the patient’s lived experiences.

Therapy exercises are organized into decks, which can be created manually or generated with ChatGPT-assisted input, with audio integrated via text-to-speech (TTS). Each deck may include multiple exercise types, such as **Complete the Sentence** (Figure 2a) and **Repeat the Sentence** (Figure 2b).

Complete the Sentence targets syntax and sentence construction through 1:1 word-to-blank matching. Patients reconstruct incomplete sentences by dragging and dropping words into blanks, with distractors included to increase task complexity. In the future, we plan to allow users to construct their own sentences based on given scenarios. To support this, we are exploring the integration of automatic parsers to detect lexical and syntactic errors, extending beyond simple rule-based matching.

Repeat the Sentence focuses on improving pronunciation and auditory comprehension. Patients listen to a sentence – either generated using Google Text-to-Speech or manually uploaded – and repeat it aloud. Real-time automatic speech recognition (ASR) analyses the spoken response and provides feedback based on a predefined similarity threshold (set at 75%).




Figure 2: Prototype 2 screenshots

To encourage daily engagement, caregivers can assign personalized tasks, which patients access via a *Daily Challenge* feature. The system selects exercises randomly from pre-set decks. All user interactions – including accuracy, duration, and frequency – are logged in MongoDB and visualized through the therapist dashboard. Performance summaries are generated at the end of each session and can be reviewed daily or weekly. This feature supports **remote monitoring** and **adaptive planning**. The interface is minimalistic and aphasia-friendly, with large, clearly labelled elements accommodating motor and cognitive impairments (Janeback and Jonsson, 2020). Therapists can register patients, link caregiver accounts, and manage therapy routines.

To promote consistent engagement, **SpeakNow** employs lightweight gamification grounded in the Octalysis framework (Gellner and Buchem, 2022). A virtual tree metaphor represents progress: completing daily tasks nurtures its growth, while missed sessions reset it, tapping into the emotional drive to avoid loss (Figure 2c). Instant feedback via encouraging messages further reinforces engagement.

Built with the MERN stack and hosted on AWS, **SpeakNow** supports media APIs, secure storage, and scalable backend infrastructure.

4 Future Directions and Conclusions

Our work represents an early but important step toward AI-enhanced assistive technologies for aphasia therapy in Singapore. A key limitation is that, although our prototypes have been conceptually and technically validated, they have not yet been tested with patients. Meaningful progress in aphasia rehabilitation depends on iterative, user-centred evaluations, which we plan to conduct once approvals are granted. Without this feedback, no technology – however sophisticated – can fully meet users’ complex needs.

Our ethnographic findings have provided valuable insights into everyday therapy challenges, helping us shape tools that reflect local linguistic diversity, cultural context, and patient motivation. These insights guided the design of our

second prototype, **SpeakNow**, which supports functional communication through personalized, adaptive, and culturally grounded therapy experiences.

Looking ahead, we plan system-level enhancements, including replacing rule-based parsing with syntactic parsers to detect malformed utterances; integrating semantic vector and visual-to-language models to suggest related words during naming tasks; and incorporating music-based therapy elements such as rhythmic repetition and singing to stimulate formulaic language patterns. We also intend to experiment with LLM-based conversational agents to train pragmatic and turn-taking skills – bridging structured exercises and spontaneous conversation. Once implemented, these enhancements will be evaluated with patients for usability, engagement, and therapeutic effectiveness.

Advancing aphasia research will require diverse and representative datasets, particularly for non-English and disfluent speech, to improve the robustness of automatic speech recognition and language modelling. Future systems should embrace multilingual and multimodal interaction – including gestures, facial expressions, and prosody – to better mirror real-world communication. Integrating neurocognitive markers, such as ERP-based recovery indicators, into adaptive therapy algorithms is another promising direction requiring close collaboration between AI researchers and neuroscientists.

Despite rapid progress, most current technologies lack real-time adaptability, personalized feedback, and cultural sensitivity. Few systems have been evaluated for long-term engagement or sustainability in clinical and home settings. Addressing these gaps requires interdisciplinary collaboration and continuous co-design with patients and clinicians. Future systems that combine clinical insight with adaptive, feedback-driven designs can more effectively help individuals with aphasia regain functional communication.

AI Use Statement

Artificial intelligence tools (ChatGPT, OpenAI) were used to assist with language editing, improve clarity and grammatical consistency, and generate an initial version of the manuscript’s abstract after the main text had been completed. All AI-generated content was reviewed, corrected, and validated by the authors prior to inclusion.

References

- Alexandra C. Salem, Robert C. Gale, Mikala Fleegle, Gerasimos Fergadiotis, and Steven Bedrick. 2023. Automating intended target identification for paraphasias in discourse using a large language model. *Journal of Speech, Language, and Hearing Research (JSLHR)*, 66(12):4949–4966.
- Ana M. Haro-Martínez, Genny Lubrini, Rosario Madero-Jarabo, Exuperio Díez-Tejedor, and Blanca Fuentes. 2018. Melodic intonation therapy in post-stroke nonfluent aphasia: a randomized pilot trial. *Clinical rehabilitation*, 33(1):44–53.

- Andrea Caria. 2025. Toward predictive communication: The fusion of large language models and non-invasive BCI spellers. *Sensors*, 25(13):3987.
- Andreea I. Niculescu, Betsy van Dijk, Anton Nijholt, Dilip K. Limbu, Swee Lan See, and Anthony H.Y. Wong. 2010. Socializing with Olivia, the youngest robot receptionist outside the lab. In S.S. Ge *et al.*, editor, *Proc. of the Int. Conference in Social Robotics*, volume 6414 of *LNAI*, pages 50–62, Berlin, Germany. Springer-Verlag.
- Andreea I. Niculescu, Bimlesh Wadhwa, and Evan Quek. 2016a. Smart city technologies: design and evaluation of an intelligent driving assistant for smart parking. *Int. Journal on Advanced Science, Engineering and Information Technology (IJASEIT)*, pages 143–149.
- Andreea I. Niculescu, Bimlesh Wadhwa, and Evan Quek. 2016b. Technologies for the future: evaluating a voice enabled smart city parking application. In *Proc. of the 4th Int. Conference on User Science and Engineering (i-USER)*, pages 46–50, Melaka, Malaysia. IEEE.
- Andreea I. Niculescu, George M. White, Swee Lan See, Ratna U. Waloejo, and Yoko Kawaguchi. 2008. Impact of English regional accents on user acceptance of voice user interfaces. In *Proc. of the 5th Nordic conference on Human-computer interaction: building bridges NordiCHI '08.*, pages 523–526, Lund, Sweden. ACM.
- Andreea I. Niculescu, Ivan Kukanov, and Bimlesh Wadhwa. 2020. DigiMo – towards developing an emotional intelligent chatbot in Singapore. In *Proc. of the 2020 Symposium on Emerging Research from Asia and on Asian Contexts and Cultures*.
- Andreea I. Niculescu, Jochen Ehnes, Minghui Dong, Changhuai You, and Paul Chan Yaozhu. 2023. Opportunities and challenges of designing assistive technologies for aphasia patients in Singapore: The case of a speech evaluation prototype. In *Proc. of the AAAI Symposium Series*, volume 1, pages 90–93.
- Andreea I. Niculescu, Mei Quin Lim, Seno A. Wibowo, Kheng Hui. Yeo, Boon Pang Lim, Michael. Popow, Dan Chia, and Rafael E. Banchs. 2015. Designing IDA – an intelligent driver assistant for smart city parking in Singapore. In *Proc. of the 15th IFIP Conference on Human Computer Interaction (INTERACT)*, LNCS, pages 510–513, Bamberg, Germany. Springer Int. Publishing.
- Andreea I. Niculescu, Yujia Cao, and Anton Nijholt. 2009. Stress and cognitive load in multimodal conversational interactions. In Constantine Stephanidis, editor, *Proceedings of the 13th International Conference on HumanComputer Interaction (HCI International 2009)*, pages 891–895. Springer. (Proceedings and Poster DVD).
- Anika Stockert, Max Wawrzyniak, Julian Klingbeil, Katrin Wrede, Dorothee Kümmerer, Gesa Hartwigsen, Christoph P. Kaller, Cornelius Weiller, and Dorothee Saur. 2020. Dynamics of language reorganization after left temporo-parietal and frontal stroke. *Brain*, 143(3):844–861.
- Antonio R. Damasio. 1992. Aphasia. *New England Journal of Medicine*, 326(8):531–539.
- AphasiaBank. 2017. Aphasiabank. <https://aphasia.talkbank.org/>. Accessed: July 2025.
- AphasiaScripts. 2021. Aphasiascripts. <https://www.sralab.org/aphasiascripts>. Accessed: July 2025.
- Basem Marie and Fadi AlSwaiti. 2024. Validity and validation in language testing: Current state and future guidelines in aphasiology. *Aphasiology*, 38(1):1–21.
- Benjamin Stahl, Bettina Mohr, Verena Büscher, Felix R. Dreyer, Guglielmo Lucchese, and Friedemann Pulvermüller. 2018. Efficacy of intensive aphasia therapy in patients

- with chronic stroke: A randomised controlled trial. *Journal of Neurology, Neurosurgery & Psychiatry*, 89(6):586–592, June.
- Bing Hong Lim. 2024. Design and implementation of web application for aphasia training. Final Year Project Report.
- Bungalow Software. 2021. Recover speech and language. <https://www.bungalowsoftware.com/> Accessed: October 2025.
- Carolyn Gellner and Ilona Buchem. 2022. Evaluation of a gamification approach for older people in e-learning. In *Proc. of INTED2022*, pages 596–605.
- Cherie Strikwerda-Brown, Siddharth Ramanan, and Muireann Irish. 2019. Neurocognitive mechanisms of theory of mind impairment in neurodegeneration: a transdiagnostic approach. *Neuropsychiatric Disease and Treatment*, 15:557–573.
- Constant Therapy. 2023. Constant therapy. <https://constanttherapyhealth.com/constant-therapy/>. Accessed: July 2025.
- Dalia Cahana-Amity and Martin L. Albert. 2014. Brain and language: Evidence for neural multifunctionality. *Behavioural Neurology*, 2014:260381.
- David Caplan. 2020. *Neurolinguistics and linguistic aphasiology*. Cambridge University Press, Cambridge.
- Frank L. Loverso, Thomas E. Prescott, and Michael Selinger. 1992. Microcomputer treatment applications in aphasiology. *Aphasiology*, 6(2):155–163.
- Ineke van der Meulen, W. Mieke E. van de Sandt-Koenderman, Majanka H. Heijenbroek-Kal, Evy G. Visch-Brink, and Gerard M. Ribbers. 2014. The efficacy and timing of melodic intonation therapy in subacute aphasia. *Neurorehabilitation and Neural Repair*, 28(6):536–544.
- iTalkBetter. 2024. <https://italk.html>. Accessed: October 2025.
- Jara Stalpaert, Sofie Standaert, Lien D’Helft, Marijke Miatton, Anne Sieben, Tim Van Langenhove, Wouter Duyck, Pieter van Mierlo, and Miet De Letter. 2022. Therapy-induced electrophysiological changes in primary progressive aphasia: A preliminary study. *Frontiers in Human Neuroscience*, 16.
- Jerrin T. Panachakel and Angarai G. Ramakrishnan. 2021. Decoding covert speech from EEG – a comprehensive review. *Frontiers in Neuroscience*, 15:642251.
- Johanna Janebäck and Ella Jonsson. 2020. Designing an activity application for people with cognitive disabilities: What should be considered when designing a UI for people with cognitive disabilities? Master’s thesis, Chalmers University of Technology.
- Julius Fridriksson, Chris Rorden, Jordan Elm, Souvik Sen, Mark S. George, and Leonardo Bonilha. 2018. Transcranial direct current stimulation vs sham stimulation to treat aphasia after stroke: A randomized clinical trial. *JAMA Neurology*, 75(12):1470–1476.
- Kamel Smaili, David Langlois, and Peter Pribil. 2022. Language rehabilitation of people with Broca aphasia using deep neural machine translation. In *Proc. of the Fifth Int. Conference on Computational Linguistics in Bulgaria (CLIB 2022)*, page 162 – 170, Sofia, Bulgaria, September. Department of Computational Linguistics, IBL – BAS.
- Katerina Hilari and Sally Byng. 2009. Health-related quality of life in people with severe aphasia. *Int. Journal of Language & Communication Disorders*, 44(2):193–205.
- Ke Su and Liang Tian. 2025. Systematic review: progress in EEG-based speech imagery brain-computer interface decoding and encoding research. *Peer J Computer Science*, 11:e2938.

- Kirrie J. Ballard, Nicole M. Etter, Songjia Shen, Penelope Monroe, and Chek Tien Tan. 2019. Feasibility of automatic speech recognition for providing feedback during tablet-based treatment for apraxia of speech plus aphasia. *American Journal of Speech-Language Pathology*, 28(2S):818–834.
- Kyoung Eun Kang, Min Jae Baek, Sangyun Kim, and Nam-Jong Paik. 2009. Non-invasive cortical stimulation improves post-stroke attention decline. *Restorative neurology and neuroscience*, 27(6):645 – -650.
- Kyran Tsapkini, Kimberly T. Webster, Bronte N. Ficek, John E. Desmond, Chiadi U. Onyike, Brenda Rapp, Constantine E. Frangakis, and Argye E. Hillis. 2018. Electrical brain stimulation in different variants of primary progressive aphasia: A randomized clinical trial. *Alzheimer's & Dementia (New York, N. Y.)*, 4:461–472.
- Laura L. Murray. 2012. Attention and other cognitive deficits in aphasia: Presence and relation to language and communication measures. *American Journal of Speech-Language Pathology*, 21(2):51–64.
- Laurin Wagner, Mario Zusag, and Theresa Bloder. 2023. Careful whisper – leveraging advances in automatic speech recognition for robust and interpretable aphasia subtype classification. In *Proc. of Interspeech 2023*, Dublin, Irland.
- Lei Mao, Jong Ho Lee, Yasmeen F. Shah, and Stephanie Valencia. 2025. Design probes for AI-driven AAC: Addressing complex communication needs in aphasia. *arXiv preprint arXiv:2504.09435*.
- Leora R. Cherney, Janet P. Patterson, and Anastasia M. Raymer. 2011. Intensity of aphasia therapy: Evidence and efficacy. *Current Neurology and Neuroscience Reports*, 11(6):560–569.
- Lucie Kytarová, Petr Nilius, and Kateřina Vitásková. 2017. Aphasia and nonlinguistic cognitive rehabilitation: Case study. In *European Proc. of Social Behavioural Sciences*, pages 402–412.
- Luis F. D’Haro, Andreea I. Niculescu, Caixia Cai, Suraj Nair, Rafael E. Banchs, Alois Knoll, and Haizhou. Li. 2017. An integrated framework for multimodal human-robot interaction. In *Proc. of Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*, Kuala Lumpur, Malaysia. IEEE (APSIPA).
- Mara Barberis, Pieter De Clercq, Bastiaan Tamm, Hugo Van hamme, and Maaïke Vandermosten. 2024. Automatic recognition and detection of aphasic natural speech. *arXiv preprint arXiv:2023arXiv230801327W*.
- Marcelo L. Berthier. 2005. Poststroke aphasia. *Drugs & Aging*, 22(2):163–182.
- Marcelo L. Berthier. 2021. Ten key reasons for continuing research on pharmacotherapy for post-stroke aphasia. *Aphasiology*, 35(6):824–858.
- Marcus Meinzer, Robert Lindenberg, Mira M. Sieg, Laura Nachtigall, Lena Ulm, and Agnes Flöel. 2014. Transcranial direct current stimulation of the primary motor cortex improves word-retrieval in older adults. *Front Aging Neurosciences*, 6:253.
- Mariacristina Musso, David Hübner, Sarah Schwarzkopf, Maria Bernodsson, Pierre LeVan, Cornelius Weiller, and Michael Tangermann. 2022. Aphasia recovery by language training using a brain–computer interface: a proof-of-concept study. *Brain Communications*, 4(1):fcac008, 02.
- Marsel M. Mesulam. 2003. Primary progressive aphasia—a language-based dementia. *The New England Journal of Medicine*, 349(16):1535–1542.

- Mayo Clinic. 2025. Aphasia: Diagnosis and treatment. <https://www.mayoclinic.org/diseases-conditions/aphasia/diagnosis-treatment/drc-20369523>. Accessed: July 2025.
- Melissa D. Stockbridge. 2022. Better language through chemistry: Augmenting speech-language therapy with pharmacotherapy in the treatment of aphasia. In *Handbook of Clinical Neurology*, volume 185, pages 261–272. Elsevier.
- Michael Biel, Hillary Enclade, Amber Richardson, Anna Guerrero, and Janet Patterson. 2022. Motivation theory and practice in aphasia rehabilitation: A scoping review. *American Journal of Speech-Language Pathology*, 31:2421–2443.
- Michael P. Towey. 2012. Speech therapy telepractice for vocal cord dysfunction (VCD): Mainecare (Medicaid) cost savings. *Int. Journal of Telerehabilitation*, 4(1):33–36.
- Mudiyansele D.P.M. Herath, Arachchilage S.A. Weraniyagoda, Thilina M. Rajapaksha, Patikiri A. D. S. N. Wijesekara, Kalupahana L. K. Sudheera, and Peter H.J. Chong. 2022. Automatic assessment of aphasic speech sensed by audio sensors for classification into aphasia severity levels to recommend speech therapies. *Sensors (Basel, Switzerland)*, 22(18):6966.
- Nina Simmons-Mackie, Meghan C. Savage, and Linda Worrall. 2014. Conversation therapy for aphasia: a qualitative review of the literature. *Int. Journal of Language & Communication Disorders*, 49(5):511–526.
- Pavel Senin. 2008. Dynamic time warping algorithm review. Technical Report CSDL-08-04, CSDL.
- Peter E. Turkeltaub, Kelly C. Martin, Alycia B. Laks, and Andrew T. DeMarco. 2025. Right hemisphere language network plasticity in aphasia. *Brain: a journal of neurology*, Advance online publication. <https://doi.org/10.1093/brain/awaf308>.
- Rajani Sebastian, Kyrana Tsapkini, and Donna C. Tippet. 2016. Transcranial direct current stimulation in poststroke aphasia and primary progressive aphasia: Current knowledge and future clinical applications. *NeuroRehabilitation*, 39(1):141–152.
- Ran Li, Nishaat Mukadam, and Swathi Kiran. 2022. Functional MRI evidence for reorganization of language networks after stroke. In *Handbook of Clinical Neurology*, volume 185, pages 131–150. Elsevier.
- Richard C. Katz. 2010. Computers in the treatment of chronic aphasia. *Seminars in Speech and Language*, 31(1):34–41.
- Sentence Sharper. 2023. An innovative approach to language software for aphasia. <https://sentenceshaper.com/>. Accessed: October 2025.
- Sol Fittipaldi, Agustin Ibanez, Sandra Baez, Facundo Manes, Lucas Sedeno, and Adolfo M. Garcia. 2019. More than words: Social cognition across variants of primary progressive aphasia. *Neuroscience and Biobehavioral Reviews*, 100:263–284.
- Sophie Roberts, Rachel M. Bruce, Louise Lim, Hayley Woodgate, Kate Ledingham, Storm Anderson, Diego L. Lorca-Puls, Andrea Gajardo-Vidal, Alexander P. Leff, Thomas M. H. Hope, David W. Green, Jennifer T. Crinion, and Cathy J. Price. 2022. Better long-term speech outcomes in stroke survivors who received early clinical speech and language therapy: What’s driving recovery? *Neuropsychological Rehabilitation*, 32(9):2319–2341.
- Susie Parr. 2007. Living with severe aphasia: Tracking social exclusion. *Aphasiology*, 21(1):98–123, January.
- Swathi Kiran and Cynthia K. Thompson. 2019. Neuroplasticity of language networks in aphasia: Advances, updates, and future challenges. *Frontiers in Neurology*, 10:295.

- Tze Yuang Chong, Kye Min Tan, Kah Kuan Teh, Chang Huai You, Hanwu Sun, and Huy Dat Tran. 2019. The I2Rs ASR system for the VOICES from a Distance Challenge 2019. In *Proc. of the 20th Annual Conference of the Int. Speech Communication Association, Interspeech 2019*, pages 2458–2462, Graz, Austria.
- VASStx. 2017. Sample. aphasia software finder. <https://www.aphasiasoftwarefinder.org/vasttx-therapy-samples>. Accessed: July 2025.
- World Health Organization. 2004. *The Atlas of Heart Disease and Stroke*. World Health Organization. Published in collaboration with the U.S. Centers for Disease Control and Prevention.
- Xiaofan Bu, Peter Hf. Ng, Ying Tong, Peter Q. Chen, Rongrong Fan, Qingping Tang, Qinqin Cheng, Shuangshuang Li, Andy Sk Cheng, and Xiangyu Liu. 2022. A mobile-based virtual reality speech rehabilitation app for patients with aphasia after stroke: development and pilot usability study. *JMIR Serious Games*, 10(2):e30196.
- Xiaoying Zhang, Zuliyaer Talifu, Jianjun Li, Xiaobing Li, and Feng Yu. 2023. Melodic intonation therapy for non-fluent aphasia after stroke: A clinical pilot study on behavioral and DTI findings. *iScience*, 26(9):107453.
- Xiaoyun Zhong. 2024. AI-assisted assessment and treatment of aphasia: a review. *Frontiers in Public Health*, 12:1401240.
- Yiting Emily Guo, Leanne Togher, and Emma Power. 2014. Speech pathology services for people with aphasia: What is the current practice in Singapore? *Disability and Rehabilitation*, 36(8):691–704, April.

WHO DROPS /ʁ/ AND WHERE? INSIGHTS FROM SPEAKER VARIATION IN FRENCH RHOTIC DELETION

ANISIA POPESCU¹, YARU WU², LORI LAMEL¹,
AND IOANA VASILESCU¹

¹*LISN, Université Paris Saclay*
anisia.popescu@lisn.fr, lamel@limsi.fr,
ioana.vasilescu@lisn.fr

²*Université de Caen Normandie*
yaru.wu@unicaen.fr

Abstract

This study presents a computational approach to analysing /ʁ/ deletion in French journalistic speech, leveraging a 150-hour annotated speech corpus processed via forced alignment with pronunciation variants. The results concur with phonological predictions, showing higher deletion rates in word-final positions. However, weak correlations at the speaker level, along with variability across professional speakers' categories, point to potential overestimations of non-canonical deletions, particularly in word-initial positions. Through manual verification, we identify higher misclassification rates in initial /ʁ/ contexts – attributable both to the fact that aligners often treat segments below their minimal duration threshold as deletions, and to biases in the ASR's baseline pronunciation dictionary. These findings underscore the need for linguistic expertise in interpreting automatic phone alignment outputs and refining phonetic analysis tools. By combining empirical phonetic inquiry, speech technology, and quality control mechanisms, our work addresses key challenges in automated speech processing, digital resource accuracy, and the integration of linguistic knowledge in computational linguistics.

Keywords: automatic alignment, pronunciation variants, phonetic reduction, natural language processing

1 Introduction

Phonetic reduction phenomena are key drivers of phonetic variation across languages and speakers. Words are often realized in forms that deviate from their canonical structures, a phenomenon that is particularly pervasive in conversational speech. For example, in American English, Johnson (2004) showed that over 60% of words in their conversational corpus deviate from their canonical form on at

least one phone. However, reductions can occur in nearly all types of speech. Byrd (1994) defines these deviations as “reduced forms of a word are simplified with respect to the canonical (underlying) form and often involve assimilations, vowel centralization and the deletion or simplifications of segments”. This paper focuses on one specific type of reduction: the complete deletion of a segment. More precisely, we focus on the deletion of the rhotic consonant in French, the most common consonant in the language (Adda-Decker, 2006), which has been shown to exhibit considerable variation (Gendrot *et al.*, 2015; Fougeron, 2007; Chafcouloff, 1980).

Word-final /ʁ/ deletion, particularly in post-obstruent contexts, has received considerable attention in the literature, with studies involving both laboratory experiments and large corpus analyses. This phenomenon is well-attested in French, closely associated with conversation speech and recognized as a sociolinguistic marker (Laks, 1977; Guiraud, 1965; Boughton, 2015; Avanzi, 2023).

In a more comprehensive corpus study, Wu and Adda-Decker (2021) examined /ʁ/ deletion across various contexts, with particular emphasis on the phonetic environment, position within the word, and speech style. Their findings show that /ʁ/ deletion occurs in all contexts, albeit at different degrees. The present study builds on this work, employing large corpora and automated methods to investigate speaker-specific variation, incorporating speaker metadata to provide deeper insights into individual speech patterns. It is well-established that phonetic variation is position-dependent, with reductions being more prevalent in word-final positions than in word-initial ones (Scheer and Ségéral, 2008). If a speaker exhibits reduction in a weak position but not in a strong position, this pattern is expected and unproblematic. However, if a speaker reduces in a strong position but not in a weak position, this discrepancy would be unusual and warrant further investigation. By analysing variation at the speaker level, this study aims to offer a more nuanced understanding of the phenomenon, while also providing a scalable method for automated analyses. Consequently, we focus specifically on /ʁ/ deletion in word-initial and word-final positions, as these represent key contexts for understanding positional variation in phonetic reduction.

2 Methods

To extract tokens containing word-initial and word-final /ʁ/, we utilized two publicly available corpora of French speech: The ESTER corpus (Galliano *et al.*, 2006), which consists of formal broadcast news speech, and the ETAPE corpus (Gravier *et al.*, 2012), which includes more informal journalistic speech, such as debates and conversations. Together, these corpora provide approximately 150 hours of recorded speech (100 hours from ESTER and 50 hours from ETAPE).

The data was processed using an automatic speech recognition (ASR) system, which performed a phoneme-level forced alignment. The dictionary specifying the

pronunciations of the words in the manual orthographic transcriptions was augmented with pronunciation variants. In particular /ʁ/ was annotated to account for both its presence (canonical pronunciation: word initial - rue [ʁy] “street”; word final - *ministre* [ministʁ] “minister”) and its possible deleted form (non-canonical pronunciation: word initial - rue [y] “street”; word final: *ministre* [minist] “minister”).

From the alignment output, we extracted additional features, including the position of the target /ʁ/, its pronunciation type (canonical: /ʁ/ present; non-canonical: /ʁ/ deleted), phonotactic complexity (whether the target /ʁ/ occurred in a cluster or as a singleton), and speaker identity. For each speaker, we also collected additional metadata for public figures, such as their profession and nationality. The final dataset was comprised solely of native French speaker for whom metadata were complete. Furthermore, only those with both word-initial and word-final /ʁ/ observations were included in the dataset. This process resulted in a final dataset comprising 817 speakers and 159,261 individual /ʁ/ tokens. Table 1 shows the counts and percentages for all /ʁ/ observations for each corpus and position in the word.

For each speaker in the dataset, the deletion rates (Del%) of word-initial and word-final /ʁ/ were computed. To quantify the difference between these two positional deletion rates we introduced a new variable:

$$\Delta_{Del\%} = Initial_{Del\%} - Final_{Del\%} \quad (1)$$

Negative values of $\Delta_{Del\%}$ indicate a higher rate of /ʁ/ deletion in word-final position, while positive values indicate a higher rate of /ʁ/ deletion in word-initial position. From a phonological perspective we expect predominantly negative values, as /ʁ/ deletion is generally more frequent in word-final positions. In addition, to account for professional background - commonly used as a proxy for social class - we defined 18 distinct occupational categories. Additionally, an “Other” category was created to group professions with a maximum of four representatives (e.g., priests, architects, military officers). This classification helped maintain a well-distributed and representative dataset for further analysis.

3 Results

This section presents the results of the analysis, focusing on the patterns of /ʁ/ deletion across different speakers and positional contexts. The findings will be outlined in detail, highlighting both overall trends and speakers’ specific variation.

Table 1: Counts and percentages for all /k/ observations per *Phonotactic complexity*, separately for Initial and Final positions across the ESTER and ETAPE corpora.

Corpus			Phonotactics	Example	Count	%
ESTER	Final	Cluster CCR		<i>ministre</i> [minist(ɤ)]	599	0.37%
ESTER	Final	Cluster CR		<i>cadre</i> [kad(ɤ)]	6,892	4.3%
ESTER	Final	Singleton R		<i>pour</i> [pu(ɤ)]	34,979	21.9%
ESTER	Initial	Cluster CCR		<i>scrutin</i> [sk(ɤ)ytɛ̃]	175	0.1%
ESTER	Initial	Cluster CR		<i>trente</i> [t(ɤ)ãt]	15,809	9.8%
ESTER	Initial	Singleton R		<i>rapide</i> [(ɤ)apid]	13,305	8.3%
ETAPE	Final	Cluster CCR		<i>administrent</i> [administ(ɤ)]	630	0.39%
ETAPE	Final	Cluster CR		<i>votre</i> [vɔt(ɤ)]	9,377	5.8%
ETAPE	Final	Singleton R		<i>faire</i> [fɛ(ɤ)]	40,919	25.6%
ETAPE	Initial	Cluster CCR		<i>strict</i> [st(ɤ)ikt]	194	0.12%
ETAPE	Initial	Cluster CR		<i>prime</i> [p(ɤ)im]	19,878	12.4%
ETAPE	Initial	Singleton R		<i>rue</i> [(ɤ)ly]	16,504	10.3%

Table 2: Counts and percentages of non-canonical pronunciation classification per *Phonotactic complexity*, separately for Initial and Final positions. Examples in French are provided

Position	Phonotactic	Example	Total Count	Non-Canonical Count	Percentage
Final	Cluster CCR	<i>ministre</i> [ministɤ]	1,229	406	33.0%
Final	Cluster CR	<i>cadre</i> [kadɤ]	16,269	4,896	30.1%
Final	Singleton R	<i>pour</i> [puɤ]	75,898	16,023	21.1%
Initial	Cluster CCR	<i>scrutin</i> [skɤytɛ̃]	369	62	16.8%
Initial	Cluster CR	<i>trente</i> [tɤãt]	35,687	4,371	12.2%
Initial	Singleton R	<i>rapport</i> [kɤapɔɤ]	29,809	4,050	13.6%

3.1 Overall Trend

Table 2 presents the counts and percentages of non-canonical pronunciations for /ɤ/ deletion, categorized by phonotactic complexity and word position (initial vs. final). The data reveal a clear positional difference, with higher deletion rates observed in final positions compared to initial positions. Final positions exhibit deletion percentages ranging from 21.1% to 33.3% while initial positions show lower rates, from 12.2% to 16.8%. Figure 1 supports this trend, illustrating the mean of the distribution of $\Delta_{Del\%}$ as - 11.7, a negative value indicating greater deletion in word-final positions. This positional effect further supported by a

logistic mixed-effects model (package lme4: (Bates *et al.*, 2015)), which accounts for position, corpus and speaker gender, with random intercepts for speaker and word. The model indicates that canonical productions are significantly more likely to occur in word-initial position than in word-final position ($p < .001$). Furthermore, there is a corpus effect, with the less formal ETAPE corpus showing higher likelihood of non-canonical productions ($p < .001$). Additionally, the gender effect reveals that males are more likely to produce canonical forms ($p < .001$).

3.2 Speaker Specific Behaviour

Figure 2 illustrates the correlation between initial and final deletion rates for each individual speaker. The results reveal a positive small-to-moderate correlation ($r = 0.35305$), as shown by the red dotted line. The observed value suggests a weaker relationship between deletion rates in these two positions than expected.

Additionally, figure 3 displays the non-canonical (blue) versus the canonical (white) production percentages across our 19 professional categories. The figure organizes non-canonical word-final productions (right-side of the graph) in decreasing order. Professions associated with less formal speech, such as performers (actors, singers, dancers) and athletes, exhibit higher levels of /ʁ/ deletion in word-final position than those associated with more formal speech, such as journalists and TV/Radio hosts. The pattern for initial non-canonical productions (left-side of the graph) was expected to mirror this decreasing trend; however, the results show irregular non-canonical percentages across professions, especially notable in categories such as arts (painters, sculptors, photographers) and members of unions. This suggests that speakers who exhibit less deletion of /ʁ/ in word-final positions do not necessarily exhibit less deletion of /ʁ/ in word-initial positions, contradicting our original prediction.




Figure 1: Distribution of deletion rates $\Delta Del\%$ observed in the corpus. The mean of the distribution is indicated by the red line, with the corresponding value displayed in white.




Figure 2: Final vs. Initial deletion rates $Del\%$ as a function of per speaker (dots).

4 Discussion

Overall, the results align with our predictions, showing lower rates of /ʁ/ deletion (i.e., non-canonical productions) in the stronger word-initial position compared to the weaker word-final position. However, deletion rates in word-initial position remain relatively high, exceeding 10%, which is consistent with previous findings reported by Wu and Adda-Decker (2021). Despite this general trend, speaker-specific behaviour reveals only a weak positive correlation between deletion rates in word-initial and word-final positions. This, along with the patterns observed across professional categories, suggests that errors in the automatic phone alignment system’s categorization of canonical and non-canonical productions may influence the results. To investigate this possibility, we manually checked a very small subset of the data, approximately 400 non-canonical tokens, including 300 word-initial and 100 word-final instances of /ʁ/. Our findings indicate a marked discrepancy in classification accuracy between the two positions. Word- final /ʁ/ tokens were found to be highly accurate, with over 98% of the non-canonical cases corresponding to actual deletions, as reported by previous studies (Wu *et al.*, 2019; Wu and Adda-Decker, 2021). In contrast, manual checks of word-initial /ʁ/ show remarkable results regarding the performance of the alignment, exhibiting a high percentage of miss-alignments. Among the non-canonical misaligned word- initial tokens, approximately 30% involved words beginning with /ʁ/ that followed a word ending in /ʁ/ (e.g., *pour* *reservoir* [pʁʁəvwaʁ], *avoir* *repris* [avwaʁvɛʁpʁi], *propre* *révision* [pʁɔʁvɛvizjɔ̃]), suggesting segmentation or transcription errors. Another 50% of these cases involved /ʁ/ segments with durations under 30ms – below the minimum phone

duration threshold used by the system – which likely led to the system eliminating rather than preserving the segment. A secondary alignment procedure that disallowed segment deletion reassigned these short segments a duration of exactly 30ms, confirming their reduced status. An additional 5% of the misclassifications could be attributed to non-native or foreign-accented speech. The remaining 10% of cases showed no clear reason for the deletion classification, highlighting possible limitations in the alignment system. The higher accuracy for word-final /ʁ/ likely results from the baseline dictionaries used to train the acoustic model. These dictionaries included phonologically motivated variants that captured the well-documented phenomenon of word-final /ʁ/ deletion, thereby enhancing recognition in that position. For instance, the word *livre* “book” was represented with three variants: the canonical [livʁ], a variant with an epenthetic schwa [livʁə] and a non-canonical reduced form [liv].




Figure 3: Non-canonical (blue) vs. canonical (white) production percentages by professional category (y- axis). Word-final deletion rates appear on the right side, word-initial on the left. Professional categories are ordered by decreasing non-canonical word-final percentages.

Limitations

One limitation of this study is its full reliance on automated processes without manual verification of the acoustic data. While manually checking nearly 160,000

tokens is impractical, incorporating a small-scale validation – such as checking 2% of the data – would be a feasible approach and is planned for future research. Another limitation is the non-inclusion of additional factors known to influence phonetic reduction, such as word frequency, phonetic context, speech rate, and more detailed speaker characteristics beyond identity (e.g., dialectal information, more precise social class identifiers). Incorporating these variables in future analyses will provide a more comprehensive understanding of the patterns observed in this study.

5 Conclusion

This study examined /ʁ/ deletion in word-initial and word-final positions in French using a 150-hour corpus of journalistic speech, with a particular focus on speaker-specific variation. Overall, the results agree with phonological predictions of phonetic reduction, confirming higher rates of deletion in word-final positions. However, speaker-level and profession-based patterns suggest potential misclassifications and overestimation of non-canonical productions by the ASR system in word-initial contexts. These findings highlight the importance of linguistic expertise in interpreting and refining the outputs of speech technology tools, particularly when dealing with complex phonetic variation in automated systems.

AI Use Statement

In this work, AI tools were used solely to rephrase and refine selected passages in order to enhance clarity and readability. All substantive content and interpretations remain our own.

References

- Bernard Laks. 1977. Contribution empirique à l'analyse socio-différentielle de la chute des /r/ dans les groupes consonantiques finals. *Langue Française*, 34:109–125.
- Cécile Fougeron. 2007. Word boundaries and contrast neutralization in the case of enchainement in French. *Papers in Laboratory Phonology IX: Change in Phonology*, 9:609–642.
- Cedric Gendrot, Barbara Kuehnert, and Didier Demolin. 2015. Aerodynamic, articulatory and acoustic realization of French. In *Proc. of the 18th International Congress of Phonetic Sciences (ICPhS2015)*.
- Dani Byrd. 1994. Relations of sex and dialect to reduction. *Speech Communication*, 15:39–54.
- Douglas Bates, Martin Machler, Ben Bolker, and Steve Walker. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48.
- Guillaume Gravier, Gilles Adda, Niklas Paulsson, Matthieu Carré, Aude Giraudel, and Olivier Galibert. 2012. The étape corpus for the evaluation of speech-based tv content processing in the French language. In *Proc. LREC 2012*, pages 114–118.
- Keith Johnson. 2004. Massive reduction in conversational American English. In

- Spontaneous speech: Data and analysis. Proceedings of the 1st session of the 10th international symposium*, pages 29–54.
- Martine Adda-Decker. 2006. De la reconnaissance automatique de la parole à l'analyse linguistique de corpus oraux. In *Actes des XXVIes Journées d'Etude sur la Parole JEP*, pages 389–400.
- Mathieu Avanzi. 2023. 'Vot' artic' est formidab': une étude multifactorielle de la chute des liquides post-obstruantes finales de mot en français. *Journal of French Language Studies*, pages 1–31.
- Michel Chafcouloff. 1980. Les caractéristiques acoustiques de /j,y,w,l,r/ en français. *Travaux de l'Institut de Phonétique d'Aix*, 7:7–56.
- Pierre Guiraud. 1965. *Le français populaire*. PUF, Paris.
- Sylvain Galliano, Edouard Geoffrois, Guillaume Gravier, Jean-Francois Bonastre, Djamel Mostefa, and Kalid Choukri. 2006. Corpus description of the ester evaluation campaign for the rich transcription of French broad- cast news. In *Proc. LREC 2006*, volume 6, pages 315–320.
- Tobias Scheer and Phillippe Segéral, 2008. *Positional factors in Lenitio a Fortition*, pages 131–172. De Gruyter Mouton, Berlin, New York.
- Yaru Wu and Martine Adda-Decker. 2021. Distribution and deletion of /ʁ/ in fluent speech. *Studii de lingvistica*, 11:39–53.
- Yaru Wu, Cedric Gendrot, Martine Adda-Decker, and Cécile Fougeron. 2019. Post-consonantal word-final /r/ realization in French: contributions of large corpora. In *Proc. of the International Congress of Phonetic Sciences ICPhS219*.
- Zoe Boughton. 2015. Social class, cluster simplification and following context: Sociolinguistic variation in word- final post-obstruent liquid deletion in French. *French Language Studies*, 25:257–286.

F5-TTS-RO: EXTENDING F5-TTS TO ROMANIAN TTS VIA LIGHTWEIGHT INPUT ADAPTATION

RADU GABRIEL CHIVEREANU, TIBERIU BOROȘ

Research Institute for Artificial Intelligence “Mihai Drăgănescu”

`rchivereanu@gmail.com, tibi@racai.ro`

Abstract

This work introduces a lightweight input-level adapter for the F5-TTS model that enables Romanian Language support. To preserve the existing capabilities of the model (voice cloning, English and Chinese support), we keep the original weights frozen, append a sub-network to the model and train it as an extension for the textual embedding matrix of the text encoder. For simplicity, we rely on the F5 specific ConvNeXt module to model the co-dependencies between these new character-level embeddings. The module serves as a “soft” letter-to-sound layer, converting Romanian text into a continuous representation that the F5-TTS model uses to produce naturally sounding Romanian utterances. We evaluate the model with a pool of 20 human listeners across three tasks: (a) audio similarity between reference and generated speech, (b) pronunciation and naturalness and (c) Romanian-English code-switching. Comparison with the mms-tts-ron model demonstrates that our approach preserves voice cloning capabilities and enables, to some extent, code-switching within the same utterance. We open-source our code and provide example audio samples at <https://github.com/racai-ro/Ro-F5TTS>.

Keywords: TTS, machine learning, flow-matching, Romanian, speech

1. Introduction

Text to speech (TTS) technology has advanced rapidly in recent years, using **a wide variety of architectures based on neural networks**. Beyond text-based conditioning, several approaches incorporate audio-based conditioning to capture and reproduce speaker-specific characteristics in synthesized speech, enabling zero-shot multi-speaker TTS (ZS-TTS) (Azzuni and El Saddik, 2025) and effective voice cloning with limited access to speaker samples (1-5 seconds of audio).

In this work, **we focus on extending the F5-TTS model** (Chen *et al.*, 2025), an open-source Non-Auto-regressive generative framework, to the Romanian language. F5-TTS builds upon E2-TTS (Eskimez *et al.*, 2024), introducing a ConvNeXt module for text encoding and a multimodal DiT module (Peebles and

Xie, 2023) as backbone. It operates with characters as tokens, padded with filler tokens to match the length of the mel spectrograms.

As a result, the model can be trained using only paired audio–text data, without requiring additional fine-grained annotations. Moreover, F5-TTS demonstrates strong capabilities in speaker identity cloning. To adapt F5-TTS for the Romanian language, we introduce a lightweight and non-invasive ConvNeXt based adapter that operates only at the input level, leaving the model’s internal layers frozen. The adapter maps Romanian characters into a continuous input space that the model can process, effectively performing a soft letter-to-sound conversion while leveraging and interpolating between the model’s previously learned letter-to-latent representations.

Keeping the core model frozen also **reduces the risk of overfitting to speakers** in the training set. An additional advantage of this approach is its natural support for code-switching, which is the ability of a model to speak two or more languages in the same utterance, an increasingly important capability given the widespread prevalence of bilingualism (Méndez Kline and Zellou, 2025).

Our contribution is the following: We propose a simple and lightweight adapter for F5-TTS that enables Romanian speech synthesis with voice-cloning capabilities. We conduct **subjective evaluations** of speaker similarity and pronunciation with a panel of 16 native speakers in three distinct tasks. We perform **objective evaluations** using the Word Error Rate (WER) on transcribed audio and measures of speaker similarity.

2. Related Work

State-of-the-art open-source zero-shot text-to-speech (ZS-TTS) systems, such as (Casanova *et al.*, 2024, Chen *et al.*, 2025), currently lack explicit support for Romanian. Existing work on ZS-TTS for Romanian typically relies on adapting large pretrained models and architectures to address data scarcity and related challenges. For example, (Răgman and Stan, 2024) extended the Fast-Pitch architecture (Łańcucki, 2021) to Romanian and enabled multi-speaker synthesis by conditioning on speaker embeddings extracted with TitaNet-L (Koluguri *et al.*, 2022).

A central difficulty in adapting pretrained models to new tasks is *catastrophic forgetting*, where finetuning degrades previously acquired knowledge (Wang *et al.*, 2024). In the context of speech, several strategies have been proposed, including model merging, experience replay, and adapter-based approaches (Hsiao *et al.*, 2025, Vander Eeckt and Van Hamme, 2023, Chen and Garner, 2024, Yang *et al.*, 2022). Replay buffers are particularly effective, but they depend on storing samples from prior training languages. In contrast, our work focuses on adaptation using only new language data, **without retaining or revisiting past examples**.

These adaptation strategies are also relevant in *code-switching* scenarios, where maintaining knowledge of multiple languages is crucial. Current methods typically leverage pretrained multilingual text encoders (Zhou *et al.*, 2020) or explicit multi-language training (Cao *et al.*, 2019). Our approach instead combines the input embedding matrix from previously learned languages with the output of a newly trained input adapter. However, since the model is not explicitly optimized to align embedding spaces, its code-switching ability remains limited.

3. Proposed Methodology

Our method involves using a ConvNext module as an input level adapter to align representations with the learned space of embeddings of the model.

3.1 Dataset

For our experiments we rely on the SWARA Speech Corpus (Stan *et al.*, 2017), which comprises over 21 hours of high-quality recordings from 17 speakers. This dataset is particularly suitable for demonstrating that our adapter retains speaker cloning capability, as its limited speaker diversity presents a heightened risk of overfitting.

For evaluation, we used examples from the Common Voice 22.0 dataset (Ardila *et al.*, 2020), the Romanian subset.

3.2 Model Architecture

To preserve the full capabilities of F5-TTS, we keep it frozen and train only a lightweight module that maps the Romanian text input into the input space of the model. Let $x = (c_1, c_2, \dots, c_T)$ denote the input Romanian text sequence of length T , where c_i is a character, and let $E \in \mathbb{R}^{|V| \times d_e}$ be the trainable embedding layer. Each character is first assigned to a continuous embedding vector $h_0 = E(x) = [E(c_1), E(c_2), \dots, E(c_T)] \in \mathbb{R}^{T \times d_e}$. To model dependencies between character embeddings, we apply a ConvNeXt architecture with **1D convolutions** along the temporal dimension (implementation from the F5-TTS repository), producing contextualized embeddings $h_{ctx} = \text{ConvNeXt}(h_0) \in \mathbb{R}^{T \times d_h}$.

Finally, the output is fed into the frozen F5-TTS model to generate speech: $\hat{y} = \text{TTS}(h_{ctx})$. In short, the entire pipeline is compactly represented as $\hat{y} = \text{TTS}(\text{ConvNeXt}(E(x)))$. Figure 1 displays a visual representation of our approach.




Figure 1: Overview of our method (right) compared to the original F5-TTS approach (left). We replace the model’s initial character level embedding matrix with a new learnable matrix with a ConvNeXt on top.

Code Switching

We also evaluated the model’s ability to generate code-switching examples between Romanian and English. Let $x = (c_1, c_2, \dots, c_T)$ be the input text sequence containing Romanian and English characters, and let m_R and m_E be masks indicating Romanian and English characters, respectively. We use the special token $\sim$ to signal language switching. The embeddings are computed as follows:

$$h_R = \text{ConvNeXt}(E_R(x) \odot m_r), h_E = E_{TTS}(x \odot m_E)$$

where E_R is the trainable embedding module for Romanian characters and E_{TTS} is the frozen embedding layer of the TTS model for English characters. The final sequence is obtained by merging the embeddings:

$$h_{cs} = h_R + h_E, \hat{y}_{cs} = \text{TTS}(h_{cs}).$$

This approach is limited, as the language module and the model itself were not explicitly optimized to handle embeddings that represent English and Romanian at the same time. We evaluated the naturalness of the approach in our experiments and found promising results.

3.3 Training

The training objective involved applying the original generative flow-matching objective of F5-TTS to the SWARA Corpus. Hyperparameters are shown in Table 1.

Table 1: Training hyperparameters

Hyperparameter	Value
Learning rate	1×10^{-4}
Steps	40500
Batch size (audio frames)	16384
Warmup updates	50
Max samples	128

The training was executed with **1 NVidia A100 GPU**, running for 12 hours.

4. Experiments

We conducted a series of experiments to evaluate the performance of our solution. The tests are divided between an subjective evaluation campaign (Section 4.1) using volunteers and an objective evaluation metric (Section 4.2), namely the automatically computed word error rate (WER) and speaker similarity.

4.1 Subjective Evaluation

For subjective evaluation, leveraging a pool of 16 native speakers we defined three tasks: (a) speaker similarity; (b) naturalness and (c) code-switching.

For comparison, we used MMS-TTS-RON, the Romanian checkpoint from the Massively Multilingual Speech (MMS-TTS) project (Pratap *et al.*, 2024), along with the base F5-TTS model. MMS-TTS-RON was chosen because it explicitly supports Romanian, being a VITS (Kim *et al.*, 2021) model trained on a Romanian corpus.

In the speaker similarity tests, we added samples from original speakers. Thus, our low anchor (a system designed for lower-bound normalization of the results) was the unmodified F5-TTS and for speaker similarity, the original voice was used for high-bound normalization.

The testing procedure for each task consisted in selecting a list of utterances that were synthesized or extracted from the dataset - for cases where natural speech samples were used. In each evaluation step, the listener would be presented with 3 (or 4) parallel samples, generated by each system (respectively natural). The listener had to evaluate the parallel samples with a score from 0 (Bad) to 100 (Excellent), not knowing the model used for synthesis.

In what follows, we detail the speaker similarity task in Subsection 4.1.1, the naturalness evaluation task in Subsection 4.1.2 and the code-switching task in Subsection 4.1.3.

Speaker Similarity

We evaluated 10 trials. Each trial consisted of the same sentence synthesized by all tested TTS models, and the original speaker’s recording was included as one of the stimuli. Participants were also provided with a reference sample of the original speaker’s voice and rated each stimulus based on how similar the voice and intonation were to the reference. This setup allowed us to observe whether participants could correctly identify the original recording among the synthesized samples.

The prompt given for this task was: *“Please rate each audio sample according to how similar the speaker sounds to the reference speaker.”*

We observed that users often found it challenging to correctly identify the reference speaker among the stimuli. The results of each trial are shown in Figure 2.




Figure 2: Results on speaker similarity task.

We observed that some of the listeners had difficulties identifying the real speaker, with overall high scores for our model, which, even in poor audio settings, still suggests that the model efficiently encompasses speaker identity in generated output, in terms of speaking style and voicetimbre.

Pronunciation and Naturalness

We tested 9 longer sentences to assess how accurately the speech was pronounced and how natural it sounded. Table 2 displays the sentences synthesized for the task.

Table 2: Sentences used for evaluating pronunciation.

Am avut astăzi o întâlnire foarte lungă cu clientul, care a durat aproape două ore, și am discutat multe detalii importante despre proiect.
Poți să-mi trimiți și mie, te rog, documentul respectiv, ca să pot verifica toate informațiile înainte de întâlnirea de mâine?
Săptămâna aceasta vreau să mă relaxez acasă, să citesc ceva și poate să vizionez un film sau să mă uit la un serial interesant.
Este extrem de enervant când oamenii nu răspund la mesajele pe care le trimiți, mai ales dacă ai nevoie urgentă de informații.
Îți voi trimite un raport complet mâine dimineață, imediat după ce reușesc să vorbesc cu toți colegii implicați în proiect.
Nu sunt sigur dacă are sens, dar hai să încercăm oricum, chiar dacă există riscul să nu funcționeze perfect.
A spus că termenul limită este vineri, dar sincer, nu cred că va termina tot ce și-a propus până atunci, așa că trebuie să verificăm împreună progresul.
Mă duc să iau o cafea de la cafenea și revin imediat, ca să putem continua discuția fără întreruperi.
Am avut o întâlnire foarte devreme dimineața și nu am avut timp nici măcar să măănânc, așa că am plecat de acasă cu stomacul gol.

The results are shown in Figure 3.




Figure 3: Results on pronunciation and naturalness task.

In this task, the MMS-TTS-RON model achieved higher overall scores for most sentences. Part of this result is due to listeners noticing an English accent in certain examples, which affected their evaluation of our model.

The prompt shown to the listeners was: *“Please rate each audio sample based on the pronunciation of the words and how natural it sounds.”*

Code Switching

We assessed the model’s ability to handle instances where both Romanian and English are used within the same utterance, using 9 examples. Table 3 displays the code-switching examples included in our evaluation.

Table 3: Code-switching sentences used for evaluation.

Am avut un call with the client today și a durat două ore.
Poți să-mi trimiți și mie un link to that te rog?
Săptămâna asta vreau să mă relaxez and maybe watch a movie or something.
E super annoying când nu răspunde lumea la mail.
Îți fac un update pe mâine dimineață după ce vorbesc cu ei.
Nu știu dacă are sens dar let's try anyway.
A zis că deadline-ul pe vineri dar actually nu cred că termină.
Mă duc să iau un coffee to go și vin imediat brother.
Am avut un meeting super early și nici nu am apucat să mănânc.

Results of the evaluation are shown in Figure 4.




Figure 4: Results on Romanian-English code-switching task.

For code-switching, listeners were instructed as follows: “*Please rate each audio sample based on how natural the transition between Romanian and English is.*” While the results indicate potential, the tests were limited in size. Further analysis is required, both from a qualitative perspective, examining perceptual naturalness and transition smoothness, and from a mechanistic perspective, to understand what occurs internally during language switching.

The system could benefit from automated language switch detection rather than manual annotation, an integration that we aim to investigate in future work.

4.2 Objective Evaluation

For our objective evaluation, we follow the approach from (Chen *et al.*, 2025) and (Răgman and Stan, 2024) and use the transcription and speaker recognition models. In our case, we applied them to a set of **1000 generated samples** from Common Voice. We also add for comparison results obtained by fully-finetuning the F5-TTS model on our training data.

Word Error Rate

The model we selected is gigant/whisper-medium-romanian, a finetuned variant of Whisper (Radford et al., 2023), available on the Huggingface platform at (Gigant, 2023). The transcriptions are then used to compute Word Error Rate (WER), Match Error Rate (MER), Word Information Lost (WIL) and Word Information Preserved (WIP). The results are displayed in Table 4.

Table 4: Evaluation metrics for Romanian TTS systems.

Metric (%)	MMS-TTS-RON	RO-F5TTS	F5-TTS-FULL-FT
WER (↓)	5.77	5.27	3.62
MER (↓)	5.63	5.15	3.52
WIL (↓)	9.52	8.23	5.39
WIP (↑)	90.48	91.77	96.92

We observe that the fully fine-tuned variant achieved better results, suggesting that the adapter is not fully able to capture all task-specific nuances, particularly the phonetic differences across languages, which likely contributed to lower performance on the pronunciation task.

Speaker Similarity

To assess speaker similarity, we employed TitaNet-L to compute speaker embeddings based on audio, following the approach of (Răgman and Stan, 2024). Cosine similarity was calculated between the reference and generated audio

samples. The results presented in Table 5 show a consistently high similarity between the reference speaker and the generated voice.

Table 5: Cosine Similarity Statistics between reference and generated audio for two models.

Statistic	RO-F5TTS	F5-TTS-FULL-FT
Mean	0.9013	0.7946
Standard Deviation	0.0319	0.0578
Minimum	0.7261	0.5446
Maximum	0.9656	0.9186
Median	0.9061	0.8033

We observe a substantial improvement over the F5-TTS fine-tuned model. The Speaker Similarity and WER results highlight the trade-off between the current adapter-based approach and full fine-tuning, reflecting a balance between maintaining voice cloning capabilities and accurate language-specific speech synthesis.

5. Conclusions & Future Work

We present a **small input-level adapter that enables Romanian text synthesis with F5-TTS** while keeping the pretrained model frozen. The adapter replaces the character embedding layer and adds a lightweight 1D ConvNeXt that **produces continuous “soft” letter-to-sound** consumable by the frozen backbone.

In subjective tests with 20 native speakers, we observed high speaker similarity. Pronunciation and naturalness were competitive, though occasionally influenced by an English accent. At this stage, it remains unclear whether this limitation arises from the adapter itself, from the underlying base model’s lack of direct training on Romanian, or from the interplay between the two, as the adapter may not be strong enough to represent patterns the model has never encountered.

Objectively, our system achieved a lower WER than MMS-TTS-RO (5.27% vs. 5.77%) and stronger speaker similarity compared to the fully finetuned F5-TTS.

Future work includes:

1. A detailed analysis of the representations in order to understand how the adapter encodes and projects the textual input with respect to the previously learned embedding space.
2. Exploring automated language switch detection and explicit bilingual training to improve code-switching performance.
3. Extending the framework to additional languages and models.

AI Use Statement

Generative AI was used only for minor stylistic edits. Hypotheses, literature review, methods, experiments and conclusions were entirely produced by the authors.

References

- Adrian Łańcucki. 2021. FastPitch: Parallel Text-to-Speech with Pitch Prediction. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6588–6592.
- Adriana Stan, Florina Dinescu, Cristina Țiple, Șerban Meza, Bogdan Orza, Magdalena Chirilă, and Mircea Giurgiu. 2017. The SWARA speech corpus: A large parallel Romanian read speech dataset. In *2017 International Conference on Speech Technology and Human-Computer Dialogue (SpeD)*, pages 1–6.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine Mcleavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 28492–28518. PMLR, 23–29 Jul.
- Chi-Yuan Hsiao, Ke-Han Lu, Kai-Wei Chang, Chih-Kai Yang, Wei-Chih Chen, and Hung-yi Lee. 2025. Analyzing Mitigation Strategies for Catastrophic Forgetting in End-to-End Training of Spoken Language Models. In *Interspeech 2025*, pages 3234–3238.
- Edresson Casanova, Julian Weber, Christopher D Shulby, Arnaldo Candido Junior, Eren Gölge, and Moacir A Ponti. 2022. YourTTS: Towards zero-shot multi-speaker TTS and zero-shot voice conversion for everyone. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 2709–2720. PMLR, 17–23 Jul.
- Edresson Casanova, Kelly Davis, Eren Gölge, Gökem Gökner, Iulian Gulea, Logan Hart, Aya Aljafari, Joshua Meyer, Reuben Morais, Samuel Olayemi, and Julian Weber. 2024. XTTS: a Massively Multilingual Zero-Shot Text-to-Speech Model. In *Interspeech 2024*, pages 4978–4982.
- Gigant. 2023. Whisper Medium Romanian. Hugging Face model repository. Fine-tuned version of openai/whisper-medium on Common Voice 11.0 & Romanian speech synthesis corpus.
- Haolin Chen and Philip N. Garner. 2024. Bayesian parameter-efficient fine-tuning for overcoming catastrophic forgetting. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32:4253–4262.
- Hussam Azzuni and Abdulmotaleb El Saddik. 2025. Voice cloning: Comprehensive survey. *arXiv preprint arXiv:2505.00579*.
- Jaehyeon Kim, Jungil Kong, and Juhee Son. 2021. Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine*

- Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5530–5540. PMLR, 18–24 Jul.
- Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. 2024. A comprehensive survey of continual learning: Theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8):5362–5383.
- Mu Yang, Shaojin Ding, Tianlong Chen, Tong Wang, and Zhangyang Wang. 2022. Towards lifelong learning of multilingual text-to-speech synthesis. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8022–8026. IEEE.
- Nithin Rao Koluguri, Taejin Park, and Boris Ginsburg. 2022. Titanet: Neural model for speaker representation with 1d depth-wise separable convolutions and global context. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8102–8106.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis Tyers, and Gregor Weber. 2020. Common voice: A massively-multilingual speech corpus. In Nicoletta Calzolari, Frédéric B  chet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, H  l  ne Mazo, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4218–4222, Marseille, France, May. European Language Resources Association.
- Sefik Emre Eskimez, Xiaofei Wang, Manthan Thakker, Canrun Li, Chung-Hsien Tsai, Zhen Xiao, Hemin Yang, Zirun Zhu, Min Tang, Xu Tan, Yanqing Liu, Sheng Zhao, and Naoyuki Kanda. 2024. E2 TTS: Embarrassingly Easy Fully Non-Autoregressive Zero-Shot TTS. In *2024 IEEE Spoken Language Technology Workshop (SLT)*, pages 682–689.
- Steven Vander Eeckt and Hugo Van Hamme. 2023. Using adapters to overcome catastrophic forgetting in end-to-end automatic speech recognition. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Teodora R  gman and Adriana Stan. 2024. Efficient Training Strategies for Natural Sounding Speech Synthesis and Speaker Adaptation Based on FastPitch. In *2024 IEEE 20th International Conference on Intelligent Computer Communication and Processing (ICCP)*, pages 1–6.
- Tyler M  ndez Kline and Georgia Zellou. 2025. The perception of code-switched vs. monolingual sentences in TTS voices. *Frontiers in Computer Science*, Volume 7 - 2025.
- Vin  el Pratap, Andros Tjandra, Bowen Shi, Paden Tomasello, Arun Babu, Sayani Kundu, Ali Elkahky, Zhaoheng Ni, Apoorv Vyas, Maryam Fazel-Zarandi, Alexei Baevski, Yossi Adi, Xiaohui Zhang, Wei-Ning Hsu, Alexis Conneau, and Michael Auli. 2024. Scaling speech technology to 1,000+ languages. *Journal of Machine Learning Research*, 25(97):1–52.
- William Peebles and Saining Xie. 2023. Scalable diffusion models with transformers. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4172–4182.

- Xuehao Zhou, Xiaohai Tian, Grandee Lee, Rohan Kumar Das, and Haizhou Li. 2020. End-to-end code-switching tts with cross-lingual language model. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7614–7618. IEEE.
- Yuewen Cao, Xixin Wu, Songxiang Liu, Jianwei Yu, Xu Li, Zhiyong Wu, Xunying Liu, and Helen Meng. 2019. End-to-end code-switched tts with mix of monolingual recordings. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6935–6939. IEEE.
- Yushen Chen, Zhikang Niu, Ziyang Ma, Keqi Deng, Chunhui Wang, JianZhao JianZhao, Kai Yu, and Xie Chen. 2025. F5-TTS: A fairytaler that fakes fluent and faithful speech with flow matching. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar, editors, *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6255–6271, Vienna, Austria, July. Association for Computational Linguistics.

CHAPTER 2

LANGUAGE VARIETIES

A REVIEW OF ADVANCES IN COMPUTATIONAL ANALYSIS OF KAZAKH POLITICAL DISCOURSE

LAURA ORYNBAY, GULMIRA BEKMANOVA,
ALTANBEK ZULKHAZHAY, AYAULYMSAIRANBEKOVA

¹*L.N. Gumilyov Eurasian National University, K. Satpayev Street,
Astana, Kazakhstan
laura.aktobe.kz@gmail.com, gulmira-r@yandex.kz,
altinbekpin@gmail.com, sairanbekova98@gmail.com*

Abstract

This review article examines advances in computer analysis of Kazakh political discourse in the context of broader research in natural language processing, knowledge engineering, and digital humanities. It addresses three main issues: the digitisation of the public sphere, the morphological complexity of the Kazakh language, and the ongoing transition to the Latin alphabet. Rather than introducing new empirical data, the paper integrates and systematises existing studies into a reproducible methodological framework. The discussion highlighted achievements in several areas, including the establishment of an ethical and legal framework for obtaining text, the use of ontological modelling to structure campaign communication and enable semantic justification, the development of preprocessing strategies that take into account spelling variations and morphological ambiguity, as well as the introduction of subsequent methods such as sentiment analysis, stance identification, semantic similarity, and framing analysis. An integrated methodological approach is proposed that combines ethical data collection, ontology development, language preprocessing, and computational interpretation into a single framework that emphasises reproducibility and comparative evaluation. Special attention is paid to the need to create annotated corpora, specialised dictionaries and reproducible workflows for the development of this developing field. By summarising methods, challenges, and opportunities, the article demonstrates how computer-based approaches can enrich the study of Kazakh political discourse and contribute to broader debates in the fields of political communication, computational linguistics, and digital humanities.

Keywords: Kazakh political discourse, computational linguistics, natural language processing, morphological analysis, sentiment analysis, stance detection, political genre classification, ontology-based modelling, low-resource languages

1 Introduction

Computer analysis of political discourse in Kazakhstan is associated with three problems: (i) a **rapidly digitising public sphere** where messages are exchanged through speeches, party documents, and social media; (ii) a **morphologically rich language with limited resources**, whose agglutinative structure complicates tokenisation, lemmatisation, and feature extraction; and (iii) an **ongoing transition of writing** in Latin, which affects the normalisation of spelling and compatibility of tools. These factors necessitate an integrated approach that combines **ethical data management** to obtain confidential political content, **knowledge engineering** (to determine the structure of the subject area), **language-specific preprocessing** (Latin morphology) and **downstream analytics** (sentiment, stance, framing, genre classification).

This paper is intended as a review and synthesis rather than a report on new experimental results. It brings together recent empirical findings from NLP projects in the Kazakh language and related computational studies on political communication, which are summarised into a coherent methodological pipeline. The goal is to provide researchers with a consolidated reference that integrates ethical, linguistic, and computational perspectives, as well as to identify reproducible methods for future empirical research in this evolving domain.

In this review, the authors summarise the latest, verifiable achievements: the legal and ethical foundations of the use of open political texts in scientific research (Bekmanova *et al.*, 2024), the OWL (Mukanova *et al.*, 2024) election advertising ontology, cognitive-computational approach to NLP for discourse interpretation (Orynbay *et al.*, 2025), Latin-morphological script formalisation for the Kazakh language (Zhetkenbay *et al.*, 2025), sentiment creation resources (Yeshpanov and Varol, 2024), scalable semantic similarity methods (Ayazbayev *et al.*, 2023), genre conclusions from the inaugural speeches of (Abdiman, 2025), and a diachronic analysis of slogans and programs of Nassimov (2022). The authors propose a concept that will turn these fragments into a holistic methodology and outline specific evaluation protocols to make the results reproducible.

2 Ethical Frameworks for Political Text Research

Compliance with ethical and legal norms when collecting data on political discourse (Bekmanova *et al.*, 2024) emphasises the important role of ethical and legal guarantees in the processing of political content from open sources, especially in the unstable sphere of social networks. They prioritise transparency, legal compliance, and careful documentation of the origin of data – principles aligned with broader guidelines for the responsible use of artificial intelligence in high-stakes environments (Floridi and Cowls, 2022).

In Kazakhstan, publicly available political statements, such as campaign speeches or official announcements, are generally legitimate for analysis.

However, content created by private individuals requires special care to prevent privacy violations and unintended harm. In the corpus of political discourse, messages from public figures (low sensitivity) are often mixed with civic content (higher sensitivity), which requires differentiated treatment strategies.

Following Bekmanova *et al.* (2024), this research recommend:

- **scope and provenance:** limit to publicly available, official content and/or publications by public figures or institutions; specify source URLs, timestamps and compliance with platform requirements;
- **data minimisation:** collect only the fields necessary for analysis (text, actor's role category, approximate time); avoid storing private individuals' IDs when it is not necessary;
- **legality and platform rules:** document the legal framework and conditions of the data collection platform/API; respect robots and speed limits; prefer datasets from official archives;
- **consent and role separation:** treat government officials/organisations differently from individuals; summarise signals rather than quoting verbatim to get information from citizens;
- **bias and harm safeguards:** audit of polarised indicators (for example, by party/region/topic); registration of false positives that may mischaracterise groups; prevention of profiling at the individual level;
- **reproducibility with privacy:** release derived functions (embeddings, thematic distributions) and share the code/score, not the source text, when there is no certainty about licensing/confidentiality.

This structured, principles-based approach ensures that political discourse research in Kazakhstan remains both methodologically sound and socially responsible, ensuring reproducibility of results without compromising ethical integrity.

3 Ontology Engineering for Campaign Discourse

Mukanova *et al.* (2024) present an ontological model for the field of election advertising, implemented in OWL using the Protégé platform. This ontology formalises the basic concepts, relationships, and properties related to Kazakhstani election campaigns, allowing for automatic analysis of election campaign texts. Their work is consistent with previous work in the development of the Schneider Political Discourse Ontology (2021) and illustrates the value of domain-specific semantic resources for analysing Kazakhstan's political rhetoric.

Ontologies of this kind serve as a semantic basis for normalising heterogeneous textual data, integrating knowledge representation with sentiment analysis and argumentation mining, which allows for a more complete interpretation of campaign messages.

Based on (Mukanova *et al.*, 2024), an extended structure is proposed in which the ontology supports both annotation and logical inference:

- **Core classes/relations:**

- Actor (party, candidate, institution),
- Message (speech, manifesto section, slogan),
- PolicyArea (economy, healthcare, etc.),
- Device (rhetorical device),
- Frame (e.g., modernisation, stability),
- TargetAudience;

With relations such as

- advocates (Actor, PolicyArea),
- frames (Message, Frame),
- targets (Message, Audience),
- employs (Message, Device).

- **Competency questions:**

- CQ1 – Which policy areas does Party X emphasise this cycle?
- CQ2 – How do frames differ across parties on the same issue?
- CQ3 – Which rhetorical devices correlate with positive sentiment?

- **Annotation interface:** a light schema to tag policy areas, frames, devices; export to RDF/OWL for reasoning and SPARQL querying.

- **Alignment and enrichment:**

- link PolicyArea to a small controlled vocabulary;
- map synonyms via distributional similarity (Ayazbayev *et al.*, 2023);
- keep an Alias property for historical party names/coalitions.

- **Reasoning use-cases:** infer under-specified relations (e.g., if a manifesto section contains a “tax reduction” pattern, infer frames(..., modernisation) and advocates(..., economy)).

4 Language-specific Preprocessing and Morphology

Zhetkenbay *et al.* (2025) solve the main problem in processing the natural Kazakh language: achieving accurate morphological analysis in the context of the country’s transition to a new Latin alphabet. They propose a hybrid approach that combines formalised rule-based grammar descriptions with a Bayesian back propagation neural network with regulation (BR-BPNN). Trained on a database of 1,000 annotated nouns, their system provides more than 91.5% accuracy in predicting suffixes, while explicitly modelling vowel harmony and consonant assimilation.

Such morphological precision is crucial for analysing political discourse, where accurate lemmatisation, dependency analysis, and keyword extraction depend on reliable preprocessing. Without this, the agglutinative nature of the Kazakh language and its vowel harmony rules can lead to the spread of errors throughout the NLP pipeline.

Based on the work of Zhetkenbay *et al.* (2025), we propose an improved morphological processing workflow adapted for the analysis of political texts:

- **orthographic normalisation** to unify legacy Cyrillic, mixed Latin variants, and punctuation styles;
- **morphological pipeline:**
 - normalisation;
 - hybrid suffix prediction using the BR-BPNN model;
 - lemmatisation conditioned on predicted suffix class;
 - heuristic checks for ambiguous allomorphs.
- **Ambiguity resolution:** use n-best suffix candidates with lattice decoding and re-rank by contextual features (neighbour tokens, policy lexicons).
- **Quality controls:** create a 1-2k token gold set of political texts (balanced by genre: slogans, manifesto, speech) to evaluate lemma accuracy and suffix accuracy; track error types (harmony error, boundary error, ambiguous proper names)

This extended structure ensures that the preliminary morphological processing for Kazakh political discourse is both linguistically accurate and computationally reliable, which ensures higher accuracy in subsequent NLP tasks.

5 Sentiment, Stance and Similarity

Measurements of mood, position, and semantic similarity are an important component of computer analysis of political discourse, especially when studying campaign rhetoric. These methods allow us to determine the emotional tone, the position of the subject in relation to a political issue, and the degree of thematic or lexical overlap between different messages. In the context of Kazakhstan’s political discourse, where short, slogan-like messages are juxtaposed with lengthy manifestos, careful adaptation of these methods is required to ensure reliable and context-sensitive interpretation.

To determine attitudes and positions, a practical starting point is to use the **KazSAnDRA** dataset (Yeshpanov and Varol, 2024), which contains annotated examples suitable for basic training. Models should be evaluated separately for political and non-political subsets to assess the change in the subject area. Threshold calibration is important when working with texts of different lengths:

short political slogans often require a higher sensitivity to decision-making, while for long political paragraphs, more subtle thresholds of polarity are more important. To avoid overestimation of indicators due to class imbalance, it is recommended to use a class-balanced sample and a macro-F1 scoring.

A more advanced approach involves frame-aware sentiment analysis, in which polarity detection is conditioned on the PolicyArea and Frame attributes from the ontology. This condition helps to avoid false assignments of polarity, such as interpreting a neutral policy description as inherently positive or negative. Aggregating sentiment at both the message and actor levels allows you to identify long-term rhetorical trends within campaigns.

Semantic similarity methods expand the possibilities of analysis beyond sentiment and stance. The Spark-based pipeline (Ayazbayev *et al.*, 2023) can be used to solve three complementary tasks:

1. **Query expansion** – for example, expanding the term “zhumys oryndary” (workplaces) to related expressions such as “enbek” (labour), “zhumyspen qamtu” (employment), and “ondiris” (production), thereby improving recall in document retrieval.
2. **Clustering** – grouping thematically similar paragraphs across different political actors, facilitating comparative thematic analysis.
3. **Near-duplicate detection** – identifying cases where campaign slogans or policy statements are reused across election cycles or between parties, which may indicate rhetorical convergence or message borrowing.

By combining the definition of sentiments and positions with ontology-based framing and large-scale semantic similarity analysis, it becomes possible to build a multidimensional representation of political discourse. This comprehensive approach not only reflects the emotional aspects and attitudes towards the election campaign, but also reflects the thematic relationships and rhetorical coincidences that shape the competitive landscape of Kazakh politics.

6 Genre and Framing: From Qualitative Insight to Features

In the analysis of political discourse, genre and framing act as two interdependent semantic layers. Genre defines the communicative format – be it a speech, slogan, or manifesto – while framing defines the interpretative prism through which key issues are presented. The identification and practical implementation of these elements are important both for understanding the rhetorical strategies of political figures and for the evolution of political narratives over time.

One of the important genres in the Kazakh political context is the **inaugural address**. As Abdiman (2025) notes, inaugurations serve legitimising, ceremonial, and programmatic purposes. They often use distinctive linguistic features such as deictic pronouns (“biz” – we), modal verbs indicating commitment, parallel syntactic structures, and calls for unity. These repetitive patterns can be encoded as

explicit signs for automatic **genre classification** and **identification of rhetorical techniques**.

Another popular genre is **slogans and election programs**. Nassimov (2022) documents recurring themes – stability, modernisation, social security, and unity – between 1999 and 2021. These frames can be implemented by creating a target dictionary of frames or by training poorly controlled frame classifiers that combine initial terms with similarity expansion techniques. Computationally tracking the **frame share** in each political party and identifying **frame shifts** during election cycles can provide valuable information about changing political priorities.

A particularly revealing aspect arises when comparing genres directly. Cross-genre analysis makes it possible to assess the correspondence between **program frames** and **speech frames**, revealing whether campaign manifestos and public appeals remain consistent in their thematic focus. This can be quantified by applying indicators such as the Jensen–Shannon divergence to the distribution of detected frames, thereby revealing cases of strong thematic continuity or a noticeable discrepancy between election promises and public rhetoric.

In general, the systematic integration of genre identification and frame analysis allows for a more complete understanding of political discourse. Moving from qualitative analysis to computational characteristics, it becomes possible to conduct a large-scale longitudinal analysis of rhetorical strategies and their evolution within the framework of the political landscape of Kazakhstan.

7 Methods and Resources for Kazakh Political Discourse Analysis

7.1 Integrated Pipeline

An integrated approach to the analysis of Kazakhstani political discourse combines computational linguistics, the development of ontologies and genre research of discourse. Based on previous research and existing tools, the proposed framework combines disparate methodologies into a single pipeline (Figure 1). This structure allows for both qualitative and quantitative analysis, while respecting ethical and legal standards for data collection and processing. The structure consists of five consecutive layers (Figure 1), each of which corresponds to a specific methodological stage:

- Data Acquisition layer – collects speeches, manifestos, slogans, debates and messages on social networks using ethical filtering protocols (Bekmanova *et al.*, 2024).
- Preprocessing layer – performs spelling normalisation, tokenisation, lemmatisation and noise removal based on morphological rules for the Kazakh language with Latin alphabet (Zhetkenbay *et al.*, 2025).
- Knowledge Representation layer – integrates ontologies (Mukanova *et al.*, 2024), connects entities and enriches vocabulary using semantic similarity methods (Ayazbayev *et al.*, 2023).

- Analysis layer – implements sentiment analysis (Yeshpanov and Varol, 2024), genre classification (Abdiman, 2025), frame definition (Nassimov, 2022) and cognitive thinking (Orynbay *et al.*, 2025).
- Output and Visualisation layer – creates dashboards, network graphs, long-term trend charts, and automated text reports for research and policy purposes.




Figure 1: Sequential methodology for Kazakh political discourse analysis.

7.2 Stage Details

Table 1 provides a detailed description of each stage, indicating specific methods and their sources.

Table 1: Description of pipeline stages for Kazakh political discourse analysis.

Stage	Description	Key Methods & Tools	References
Data Acquisition	Collect open political texts with legal–ethical screening	Web scraping/APIs, archival retrieval, curation	Bekmanova et al., 2024
Preprocessing	Normalise (Cyrillic→Latin), tokenise, lemmatise, denoise	Script conversion, custom tokenisers, morphological rules	Zhetkenbay et al., 2025
Knowledge Representation	Ontology integration; entity linking; semantic enrichment	OWL/Protégé; embeddings; Spark pipelines	Mukanova et al., 2024; Ayazbayev et al., 2023
Analysis	Sentiment/stance, genre, frames, cognitive reasoning	ML/DL classifiers; topic/argument mining	Yeshpanov and Varol, 2024; Abdiman, 2025; Nassimov, 2022; Orynbay et al., 2025
Output & Visualisation	Present insights and longitudinal trends	Dashboards, networks, time series	Yeshpanov and Varol, 2024; Abdiman, 2025; Nassimov, 2022; Orynbay et al., 2025

7.3 Comparative Summary of Reviewed Studies

Table 2 compares core works covered in the review.

Table 2: Comparative summary of reviewed studies.

Author(s) & Year	Focus	Methodology / Tools	Dataset	Key Findings	Relevance
Bekmanova et al., 2024	Ethics of analyzing Kazakh political discourse	Legal–ethical analysis	Public speeches, campaigns, social media	Transparency and legal compliance	Sourcing/processing guidelines
Mukanova et al., 2024	Ontology for pre-election advertising	OWL/ Protégé	Campaign texts	Structured domain model	Semantic search/ reasoning
Orynbay et al., 2025	Cognitive computing in NLP	Taxonomy of techniques	Mixed corpora	Multimodal/ context-aware interpretation	Contextual inference
Zhetkenbay et al., 2025	Latin-script morphology	Hybrid rules + BR-BPNN	1,000 nouns	>91% suffix prediction	Robust pre-processing

Author(s) & Year	Focus	Methodology / Tools	Dataset	Key Findings	Relevance
Yeshpanov and Varol, 2024	KazSAnDRA sentiment dataset	Annotation; ML baselines	Thousands of sentences	Strong baselines	Enables sentiment / stance
Ayazbayev <i>et al.</i> , 2023	Semantic similarity (Spark)	Embeddings; cosine similarity	Large Kazakh corpus	Semantically close word sets	Query expansion; clustering
Abdiman, 2025	Inaugural addresses genre	Comparative rhetorical/ functional analysis	KZ/US inaugurals	Legitimising strategies; unity appeals	Genre features
Nassimov, 2022	Programs & slogans (1999-2021)	Thematic analysis	Party programs & slogans	Frames: stability, modernisation, unity, welfare	Frame features

Despite the fact that each of the considered resources and methods has been independently tested in previous publications, a quantitative assessment of their joint work within the framework of an integrated approach has not yet been carried out. Evaluating the end-to-end process of working with a representative body of political texts is a clear direction for future research. Such a study would make it possible to measure the scope of the corpus, the accuracy of preprocessing, and the consistency of ontologies at different stages, thereby demonstrating the practical effectiveness of the proposed structure.

8 Limitations and Risk Management

Domain change: models trained on common sentiment datasets can adjust to non-political patterns; mitigation involves fine-tuning on domain specific political subsets.

Ambiguity and irony: slogans and tweets compress meaning; integrate rhetorical device detectors and contextual windows.

Bias: Frame patterns and tonality may reflect the bias of the annotators; diversify the annotators and communicate agreement between the annotators.

Script variability: Expect Cyrillic and Latin letters to mix frequently; use reliable mapping tables and unit tests for normalisation.

9 Conclusion

This study highlights that computer analysis of Kazakhstani political discourse requires careful integration of ethical standards, language-specific processing tools, and structured knowledge representation. Advances in this field have provided new resources such as ontological models, databases on feelings and positions, morphological analysers, and frameworks for analysing rhetoric and

frameworks. Together, these developments form the basis for a methodology that can solve problems unique to the Kazakh language, including its agglutinative morphology, writing variability, and rich rhetorical diversity. At the same time, significant obstacles remain: short texts often contain ambiguity, changing the subject area affects the reliability of the model, annotations can introduce distortions, and the inconsistency of scenarios complicates processing. Overcoming these challenges will depend on transparent assessment protocols, balanced corpora, and practices that protect data confidentiality and integrity. Overall, current research demonstrates a shift from isolated innovations to a consolidated methodological framework. With continuous improvement, these tools and methods can provide meaningful insights into political communication, contributing to global advances in resource-limited natural language processing.

As a review document, this study is more conceptual and integrative than experimental. In future work, this synthesis will be expanded by creating a large-scale corpus and empirical comparative analysis of the entire project.

AI Use Statement

Generative artificial intelligence tools were used to a limited extent in the preparation of this manuscript. In particular, artificial intelligence-based language models were used for proofreading, improving grammar, and enhancing the stylistic quality and clarity of the text. The main ideas, summary of the analysed literature, methodological basis, and final conclusions were developed exclusively by the authors.

Acknowledgements

This research is funded by the Science Committee of the Ministry of Science and Higher Education of the Republic of Kazakhstan (Grant No. AP19679847).

References

- Assel Mukanova, Banu Yergesh, Gaziza Yelibayeva, Gulmira Bekmanova, and Bibigul Razakhova. 2024. Developing an ontological model for pre-election advertising. In *Proceedings of the 9th International Conference on Information Systems Engineering (ICISE 2024)*, pages 1–6. ACM.
- Daniyar Ayazbayev, Andrey Bogdanchikov, Kamila Orynbekova, and Iraklis Varlamis. 2023. Defining semantically close words of Kazakh language with distributed system apache spark. *Big Data and Cognitive Computing*, 7(4):160.
- Gulmira Bekmanova, Banu Yergesh, Assel Omarbekova, Yerkyn Ongarbayev, and Altanbek Zulkhazhav. 2024. Ethical aspects of analyzing Kazakh political discourse. In *ICCSA 2024 Workshops, Lecture Notes in Computer Science*, volume 14819, pages 138–145. Springer.
- Laura Orynbay, Gulmira Bekmanova, Banu Yergesh, Assel Omarbekova, Ayaulym

- Sairanbekova, and Altynbek Sharipbay. 2025. The role of cognitive computing in nlp. *Frontiers in Computer Science*, 6:1486581.
- Lena Zhetkenbay, Altynbek Sharipbay, Bibigul Razakhova, Gulmira Bekmanova, Alibek Barlybayev, Aizhan Nazyrova, and Banu Yergesh. 2025. Formalization of morphological rules for Kazakh nouns in the new Latin alphabet. *Journal of Applied Data Sciences*, 5(3):1999–2019.
- Luciano Floridi and Josh Cowls. 2022. A unified framework of five principles for ai in society. *Machine learning and the city: Applications in architecture and urban design*, pages 535–545.
- Murat Nassimov. 2022. Election programs and political slogans in Kazakhstan. *Codrul Cosminului*, 28(1):115– 144.
- Rustem Yeshpanov and Huseyin Atakan Varol. 2024. Kzsandra: Kazakh sentiment analysis dataset of reviews and attitudes, *arXiv preprint arXiv:2403.19335*.
- Zhaniya Omirkhankyzy Abdiman. 2025. *Inaugural Address as Part of Political Discourse*. Ph.D. thesis, Al-Farabi Kazakh National University. Doctor of Philosophy (PhD). URL (abstract): [https://farabi.university/storage/files/2030243563683989b2ac6e1925763992_Abtract %20Abdiman%20Zh.pdf](https://farabi.university/storage/files/2030243563683989b2ac6e1925763992_Abtract%20Abdiman%20Zh.pdf).

GEOGRAPHIC DISTRIBUTION OF DACO-ROMANIAN SUBDIALECTS THROUGH COMPUTATIONAL METHODS

SILVIU-IOAN BEJINARIU¹, FLORIN-TEODOR OLARIU²,
RAMONA LUCA¹

¹ *Institute of Computer Science,*

Romanian Academy Iași Branch, 2, T. Codrescu Street, Iași, Romania
{silviu.bejinariu, ramona.luca}@iit.academiaromana-is.ro

² *“Alexandru Philippide” Institute of Romanian Philology,*

Romanian Academy Iași Branch, 2, T. Codrescu Street, Iași, Romania
olariuft@yahoo.com

Abstract

Linguistic atlases document, map, and compare language variations across space and time. They help to identify dialectal variations in terms of vocabulary, pronunciation, grammar or regional expressions in a language. Linguistic atlases also help to study the evolution of language by comparing regional and historical variants to identify phonetic, lexical or morphological changes. In comparative linguistic research, such atlases also help classify dialects and identify cross-language influences. It is well known that the Daco-Romanian dialect has a large number of subdialects. This paper proposes a computerized method for identifying its dialect areas using lexical criteria. Dialect data were taken from map legends in *Romanian Linguistic Atlas by Regions – Synthesis, vol. III*, and several unpublished maps from *vol. IV*. A clustering of the survey points is performed using two approaches: hierarchical clustering using dialect distances computed between all survey points and K-Means clustering using semantic vectors (embeddings) extracted using pre-trained Sentence Transformers models. The results obtained by the two methods are discussed and compared in the case of three areas of the Daco-Romanian dialect: Romanian Moldova and Bukovina and Republic of Moldova. The dialect areas obtained using the first approach are closer to those resulting from traditional dialect research. The clustering results are presented in graphical format by overlaying on interactive dialect maps.

Keywords: linguistic atlases, dialects, clustering, lexical distance, sentence transformer.

1 Introduction

Work on information systems for publishing Romanian linguistic atlases began over two decades ago in Iași and Cluj-Napoca. It continues today through collaboration between several institutes of the Romanian Academy, including the Institute of Computer Science, the “Iorgu Iordan – Al. Rosetti” Institute of Linguistics, and the “Alexandru Philippide” Institute of Romanian Philology. Earlier research mainly focused on creating tools (e.g., ALRMaps, ALR-Word-AddIn) for digitizing and publishing dialect data in different atlas formats: analytical and synthetic maps, un-cartographed materials, indexes, dialect texts (Bejinariu *et al.*, 2022, 2024), and online as analytical maps with multimedia capabilities (Olariu *et al.*, 2023). Some dialectal analysis tools were developed but their main goal was to improve the graphic quality or to create new layouts of the published maps. Similarly, the dialectometric studies developed within the framework of the collaboration between the “Sextil Pușcariu” Institute of Linguistics and Literary History, Romanian Academy Cluj-Napoca Branch and the York University, Toronto, were used in the Romanian Online Dialect Atlas (RODA) systems for the creation and publication of dialectometric interpretive maps (Embleton *et al.*, 2013).

The Romanian language has four dialects: Daco-Romanian, Aromanian, Megleno-Romanian and Istro-Romanian spoken in Romania, Republic of Moldova, Ukraine, Serbia, Hungary, Albania, Bulgaria, Republic of North Macedonia, Greece and Croatia. The Daco-Romanian dialect is spoken in Romania, Republic of Moldova, Ukraine (Transcarpathia, northern Bukovina) but also in several localities in Hungary, Serbia and Bulgaria that are located in the border areas.

The interpretation of the maps of the Romanian Linguistic Atlas (ALR), first published in 1938, revived the issue of classifying dialectal varieties of Daco-Romanian from a synchronic perspective. On the basis of the research undertaken by the Cluj School of Linguistics, represented by Sextil Pușcariu, Sever Pop, and Emil Petrovici, later joined by Radu Todoran, five dialectal units within Daco-Romanian (also referred to as subdialects) were identified (Saramandu and Nevaci, 2009):

- the Wallachian subdialect: spoken in Wallachia, Dobruja, Oltenia (except for the north-western area), as well as in the southern and south-eastern parts of Transylvania;
- the Moldavian subdialect: spoken in Moldavia, Bessarabia (extending beyond the Dniester), Bukovina, and the north-eastern region of Transylvania;
- the Banat subdialect: spoken in Banat, in north-western Oltenia, and in south-western Transylvania;

- the Crișana subdialect: spoken in Crișana and in the western region of Transylvania;
- the Maramureș subdialect: spoken in the northern part of the country.

Some authors (e.g. Alexandru Philippide) consider other classifications of subdialects which are usually divided into two large groups: the Northern (Moldavian–Transylvanian) group and the Southern (Meridional) group. The various classifications of dialects and the dialectal areas are determined by the chosen criteria, which can be phonetic, morphological, and lexical.

This paper proposes a lexical based computerized method for identifying the dialect areas of the Daco-Romanian dialect. The study is focused on the survey points network used in the *Romanian Linguistic Atlas by Regions. Synthesis* (ALRR. *Synthesis*) and *Moldovan Linguistic Atlas* (ALM).

2 The Romanian Linguistic Atlas by Regions. Synthesis

ALRR. *Synthesis* includes a network of 1203 localities, of which 998 in Romania and 205 in the Republic of Moldova and Ukraine. It brings together dialect material from the seven regional linguistic atlases in Romania (Banat, Crișana, Maramureș, Transylvania, Moldavia and Bukovina, Wallachia and Dobruja, Oltenia) and the *Moldovan Linguistic Atlas* published in the Republic of Moldova (Saramandu *et al.*, 2018).

In fact, the dialect material included in ALM was collected from 240 localities in Bessarabia, Transnistria, northern Maramureș (Transcarpathia), northern Bukovina, Hertsa region, south-eastern Ukraine, Caucasus (Abkhazia and Krasnodar Krai in the Russian Federation), and the Asian part of the former USSR (Kazakhstan, Kyrgyzstan, and Primorye in the Russian Far East). Yet, only 205 localities, those from Republic of Moldova and partially from Ukraine are considered in ALRR. *Synthesis*.

2.1 Input dataset selection

Until now, three volumes of ALRR. *Synthesis* were published and the fourth volume is in preparation. From these, only the third volume was prepared using a computer application which means the legends are available in electronic format and these only partially. Some legends from the upcoming volume are also available. As noted by Bejinariu *et al.* (2022), these legends are stored in Excel files listing each survey point's forms and corresponding symbols (Fig. 1). The actual forms (with phonetic phenomena) are prepared in text documents which will be drawn on maps (Fig. 2).

The third volume of ALRR. *Synthesis* contains 155 maps. For certain words, when the aspect of interest was not only lexical but phonetic or morphologic, multiple maps were included. From this collection of legends only 33 of them were selected from those which presents the lexical variation. From these only 18 maps include the forms used in the Republic of Moldova. Therefore, in the

following sections 33 legends are used if the analysis is performed for survey points in Romania and only 18 legends if the analysis considers survey points in the Republic of Moldova. Also, in some survey points the lexical forms are missing, either because the interviewee does not know the answer, or a different lexical form than those included in the map's legend is used (see the last three lines in the legend of the map in Fig. 2). Thus, for some words the number of forms is lower than 998 for Romania (RO) and lower than 205 for Republic of Moldova (RM). More details are presented in Table I.

Table 1: Number of available forms for each notion

Points		<i>feastă</i>	<i>tencuială</i>	<i>muruiesc</i>	<i>jgheab</i>	<i>căpriori</i>	<i>olane</i>	<i>leațuri</i>	<i>murdară</i>	<i>tălpei</i>	<i>vas de mămăligă</i>
RO	998	611	951	895	966	986	956	984	971	976	901
RM	205	116		194	203						202

(continued)

Points		<i>lighean</i>	<i>față de masă</i>	<i>ladă</i>	<i>mă uit</i>	<i>fedeleș</i>	<i>unchi</i>	<i>bunic</i>	<i>bunică</i>	<i>flăcău</i>	<i>însărcinată</i>	<i>naște</i>	<i>steapă</i>
RO	998	993	978	980	998	825	994	994	997	984	981	987	970
RM	205	203	202	201	203			202	202		200	203	199

(continued)

Points		<i>tresare</i>	<i>de-a bușilea</i>	<i>nun</i>	<i>cumătru</i>	<i>cumătră</i>	<i>leagăn</i>	<i>cumnat</i>	<i>văduv</i>	<i>văduvă</i>	<i>sicriu</i>	<i>cimitir</i>
RO	998	990	966	995	988	981	995	989	988	990	993	985
RM	205		194	189					201		200	201

Another problem is related to the different keyboard layouts used for writing Romanian text with diacritics. In these layouts, the encodings of the characters ‘ș’, ‘ț’, ‘Ș’ and ‘Ț’ are different. To avoid confusions and mistakes, all characters written in legends using the old encoding (which actually corresponds to the Romanian Legacy keyboard layout) are replaced by the new encoding (corresponding to Romanian Standard keyboard layout). However, there are some

applications, e.g. QGIS which is used to prepare map's template (Bejinariu *et al.*, 2022) that still use the old encoding to draw text on maps.


2.2 Input dataset preparation

The Excel legends were merged into a single dataset containing all necessary information for the experiments (Fig. 3). First, the indices of the forms used (column D in Fig. 1) are replaced with the forms themselves and added to the final table of forms.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X
					GSIDS_043	GSIDS_044	GSIDS_046	GSIDS_047	GSIDS_040	GSIDS_041	GSIDS_042	GSIDS_112	GSIDS_080	GSIDS_089	WEEDN_071	WEEDN_189	WING1_231	GSIMS_033	TIMES_042	TIMES_065	TIMES_063			Forma
1	Sat	Nord	Est	Forma																				
2	1	2206.12	1724.78	6						1														6
3	2	2179.68	1676.31	3			1																	3
4	3	2131.21	1658.69	6						1														6
5	4	2054.10	1546.33	6						1														6
6	5	1820.57	1649.88	6						1														6
7	6	1725.84	1704.95	3			1																	3
8	7	1587.05	1698.34	3			1																	3
9	8	1485.71	1821.72	3			1																	3
10	9	1635.52	1867.98	3			1																	3
11	10	1752.28	1863.58	3			1																	3
12	11	1820.57	1920.86	6						1														6
13	12	1666.36	1945.09	3			1																	3
14	13	1494.52	1905.43	3			1																	3
15	14	1595.86	1980.34	3			1																	3
16	15	1917.51	1923.06	3			1																	3
17	16	2223.74	1898.82	6						1														6
18	17	1919.71	1984.75	7							1													7
19	18	1401.99	2083.88	3			1																	3

Figure 1: Example of original legend file from *ALRR. Synthesis vol. III*

In the original Excel files, the survey points (column A in Fig. 1) are indexed using numeric keys as follows: points in Romania from 1 to 998 and points from Republic of Moldova from 01 to 0224. To avoid confusion during numerical processing (e.g., between “1” and “01”), survey point keys were renamed as R1–R998 and M1–M224 respectively (column A in Fig. 3). The coordinates specified in the input Excel files (columns B and C in Fig. 1) are expressed in pixels and are replaced with the geographical coordinates of the survey points to allow their correct positioning on the map in GIS systems. Finally, a new column that contains the name of the region in which the survey point is located is added (column D in Fig. 3). The resulted table is saved in CSV format with UTF-8 encoding.

Figure 2: Synthetic map from ALRR. *Synthesis vol. III*, re-digitalized

	A	B	C	D	E	F	G	H	I	J	K	L
1	Localitate	Lat	Lon	Regiune	țeastă	tencuiesc	tencuială	cercuiesc	muruiesc	lacăte	încui	jgheab
2	R1	44.76187	22.39847	Banat	troacă	văcăluiesc	maltăr	nuelesc	spoiesc	lacăte	încui	rinăr
3	R2	44.72223	22.37047	Banat	scăfirile	măltăresc	maltăr	bat scuarță de cei	-	lacăte	încui	rinăr
4	R3	44.70651	22.31685	Banat	-	văcăluiesc	maltăr	bat scuarță de cei	spăl vatra	lacăte	încui	rinăr
5	R4	44.61374	22.23468	Banat	-	văcăluiesc	maltăr	bat scuarță de cei	netezesc	lacăte	încui	vălău
6	R5	44.68841	21.97074	Banat	-	văcăluiesc	-	-	spoiesc	lacăte	încun	dudă
7	R6	44.72909	21.8631	Banat	troacă	măltăresc	maltăr	-	spoiesc	lacăte	încun	vălău
8	R7	44.71838	21.70927	Banat	-	măltăresc	maltăr	bat nuele	netezesc	lacăte	încui	vălău
9	R8	44.81364	21.59111	Banat	-	măltăresc	maltăr	bat nuele	spoiesc	lacăte	încun	vălău
10	R9	44.85661	21.75549	Banat	troacă	măltăresc	maltăr	bat scuarță de cei	fac vatra	lacăte	încui	vălău
11	R10	44.85741	21.88581	Banat	-	măltăresc	maltăr	ciercuiesc	-	lacăte	încun	vălău
12	R11	44.90575	21.95966	Banat	-	văcăluiesc	maltăr	-	spoiesc	lacăte	încun	vălău
13	R12	44.91962	21.78641	Banat	glavă	măltăresc	maltăr	bat scuarță de cei	ung	lacăte	încun	vălău
14	R13	44.88129	21.59689	Banat	-	măltăresc	maltăr	bat lăceș	netezesc	lacăte	încun	dudă
15	R14	44.94528	21.70619	Banat	glavă	măltăresc	maltăr	bat scuarță de cei	lipesc	lacăte	încun	vălău
16	R15	44.91087	22.06802	Banat	-	măltăresc	maltăr	ciercuiesc	-	lacăte	încun	vălău
17	R16	44.90145	22.41245	Banat	-	văcăluiesc	maltăr	bat nuele	-	-	încui	rinăr
18	R17	44.96028	22.06804	Banat	-	văcălui	maltăr	ciercuiesc	-	lacăte	încun	vălău
19	R18	45.02104	21.48521	Banat	-	măltăresc	maltăr	-	spoiesc	lacăte	încun	vălău
20	R19	45.0245	21.64	Banat	-	tencuiesc	-	-	ung	lacăte	încun	vălău
21	R20	45.00893	22.16471	Banat	-	măltăresc	maltăr	bat nuele	-	lacăte	încun	vălău
22	R21	45.00403	22.36837	Banat	-	văcălui	maltăr	ștucăturiesc	-	lacăte	încun	vălău
23	R22	45.04934	22.22479	Banat	-	tencuiesc	-	ciercuiesc	-	lacăte	încun	dudă

Figure 3: Formatted content of the input dataset

3 Dialect areas identification using lexical criteria

Dialectometry is a subfield of dialectology introduced by Jean Séguy (1971). Hans Goebel (1981) developed and refined the theory in the 1970s and 1980s. Dialectometry is a computational approach to dialectology that studies language variation through quantitative and statistical analysis. Mainly, it is based on linguistic distances which are computed to evaluate the differences between

languages or dialects. More modern methods take into account the semantic domain which is important when analysing the geographical distribution of lexical variation.

3.1 Dialect areas identification by lexical distances

Linguistic distance measures the degree of difference or similarity between two languages or dialects. Linguistic distance evaluates morpho-syntactic, phonological and phonetic features which are considered by some authors (Séguy, 1971) the most relevant. Therefore, when only lexical variation is considered, the recommended term is *lexical distance*. Thus, from now on in this paper we will use the term *lexical distance*. In practice several distances are used: *Cosine*, *Jaccard*, *Levenshtein*, etc. (Goebel, 2018).

Given a dataset that contains the forms used for n words (notions) in m survey points, the lexical distances are computed as follows, using data already prepared in the Excel file described above.

- compute $Sim_{jk}^i, i = \underline{1, n}; j, k = \underline{1, m}$, similarity between the forms of word i in survey points j, k .
- compute the matrix of lexical distances LD between the survey points. It is a square and symmetric matrix with zero on the main diagonal:
- $LD_{jj} = 0; LD_{jk} = LD_{kj} = 1 - \frac{\sum_i CS_{jk}^i}{n^*}$
for $i = \underline{1, n}; j, k = \underline{1, m}, len(form_j^i) > 0, len(form_k^i) > 0$ where n^* is the number of nonempty forms for both words j and k .

The *Cosine* similarity CS between two forms is computed as follows: both forms are transformed into the frequency vectors X, Y of their elements (characters) and $CS(X, Y) = \frac{(X \cdot Y)}{\|X\| \cdot \|Y\|}$ where $X \cdot Y$ is the dot product. CS has values in range $[0; 1]$ where a value of 1 means full similarity but not complete matching of the two forms, e.g. *Cosine* similarity between a word and its anagram is 1. Yet it is used in dialectology for its low complexity. The *Jaccard* similarity JS does not consider the number of occurrences of each character. Let X and Y the sets of characters in the two forms. $JS(X, Y) = \frac{|X \cap Y|}{|X \cup Y|}$. JS also has values in range $[0; 1]$. The *Levenshtein* distance between two words is the minimum number of single-character edits (insertions, deletions or substitutions) required to change one word into the other. It is very sensitive to length and small spelling differences, but it is recommended for lexical comparisons (word to word).

To determine the dialect areas, a hierarchical agglomerative clustering algorithm is further applied on the distance matrix. It starts with a set of clusters each containing a single point from the input set. The clusters between which the distance is minimal are repeatedly unified and then the distances between the clusters are recalculated. The process stops when the set contains only one cluster.

The cluster's merging process is represented using a tree diagram (dendogram). By analysing the dendogram, clusters can be extracted in a specific step when their number is the desired one.

3.2 Dialect areas identification using semantic vectors

The second approach to dialect clustering is based on sentence transformers. They transform sentences in fixed-length vector representations (embeddings). Some models analyse the context in which the words are used increasing the quality of the embeddings. The generated embeddings can be used for semantic search, clustering, information retrieval and many other applications. In the experiments described in this paper, the following models are used:

- *distiluse-base-multilingual-cased*. This sentence-transformer model maps sentences and paragraphs to a 512-dimensional dense vector space and can be used for tasks like clustering or semantic search. This model is based on the DistilBERT multilingual model which was trained on the concatenation of Wikipedia in 104 different languages. *Distiluse-base-multilingual-cased* was trained for semantic similarity and finally a pooling was applied to reduce dimensionality (Reimers and Gurevych, 2019).
- *bert-base-romanian-cased-v1*. This is the first purely Romanian transformer-based language model, based on the Multilingual BERT model, it maps sentences to 768-dimensional vectors and it was trained using about 15GB data from OSCAR, OPUS and Wikipedia corpora (Dumitrescu *et al.*, 2020). Other BERT models for the Romanian language are: RoBERT (Masala *et al.*, 2020) and Distil-BERT (Avram *et al.*, 2022).

The semantic vectors for the survey points are extracted from "sentences" obtained by concatenating the forms of all words in the respective survey point. Then, the well-known K-means clustering algorithm is applied to group unlabelled vectors into clusters.

3.3 Implementation and delivery of results

The dialect areas identification procedures are implemented in Python. Among the packages used, we mention the most relevant ones: `pandas`, `scipy`, `numpy`, `sklearn`, `transformers`, `sentence transformers` and `folium`. The `folium` package makes it easy to visualize data on interactive leaflet maps. It enables both the binding of data to a map for visualization, and the use of overlay markers as vector or raster images, as well as HTML elements. In our implementation, the results are saved in HTML files that contain maps with markers of different colours (Fig. 4). The name of the survey point, region, country as well as the cluster number and the list of the used forms are displayed

using a tooltip control for each point. The identifier of the cluster to which each survey point belongs is saved in a new column entitled ‘cluster’ in the CSV file.

Five clustering procedures were applied. Three of these are based on the *Cosine*, *Jaccard* and *Levenshtein* distances and Agglomerative clustering. In the other two procedures, the K-means algorithm was applied in the multi-dimensional space of the embeddings extracted using the *distiluse-base-multilingual-cased* and *bert-base-romanian-cased-v1* models. In Table II is shown the list of the areas on which the clustering procedure was applied as well as the number of word and the number of clusters in which the survey points were grouped.




Figure 4: Interactive map with clustering results

Table 2: Areas in which the distribution of dialects was studied

	Areas	# words	# clusters
RO	Romania	33	7 ÷ 12
RM	Republic of Moldova	18	2 ÷ 5
MB	Moldavia and Bukovina	33	2 ÷ 5
RO+RM	Romania + Republic of Moldova	18	7 ÷ 15
MB+RM	Moldavia and Bukovina + Republic of Moldova	18	2 ÷ 7

4 Results and discussion

Dialectologists define dialectal areas by analysing morpho-syntactic or lexical variations across selected words and expressions, and these areas may vary depending on the chosen criteria. Also, in this experiment in which dialect areas are detected

based on lexical criteria, a relatively small number of words are used and these may not necessarily be the most representative. These experiments represent a starting point; further work will expand the analysis to a larger set of words.

4.1 Detection of dialect areas in Romania


Figure 5 shows the maps obtained by grouping the survey points in Romania in 10 clusters using four of the above-mentioned procedures. The colour assigned to each area differ from one map to another because they are assigned according to the order in which the clusters are determined. From the analysis of the four maps, the following conclusions can be noted:

- the dialect areas only partially correspond to the historical regions of Romania which are relatively similar to the dialect areas established by dialectologists;
- the areas differ the most in Transylvania, where the dialect areas are divided differently even by dialectologists;
- the cluster points determined with the help of semantic vectors are quite mixed from the point of view of their geographical position.


Lexical distances measure differences at the character level, whereas embeddings focus on semantic relationships through tokenization. In the example presented, for the *bert-base-romanian-cased-v1* model, which was trained specifically for the Romanian language, only 14193 (~40%) of the words (counted without removing duplicates) were covered by the vocabulary. The remaining 20955 words were broken into 50092 tokens. Therefore, in the following paragraphs we will refer especially to lexical distances-based clustering results.

Other two examples of dialect areas detections are displayed in Fig. 6. The subdialect spoken in Maramureş appears as a distinct area starting from the clustering in 7 areas, even though it appears to be close to the one spoken in southern Crişana (Fig. 6.b) and southern Transylvania (Fig. 6.a). Also, different dialect areas are identified in the southern Banat and another in the transition area between Oltenia and Wallachia. In contrast, the two dialect areas of Moldavia and Bukovina remain constant even though there are differences here too, as we will see below.


GEOGRAPHIC DISTRIBUTION OF DACO-ROMANIAN SUBDIALECTS




a. Romania, 10 clusters, *Cosine* distance



b. Romania, 10 clusters, *Levenshtein* distance



c. Romania, 10 clusters, *bert_base_romanian*



d. Romania, 10 clusters, *distiluse-base-multilingual*

Figure 5: Results comparison of dialectal areas detection in Romania

4.2 Detection of dialect areas in Moldavia and Bukovina

Some significant results for dialectal areas detection in Moldavia and Bukovina are shown in Fig. 7. The number of clusters was varied in the range [2; 5]. The subdialect spoken in Bukovina was highlighted in all clustering variants. Starting from the grouping into 3 clusters, the subdialect used in the north-eastern part of Moldova (Iași and Botoșani counties) is also highlighted (Fig. 7.a). In the case of the grouping into 5 dialect areas, a separation occurs between the central and southern parts, probably due to the Wallachian dialect that has influences in this area. There is also a separation of the Bucovina area, in which a dialect area is highlighted in the mountain area of Suceava county (Fig. 7.b).




Figure 6: Dialectal areas detection results for Moldavia and Bukovina




Figure 7: Dialectal areas detection results for Moldavia and Bukovina

4.3 Detection of dialect areas in Republic of Moldova


For the Republic of Moldova, Fig. 8 presents two interesting results. We recall that in this case only 18 words were used and the set of survey points also includes localities from Ukraine. In the case of the 4-groups clusterization (Fig. 8.a), the areas from Northern Bukovina (including localities from Transcarpatia) on the one hand and Transnistria with the Odessa region (where Russian speakers are in the majority) on the other hand are clearly highlighted. The other two areas are in the centre of Bessarabia (with Romanian speakers in the majority) and the northern and southern parts (with strong Russian influences) respectively.

An apparently unexpected situation is presented in the Fig. 8.b in which the survey points from both Republic of Moldova and Moldavia and Bukovina were considered. In the case of 7-groups clustering, the grouping mode remains unchanged: the four dialect areas mentioned above in Republic of Moldova and the rest in Moldavia and Bukovina. The only linguistic similarity occurs in the


Bukovina area of Romania and Northern Bukovina in Ukraine. It would have been expected that similarities would appear between points on the two sides of the Prut River, but these are missing. These similarities do not appear even if we reduce the number of clusters to 3.

4.4 Detection of dialect areas in Romania and Republic of Moldova

In the last experiment we considered all survey points from Romania and Republic of Moldova together (Fig. 9). The number of clusters was varied in range [7; 15]. It is obvious that a reduced number of 7 groups is not enough (Fig. 9.a): points in Oltenia and Wallachia are grouped together and similarly the points from Bukovina and Transylvania.




a. Republic of Moldova,
4 dialect areas, *Cosine* distance




b. Republic of Moldova + Moldavia and
Bukovina,
7 dialect areas, *Cosine* distance

Figure 8: Dialectal areas detection results for Republic of Moldova



a. 7 dialect areas, *Cosine* distance



b. 14 dialect areas, *Jaccard* distance

Figure 9: Results comparison of dialectal areas detection in Romania and Republic of Moldova

If the number of clusters is doubled (Fig. 9.b), the dialect areas generated for Romania and Republic of Moldova respectively are similar to those generated for

each one separately. Some changes occur regarding the Maramureş area whose points are grouped with those from Transcarpathia. Moreover, in Republic of Moldova the points from Transnistria are grouped with those from northern and southern Bessarabia, obviously due to the influence of the Russian language. Also, the lexical separation between Moldavia and the Republic of Moldova continues to manifest itself. Although a closer similarity between the two lexicons might be expected, it is important to recall that these surveys date from the 1970s and that, except for the interwar period, the Republic of Moldova has long been under strong Russian influence in its various forms.

4.5 Objective evaluation of clustering results

In the above paragraphs, the number of clusters for determining dialect areas was chosen so as to be as close as possible to the dialect areas resulting from dialectological publications. Next, we will make an objective evaluation of the results obtained from clustering. In the literature, different measures are used for this evaluation, the most known being Silhouette Score and Davies-Bouldin Index (Liu *et al.*, 2010).

The Silhouette Score (SS) is a commonly used metric that measures how similar an object is to the other objects in its own cluster compared to objects in other clusters. The score for each point ranges from -1 (assigned to wrong cluster) to 1 (the point is well matched to its cluster). By averaging these scores across all points, the Silhouette Score provides an overall measure of how well-defined and separated the clusters are. The Davies-Bouldin Index (DBI) is an evaluation measure that assesses both the compactness of clusters and their separation from one another. For each cluster, it computes the ratio between the average distance of points within the cluster (intra-cluster similarity) and the distance to the nearest other cluster (inter-cluster separation). The final DBI is the average of these ratios across all clusters. Lower DBI values indicate better clustering quality. Also, the Multidimensional Scaling (MDS) was applied to embed the pairwise lexical distances matrix of survey points into a 5-dimensional Euclidean space, preserving the original dissimilarities for cluster evaluation.

A dialect area clustering assessment was carried out for all 5 areas mentioned above: Moldavia and Bukovina; Moldavia and Bukovina + Republic of Moldova; Republic of Moldova; Romania; Romania + Republic of Moldova. The number of clusters was varied in range [2; 20] and all three distances were used: *Cosine*, *Jaccard* and *Levenshtein*. In Figs. 10-14 are shown the charts with values of the two evaluation measures and the resulting maps for the best clustering evaluation values.

For the Moldova and Bukovina area (Fig. 10) the best values are obtained for clustering into two dialect areas. These are generally accepted even though the Moldavian area can be divided into several subareas. In case Moldavia and Bukovina and Republic of Moldova are analysed together the best results are

GEOGRAPHIC DISTRIBUTION OF DACO-ROMANIAN SUBDIALECTS

obtained for 3-dialect areas clustering. The two areas remain separate with the exception of Bucovina and Northern Bucovina, which are unified in the same cluster (Fig. 11).




Figure 10: Evaluation measures and map of dialect areas in Moldavia and Bukovina




Figure 11: Evaluation measures and map of dialect areas in Moldavia and Bukovina + Republic of Moldova




Figure 12: Evaluation measures and map of dialect areas in Republic of Moldova

The dialect areas clustering results for Republic of Moldova the results are not consistent. When using the *Levenshtein* distance, clustering into 4 areas is the best,

but for the other distances and measures the best values obtained are for 6, 8 and 11 clusters. Also, the dialect areas in Fig. 12 differ significantly from those obtained using *Cosine* distance for the same number of clusters (Fig. 8.a). For determining the dialect areas in Romania, the best values were obtained in the case of clustering into 4 groups. As can be seen on the map in Fig. 13, the dialects from Wallachia, Dobruja and Oltenia were unified, as well as those from Transylvania, Crişana and Maramureş. The same 4 clusters are obtained if the dialects from Romania and Republic of Moldova are analysed together. In this case, the Banat, Bukovina, Northern Bukovina and Transcarpatia areas are grouped with Transylvania, while Republic of Moldova remains separate as dialect area.




Figure 13: Evaluation measures and map of dialect areas in Romania




Figure 14: Evaluation measures and map of dialect areas in Romania + Republic of Moldova

5 Conclusions and future work

The paper addresses the problem of automatic determination of subdialect areas using only the lexical features of the spoken language. Between the two approaches, with lexical distances and semantic vectors, the first one yielded more

consistent results aligning with Romania's and Republic of Moldova's regional geography. However, the number of words whose forms were analysed and grouped has been relatively small. This is also the reason why, when evaluating clustering, the optimal number of dialect areas differs significantly from that obtained by classical methods. We consider this to be an initial approach that will be developed by increasing the number of words as well as analysing the variation of morpho-syntactic features.

AI Use Statement

In the final editing stage of the paper, ChatGPT (<https://chatgpt.com/>) was used to check spelling and detect overly complex or unclear phrases that required splitting into multiple sentences. All corrections and modifications were made by the authors of the paper.

Acknowledgement

We thank Manuela Nevaci and Nicolae Saramandu for allowing the use of legends of the Romanian Linguistic Atlas by Regions. Synthesis. We thank Daniela Gîfu for the information about sentence transformers and especially Romanian transformer models.

The work related to the Moldovan Linguistic Atlas was supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS–UEFISCDI, project number PN-IV-P8-8.3-ROMD-2023-0318, within PNCDI IV. All other results were obtained within the framework of the fundamental research projects of the Romanian Academy.

References

- Andrei-Marius Avram, Darius Catrina, Dumitru-Clementin Cercel, Mihai Dascălu, Traian Rebedea, Vasile Păiș, and Dan Tufiş. 2022. Distilling the Knowledge of Romanian BERTs Using Multiple Teachers. In N. Calzolari, F. Béchet, P. Blache, K. Choukri, C. Cieri, T. Declerck, S. Goggi, H. Isahara, B. Maegaard, J. Mariani, H. Mazo, J. Odijk, and S. Piperidis (eds.) *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 374–384, Marseille, France. European Language Resources Association.
- Florin-Teodor Olariu, Alexandru Laurențiu Cohal, Luminița Botoșineanu, Veronica Olariu, Ramona Luca, and Silviu-Ioan Bejinariu. 2023. Making Linguistic Resources Accessible. The Audio Archive of the New Romanian Linguistic Atlas by Regions. Moldova and Bukovina. In A. Vasilescu, D. Gîfu, D. Tufiş, and D. Cristea (eds.) *Proceedings of the 18th International Conference on Linguistic Resources and Tools for Natural Language Processing – ConsILR-2023*, Brașov, 11-14 dec. 2023, pages 199-214.
- Hans Goebl. 1981. Éléments d'analyse dialectométrique (avec application à l'AIS). In *Revue de linguistique Romane*, 45, pages 349–420.

- Hans Goebel. 2018. Dialectometry. In C. Boberg, J. Nerbonne and D. Watt (editors), *The Handbook of Dialectology*, John Wiley & Sons, Inc., pages 123-142.
- Jean Séguy. 1971. La relation entre la distance spatiale et la distance lexicale. In *Revue de Linguistique Romane*, 35. pages 335–57.
- Mihai Masala, Stefan Ruseti, and Mihai Dascalu. 2020. RoBERT – A Romanian BERT Model. In D. Scott, N. Bel, and C. Zong (eds.) *Proceedings of the 28th International Conference on Computational Linguistics*, Barcelona, Spain (online), International Committee on Computational Linguistics, pages 6626–6637.
- Nicolae Saramandu (coord), Mihaela Morcov, Manuela Nevaci, Daniela Răuțu, Irina Floarea, Carmen-Ioana Radu, Ionuț Geană, and Mara Iuliana Manta. 2018. *Romanian Linguistic Atlas by Regions. Synthesis* (original title in Romanian: *Atlasul lingvistic român pe regiuni. Sinteză*). vol. III, Publishing House of the Romanian Academy, <http://www.lingv.ro/images/ALRR%20sinteza.pdf>, last accessed at 20.08.2025.
- Nicolae Saramandu and Manuela Nevaci. 2009. *Multilingualism and minority languages in Romania* (original title in Romanian: *Multilingvism și limbi minoritare în România*). „Iorgu Iordan - Alexandru Rosetti” Institute of Linguistics, Bucharest, <https://rgnpress.ro/wp-content/uploads/2018/07/Multilingvism-si-limbi-minoritare-in-Romania.pdf>, last accessed at 1.11.2025.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In K. Inui, J. Jiang, V. Ng, and X. Wan (eds.) *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, Hong Kong, Nov. 3–7, 2019, pages 3982-3992, Association for Computational Linguistics.
- Sheila Embleton, Dorin Uritescu, and Eric S. Wheeler. 2013. Defining dialect regions with interpretations: Advancing the multidimensional scaling approach. In *Literary and Linguistic Computing*, 28(1), pages: 13–22.
- Silviu-Ioan Bejinariu, Ludmila Malahov, Tudor Bumbu, Florin-Teodor Olariu and Ramona Luca. 2024. New Achievements in the Digitization of Linguistic Atlases of the Romanian Language. In A. Baci, D. Nagy-Onița, D. Tufiș, and D. Cristea (eds.) *Proceedings of the 19th International Conference on Linguistic Resources and Tools for Natural Language Processing – ConsILR-2024*, Alba Iulia, 14-16 nov. 2024, pages 23-41.
- Silviu-Ioan Bejinariu, Ramona Luca, Florin Iftene, Manuela Nevaci, Carmen Irina Floarea and Nicolae Saramandu. 2022. GIS Based Publishing System for the Romanian Linguistic Atlas by Regions – Synthesis. In S. Cojocaru, V. Bobicev, T. Verlan, D. Tufiș, and D. Cristea (eds.) *Proceedings of the 17th edition of The International Conference on Linguistic Resources and Tools for Natural Language Processing – ConsILR-2022*, Chișinău, 10-12 November 2022, pages 81-92.
- Ștefan Dumitrescu, Andrei-Marius Avram, and Sampo Pyysalo. 2020. The birth of Romanian BERT. In T. Cohn, Y. He, and Y. Liu (eds.) *Findings of the Association for Computational Linguistics: EMNLP 2020* (online), pages 4324–4328, Association for Computational Linguistics.
- Yanchi Liu, Zhongmou Li, Hui Xiong, Xuedong Gao, and Junjie Wu. 2010. Understanding of Internal Clustering Validation Measures. In G.I. Webb, B. Liu, C. Zhang, D. Gunopulos, and X. Wu (editors). *Proceedings of the 2010 IEEE International Conference on Data Mining*, Sydney, pages 911-916.

CHAPTER 3
DICTIONARY-BASED METHODS
AND LEXICOGRAPHIC
RESOURCES

FROM TRANSITIONAL SCRIPT TO CRITICAL DIGITAL EDITION: DIGITIZING *DISIONĂRAȘ ROMÂNESC DE CUVINTE TEHNICE ȘI ALTELE GREU DE ÎNȚELES* (1851)

TUDOR BUMBU¹, ELENA ISABELLE TAMBA², ANA VULPE³,
MARIUS-RADU CLIM², MARIANA NASTASIA²,
LIVIA CĂRUNTU-CARAMAN³, LILIANA BOTNARU³,
CRISTINA CATEREV³, CLAUDIUS TEODORESCU⁴

¹ *Institute of Mathematics and Computer Science “V. Andrunachievici”,
State University of Moldova
tudor.bumbu@math.md*

² *Institute of Romanian Philology “Alexandru Philippide”,
Romanian Academy – Iași Branch
isabelle.tamba@gmail.com
mariusclim@gmail.com
nastasiamarianal@gmail.com*

³ *Institute of Romanian Philology “Bogdan Petriceicu-Hasdeu”
State University of Moldova
vanet1854@gmail.com
liviacaraman88@gmail.com
botnari.liliana86@gmail.com
cristinacaterev@gmail.com*

⁴ *Academy of Sciences and Literature, Mainz
claudius.teodorescu@gmail.com*

Abstract

The paper focuses on the development of the process of digitising lexicographical research in Romania, with particular emphasis on a challenging aspect, the digitisation of dictionaries in Romanian Transitional alphabet (RT), which pose rather specific problems and, consequently, the development of the Romanian Lexicographical Corpus (CLRE – <https://clre.solirom.ro/>). We present the workflow and tools developed to produce a digital critical edition of Theodor Stamati's *Disionăraș românesc de cuvinte tehnice și altele greu de înțeles* (Iași, 1851), a dictionary of about 4,200 entries written in the RT. Our pipeline combines: (1) a customised OCR model trained in ABBYY FineReader 16 on representative pages (about 13,000 characters), achieving below 5% CER, (2) a web-based lexicographic interface for post-OCR verification and semantic markup (etymology, analytical definition, synonym/antonym-based definitions,

examples, variants, norms, stable expressions), and (3) an assisted transliteration stage from RT to Latin Romanian (via the HeDy platform), designed to *preserve all prior annotations*, the OCR output, and page scans. The interface adds dedicated editor tools, and a case-sensitive search engine that supports headword lookup in both alphabets, direct page navigation, result previews etc. We outline a font development plan to reproduce key historical ligatures (e.g., *ib*) for a faithful visual rendering. The outcome will be an online, queryable edition aligned with CLRE and a print-ready republication with ISBN that presents the transcription, the transliteration, and the critical part. This work is undertaken within the project PHILORD (PN-IV-P8-8.3-ROMD-2023-0318) by a joint team of lexicographers and IT specialists from Romania and the Republic of Moldova.

Keywords: lexicography, Romanian Transitional alphabet, OCR, dictionary annotation, transliteration, digital edition, CLRE

1 Introduction

Romanian was written over different periods with Cyrillic, Latin, and during the nineteenth century with a transitional alphabet that mixes elements from both. This layered history complicates digitisation, search, and scholarly editing. Our paper addresses these challenges through the case of Theodor Stamati's *Disionăraș românesc de cuvinte tehnice și altele greu de înțeles* (Iași, 1851), a dictionary that is important for the modernisation and the enrichment of Romanian vocabulary in the nineteenth century.

Our goal is to produce a critical digital edition that preserves the original script and page images, supports rich annotation and transliteration, and integrates with the *Corpus lexicografic românesc electronic* (CLRE).

We work within the PHILORD project (PN-IV-P8-8.3-ROMD-2023-0318), in a joint team of lexicographers and IT specialists from research institutes in Iași and Chișinău. The edition is designed for two audiences. It offers philologists a faithful view of the source through paired text and scan and through segment-level markup. It offers digital-humanities users a searchable resource with stable identifiers in a modern Romanian alphabet.

Our approach is practical and incremental. First, we prepare page images and extract text with optical character recognition. Second, we verify and annotate the text in a web interface that supports different markup elements and revision modules. Third, we transliterate the RT into Latin Romanian while keeping previous annotations. Finally, we plan typography that reproduces key historical features such as ligatures and diacritics.

We report five concrete results:

1) An OCR setup in ABBYY FineReader 16, trained on 20 representative pages (about 13,000 characters), using a letter set that combines Latin and RT and a lexicon of more than 10,000 historical lemmas. The model reaches a character error rate below 5 percent, with the most frequent confusion between similar looking letters such as *ō* and *ø*.

2) A browser-based post-OCR interface where editors verify entries against scans and apply markups such as analytical definition, etymology, sense, definition by synonyms, definition by antonyms, examples, variants, norms, and stable expressions. Formatting options include bold, italics, and letter-spacing to preserve historical spacing in headwords and other salient forms. All annotations are saved to XML.

3) A case-sensitive search engine that supports headword lookup in both scripts, automatic transliteration from Latin input to RT, direct page search, previews with highlighting, automatic opening of the selected article, and a mini keyboard for Latin Romanian and RT.

4) An assisted transliteration stage that uses the HeDy service¹. The verification interface mirrors the post-OCR editor and shows the original Cyrillic text and the scan next to the transliterated text. The process preserves all annotations created earlier.

5) A plan for a dedicated OpenType font to render historical ligatures and diacritics accurately. The ligature *ix* will be supported, and the breve will be correctly centred above *ŷ*. This improves fidelity in both screen and print outputs.

The next sections present related work, the methodology (preprocessing with ScanTailor, OCR training and evaluation, post-OCR verification, annotation and search, transliteration, data management, first steps in the development of a specialized font), conclusions with future work, including the online publication aligned with CLRE and a print-ready edition with ISBN.

2 Related Work

The work draws on three strands: (a) the digitisation of Romanian lexicography and the development of shared resources; (b) OCR and text processing for historical Romanian in Cyrillic and transitional alphabets; (c) transliteration methods and editorial models for critical digital editions.

Projects of the Romanian Academy and partner institutions have prepared the ground for large-scale electronic lexicography and for interoperable resources. The *Corpus lexicografic românesc electronic (CLRE)* provides alignment and access across dictionaries and supports editorial workflows and reuse (Tamba, 2014; Tamba Dănilă *et al.*, 2012; Haja *et al.*, 2023; Tamba, 2022; Tamba, 2024; Cristea *et al.*, 2024). These works describe both the institutional context and the practical constraints of digitisation in Romanian, including data formats, editorial policies, and dissemination.

A foundational strand for Romanian lexicography is the digitisation and online dissemination of the Academy's thesaurus (*DA/DLR*). These resources establish

¹ Digitisation platform with tools and resources for processing Cyrillic Romanian documents: <https://digitizare.math.md/>

editorial conventions, data formats, and discovery patterns that our work aligns with (e.g. TEI-based exports, headword anchoring, and corpus interoperability).

Romanian materials printed in the nineteenth century frequently use RT. OCR on these sources requires customised alphabets and historical lexica. Studies report pipelines and tooling for documents printed in RT, including character inventories and pre-/post-processing steps (Cojocaru *et al.*, 2016; Demidova and Burţeva, 2017). General-purpose OCR and layout tools such as ABBYY FineReader, Transkribus and Tesseract are often combined with document pre-processing to handle non-modern ligatures and spacing. For transitional print, published accounts emphasise the need to train lettersets that cover both scripts and to evaluate recurring confusions at character level. Our OCR setup follows this line and contributes a letterset and lexicon targeted at RT.

A recent work (Bumbu *et al.*, 2023) analyses OCR for RT within the HeDy platform and compares two practical strategies. The first reproduces the original variant in the target alphabet and, with ABBYY FineReader prepared for that variant, yields about 7 per cent word error but requires careful setup. The second uses an intermediate Latinised “ligature” alphabet in FineReader to normalise mixed Latin and Cyrillic forms; this lowers the error to about 4.8 per cent and avoids Unicode support issues in training interfaces. On a test page of 1,172 characters, accuracy rose from 92.1 per cent to 97.4 per cent after adding a user dictionary. The study lists typical confusions, notably **Ѣ** vs **B**, **н** vs **II**, and **III**, and proposes a working definition of the RT with 36 letters including **Ѣ ѣ, ІЄ іє, А а, Ӧ ӧ, Ӣ ӣ, Ӥ ӥ, Ӧ ӧ, Ӣ ӣ, Ӥ ӥ, Ӧ ӧ, Ӣ ӣ, Ӥ ӥ**. These observations informed our letterset and dictionary choices and our focus on entry-level review of systematic errors.

A complementary line of research (Frincu *et al.*, 2023) proposes machine-learning pipelines for transliteration using Tesseract with small training sets. Frincu *et al.* (2023) compare training from scratch with transfer learning from Latin or Cyrillic models and evaluate both Latin and RT targets under repeated cross validation. They report character error rates as low as 1.8 per cent for Latin output and around 17.7 per cent when targeting RT directly. The gap is attributed to frequent misclassification of visual similar Latin and Cyrillic vowels, as well as challenges posed by mixed-script layouts. The authors recommend a two-phase workflow that first performs OCR into Latin or RT and then applies context rules and corpus-based correction.

The Romanian transitional alphabet requires rules that are sensitive to local context. Methodological overviews for Romanian digital resources discuss transliteration as a necessary step for access and integration, while acknowledging that philological use still needs the original script (Tamba, 2014; Tamba, 2022; Tamba, 2024; Cristea *et al.*, 2024). Technical work from Chişinău and Iaşi explores algorithmic approaches and services for converting Romanian Cyrillic to Latin and for handling historical scripts (Cojocaru *et al.*, 2016). Our project uses

the HeDy platform for assisted transliteration and integrates verification in an interface that preserves both the Cyrillic text and the page image.

3 Materials and Methodology

This section describes the source materials, the preparation of page images, OCR, post-OCR verification and article-level markup, search, assisted transliteration from RT to Latin, data management, and the design of a dedicated font. Most scans are two-page openings; we split them into single pages during preprocessing so that OCR and later alignment work at page level.

3.1 Source and Scans

We work on the first edition of Theodor Stamati, *Disionăraș românesc de cuvinte tehnice și altele greu de înțeles (Iași, 1851)*. The dictionary includes about 4,200 entries and is printed in the RT. The image set reflects typical nineteenth-century typography: mixed Latin and Cyrillic forms, variable diacritics, headwords with wider spacing, and display capitals. Figure 1 shows a two-page scan (pp. 4-5). For OCR we prepared single-page inputs and then applied the other preprocessing steps described below.

3.2 Preprocessing with ScanTailor

We used ScanTailor to process the scans. Because most scans were two-page openings, we first split each opening into verso and recto pages, then applied the steps below to each page: deskewing, binarisation, margin and content-box detection, and contrast adjustment. Where the gutter produced curvature or shading, we adjusted the content box so that recognition was not affected by the fold. A few pages (1, 2, 35, 36) were missing in our copy. For these we consulted a copy of the dictionary and extracted the missing page scans. No additional preprocessing was required as the pages in this copy were already pre-processed. The next subsection presents the OCR stage with ABBYY FineReader.




Figure 1: Two-page scan (pp. 4–5) from Stamati, *Disionăraș românesc de cuvinte tehnice și altele greu de înțeles (Iași, 1851)*.

3.3 OCR with ABBYY FineReader 16

We trained a customised FineReader model for the RT used in Stamati's dictionary. The training set contains 20 representative pages at roughly 650 characters per page (about 13,000 characters in total). The model uses a letter set that mixes Latin and RT and a vocabulary of more than 10,000 lemmas compiled from an RT dataset² and from documents mentioned in the work (Bumbu et al., 2023). On a held-out set the model achieves a character error rate below 5 per cent. The most frequent systematic confusion is **б** with **в**.

3.4 Post-OCR Verification and Annotation

Post-OCR work was carried out in a dedicated web interface for transcription and annotation developed within the CLRE framework³.

Before uploading the recognised text together with the page images, we added simple inline markers to the text stream to ease import and indexing. Page breaks are marked by ==X==, which indicates the end of page X; text before ==X== belongs to page X, text after it to X+1. Figure 2 illustrates this import format.

```

А фр. кяте, а треї леї ока.
===
Аbate, авъ; пъринте бърънѣ; фъуарникѣ.
===
Абдікацїе, л. абізчере, депѣчере де бѣнъ воіе а љнеї дерегъторїї, а љнеї вреднічїї, а љнеї короне.
===
Абеуедѣ, абеуедарїѣ, алфавитарїѣ, къртічікѣ кѣпінзетоаре де а, б. в.
Абісѣ, гр. аѣнклѣ, пренаетіе.
===
Аблатівѣ, л. љнѣл дін казѣріле граматічеї.
===
Абнормѣ, л. абѣтѣтѣ де ла перѣлѣ, афарѣ де фіре.
===
А б о л е с к ѣ, л. оборѣ. німіческѣ. десфіінѣзѣ, (о леѣе, љнѣ обічеїѣ).
===
Аболіѣе, аболічіѣне, оборѣре, німічі-ре, десфіінѣаре (леѣе).
Абомінабілѣ, л. ѣрѣтѣ.
===
Абонаментѣ, фр. пренѣмераѣе ла љнѣ театрѣ орї ла љнѣ алѣтѣ лѣкрѣ фолосіторїѣ;
abonnement suspendu, абонаментѣ невалабілѣ пентрѣ оѣатѣ.
===
==1==
Абонатѣ, фр.чел че сѣбѣтѣнсеамѣнѣ (іскѣлеѣе), орї респѣнде (плѣтеѣе) ѣнаїнте ла абонаментѣ, пренѣмерѣторїѣ (ла о карте).
===

```

Figure 2: Recognised page import format:
page breaks marked as ==I==, which indicates the end of page 1,
and article boundaries as ===.

² <https://www.kaggle.com/code/mariuspenteliuc/analyze-rt-dataset>

³ <https://sollirom-clre.gitlab.io/texts/stamati1851/admin-site/transcription/>
("Alexandru Philippide" Institute for Romanian Philology – Romanian Academy,)

In this interface the recognised text and the scan are shown side by side. Editors verify each entry and apply semantic markup at segment level: analytical definition, etymology, sense, definition by synonyms, definition by antonyms, examples, variants, norms, and stable expressions. Formatting controls include bold, italics, and letter-spacing. Letter-spacing preserves the wider spacing used in headwords and other salient strings in the original print. All markup is saved to XML using TEI standard, so that the edition can be queried, exported, and aligned with CLRE.

We also implemented a case-sensitive search layer with the following functions: headword lookup in Latin and in RT with automatic transliteration from Latin input to RT; direct page search by number; inline previews with highlighting; automatic opening of the selected entry. These tools reduce navigation time. Figure 3 shows the editing and search interface.




Figure 3: CLRE web interface for transcription, annotation, and search.

3.5 Transliteration by HeDy

For transliteration we use the service provided by the HeDy platform in a web interface similar to that described above. Transliteration is performed after markup, so that annotations are preserved. The transliteration interface mirrors the post-OCR editor: the original Cyrillic text is displayed below the transliterated text, and editors can confirm the result of transliteration entry by entry. The implementation keeps capitalisation and does not break segment boundaries or styles. Figure 4 shows the transliteration view.




Figure 4: Transliteration view: original script and scan on the left, Latin transliteration on the right.

3.6 Data Management

All annotations, offsets, and attributes are stored in XML. We validate load–edit–save integrity and version all exports in GitLab. The prepared bundles include XML files for transcriptions as well as for transliterations, and links to page scans. These bundles support alignment in CLRE and preparation of the print-ready edition with ISBN.

3.7 Font Design

We are developing a manuscript-specific OpenType font to render historical ligatures and diacritics. For example, the sequence *ix* will be rendered as a contextual ligature with a tying line, and *ix* as a separate ligature. The goal is to reproduce the behaviour of the original type while keeping the text layer searchable on the web and in print. Table 1 sketches the target glyphs for some letters and ligatures.

Table 1: Selected Unicode sequences and target OpenType behaviour.

Unicode sequence	Current font	Target font
U+0069	i	ı
U+A64B	ȝ	ȝ
U+A64B U+0306	ȝ̃	ȝ̃

Unicode sequence	Current font	Target font
U+0069 U+A64B	îȚ	İȚ
U+0069 U+A64B U+0306	îȚ	İȚ

4 Conclusion and Future Work

The scientific lexicographical component of the PHILORD project envisages the creation of a critical digital edition of Theodor Stamati's work, *Disionăraș românesc de cuvinte tehnice și altele greu de înțeles (1851)*, a dictionary (containing c. 4200 entries) that was very important for the modernisation and neologisation of the Romanian language (but also for terminological elements) in the 19th century, thus creating a model that can be used to compile digital critical editions of dictionaries in the transitional alphabet of the Romanian language (spoken in Romania or Moldova). We deliver a methodology that starts from scan to critical digital edition. Next steps are to finalize the transliteration verification interface; complete the manuscript-specific font; enrich the search engine; publish the online edition aligned with CLRE; and prepare the print-ready ISBN edition of this dictionary which supposes the next structure:

Foreword; Philological study – cultural, social, etc. contextualization; Linguistic study – phonetics, morphology, lexicon, etymology, etc. Study on the digitization of the dictionary – contextualization (methodologies for the digital transliteration of texts into the transitional alphabet; digitized old dictionaries – LBe, eRomLex), the CLRE method, etc. Note on the edition; List of abbreviations; Text in the transitional alphabet; Text in Latin script (in interpretative transcription); Index of words, forms, and idioms in the transitional alphabet; Index of words, forms, and idioms in Latin script (in interpretative transcription); Bibliography; General bibliography; Bibliography on Stamati; Facsimiles; Abstracts in French, English, German.

Although this report documents one dictionary (Stamati 1851), the pipeline can be adapted to other dictionaries and prints from 1830–1860.

AI Use Statement

AI tools were used for limited editorial assistance, including grammar and clarity improvements, preliminary summaries of literature abstracts. All AI-assisted suggestions were reviewed, corrected, and validated by the authors. No AI tools were used to generate original research ideas, data, analyses, or interpretations.

Acknowledgements

This work was supported by a grant from the Romanian Ministry of Research, Innovation and Digitalisation, UEFISCDI, Strategii de convergență epistemică și

metodologică în cercetarea filologică din România și Republica Moldova [Strategies for epistemic and methodological convergence in philological research in Romania and the Republic of Moldova] (PHILORD), code PN-IV-P8-8.3-ROMD-2023-0318; 2024–2026; project director: dr. Florin-Teodor Olariu.

References

- Bumbu, T., Burțeva, L., Cojocaru, S., Colesnicov, A., Malahov, L. 2023. Distinctive features of recognition for documents printed in the Romanian transitional alphabets. *Computer Science Journal of Moldova*, 31(3), 340–350. <https://doi.org/10.56415/csjm.v31i.17>
- CLRE – Corpus lexicografic românesc electronic. <https://clre.solirom.ro/>
- Cojocaru, S., Burtseva, L., Ciubotaru, C., Colesnicov, A., Demidova, V., Malahov, L., Petic, M., Bumbu, T., Ungur, Ș. 2016. On Technology for Digitization of Romanian Historical Heritage Printed in the Cyrillic Script. In *Proceedings of the Conference on Mathematical Foundations of Informatics (MFOI 2016)*, Chișinău, 160–167.
- Cristea, D., Haja, G., Tamba, E. I., Tufiș, D. 2024. Digitalization of Cultures. Preparations for the Adoption of a Strategy for the Romanian Language. In *Proceedings of the 18th International Conference on Linguistic Resources and Tools for Natural Language Processing*, Brașov, 11–14 December 2023, Iași, Editura Universității „Alexandru Ioan Cuza” din Iași, 23–40.
- Demidova, V., Burțeva, L. 2017. *Digitizarea și vizualizarea alfabetului chirilic românesc de tranziție (1830–1860)*. Akademos, 1(44), 24–29.
- Frincu, M., Frincu, S., Penteliuc, M. E. 2023. Challenges and Solutions in Transliterating 19th-Century Romanian Texts from the Transitional to the Latin Script. In *Proceedings of LDK 2023*, Vienna, 226–231, NOVA CLUNL.
- Haja, G., Tamba, E., Teodorescu, C. 2023. L'étape actuelle dans la numérisation de la recherche linguistique / lexicographique de l'Académie Roumaine, In Dolores Corbella, Josefa Dorta et Rafael Padrón (eds.). *Perspectives de recherche en linguistique et philologie romanes, textes choisis par la Société de linguistique romane*, Strasbourg, Éditions de linguistique et philologie/Société de Linguistique Romane, p. 1327–1338.
- Seche, M. 1966. *Schiță de istorie a lexicografiei române*, vol. I: *De la origini până la 1880*. București: Editura Științifică.
- Stamati, T. 1851. *Disionăraș românesc de cuvinte tehnice și altele grele de înțeles*. Iași: Tipografia Buciumului Român.
- Tamba, E. I. 2022. La lexicografia Rumana. Historia y Actualidad, in *Lexicografia de las lenguas románicas. Perspectiva histórica*. Vol. I, Edited by Félix Córdoba Rodríguez, Ernesto González Seoane, María Dolores Sánchez Palomino, De Gruyter, p. 265–282.
- Tamba, E. I. 2024. Le Processus de Numérisation de la Lexicographie en Roumanie: Présent et Perspectives, in „Lexikos” (AFRILEX Series), 2024, 34(1), 1-20; <https://lexikos.journals.ac.za/pub/article/view/1857>; DOI: <https://doi.org/10.5788/34-1-1857>.
- Tamba, Elena Isabelle. 2022. The Role of the Institutes of the Romanian Academy in the digitalization process of the linguistic research, In Željko Jozić and Sabine Kirchmeier (eds.). *The Role of National Language Institutions in the Digital Age. Contributions to*

- the EFNIL Conference 2021 in Cavtat*. Budapest: Nyelvtudományi Kukatóközpont, p. 91-100 (<https://efnil.org/wp-content/uploads/2023/11/EFNIL-Cavtat-Book.pdf>)
- Tamba Dănilă, E., Clim, M.-R., Catană-Spenchiu, A., Patrașcu, M. 2012. The Evolution of the Romanian Digitalized Lexicography. The Essential Romanian Lexicographic Corpus, In R. Vatvedt Fjeld, and J. Matilde Torjusen (eds.) *Proceedings of the 15th EURALEX International Congress*, 7–11.08.2012, Oslo, Press Representrals, UiO; p. 1014–1017; http://www.euralex.org/proceedings-toc/euralex_2012/
- TEI Lex-0: A baseline encoding for lexicographic data. <https://dariah-eric.github.io/lexicalresources/pages/TEILex0/TEILex0.html>
- Ursu, N. A. 1962. *Formarea terminologiei științifice românești*. București: Editura Științifică.
- Variot, Estelle. 2009. Lexicul, martorul contactelor între popoare: Disionărașul românesc de cuvinte tehnice și altele greu de înțeles al lui Teodor Stamati (Le lexique, témoin des contacts entre les peuples: Le petit dictionnaire de termes techniques et d'autres difficiles à comprendre de Teodor Stamati), dans G.G. Neamțu, Ștefan Gencărașu, Adrian Chircu (éds.) *Actes du colloque « Limba română – abordări tradiționale și moderne » (La langue roumaine – approches traditionnelles et modernes)*, 19 - 20 October 2007, Cluj-Napoca, Presa Universitară Clujeană, 2009, p. 275-283.
- Variot, Estelle. 1997, Un moment significatif de l'influence française sur la langue roumaine: Le dictionnaire de Teodor Stamati, Iassy, 1851. Thèse de doctorat, <http://diffusiontheses.fr/22635-these-de-variote-estelle.html>.
- Vulpe, A., Tamba, E. I., Botnari, L., Caruntu-Caraman, L. 2025. Din istoria lexicografiei explicative românești: Teodor Stamati, „Disionăraș românesc de cuvinte tehnice și altele greu de înțeles (1851)” – recuperări, reconsiderări, in *Lucrările Conferinței Științifice Naționale „Lecturi în memoriam acad. Silviu Berejan”* (ed. a VII-a), coord.: Ecaterina Brăguță, Chișinău, CEP USM, Institutul de Filologie Română „Bogdan Petriceicu-Hasdeu” al USM, 79–93.

EDTLR – A NEW BEGINNING

PETRU REBEJA¹, DAN CRISTEA^{1,2}

¹ *Institute of Computer Science, Iași branch of the Romanian Academy*
petru.rebeja@gmail.com

² *Faculty of Computer Science, "Alexandru Ioan Cuza" University of Iași*
danu.cristea@gmail.com

Abstract

This paper introduces a new software system that represents the next step in the evolution of the eDTLR project (the Digital Thesaurus Dictionary of the Romanian Language). While the new system aims to fully preserve the functionality of the previous systems, the novelty of the new design is to split the system into three separate components: (i) an application for annotating the entries of the dictionary, (ii) an application for browsing the dictionary entries, and (iii) an application for structuring the dictionary entries. This new architecture is robust and requires low maintenance while catering to the specific needs of the three groups of users: annotators, linguists, and wide public.

Keywords: Digital Romanian Thesaurus Dictionary, modular architecture, system design, linguistic annotation

1 Introduction

The *Digital Thesaurus Dictionary of the Romanian Language* shortened as *eDTLR* — *Dicționarul de Tezaur al Limbii Române* în format electronic is a project that has been in the making for numerous years, and is yet to be completed. During its history, the project has had to overcome a multitude of challenges ranging from handling intellectual property rights to technical issues.

Over the years, several attempts at moving the eDTLR project closer to its finish line have been made, out of which two major achievements are noteworthy: (i) a processing chain that works on the OCR results of the dictionary pages (Curteanu *et al.*, 2008; Curteanu *et al.*, 2009) and manages to determine the structure of dictionary entries, and (ii) an interface that allows searching dictionary entries, and viewing the scans of the dictionary page where the entry is located (Bursuc *et al.*, 2021; Bursuc *et al.*, 2022).

However, although the aforementioned attempts have achieved great results, they are still facing difficulties: the CLRE interface (Bursuc *et al.*, 2021; Bursuc *et al.*, 2022) does not allow copying the text of the dictionary entries, and the processing chain (Curteanu *et al.*, 2008; Curteanu *et al.*, 2009) suffers due to lack of quality data. This paper proposes a technical solution that combines the

strengths of the previous attempts at building the eDTLR project. The new solution contains a dedicated module that allows annotating the dictionary entries before determining their structure, and based on the determined structure, display the entries through a Web interface that will cater to the needs of both the researchers (linguists and others), and the wide public.

2 Implementation

This section provides detailed information about the implementation of the new eDTLR system: the Section 2.1 provides a birds' eye view of the solution, the Section 2.2 contains details about why some technical decisions were made during the design phase, and Sections 2.3 through 2.5 describe each of the components of the system.

2.1 Solution overview

As Figure 1 shows, the new system consists of three main components, namely:

1. An application for collaborative correction and annotation of the entries of the DA (Academia Română, 2021 2022) and DLR (Academia Română, 2022) corpora which were created by passing the page scans through an OCR¹ engine — *eDTLR Annotator*. The role of this application is, as the name implies, to facilitate the correction and annotation of the texts resulted from the OCR process.
2. A processing chain — *eDTLR Structurizer* — that is responsible for parsing both the annotated entries, and the entries that have already been digitised, determine the structure of each entry, and save the structure into an XML file in TEI (TEI Consortium, 2007) format.
3. An application for browsing the entries of eDTLR — *eDTLR Browser*, which is responsible for parsing the TEI entries created by the *eDTLR Structurizer* application, importing them into the database, and displaying the entries to the user through a Web interface.

¹ Optical Character Recognition




Figure 1: The final data flow diagram of the new eDTLR system. The page scans of the printed editions are processed by an OCR engine, then the images alongside the texts from the OCR are imported into the eDTLR Annotator. The eDTLR Structurizer then transforms either the output of the eDTLR Annotator, or the already digitised entries of DLR into files in TEI format. Then, the TEI files are imported into the eDTLR Browser application. The dashed arrow from the page scans to the OCR engine represents an independent processing that is not part of the system.

Sections 2.3 — 2.5 describe each of the components listed above in detail; the code of the systems presented herein is available on Github².

2.2 Technical considerations

The design phase of every system is where the fundamental decisions regarding the development and evolution of the said system are taken. Of course, such decisions are not taken lightly; the multitude of factors that affect the design decisions are thoroughly scrutinised before taking a specific decision.

The system presented herein underwent the same process. The design phase was mainly concerned with deciding the software stack of the system, its overall architecture and data flow (presented in Figure 1), performance, and user base. The technical considerations that resulted from the design phase are presented in the remainder of this section.

Software stack. To pick the software stack for the new eDTLR system, the following factors were considered: cost, productivity, and experience. Whereas all the modern software frameworks provide a high degree of productivity, the cost of development and running the system is different from one framework to another. As a result of this factor, all the frameworks and software products that were not free were excluded from the candidate list. Lastly, the next factor applied to the choice of the software stack was the degree of experience each member of the

² <https://github.com/arfi-iit/>

development team had with each piece of technology, which favoured the software systems that were already used prior by the entire team, or by individual members. As a result of this selection process, the following software stack was chosen: (i) CentOS Stream(Red Hat, Inc., 2025) as the operating system, (ii) Python(Van Rossum and Drake, 2009) as the development language of the application back-ends, and for the processing of the eDTLR entries, (iii) Django(Django Software Foundation, 2025) as the development framework for Web applications, (iv) plain Javascript(Eich, 1995) for enriching the user experience of the Web applications, (v) the Codemirror (Haverbeke, 2025) editor for editing eDTLR entries, and (vi) PostgreSQL(PostgreSQL Development Team, 2023) as the database engine.

Modularity. Although the entire eDTLR system can be developed as a single software pack-age, the components presented in the previous section will be developed and deployed as separate applications due to the following considerations: (i) by default the *eDTLR An-notator* application will have a limited number of users which means lower demand on computational resources, (ii) in contrast, the *eDTLR Browser* will have a larger user base, and subsequently a high demand for resources, and (iii) given their purpose the applications will evolve independently based on user feedback.

Inter-application communication. As a consequence of the decision to isolate the applications no direct communication channel (e. g.: HTTP Endpoints) will be available. The reasons for this decision are twofold: security considerations, and implementation and maintenance cost. However, since the applications need to interchange data the data exchange will be implemented through file transfers where each file adheres to a structure that is well-known for both the sender, and the receiver applications. Given that one of the objectives of the eDTLR Structurizer application is to export data in TEI (TEI Consortium, 2007) format, which is an XML file format, then it makes sense to use the XML format for inter-application communication. Furthermore, the XML format bears the advantages that for the specific case of dictionary entries it is easier to read than more compact formats like JSON.

User experience. In order to provide a great user experience, the user-facing applications (eDTLR Annotator, and eDTLR Browser) will use well-known idioms to handle user inter-action. Furthermore, the eDTLR Browser application will offer a simplified view of the dictionary entries for the large public, and a complex view for linguists that will incorporate a search by specific criteria. The point of reference for the interface design of the eDTLR Browser application will be the electronic version of the Spelling, Pronunciation, and Morphology Dictionary of the Romanian Language (Institutul de Lingvistică „Iorgu Iordan –

Alexandru Rosetti”,); this will ensure a uniform interface across all the electronic dictionaries of the Romanian Academy.

2.3 The eDTLR Annotator Application

From the technical point of view, the *eDTLR Annotator* application is responsible for the following aspects: (i) import of the page scans and OCR results into the database, (ii) facilitate the annotation of dictionary entries, (iii) implement a workflow through which to pass all the annotations, and (iv) export the data for the *eDTLR Structurizer* application (Section 2.4).

Data import

The import of the data into the *eDTLR Annotator* application is done through a custom Django management command³. This command, named suggestively `import-data`, receives as command-line arguments: (i) the path of the directory containing the files with the dictionary entries that need to be annotated, (ii) the path of the directory containing the page scans of a specific volume of the dictionary, (iii) the path of a file that maps each dictionary entry to the pages (i. e. page scans) of the imported volume, and (iv) the path of the directory from which the *eDTLR Annotator* application serves its static files. When the command is started, it executes the Algorithm 1 which involves iterating through the entries to be imported, checking if they are mapped to the dictionary pages they appear in, replacing the markup, and inserting or updating the database data.

Algorithm 1 Import of dictionary page scans and OCR results into the *eDTLR Annotator* application

1: **Input parameters:**

2: `entries_directory`: directory containing dictionary entries

3: `scans_directory`: directory containing page scans

4: `mappings_file`: file containing entry-page mappings

5: `static_directory`: directory used to host Web application static files

6: **Begin**

7: search the `input – directory` for files with the `.xml` extension

8: **if** no files found **then**

9: `exit`

10: **end if**

11: copy images from `scans_directory` into `static_directory`

12: **for** each `file` with the `.xml` extension **do**

13: read the contents of the file into `text`

³ <https://docs.djangoproject.com/en/5.2/howto/custom-management-commands/#command-objects>

```

14:     replace XML tags in text with custom markup symbols
15:     extract title_word from text
16:     if title_word not in mappings_file then
17:         print error message
18:     else
19:         insert or update title_word and text in the database
20:     end if
21: end for
22: End

```

The data import procedure needs to be executed for each volume of the dictionary that needs to have its entries annotated. The reason for this constraint is the repetition of page numbers in each volume of the dictionary. During the import, the application maps a page of the volume to the image file that contains the scan of the page in order to be displayed in the application interface. Since the only way to identify a page is by its page number, the pages need to be grouped by volume; otherwise, the application may display an entry from one volume of the dictionary alongside the image of a page from another volume, which renders the annotation impossible.

However, if the import procedure is executed multiple times for the same volume, this will only result in updating the text of the existing entries and inserting entries that were not present before. Likewise, the page scans will be replaced with the images from the import directories.

Application interface

The interface of the *eDTLR Annotator* application is separated into two areas: the admin area⁴, and the area for entry annotation. The admin area is provided out of the box by the Django framework, and offers an uniform way of administering the data of the application, e. g.: adding/removing user accounts, making small edits to the dictionary entries etc. functionality which, while important does not present significant interest in regards to annotation of dictionary entries; as such, the remainder of this section will resort to presenting the area for entry annotation.

⁴ <https://docs.djangoproject.com/en/5.2/ref/contrib/admin/>

After the annotator user logs into the application, he/she is presented with the homepage shown in Figure 2. The homepage provides an overview of the annotation efforts of the user. It displays, in the top centre region, a list of annotations that have been assigned to the user alongside the status of each annotation, such that, by clicking on an annotation from the list that does not have the status Complete, the user will be directed to the annotation editor where he/she will be able to continue the annotation process from where it was left off.




Figure 2: A snapshot of the *eDTLR Annotator* home page. The page displays the list of the entries assigned for annotation to the current user, annotation statistics, and, at the bottom, the button for requesting to annotate another entry.

On the top-right part of the homepage, the user is presented with some statistics about his/her annotations: the total number of annotations (excluding the ones with the status In progress), and a total sum of the symbols from each annotation. Below these total counts, the user is presented with a distribution of the annotations per status.

Last but not least, at the bottom of the homepage is located the button that requests to assign the user another entry to annotate. Whether the user clicks on the button to request a new annotation, or he/she clicks on an annotation from the list of already assigned annotations, the application will navigate to the annotation editor page shown in Figure 3.

The annotation editor is the page that encapsulates the workflow needed to complete the annotation of an entry which is presented in Section 2.3.3. This page is split into two horizontal panes that allow viewing the dictionary page(s) containing the entry while annotating its text.

The left pane displays the page of the dictionary where the entry is located. If the entry spans on multiple pages, the user can navigate back and forth between the pages where the entry is displayed by clicking on the navigation buttons that are displayed when hovering with the cursor over the page image. Depending on the resolution of the screen, the user may choose to zoom into the image to avoid eye strain. He/she can do so by using the zoom bar displayed at the top of the page. After zooming in, the user is also able to adjust the area that is being displayed by scrolling the image both horizontally, and vertically.

The right pane contains the annotation editor, and at the bottom two buttons which, in left to right order, allow the user to request marking the annotation as complete, and respectively to save the work in progress. The annotation editor is a CodeMirror (Haverbeke, 2025) that was adapted to use a custom Markdown syntax (Gruber, 2004) for annotating entries. The editor displays a panel that contains buttons for all custom annotations that are possible on the entry text. As shown in Figure 3, from left to right these operations are: (i) mark text as bold, (ii) mark text as italic, (iii) mark text as superscript, (iv) mark text as subscript, and (v) mark text as acronym.

Annotation flow

The annotation workflow of a dictionary entry starts from an explicit user request. To submit such a request, the user needs to click a button in the homepage of the *eDTLR Annotator* application displayed in Figure 2. Once the request is submitted, the application proceeds to select an entry to annotate.

The selection procedure tries to reduce the time it takes for an entry to become fully annotated (i. e. to be exported to the *eDTLR Browser* application). As such, when the user requests to annotate a new entry, the application will first look for entries which either are already annotated, or are being annotated by a different annotator than the current user. If such entries exist, the user gets assigned the first one from the collection. If no such entries exist, the application will try to pick the first entry that is not assigned to any user; otherwise, it will display a message that there are no more entries to annotate.

Once an entry is selected, the user gets a copy of the entry to annotate; each such copy is named an Annotation in the system, and it is saved in a separate database table from the entries. Having separate Annotation objects for the same dictionary entry allows the users to work independently of each other.

At this moment, the user can start to annotate the entry. Depending on the dictionary entry, the annotation process can be as easy as just stating that the entry is correct, or as complex as rewriting the entire entry. In either case, the user can choose at any time to save his/her work in progress without overwriting the work of the other annotator; every time the user saves an annotation the annotation will keep the status it was assigned since it was created — In progress.

While working on an annotation, the user may encounter a situation that needs to be discussed with an expert. In such case, the user will save the annotation and requests to annotate another entry until he/she receives the answer from the expert.

When the user has finished annotating the entry, he/she will request to mark the annotation as complete. When such request is submitted, the application will determine the actual status of the annotation based on how many annotations the entry has, and what are the contents of those annotations. If the entry has no other annotations, then the annotation will be assigned the status Complete. Otherwise, the contents of the entries are compared pairwise, and if a mismatch is detected, all of the annotations will be assigned the status Conflict; if all the annotations have the same text, then they will receive the status Complete.

Data export

The export of the data from the *eDTLR Annotator* application is also done through a custom Django management command named `export-complete-annotations`⁵. As the name suggests, the command exports only the entries that have the status Complete, and are annotated by two annotators.




Figure 3: A snapshot of the *eDTLR Annotator* annotation editor. The left pane of the application displays the page scan of the dictionary zoomed at 1.4× to avoid eye strain of the annotator, and centred on the entry being annotated. The right pane contains the entry editor with the action buttons, the text of the entry, and on the bottom — buttons for marking the annotation complete (on the left), and for saving the work in progress (on the right).

⁵ The actual name is `exportcompleteannotations` — same as in text but without dashes; the lack of dashes in the name is dictated by the naming conventions used in the development of the application.

The logic of the export procedure is trivial: it loads from the database entries with two annotators, and with status Complete in batches of 1000 items, and saves the text of each item into an XML file.

Table 1: Overview of the data that is currently stored into the database of the *eDTLR Annotator* application.

Total entries	15783
Complete entries	1969
In progress entries	274
Conflicting annotations	5604
Not annotated entries	13814

As Table 1 shows, currently there are less than 2000 complete entries in the database of the *eDTLR Annotator* application. As such, the execution time of the data export command is insignificant; however, as the number of complete annotations grows, the total execution time of the export procedure is also expected to grow proportionally, but the memory footprint of the export procedure should have only small variations given that the entries are exported in batches.

2.4 The eDTLR Structurizer Application

The main work of the *eDTLR Structurizer* application is to port the existing system from (Curteanu *et al.*, 2008; Curteanu *et al.*, 2009), which was developed in Java to Python. Porting the entry processing logic from Java to Python is a direct consequence of the chosen software stack mentioned in Section 2.2. Although the process of porting the code from Java to Python requires a significant amount of development effort, it has some inherent advantages. The main advantage of undergoing through this effort is that rewriting the code in another programming language while keeping the processing logic means gaining a deeper knowledge and understanding of the system which will result, in the future, in lower maintenance and debugging effort.

However, although the main objective for the *eDTLR Structurizer* application is to be ported from Java to Python, the processing logic needs to undergo some adjustments. A handful of these adjustments are concerned to accommodating the code to the mixed programming paradigm supported by Python.

Another important change to the *eDTLR Structurizer* application is the introduction of a discriminator module at the beginning of the processing flow. The discriminator module is responsible for determining the type of entry being processed (e. g.: an entry that references another entry, an entry with just one meaning etc.), and selecting the proper processing strategy.

2.5 The eDTLR Browser Application

The main responsibilities of the *eDTLR Browser* application are to: (i) allow the users to search for specific dictionary entries, (ii) display the entries that are stored in the application database, and (iii) import the data into the application database in order to update existing entries, or to increase the number of stored entries.

Data import

As mentioned in Section 2.2, the application does not expose any endpoints where the data to be imported can be directed to. As a consequence of this decision, the data is imported into the database based on a cron⁶ schedule.

As with the *eDTLR Annotator* application described in Section 2.3, the logic of the data import mechanism is encapsulated in a custom Django management command class that executes the Algorithm 2. The value of the *input – directory* parameter is specified as a command-line argument.

Algorithm 2 Import of dictionary entries into the *eDTLR Browser* application

```

1: Input parameters:
2:   input_directory: directory containing input files
3: Begin
4: search the input_directory for files with the .xml extension
5: if no files found then
6:   exit
7: end if
8: for each file with the .xml extension do
9:   convert the contents of file into internal representation
10:  read the internal id of the entry
11:  from file contents if entry with
12:  same id exists in the database then
13:    update the entry
14:  else
15:    insert the entry
16:  end if
17: end for
18: End

```

The chosen interval for triggering the data import mechanism is 30 minutes. The reason for choosing this value was that it should allow importing all of the data in a single run. However, this value is far from ideal, and makes an annotator

⁶ cron is a job scheduler in Unix-like operating systems which allows execution of jobs at specified intervals of time.

wait for half an hour to see an entry marked with status Complete to be imported into the *eDTLR Browser* application. The optimal value for the import interval will be identified empirically at a later stage in the project.

Application interface

The interface of the *eDTLR Browser* application has two main responsibilities: search, and display. The homepage of the application presents the user with a search box which allows searching for a specific entry by the title-word of the entry. The text entered by the user in the search box is then transformed into a canonical form, and only then it is matched against existing entries in the database. The interface then returns either a list of entries that match the search string, or displays a message that no entries are found. Upon clicking on a result from the list of matching entries, the application displays the whole entry in a separate page.

In near future, this rudimentary search mechanism will be expanded to allow searching entries by complex criteria. In this way, the *eDTLR Browser* application will provide an enhanced experience for the expert users (linguists and researchers) while also keeping a simple search for the wide public.

Application server

While in development, the *eDTLR Browser* application is not yet available to the large public but rather to a small group of people for testing and validation purposes. Given the small number of users thus small resource requirements, the *eDTLR Browser* application is hosted on the same server as the *eDTLR Annotator* application.

Once made available to the general public, the *eDTLR Browser* application will be hosted on a dedicated server to handle a high volume of user requests efficiently.

3 Future work

At the time of writing, due to time and various other types of constraints, the eDTLR system is far from being feature-complete. However, looking at the backlog of the project, the remaining work on the eDTLR system can be split into two categories: (i) work in progress that is not yet complete, and (ii) work planned for future.

3.1 Work in progress

Out of the multiple in-progress tasks that are on the board of the project⁷, the most important one is the implementation of a conflict resolution workflow. This

⁷ Out of precaution to avoid exposing nonpublic data, the board of the project is not made public.

feature, shown in its current implementation state in Figure 4, is as important for the annotators as it is for the quality of the data resulted from the annotations. Once finished, the conflict resolution will allow annotators to reach the inter-annotator agreement faster, and most importantly *without the mediation of an expert user*, all while ensuring a proper user-experience, avoiding errors due to annotators' fatigue, and increasing the productivity (and implicitly decreasing the overall time that it takes for an entry to be fully annotated). However, the implementation of the conflict resolution feature requires a complete rewrite of the annotation editor for two main reasons: (i) the library that sits at the core of the annotation editor (Haverbeke, 2025) needs to be updated to the latest version, and (ii) the custom markup language used in the annotation editor needs to have its underlying grammar changed.

1	**Z** s. m. invar. A treizeci și una ¹ , respectiv ultima 2. literă a	1	**Z** s. m. invar. A treizeci și una ¹ , respectiv ultima 2. literă a
2	alfabetului limbii române, și sunetul corespunzător (consoană	2	alfabetului limbii române, și sunetul corespunzător (consoană
3	fricativă, dentală, sonoră). @cf. DORF, ȘĂINEANU ² , RESMERIȚĂ, D.,	3	fricativă, dentală, sonoră). Cf. DORF, ȘĂINEANU ² , RESMERIȚĂ, D.,
4	CADE@ *Sunt risipiți – pe toată întinderea alfabetului, de la „a”	4	CADE@ *Sunt risipiți – pe toată întinderea alfabetului, de la „a”
5	până la „zed”. * @VINEA, L. I, 137, cf. DL, DM, DER.@ *A fost dezolat	5	până la „zed”. * @VINEA, L. I, 137, cf. DL, DM, DER.@ *A fost dezolat
6	intr-o zi când i se atrăsese atenția despre o altă literă că este	6	intr-o zi când i se atrăsese atenția despre o altă literă că este
7	„zet”. * (ROMANIA LITERARĂ, 1978, nr. 72, 183, cf. M. D. ENC., DEX.@	7	„zet”. * (ROMANIA LITERARĂ, 1978, nr. 72, 183, cf. M. D. ENC., DEX.@
8	**Ț** \$loc. adj. și adv.\$ **De la A la Z(ied)** (care este) de la un	8	**Ț** \$loc. adj. și adv.\$ **De la A la Z(ied)** (care este) de la un
9	capăt la altul. @cf. DEX.@	9	capăt la altul. @cf. DEX.@
10	- Pronunțat: *ze* și (în expresii) *zet*, (ieșit din uz) *zed.*	10	- Pronunțat: *ze* și (în expresii) *zet*, (ieșit din uz) *zed.*
11		11	

Figure 4: A snapshot of the new annotation editor in the conflict resolution mode. The two panes of the editor contain the conflicting annotations of the same entry. Each pane indicates on the fringe where the conflict occurs, and underline the conflicting characters.

While the update of the library does not bear any urgency, changing the grammar for the annotation markup is of high importance because not only will it provide a better user experience but will effectively prevent the annotators from inserting erroneous markup symbols by providing proper visual cues in the text.

3.2 Planned work

In order to shorten the time it takes for the complete entries to be imported into the eDTLR Browser application, the implementation of the export and import of data to and from XML files in TEI format was postponed in favour of using a temporary XML structure which only caters to import into the eDTLR Browser application. Of course, the authors are aware that using yet another custom format instead of a standard structure hinders the data interchange, which in itself has a negative impact on the research community. To alleviate this downside, the authors are committed to implement the data interchange between the eDTLR Annotator, and the eDTLR Browser applications using the TEI format, as soon as the conflict resolution feature is finished. Furthermore, as mentioned in Section 2.5.2 the eDTLR Browser application is not yet feature-complete. As such, raising the application to a level of maturity that is suitable for the large public is the next direction of the planned work.

4 Conclusion

This paper presents the technical solution designed to replace the existing implementations of the eDTLR project. In order to provide high degree of maintainability, the system is split into three parts that can evolve independently of each other: an application for annotating the dictionary entries, an application for displaying the dictionary entries, and an application for structuring the dictionary entries in a well-known format. Since the applications are independent of each other, they communicate the data through TEI encoded files. Communicating using files in TEI format bears a significant advantage — such files can be published to multiple research communities.

The user-facing applications, namely *eDTLR Annotator*, and *eDTLR Browser* were designed to provide the best possible user-experience for their target demographics. The *eDTLR Annotator* application is intended to be used by linguists, and as such implements a workflow for annotating dictionary entries. On the other hand, the *eDTLR Browser* application targets both expert users (linguists), and the wide public.

Out of the three aforementioned applications, only *eDTLR Annotator* is near being feature-complete; the other two — *eDTLR Structurizer*, and *eDTLR Browser* are still under development.

AI Use Statement

No generative artificial intelligence tools were utilized at any stage in the preparation, drafting, or finalization of this manuscript.

Acknowledgements

The authors would like to acknowledge the people that contributed to the system presented herein. For providing the data that fed the *eDTLR Annotator* application, the authors are grateful to: First Degree Scientific Researcher, PhD Gabriela Haja, Senior Researcher, PhD Elena Isabelle Tamba, and Lecturer, PhD Alex Moruz. The authors also express their gratitude for the feedback and suggestions to: Scientific Researcher, PhD Alexandra Ilie, Scientific Researcher Cecilia Bolea, and Research Assistant Laura Pistol.

Last but not least, the authors would like to thank the annotators for their dedication and patience throughout the annotation process. Special thanks also go to Cecilia Bolea, Laura Pistol, and Alexandra Ilie for supervising and guiding the annotators.

References

- Academia Română. 2021–2022. DA în CLRE: Academia Română, dicționarul limbii române [DA] (1913–1949), ediție digitală anastatică în CLRE. Corpus lexicografic românesc electronic. <https://dlr1.solirom.ro/>. [Online; accessed 2025-10-01].
- Academia Română. 2022. DLR în CLRE: Academia Română, dicționarul limbii române [DLR] (1965–2010), ediție digitală anastatică în CLRE. Corpus lexicografic românesc electronic. <https://dlr1.solirom.ro/>. [Online; accessed 2025-10-01].
- Alina-Mihaela Bursuc, Marius-Radu Clim, Mioara Dragomir, Cristina Florescu, Gabriela Haja, Laura Manea, Mariana Nastasia, Elena Isabelle Tamba, and Claudius-Marian Teodorescu. 2021. Clre. Corpus lexicografic românesc electronic, realizat în cadrul Departamentului de lexicologie și lexicografie de la Institutul de Filologie Română “Alexandru Philippide” al Academiei Române - Filiala Iași. <https://clre.solirom.ro/>.
- Alina-Mihaela Bursuc, Marius-Radu Clim, Mioara Dragomir, Cristina Florescu, Gabriela Haja, Laura Manea, Elena Isabelle Tamba, and Claudius Teodorescu. 2022. Lansarea primului volum din corpus lexicografic românesc electronic (CLRE): Da. literele a–b. In Ofelia Ichim, Alina-Mihaela Bursuc, Marius-Radu Clim, Ofelia Ichim, Maricica Munteanu, Veronica Olariu, Ana-Maria Prisacaru, and Ioana Repciuc, editors, *Provocări trecute și prezente în evoluția limbii, literaturii și culturii române*, pages 15–34. Editura Universității de Vest, Timișoara.
- Brendan Eich. 1995. Javascript programming language. Online. Initial release by Netscape.
- Django Software Foundation. 2025. Django (version 5.2).
- Guido Van Rossum and Fred L. Drake. 2009. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA.
- Institutul de Lingvistică „Iorgu Iordan – Alexandru Rosetti”. Dicționarul ortografic, ortoepic și morfologic al limbii române, ediția a iii-a revăzută și adăugită. online version <https://doom.lingv.ro/>.
- John Gruber. 2004. Markdown: Syntax documentation. <https://daringfireball.net/projects/markdown/>. Original specification of the Markdown language.
- Marijn Haverbeke. 2025. Codemirror (version 6).
- N. Curteanu, A. Moruz, D. Trandabăț, C. Bolea, M. Spătaru, and M. Husarciuc. 2009. Sense tree parsing and definition segmentation in edtlr thesaurus. In D. Trandabăț, D. Tufiș, and D. Cristea, editors, *Proceedings of the Workshop "Linguistic Resources and Instruments for Romanian Language Processing"*, pages 65–74, Iași, Romania, November. "Al.I.Cuza" University Publishing House. In Romanian. Workshop held November 19–21, 2008.
- N. Curteanu, D. Trandabăț, A. Moruz, G. Pavel, and C. Vereștiuc. 2008. Dictionary sense segmentation & dependency – a flexible strategy for thesauri parsing. In *Proceedings of the 5th European Conference on Intelligent Systems and Technologies (ECIT 2008)*. CD edition.
- PostgreSQL Development Team. 2023. PostgreSQL (version 16).
- Red Hat, Inc. 2025. Centos stream. Rolling-release Linux distribution.
- TEI Consortium, 2007. *TEI P5: Guidelines for Electronic Text Encoding and Interchange*. Text Encoding Initiative Consortium, latest release edition. Version P5, updated regularly.

ENHANCING THE PARSEME-RO CORPUS WITH NOMINAL, MODIFIER AND FUNCTIONAL MULTIWORD EXPRESSIONS

VERGINICA BARBU MITITELU¹, IOANA-MARIA BIOLAN²,
IOANA BUHNILĂ^{3,6}, MIHAELA CRISTESCU², ELENA IRIMIA¹,
CĂTĂLIN MIHĂILĂ¹, ISABELLA ȘINCA², AMALIA TODIRAȘCU⁴,
CARMEN VASILE^{2,5}

¹*Romanian Academy Research Institute for Artificial Intelligence*
{vergi, elena, catalin}@racai.ro

²*University of Bucharest, Faculty of Letters / “Solomon Marcus” Centre for
Computational Linguistics*

ioanabiolan7@gmail.com
mihaela.cristescu@litere.unibuc.ro
isabella.sinca@s.unibuc.ro
carmen.vasile@unibuc.ro

³*Centre for Data Science in Humanities, Chosun University*
ioana.buhnila@chosun.ac.kr

⁴*LiLPa, Université de Strasbourg*
todiras@unistra.fr

⁵*Romanian Academy, “Iorgu Iordan – Al. Rosetti” Institute of Linguistics*

⁶*ATILF, University of Lorraine-CNRS*
ioana.buhnila@univ-lorraine.fr

Abstract

PARSEME-Ro, the Romanian corpus annotated with verbal multiword expressions, has recently been further enhanced to include annotated multiword expressions of all the other parts of speech. The paper presents the methodology we adopted for this endeavour, which consists of an automatic phase based on knowledge extraction from Romanian lexicographic resources and their automatic annotation in the corpus, followed by manual validation of these annotations, as well as consistency check of this manual work. The corpus is developed in a multilingual context, following the same annotation principles for all languages involved, which makes it possible to compare this language phenomenon across languages. Besides constituting a language resource per se, the corpus is also usable in shared tasks concerning identification of multiword expressions.

Keywords: multiword expressions (MWE), Romanian, annotation, corpus

1 Introduction

Multiword expressions (MWE) is a linguistic phenomenon for which there is not a unique definition in the fields of linguistics, in any of its subfields such as applied linguistics, lexicology, or in Natural Language Processing (NLP). Generally speaking, MWEs are sequences of two or several words, sometimes discontinuous, which display some phonetic, lexical, morphological, syntactic, semantic or statistical idiosyncrasies (Baldwin and Kim, 2010). According to the literature and to empirical data, all parts of speech can be represented by both words and MWEs.

The label of MWE might be associated with a broad set of word combinations, including idioms, compounds, named entities, fixed expressions, terms or collocations. Such a large diversity of MWEs is due to the lack of a generally accepted definition of MWEs, to the lack of a unique name for this phenomenon and to the variety of properties such word combinations have. A defining property of idioms is their non-compositional meaning, e.g. *a pune bețe în roate (cuiva)* ‘to put a spoke into [somebody]’s wheel’ has a specific, non-compositional meaning: *a obstrucționa* ‘to obstruct’. Collocations, composed by a base and a collocative, accept morphosyntactic or syntactic variability, the sense is more compositional (*unt ranced* ‘rancid butter’), but the collocative slightly changes its sense. In the PARSEME¹ and UniDive² COST Actions, the term collocation is used to design statistically frequent word associations and are not taken into account, while other projects, such as SimpleApprenant (Todirascu and Cargill, 2019) or PREFABS (Wang and Tutin, 2025) use a different definition: specific lexical association, with compositional sense, sometimes opaque, accepting syntactic transformations. Differences between categories and definitions imply that NLP tools will not identify the same MWE and will not apply the same classification criteria.

MWEs are sources of errors or difficulties for NLP applications and tools, such as machine translation, parsing, text simplification and adaptations, due to their specific properties. Correct identification of MWEs is crucial for basic NLP tools such as tagging, parsing or named entity recognition. Wrong tokenisation might propagate errors among all the NLP annotation steps.

NLP applications have difficulties in identifying MWEs due to their specific properties. A word-for-word translation of an idiom generally fails, due to non-compositional meaning. Sometimes the MWEs are ambiguous: they can have a literal meaning, alongside the non-compositional one, in different contexts: *a închide ochii* ‘to close the eyes’ (non-compositional meaning: *a mușamaliza* ‘to hide’ in example 1), but in a specific context the literal meaning is possible (example 2):

¹ <https://typo.uni-konstanz.de/parseme/>

² <https://unidive.lisn.upsaclay.fr/>

- (1) A închis ochii la aranjamentele colegilor.
 ‘He closed his eyes to his colleagues’ arrangements.’
 “He pretended not to be aware of his colleagues’ arrangements.”
- (2) A închis ochii repede din cauza luminii soarelui.
 ‘He closed his eyes quickly because of the sunlight.’
 “He quickly closed his eyes because of the sunlight.”

Syntactic variants of MWEs might be difficult to identify by NLP applications, due to modifiers included between the words. This assertion is particularly true for verbal MWE. For these reasons, it is difficult to annotate discontinuous MWE, due to long-distance dependencies. Moreover, odd syntactic combinations form functional or modifiers MWE, but are difficult to predict following syntactic constraints: *în timp ce* “while” (functional MWE) is composed of one preposition (*în* “in”), a noun (*timp* “time”) and a pronoun (*ce* “what”). Morphosyntactic rules are not always predictable: for example, the noun *punct* “dot” in the MWE *a pune la punct* “to set up” cannot occur with either a definite or an indefinite article: **a pune la punctul* ‘to put in point-the’, **a pune la un punct* ‘to put in a point’.

For this study, we annotated MWEs using the PARSEME guidelines (Savary *et al.*, 2018). PARSEME is an international community that is developing and maintaining a multilingual corpus annotated with MWEs in those respective languages. This research community originated from the PARSEME COST Action (2013-2017). The project is being developed further within the UniDive COST Action (2022-2026), exploring the universality, diversity and idiosyncrasy in language technology, across multiple languages.

The objective of our endeavour was to augment the PARSEME corpus focusing on nominal, functional and modifier MWEs for the Romanian language. The same procedure as previously developed for verbal MWEs was followed, using a decision-tree structure to establish whether a group of words identified in the corpora constituted a MWE.

The paper is structured as follows: Section 2 contains the types of MWEs annotated in the PARSEME-Ro corpus. Section 3 presents related work with respect to the context in which we have carried out our annotation. Section 4 describes our work methodology: the corpus we have worked with, the extraction of MWEs from several Romanian lexicographic resources, the curation of this list of MWEs, the automatic annotation of the curated list of MWEs in the corpus, its manual validation and consistency check. Section 5 offers some quantitative data that illustrates the scale of our effort, after which we conclude the paper.

2 Types of MWEs in PARSEME-Ro

The types of MWEs that have been annotated in our corpus (previously and in the work reported here) are: verbal, nominal, adjectival, adverbial, functional

MWEs. **Verbal MWEs** include various types of MWEs with verbal behaviour. **Verbal idioms (VID)** are fixed and conventionalized expressions whose global meaning cannot be derived from the literal meanings of the words, such as *a pune bețe în roate* ‘to put sticks in the wheels’ meaning “to obstruct”. **Inherently reflexive verbs (IRV)** are verbal expressions that obligatorily occur with a reflexive marker, for example, *a se teme* ‘to fear oneself’ meaning “to be afraid”. **Light verb constructions** are combinations of a semantically light verb with a predicative noun. In their full version (**LVC.full**), the verb loses its original meaning, as in *a lua o decizie* ‘to take a decision’, meaning “to make a decision”. In the causative subtype (**LVC.cause**), the verb contributes a causative sense, as in *a pune în mișcare* ‘to put in motion’, meaning “to set in motion”. Another subtype is **inherently adpositional verbs (IAV)**, which obligatorily select a specific preposition, for instance, *a aspira la* ‘to aspire at’, meaning “to aspire to”.

Nominal MWEs function as nouns. **Nominal idioms (NID)** are noun phrases with figurative meanings that cannot be understood literally, such as *beton armat* ‘reinforced concrete’, which, besides the compositional meaning occurring in the context of building materials and constructions, also means “extremely strong or resilient” in other contexts in which it is analysed as a MWE. **Pronominal idioms (PronID)** behave like pronominal words and convey idiomatic meaning, for example, *toată suflarea* ‘all breath’, meaning “everyone”. **Deverbal nominal MWEs (NV)** are nouns derived from verbal constructions. Inside this group, **NV.VID** corresponds to nouns derived from verbal idioms, like *punere la punct* ‘putting in point’, derived from *a pune la punct* ‘to put in point’, meaning “to fine-tune/to adjust”. **NV.LVC.full** refers to nouns derived from full light verb constructions, such as *luarea unei decizii* ‘the taking of a decision’, from *a lua o decizie* ‘to take a decision’, meaning “to make a decision”. **NV.LVC.cause** corresponds to nouns derived from causative light verb constructions, for instance, *intrare în vigoare* ‘entry into force’, from *a intra în vigoare* ‘to enter into force’, meaning “to come into effect”.

Adjectival and adverbial MWEs cover expressions that function as modifiers. Adjectival and adverbial MWEs can be deverbal or idiomatic. **AV.VID** refers to adjectival or adverbial forms derived from verbal idioms, such as *fără tragere de inimă* ‘without pulling of heart’, meaning “reluctantly/unwillingly”. **AV.LVC.full** represents adjectival forms derived from full light verb constructions, for example, *aflați în dificultate* ‘found in difficulty’, from *a se afla în dificultate* “be in trouble/in difficulty”. **AV.LVC.cause** corresponds to adjectival forms derived from causative light verb constructions, such as *pus la dispoziție* ‘put at disposal’, from *a pune la dispoziție* ‘to put at disposal’, meaning “made available”. **Independent adjectival idioms (AdjID)** are fixed expressions with idiomatic adjectival meaning, such as *de mâna a doua* ‘of the second hand’, meaning “second-hand/used”. **Adverbial idioms (AdvID)** are fixed multiword adverbs, for

example, *cu sufletul la gură* ‘with the soul at the mouth’, meaning “breathlessly/eagerly/anxiously”.

Finally, **Functional MWEs** behave like function words – determiners, adpositions, conjunctions, or interjections. These expressions typically form relatively short, closed lists of cases within each specific language. **Determiner idioms (DetID)** are MWEs whose neutral form has the distribution of a determiner, such as *tot soiul de* ‘every kind of’, meaning “all sorts of/a wide variety of”. **Adpositional idioms (AdpID)** are MWEs whose neutral form functions as an adposition, for instance, *în fața* ‘in the face of’, meaning “in front of”. **Conjunction idioms (ConjID)** function as MWEs with the distribution of a conjunction, with examples like *în caz că* ‘in case that’, meaning “in case/if”. Lastly, **Interjection idioms (IntjID)** are MWEs whose neutral form has the distribution of an interjection, such as *Doamne ajută* ‘Lord help’, meaning “God help us!/May God help!”.

3 Related Work

The PARSEME initiative has organized several multilingual annotation campaigns and shared tasks, focusing primarily on VMWEs. The first edition of the shared task in 2017, PARSEME 1.0, covered 18 languages (Romanian included) and produced a large, manually annotated multilingual corpus of VMWEs, establishing a foundational typology and addressing challenges such as syntactic discontinuity and long-distance dependencies (Savary *et al.*, 2017; Ramisch *et al.*, 2018). Subsequent campaigns, PARSEME 1.1 and 1.2, between 2018 and 2020, expanded the corpus for 19, respectively 14 languages, refined annotation guidelines, and introduced semi-supervised evaluation settings, producing more than 5 million tokens and over 60,000 annotated VMWEs. These corpora have become benchmarks for supervised MWE identification systems and have fostered the development of robust, cross-lingual VMWE identification methods (Ramisch *et al.*, 2020).

The PARSEME annotation guidelines formalized a universal procedure for identifying MWEs across languages, combining a decision-tree methodology with cross-lingual consistency tests and language-specific clarifications (Savary *et al.*, 2018; Savary *et al.*, 2023). This dual strategy has been essential to achieving both consistency and flexibility in multilingual annotation. PARSEME 1.3 (Savary *et al.*, 2023) further extended the multilingual corpus to 26 languages, aligning annotations with Universal Dependencies³ (de Marneffe *et al.*, 2021) v.2, improving linguistic coverage, and making the corpus a versatile resource for multilingual verbal MWE research.

³ <https://universaldependencies.org/>

Building on this foundation, the current UniDive COST Action integrates the PARSEME methodology to extend annotation coverage to nominal, modifier, and functional MWEs. The current annotation campaign introduces version 2.0 of the guidelines⁴ and includes pilot annotations from several language teams. These efforts highlight the importance of inter-annotator agreement, the use of complementary tools to ensure consistency, and language-specific adaptations for typologically diverse languages (Savary *et al.*, 2023).

Our contribution continues this trajectory by enhancing the Romanian PARSEME corpus (PARSEME-Ro) (present in the PARSEME corpus from its very first edition) (Barbu Mititelu *et al.*, 2019) with systematically annotated nominal, modifier, and functional MWEs. We build on the established decision-tree annotation methodology and cross-lingual testing, introducing Romanian-specific clarifications where necessary.

4 Methodology

As mentioned, the previous editions of the PARSEME-Ro corpus were manually annotated. However, for expanding its annotation with MWEs of other parts of speech, we decided to automate the process, given that the load of these new MWEs seemed quite high and some of them are very redundant (e.g. the functional ones). This automatic process has always been doubled by manual intervention, precisely because of the challenges this phenomenon raises, some of which were explained in the Introduction. All the steps towards the final annotation of PARSEME-Ro corpus are described in the rest of this section.

4.1 The Corpus

The PARSEME-Ro corpus contains only journalistic texts from the period 2003-2006 and is made up of 56,703 sentences with 1,015,623 tokens. In this annotation campaign, we added 17,168 sentences to the corpus (containing 486,141 tokens), which were also journalistic texts, from the period 2010-2016, so that we kept the corpus homogeneous as far as its genre is concerned.

4.2 The Automatic Extraction of MWEs from Dictionaries

The first stage in the annotation of this corpus involved identifying dictionaries of idioms and expressions available either in an online version with search and/or download functionality, or in digital formats (.DOCX, .XLS, .PDF, etc.). Our focus was on resources that allow automatic processing, whose information can be standardized into a common format, and on the basis of which a list of expressions could be constructed and labelled in accordance with the PARSEME

⁴ See PARSEME Shared Task 2.0 Guidelines: <https://parsemefr.lis-lab.fr/parseme-st-guidelines/2.0/>

conventions (see Section 2). Accordingly, we identified the following data sources:

- *Dicționarul ortografic, ortoepic și morfologic al limbii române* (DOOM⁵, The Orthographic, Orthoepic, and Morphological Dictionary of the Romanian Language)⁵; using specific queries to interrogate the web interface, we were able to retrieve and download all idiomatic expressions corresponding to a given part of speech (nouns, adjectives, adverbs, prepositions, conjunctions, interjections).
- *Dicționarul explicativ al limbii române* (DEX, The Explanatory Dictionary of the Romanian Language, the online version⁶): although it does not permit downloading lists of idiomatic expressions, it was nonetheless consulted at all stages of annotation.
- *Dicționarul de expresii, locuțiuni și sintagme ale limbii române* (DELS, Dictionary of expressions and locutions, Mărânduc, 2010): with the author's kind permission, the dictionary's .PDF format was used for the automatic extraction of expressions into an .XLS format: for each expression, the type assigned by the author (expression, idiom, phrase) is specified.
- *Dicționar de expresii românești în contexte, Vol. 1-4* (Dictionary of Romanian Expressions in Context, Ilinca, 2015, 2016, 2021a, 2021b), in .PDF format; the information was automatically extracted from the .PDF files using a dedicated script; the resulting text file was processed both automatically and manually to remove verbal expressions (which are not of interest at this stage of annotation), to correct errors inherent to the automatic extraction process, and to convert the expressions into the established common format.
- *Dicționar frazeologic al limbii române* (Phraseological Dictionary of the Romanian Language, Tomici, 2009): it was processed similarly to the Ilinca dictionary.

The second step was the compilation of the expressions from the aforementioned resources into a common list, from which duplicates were removed. The resulting list was then automatically matched against the corpus to be annotated, using a specially created script, in order to extract a list of expressions that are known to occur in the corpus and are intended for automatic annotation.

4.3 The Curation of the Automatically Extracted MWE List

The list that resulted from the five lexicographic resources was limited to 2,034 expressions that occur in the corpus and were to be annotated. This list was

⁵ <https://doom.lingv.ro/>

⁶ <https://dexonline.ro/>

labelled with PARSEME MWE categories and validated in two stages by the team of linguists. A Google Spreadsheets document was used to ensure simultaneous access for the team to the same working file.

Each expression was manually assigned one or more possible labels from the label set used in the current version of the PARSEME project (see Section 2 above). Beyond remedying errors that resulted from automatic extraction (missing segments from expressions, incorrect lemmatization, etc.), several operations were performed, of which we enumerate a few:

- Expressions that were recorded in lexicographic resources with the traditional label *expression*, *idiom* (Rom. “expresie”, “locuțiune”), but did not pass the MWE tests from the PARSEME test battery were removed.
- If an entry had two different morphosyntactic values, a second entry homonymous with the first was introduced. E.g., if the list contained only a single entry *în zadar* “in vain” this was duplicated because it can be assigned two labels: *în zadar* (AdjID), as in: *speranțe în zadar* ‘hopes in vain’ “vain hopes”, and *în zadar* (AdvID), as in: *plânge în zadar* “cry in vain”.
- If an entry with a certain morpholexical label had multiple different meanings, these were not marked differently, but remained under the same entry. For example, *de sus* (AdvID) ‘from above’ means both “coming from a superior/from authorities”, as well as “from heaven or from God”.

The list of labelled entries was cross-validated using a dual validation-annotation approach, with each annotation entry independently validated by two of the six annotators working in parallel, with no access to each other's decisions. Subsequently, the two lists were automatically processed and, through cross-comparison, expressions showing inter-annotator agreement across both annotation phases and those exhibiting disagreement were extracted into separate lists.

The validation process yielded inter-annotator consensus for 1,178 entries and disagreement for 856 entries. The overall inter-annotator agreement rate, calculated from the final lists, was 57.9% (1,178 out of 2,034 total entries). However, it should be noted that entries were added and removed during both cross-validation phases, and these modifications were not systematically tracked.

Among the agreed-upon expressions, 552 out of 1,178 (46.9%) failed to meet PARSEME guidelines for MWE identification and were classified as NOT MWE in the first phase of labelling or in the cross-validation/annotation phase. The consensus dataset was subjected to a third validation by an expert linguist.

The disagreement dataset comprised 856 expressions, of which 264 were classified as NOT MWE. All disagreement cases underwent discussion in team meetings to establish consensus decisions. The resolution outcomes varied across two categories: retention of both labels and selection and/or rejection of label(s).

In some instances, both conflicting annotations were deemed valid and preserved. This situation occurred when one annotator had assigned an additional morpholexical label during the cross-validation, creating a homonymous entry alongside the existing one (labelled in the initial annotation phase). For example, for the MWE *de moarte* “to death, deadly”, both labels were retained (AdjID: *dușman de moarte* “mortal enemy” and AdvID: *el urăște de moarte matematica* “he hates mathematics to death”). In other cases, when conflicting labels were assigned to the same entry, one was selected based on team consensus. For example, for *în caz de* “in case of”, where one annotator validated the AdpID label (prepositional MWE) and another the ConjID label (conjunctive MWE), the latter label was eliminated as incorrect. Some expressions were reclassified as non-MWEs following this reevaluation, as they do not pass any of the PARSEME tests. Examples include *fără limită* “without limit” (initially labelled as AdjID and as NOT MWE, ultimately classified as NOT MWE) and *cu putere* “with force” (initially tagged as AdvID and as AdjID, subsequently reclassified as NOT MWE in the third validation round).

Following these discussions, the dataset was refined. Expressions deemed invalid MWEs were removed (totaling 816 NOT MWEs), while new expressions were incorporated, yielding a final dataset of 2,010 expressions, which were used for the automatic annotation of the PARSEME-Ro corpus.

4.4 The Automatic Annotation of the Corpus with MWEs

The resulting list was used to automatically annotate the corpus of interest (PARSEME-Ro). In the process of identifying expressions within the corpus, we allowed a maximum of two intervening words between the words of an expression, in order to capture cases where an expression allows for the insertion of other words between its components. Although this measure does not guarantee that all expressions are captured, it is a good enough compromise, since with the rising number of intervening words, the number of possible errors is also increasing. Special attention was given to numbering the expressions, continuing from the number immediately following that used for identifying verbal expressions (annotated in a previous annotation stage), to ensure unique assignment of expression numbers within a sentence (any incorrect assignment of expression numbers leads to errors in the rendering of expressions in the validation tool and in future intents of processing or searching the corpus). The corpus was segmented into files of a maximum of 500 sentences each (for easier management in the validation stage), with their naming and encoding in accordance with the instructions provided in the PARSEME Language Leader Guide⁷.

⁷ <https://gitlab.com/parseme/corpora/-/wikis/PARSEME-Language-Leader-Guide>

4.5 The Validation and Correction of the Automatic Annotation with MWEs

Seen as “proof that something is correct” (Cambridge Dictionary⁸, 2003), validation is an essential step in considering a product to be final. In relation to the PARSEME-Ro project, automatic annotation enabled the corpus volume to be managed efficiently. However, the validation of the automatic annotation involved a stage of error correction, which was done by the entire PARSEME-Ro team.

The automatically annotated files were uploaded to the FLAT platform instance dedicated to the PARSEME project, where all manual annotation activities within the project take place. Each annotator has an individual FLAT account, where they only have access to the files for which they need to validate the automatically identified expressions, which meant doing one of the following actions:

- If the expression has two annotations, one is removed. These are the cases when an entry in the list of expressions had two different morphosyntactic values and the proper label can be identified only in context: see the discussion about the MWE *de moarte* “to death, deadly” in subsection 4.3.
- If the expression is not correctly annotated, one of the following action is done:
 - change the category label,
 - eliminate components, or
 - add components.
- If the respective string is not an expression (in that respective occurrence), its annotation is removed. For example, the construction *de aici* “from here” is considered AdvID only in contexts such as *De aici rezultă că negocierile au eşuat*. “From this it follows that the negotiations have failed.”, therefore, followed by the verb *a rezulta* “to result” or *a reieşi* “to emerge”, with the same meaning. However, the automatic annotation marked *de aici* as AdvID even in contexts such as *De aici, din Alaska, au început negocierile*. “From here, in Alaska, negotiations begin.”, in which case the annotator removed the label.

Besides correcting automatically annotated MWEs, we also had to manually annotate any expression occurring in the corpus, but not identified automatically (because it is not in the list of MWEs extracted from the lexicographic sources or because more than two words were intervening in the expression).

⁸ https://www.google.com/url?q=https://dictionary.cambridge.org/dictionary/english/validation&sa=D&source=docs&ust=1762729098442856&usg=AOvVaw1kgNkUmsL_EF2yRdw6bKMu

4.6. Ensuring Annotation Consistency

PARSEME has been preoccupied with the quality of the produced datasets, given that they were always meant to be used in shared tasks dealing with the identification of MWEs in corpora. Thus, a system of ensuring consistent annotation was set in place. A suite of Python scripts helps the users to ensure their corpora consistency. For each MWE annotated in the corpus, all its annotations are extracted together with other contexts in which the same words occur but that were not annotated (counting as skipped occurrences). The human validator carefully reads all of them and decides upon the MWE status of the respective string, upon its correct label in each context. This offers a better way of spotting similar contexts for which the same decision must be made, as well as differences in the possible usages/values of the respective string.

5 Results

At the moment of this writing, the quantitative data about our corpus look as follows:

TOTAL: 31294 MWEs → {'NID': 2751, 'AdpID': 8489, 'IRV': 2950, 'AdjID': 4003, 'AdvID': 5554, 'NV.LVC.full': 18, 'VID': 1359, 'TODO': 3, 'ConjID': 1703, 'IAV': 3166, 'PronID': 317, 'ConjID': 175, 'LVC.full': 411, 'LVC.cause': 123, 'NV.VID': 39, 'NOT': 2, 'AV.LVC.cause': 6, 'DetID': 35, 'IntjID': 80, 'AV.LVC.full': 5, 'VPC.full': 7, 'AV.VID': 14, 'AV.IRV': 2, 'AV.IAV': 58, 'NV.IAV': 11, 'NV.LVC.cause': 13}. These figures are provisional, as the corpus will undergo a final cleaning phase, as some sentences (even ones in which MWEs are annotated) are too short to be relevant for the occurrences of MWEs and we intend to remove them.

6 Conclusions

The enhanced PARSEME-Ro 2.0 corpus represents the most comprehensive resource for Romanian MWEs to date, providing a crucial foundation for advancing NLP applications for the Romanian language. We presented here the way in which we enhanced the PARSEME-Ro corpus (which was initially annotated with verbal MWEs only) with MWEs of all the other parts of speech. The annotation process is both lengthy and challenging. The type of linguistic phenomenon we considered here makes it even more problematic due to the idiosyncrasies MWEs display. On the other hand, the automatic annotation increases the burden as it proposes multiple labels to MWE. Many MWEs (adverbial or adjectival) have several possible values and should be manually disambiguated. The PARSEME-Ro 2.0 will be released together with the other corpora in this initiative and made available, just like all its previous editions, under an open license.

7 Acknowledgments

This research was partially supported by LLMs4EU, co-funded by the Digital Europe Programme under GA 101198470. Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

AI Use Statement

Generative-AI tools were used solely for minor stylistic refinements and copyediting. All content, analysis, and conclusions in this paper are the authors' own work.

References

- Agata Savary, Cherifa Ben Khelil, Carlos Ramisch, Voula Giouli, Verginica Barbu Mititelu, Najet Hadj Mohamed, Cvetana Krstev, Chaya Liebeskind, Hongzhi Xu, Sara Stymne, Tunga Güngör, Thomas Pickard, Bruno Guillaume, Eduard Bejček, Archana Bhatia, Marie Candito, Polona Gantar, Uxo Iñurrieta, Albert Gatt, Jolanta Kovalevskaitė, Timm Lichte, Nikola Ljubešić, Johanna Monti, Carla Parra Escartín, Mehrnosh Shamsfard, Ivelina Stoyanova, Veronika Vincze, Abigail Walsh. 2023. PARSEME corpus release 1.3. In *Proceedings of the 19th Workshop on Multiword Expressions (MWE 2023)*, pages 24–35, Dubrovnik, Croatia, Association for Computational Linguistics.
- Agata Savary, Marie Candito, Verginica Barbu Mititelu, Eduard Bejček, Fabienne Cap, Slavomir Čéplö, Silvio Ricardo Cordeiro, Gülşen Eryiğit, Voula Giouli, Maarten van Gompel, Yaakov HaCohen-Kerner, Jolanta Kovalevskaitė, Simon Krek, Chaya Liebeskind, Johanna Monti, Carla Parra Escartín, Lonneke van der Plas, Behrang QasemiZadeh, Carlos Ramisch, Federico Sangati, Ivelina Stoyanova, Veronika Vincze. 2018. PARSEME multilingual corpus of verbal multiword expressions. In Stella Markantonatou, Carlos Ramisch, Agata Savary, and Veronika Vincze (eds.), *Multiword expressions at length and in depth: Extended papers from the MWE 2017 workshop*, pages 87-147. Language Science Press, Berlin.
- Amalia Todirascu and Marion Cargill. 2019. SimpleApprenant: a platform to improve French L2 learners' knowledge of multiword expressions. In Fanny Meunier, Julie Van de Vyver, Linda Bradley, and Sylvie Thouésny (Eds.), *CALL and complexity - short papers from EUROCALL 2019*, pages 356-361.
- Carlos Ramisch, Agata Savary, Bruno Guillaume, Jakub Waszczuk, Marie Candito, Ashwini Vaidya, Verginica Barbu Mititelu, Archana Bhatia, Uxo Iñurrieta, Voula Giouli, Tunga Güngör, Menghan Jiang, Timm Lichte, Chaya Liebeskind, Johanna Monti, Renata Ramisch, Sara Stymne, Abigail Walsh, Hongzhi Xu. 2020. Edition 1.2 of the PARSEME Shared Task on Semi-supervised Identification of Verbal Multiword Expressions. In *Proceedings of the Joint Workshop on Multiword Expressions and Electronic Lexicons*, pages 107–118, Association for Computational Linguistics.

- Carlos Ramisch, Silvio Ricardo Cordeiro, Agata Savary, Veronika Vincze, Verginica Barbu Mititelu, Archana Bhatia, Maja Buljan, Marie Candito, Polona Gantar, Voula Giouli, Tunga Güngör, Abdelati Hawwari, Uxoa Iñurrieta, Jolanta Kovalevskaitė, Simon Krek, Timm Lichte, Chaya Liebeskind, Johanna Monti, Carla Parra Escartín, Behrang QasemiZadeh, Renata Ramisch, Nathan Schneider, Ivelina Stoyanova, Ashwini Vaidya, Abigail Walsh. 2018. Edition 1.1 of the PARSEME Shared Task on Automatic Identification of Verbal Multiword Expressions. In *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)*, pages 222–240, Santa Fe, New Mexico, USA, Association for Computational Linguistics.
- Cătălina Mărănduc. 2010. *Dicționar de expresii, locuțiuni și sintagme ale limbii române*, DELS. Corint, Bucharest.
- DEX – *Dicționarul explicativ al limbii române*. 2009. Academia Română, Institutul de Lingvistică. Univers Enciclopedic Gold, Bucharest.
- DOOM³ – Ioana Vintilă Rădulescu (coord.). 2021. *Dicționar ortografic, ortoepic și morfologic al limbii române*. Univers Enciclopedic, Bucharest.
- Marie-Catherine de Marneffe, Christopher Manning, Joakim Nivre, and Daniel Zeman. 2021. Universal Dependencies. *Computational Linguistics* 47(2): 255–308.
- Mile Tomici. 2009. *Dicționar frazeologic al limbii române*. Editura Saeculum Vizua, Bucharest.
- Timothy Baldwin and Su Nam Kim. 2010. Multiword expressions. In Nitin Indurkha and Fred J. Damerau (Eds.), *Handbook of Natural Language Processing*. Boca Raton, FL: CRC Press.
- Vasile Ilincan. 2015. *Dicționar de expresii românești în contexte* [DERC]. A-C. Volume I; 2016, D-N. Volume II; 2021a, O-R. Volume III; 2021b, S-Z. Volume IV. Presa Universitară Clujeană, Cluj-Napoca.
- Verginica Barbu Mititelu, Mihaela Cristescu, and Mihaela Onofrei. 2019. The Romanian Corpus Annotated with Verbal Multiword Expressions. In *Proceedings of the Joint Workshop on Multiword Expressions and WordNet (MWE-WN 2019)*, pages 13–21, Florence, Italy. Association for Computational Linguistics.
- Yinjie Wang and Agnès Tutin. 2025. *Tu plaisantes? Ma parole! Étude de quelques motifs récurrents exprimant la surprise dans les dialogues romanesques contemporains français, Corpus [En ligne]*, 27 | 2025.

THE AUTOMATIC DIGITALIZATION OF THREE ROMANIAN DICTIONARIES: DELR, DOOM³ AND DLR

CĂTĂLIN MITITELU¹, HORIA CUCU²

¹*Romanian Academy,
"Iorgu Iordan – Alexandru Rosetti" Institute of Linguistics
catalinmititelu@yahoo.com*

²*Zevo Tech and POLITEHNICA Bucharest
horia.cucu@zevo-tech.com*

Abstract

This article examines the automatic digitalization of three major Romanian dictionaries: the *Dicționarul Etimologic al Limbii Române* (DELR), the *Dicționarul Ortografic, Ortoepic și Morfologic al Limbii Române* (DOOM³), and the *Dicționarul Limbii Române* (DLR). Unlike projects carried out with specialized lexicographic platforms, these resources were digitized through a unified two-phase pipeline: parsing with the DEPAR tool and subsequent enrichment through tailored technologies. DEPAR automatically interprets lexicographic conventions and converts dictionary entries into structured XML-TEI, balancing shallow and deep parsing, while golder master testing ensures schema evolution and long-term consistency. The enrichment phase varies across dictionaries: DELR remained a proof of concept with minimal structuring; DOOM³ underwent deep parsing and integration of additional functionalities, notably bimodal enrichment through the ZEVO Live TTS API, which provides high-quality audio pronunciations alongside automated correction of wrong forms; and DLR is in the process of advanced shallow parsing combined with a redesigned graphical interface and dynamic functionalities. By comparing the three projects, the article demonstrates the scalability of the automatic pipeline, underlines the role of speech technologies in lexicography, and situates the Romanian experience within broader international efforts on dictionary parsing and digitalization.

Keywords: dictionary entry parsing, digitalization, audio generation, *Dicționarul Etimologic al Limbii Române* (DELR), *Dicționarul Ortografic, Ortoepic și Morfologic al Limbii Române* (DOOM³), *Dicționarul Limbii Române* (DLR)

1. Introduction

The digitalization of linguistic resources represents one of the most transformative processes in contemporary lexicography. Beyond granting universal access, digital dictionaries offer advanced possibilities for linguistic

research, natural language processing (NLP), and educational applications. The Romanian Academy's "Iorgu Iordan – Alexandru Rosetti" Institute of Linguistics has initiated an ambitious program to automatically digitize their lexicographic works, and in this paper we focus on three of them: the *Dicționarul Etimologic al Limbii Române* (DELR) (En. "The Etymologic Dictionary of Romanian"), the third edition of the *Dicționarul Ortografic, Ortoepic și Morfologic al Limbii Române* (DOOM³) (En. "The orthographic, Orthoepic and Morphologic Dictionary of Romanian"), and the *Dicționarul Limbii Române* (DLR) (En. "The Dictionary of the Romanian Language").

The approach differs fundamentally from the other institute's projects carried out in proprietary environments such as Tlex, e.g., *Dicționarul de cadre verbale* (DCV) (En. "The Valence Lexicon of Romanian Verbs"), volumes of *Inventar de cuvinte și sensuri noi atestate în mediul online* (ICSO) (En. "The Inventory of new words and senses occurring online"). Our approach relies entirely on open-source technologies and automatic parsing. TLex-based projects require manual data structuring within an interface, while our method interprets the lexicographers' metalanguage through formal grammars, ensuring consistency across all data, scalability, and minimal human intervention. This methodological choice allows the dictionaries to evolve automatically as new linguistic conventions or XML schemas are introduced.

The three dictionaries underwent a **two-phase automatic digitalization pipeline**:

Parsing with DEPAR, a tool that transforms legacy print-oriented data into structured XML-TEI.

Enrichment and indexing, which enable advanced search, filtering, and – where relevant – additional functionalities such as error correction and audio integration.

This article describes the methodology of this transformation, presents the results obtained for each dictionary, and concludes with a comparative overview.

2. The digitalization process

2.1. Motivation

The three dictionaries are developed at "Iorgu Iordan – Alexandru Rosetti" Institute of Linguistics mainly to be printed on paper. They are written in DOC format which is also offered as basis for the online version, thus the necessity to develop an automatic processing pipeline to extract and convert the dictionary content in a format that facilitates the searching not only by headword, but also by variants, part of speech, usage and other lexicographical component of an entry.

2.2. Related work

Out of the several attempts to offer access to Romanian dictionaries in electronic format (see dexonline.ro¹, the CLRE project, Pătraşcu *et al.*, 2013), we will focus here only on those that involved parsing of the existing electronic material, that was not initially subject to scanning and OCR-isation of printed material.

The automatic parsing and digitalization of dictionaries has been an active research field in both Romanian and international lexicography. A series of works have established some of the most influential models for dictionary entry parsing: the *Dictionary Sense Segmentation & Dependency (DSSD)* method (Curteanu, Trandabăţ, & Moruz, 2008) introduced a two-level parsing strategy, separating sense-tree construction from the parsing of individual definitions. Later, the *Segmentation–Cohesion–Dependency (SCD)* configurations (Curteanu, 2010) provided an optimal and portable parsing method applicable across Romanian, French, and German dictionaries, demonstrating that parsing strategies can generalize across languages and lexicographic traditions. Building on this, Curteanu and Moruz (2012) proposed a procedural DTD approach, showing how parameterized grammars can serve as a basis for automatic conversion of dictionary entries into structured XML formats. These approaches are similar to the Romanian DEPAR tool, which also relies on formal grammars to capture the lexicographers’ metalanguage and ensure consistency across large corpora.

2.3. Parsing with DEPAR

The **Dictionary Entry PARser (DEPAR)** is a dedicated tool designed to automatically interpret the metalinguistic conventions of lexicographic articles. DEPAR converts entries, originally stored in DOC format, into valid XML-TEI structures by applying:

- a **formal grammar** of dictionary articles;
- a set of **transformation rules** of the parser tree obtained.

The DEPAR tool (Dictionary Entry PARser) was first presented by Mărânduc, Mititelu, and Pérez (2014) in the context of developing automatic parsers for Romanian lexical resources. In that study, the authors proposed a flexible parsing framework capable of transforming lexicographic text into structured TEI-XML representations. DEPAR was designed to handle the morphological and derivational complexity of Romanian while remaining adaptable to dictionaries organized alphabetically or by lexical nests. This early implementation laid the groundwork for the current large-scale digitalization of DELR, DOOM³, and DLR.

¹ <https://dexonline.ro>

The parser recognizes elements such as headwords, inflected forms, grammatical markers, senses, domains, etymologies, and bibliographical references. Crucially, DEPAR relieves lexicographers of the need to master complex XML encoding.

The core module, **depar-core**, leverages ANTLR4², allowing each dictionary to have a dedicated grammar tailored to its specific entry structure. ANTLR4 was selected due to its support for LL(*) grammars, which offer greater flexibility and intuitiveness compared to LR grammars.

The processing pipeline begins with Microsoft Word documents (.doc), which are converted to well-formed XHTML using **Apache Tika**³, ensuring normalized markup and removal of extraneous Word-specific structures. The XHTML files are subsequently transformed into two files: a content file (.txt) containing raw text and a feature file (.feat) capturing style information (see Figure 1 for a detailed parsing diagram).




Figure 1: DEPAR diagram parsing

This separation allows the generated parser to focus primarily on content parsing, while style information is queried selectively in the dictionary specific grammar, through semantic predicates when necessary for structural disambiguation. Bellow is an example of a fragment from the DOOM³ grammar describing the recognition of references to other entries, being marked with style **bold**:

```
referenceForm: {bold() }? word;
word:         WORD;
```

The list of recognized styles can be easily extended and it contains the following: bold, italic, superscript, subscript, small caps, footnote anchor, footnote index, colour. A predicate (a function that returns a boolean value) is implemented for each style to be used in any moment in the grammar (see {bold() }? from above), among several additional useful predicates:

² <https://wwwantlr.org/>

³ <https://tika.apache.org/>

boundary returns true if there is a change in the current position regarding style; it proves very useful to identify sequences of tokens having the same style;

hasFeature(*Feature feature*) returns true if the current token has the corresponding feature or style;

hasExactFeatures(*Feature[] features*) returns true if the current token has exactly the list of features passed as argument.

A distinctive good practice when implementing the dictionary specific parsers is the use of **golden master testing** (also called characterization or approval testing). Each new parsing rule is validated against a gold-standard reference; only if the new version provides improvements does it become the new baseline. This mechanism ensures both **schema evolution** and consistency. For instance, when DOOM³ introduced usage indications advanced filtering, DEPAR adapted seamlessly, avoiding manual intervention across the structure of the corpus. See below an example of the schema evolution, in which the “after” part of the examples contains the new attributes introduced for coping with usage indications:

before:

```
<entry type="new"><form type="lemma"><orth
rend="____&apos;" norm="abule">abul<emph style="text-
decoration:underline">e</emph></orth><usg>arg.</usg><g
ramGrp><gram type="pos"
norm="interjecție">interj.</gram></gramGrp></form></en
try>
```

after:

```
<entry type="new"><form type="lemma"><orth
rend="____&apos;" norm="abule">abul<emph style="text-
decoration:underline">e</emph></orth><usg
type="register"
norm="argotic">arg.</usg><gramGrp><gram type="pos"
norm="interjecție">interj.</gram></gramGrp></form></en
try>
```

In practice, parsing lies between **shallow parsing** (minimal structuring, formatting preserved) and **deep parsing** (full decomposition). The balance depends on the dictionary’s internal structure and intended functionalities, and this is obvious in the way in which each of the three mentioned dictionaries was parsed, as will be shown below.

2.4. Enrichment and indexing

After parsing, dictionary entries are stored and made searchable through dedicated indexing solutions tailored for each dictionary. **In the case of DOOM³, indexing was complemented by an enrichment phase:**

a **list of incorrect word forms** and their corresponding entries;

audio files illustrating pronunciation for headwords.

The list of incorrect word forms comprises frequent misspellings, colloquial or obsolete variants, and orthographic errors that users commonly search for when consulting normative dictionaries. These forms were compiled from corpus evidence and web query logs, then mapped to their corresponding canonical entries. The inclusion of such mappings improves search tolerance and helps guide users toward standard Romanian usage.

The generation of audio files was facilitated by Zevo Live TTS API which is a speech synthesis system which generates short spoken utterances from corresponding text. The system supports Romanian and English and is designed to be extensible to other languages. It includes multiple male and female voices pre-trained in the application, and it allows for adaptation to new target voices using client-specific data. Integration into client systems is achieved via WebSocket or HTTP interfaces.

Technically, Zevo Live TTS API is a deep-learning based framework, employing Transformer neural networks across its main modules. The architecture comprises three core components:

The Text Preprocessing Module, responsible for normalizing non-standard text elements (e.g., numerals, email addresses, acronyms) is developed from scratch by Zevo Tech.

The Speech Synthesis Module generates intermediate acoustic features (e.g., mel-spectrograms) from the normalized text. This module is adapted on FastPitch: Parallel Text-to-Speech with Pitch Prediction⁴. FastPitch is a fully-parallel Transformer-based model (an evolution of FastSpeech) which models the pitch (fundamental frequency) contour per input symbol. This pitch conditioning improves expressiveness, convergence, and quality by allowing prosody manipulation, better matching of semantic emphasis, and eliminating the need for teacher-student distillation of spectrogram targets in certain settings.

The Vocoder Module converts the acoustic features produced by the synthesis module into waveform audio. This module is based on HiFi-GAN: Generative Adversarial Networks for Efficient and High Fidelity

⁴<https://arxiv.org/abs/2006.06873>

Speech Synthesis⁵. HiFi-GAN is a GAN-based vocoder that achieves high fidelity audio (≈ 22.05 kHz) with efficiency: it generates output many times faster than real-time on modern GPUs, and it generalizes well to unseen speakers.

The process of including audio generated files is under development and presupposes a manual validation. Four voices were selected as possible candidates, for each headword four audio files corresponding to each voice being generated. A group of linguists are validating the pronunciations keeping only those that are closer to the natural utterance. Only after this phase the link between dictionary entry and audio files is established. At the moment of this writing, more than 1,200 headwords have been associated their pronunciation on the website.

At present, the DOOM³ online platform integrates a dedicated **audio layer** accessible directly from each lemma page through a speaker icon placed next to the headword. The set of audio-enabled entries is continuously expanding and can be consulted on the dedicated page: <https://doom.lingv.ro/cautare/cuvinte-pronuntare>.

2.5. The whole picture of a dictionary digitalization

The digitalization process is, in general, fully automated, except for the cases where manual validation or external resources are necessary. The DOC files are automatically converted into XML (TEI) format, the indexer can bring additional information from external resources where necessary and send the content to a database. An user interface application can access the database and offer the content to the user as depicted in Figure 2.




Figure 2: Dictionary digitalization diagram

⁵ <https://arxiv.org/abs/2010.05646>

3. Results of the digitalization

3.1. DELR

For **DELR** (available at <https://delr.lingv.ro>) we started the digitalization in 2022 and it is under development as far as its linguistic content is concerned. At present it contains over 32,000 of words, ranging from A to D, with over 17,000 variants, grouped in over 9,000 of lexical nests. The parsing applied to DELR is deliberately shallow: only headwords, variants, and family words are explicitly marked, while the remainder of each entry is retained largely in its original formatting without deep structural decomposition. As a result, the functionalities currently offered are limited to searches by headword and family words, with the optional inclusion of variants. This restricted functionality reflects the project's role as a proof of concept, designed to demonstrate the feasibility of automatically parsing traditional dictionaries prepared for print. A representative example is the entry *ABÁ*, in which headwords are displayed in bold and variants in italics, while family relations are identified, but etymological details and citations remain unstructured.

```
<?xml version="1.0" encoding="UTF-8"?>
<div type="entry" xml:id="ABÁ.1">
  <p type="header">
    <index indexName="headers">
      <term key="ABÁ.1">
        <emph style="font-weight:bold">ABÁ</emph>
        <emph style="font-weight:bold;vertical-align:super">1</emph>
      </term>
    </index>
    s.f. „un fel de postav – (vêtement de) bure”
    1621 (DIR B IV, 63), var. (Ban., Trans.)
    <index indexName="variants">
      <term key="habá.1">
        <emph style="font-style:italic">habá</emph>
        <emph style="vertical-align:super">1</emph>
      </term>
    </index>,
    <index indexName="variants">
      <term key="ábă | hábă">
        <emph style="font-style:italic">(h)ábă</emph>
      </term>
    </index>.
  </p>
  <p type="other">
```

Tc.
 <emph style="font-style:italic">aba</emph>
 (< pers. < arab
 <emph style="font-style:italic">'abă</emph>
), dial.
 <emph style="font-style:italic">haba</emph>
 . Var.
 <emph style="font-style:italic">ábă</emph>
 < scr.
 <emph style="font-style:italic">aba</emph>
 , magh.
 <emph style="font-style:italic">aba</emph>
 . Var.
 <emph style="font-style:italic">hába</emph>
 < scr.
 <emph style="font-style:italic">haba</emph>
 . Var. cu
 <emph style="font-style:italic">h-</emph>
 , posibil și < ucr.
 <emph style="font-style:italic">haba</emph>
 .
 </p>
 <p type="family">
 Fam.:
 <index indexName="family">
 <term key="abagér">
 <emph style="font-weight:bold">abagér</emph>
 </term>
 </index>
 s.m. (Mold.) „fabricant sau negustor de aba”
 1673 (var.
 <index indexName="variants">
 <term key="abăgeăr">
 <emph style="font-style:italic">abăgeăr</emph>
 </term>
 </index>, doc. apud Pavlescu, E. 510) ←
 <emph style="font-style:italic">abagiu</emph>, încadrat în clasa derivatelor în -
 <emph style="font-style:italic">ar</emph>
 HEM I, 62, DA sau ←
 <emph style="font-style:italic">abagerie</emph>
 TDRG, DEX, Suciu, IT II, 29;
 <index indexName="family">
 <term key="abagerie">

<emph style="font-weight:bold">abageríe</emph>
 </term>
 </index>
 s.f. „loc în care se produce sau se vinde aba; meseria de abagiu” 1705-1732 (var.
 <index indexName="variants">
 <term key="abăgeríe">
 <emph style="font-style:italic">abăgeríe</emph>
 </term>
 </index>, Cron. anon. racov. 9) ←
 <emph style="font-style:italic">abagiu</emph>
 TDRG, CADE, SDLR, DEX, Suceiu, IT II, 29 sau ←
 <emph style="font-style:italic">abager</emph>
 DA;
 <index indexName="family">
 <term key="abagerítă">
 <emph style="font-weight:bold">abagerítă</emph>
 </term>
 </index>
 s.f. „femeie care se ocupa cu fabricarea sau vânzarea abalei sau a hainelor de aba” 1742 (Cond. Mav. II, 470) ←
 <emph style="font-style:italic">abager</emph>;
 <index indexName="family">
 <term key="abagíu">
 <emph style="font-weight:bold">abagíu</emph>
 </term>
 </index>
 s.m. „fabricant sau negustor de aba” 1642 (Doc. Ist. Buc. 28, apud Suceiu, IT II; atestarea 1582-1591, antrop.
 <emph style="font-style:italic">Abagi</emph>;
 Doc. Î. 156, dată de TDRG, este suspectă, apărând într-o listă ce conține multe nume proprii străine),
 var. (înv.)
 <index indexName="variants">
 <term key="abagéu">
 <emph style="font-style:italic">abagéu</emph>
 </term>
 </index>,
 <index indexName="variants">
 <term key="habagíu">

```

      <emph style="font-
style:italic">habagíu</emph>
    </term>
  </index> &lt; tc.
  <emph style="font-style:italic">abacı</emph>;
  <index indexName="family">
    <term key="abăgerésc">
      <emph style="font-
weight:bold">abăgerésc</emph>
    </term>
  </index>
  adj. (Olt.) „de abagerie, de abager” 1886 (HEM)
←
  <emph style="font-style:italic">abager</emph>.
</p>
<p type="confer">
  Cf. <emph style="font-weight:bold">abacı</emph>.
</p>
</div>

```

3.2. DOOM³

For **DOOM**³ (available at <https://doom.lingv.ro>) we started the digitalization in 2022, with the public launch in August 2023. It contains over 65,000 entries and over 200,000 searchable word forms. In contrast to DELR, the parsing depth here is considerably greater: entries are decomposed into their constituent components, including headword, orthographic representation, inflected forms, and detailed grammatical categories, while much of the original print formatting has been eliminated in favour of structural clarity. This deeper parsing supports advanced functionalities, enabling users to perform complex queries by combining filters such as part of speech, gender, number, register, and language of origin; additionally, the entire theoretical content, *Studiu introductiv* (En. “Introductory study”), is made available for online consultation. The dictionary is further distinguished by an enrichment layer that provides suggestions for wrong or obsolete forms, integrates audio pronunciations through speech synthesis, and thereby enhances both usability and accessibility. A representative case is the entry *abandon*, which is fully decomposed into lemma form with grammatical tags, plural forms, and orthographic variants, illustrating the fine-grained level of linguistic annotation achieved in this project.

```

<entry xml:id="AAAA-00001-AJML-95718"
sortKey="abandon">
  <form type="lemma" xml:id="AAAA-00000-EXKF-39616">
    <orth rend="_____ ' _" norm="abandon">

```

```

      aband<emph style="text-
decoration:underline">o</emph>n
    </orth>
    <gramGrp>
      <gram type="pos" norm="substantiv">s.</gram>
      <gram type="gender" norm="neutru">n.</gram>
    </gramGrp>
  </form>
  <form type="inflected" rend="prefix (comma) "
xml:id="AAAA-00003-CJEM-44164">
    <gramGrp type="partial">
      <gram type="number" norm="plural">pl.</gram>
    </gramGrp>
    <orth rend="_____'____'" norm="abandonuri">
      aband<emph style="text-
decoration:underline">o</emph>nuri
    </orth>
  </form>
</entry>

```

3.3. DLR

The digitalization of the new edition of **DLR**, accessible through its redesigned interface at <https://dlr-test.lingv.ro>, was initiated in 2024 and remains an ongoing project. At the present stage, the resource comprises approximately 880 entries, 1,590 senses, and more than 7,800 citations. The parsing approach applied to DLR may be described as advanced shallow parsing: while a significant portion of the typographic style is preserved, the parser identifies structural components such as headwords, senses, citations, semantic domains, and etymologies. This structural marking allows for enhanced functionalities, including headword searches (exact or tolerant of diacritics), as well as prefix and suffix searches employing wildcard characters (*, ?). Moreover, users can dynamically show or hide citations in order to obtain either compact or extended views of dictionary entries. A central innovation of this edition is its newly designed interface, which emphasizes minimalist navigation, dynamic query execution, and the possibility of cross-linking with other lexicographic resources; indeed, content from DELR can already be displayed within this environment upon request. A representative case is the entry *AALENIÁN*, which demonstrates explicit structural annotation of senses, grammatical information, citations, and etymology, while still maintaining fidelity to the lexicographers' stylistic conventions.

```

<entry xml:id="AAAA-00000-AXNC-12470" type="lemma">
  <form type="headword">
    <orth>AALENIÁN</orth>,
  <form type="inflected">

```

```

        <orth>-Ă</orth>
    </form>,
    <gramGrp>
        <gram type="pos" norm="substantiv">s.</gram>
        <gram type="gender" norm="neutru">n.</gram>
    </gramGrp>,
    <gramGrp>
        <gram type="pos" norm="adjectiv">adj.</gram>
    </gramGrp>
</form>
<dictScrap xml:id="global-usg">
    <usg type="dom">
        <emph style="font-style:italic">GEOL.</emph>
    </usg>
</dictScrap>
<dictScrap xml:id="senses">
    <sense type="number" xml:id="AAAA-00003-JAFT-41277">
        <def>
            <emph style="font-weight:bold">1.</emph>
S. n. Primul etaj al jurasicului mediu.
        </def>
        <dictScrap>
            <cit>
                <emph style="font-weight:bold">1949</emph>
                <emph style="font-style:italic">LTR</emph>, I.
            </cit>
            <cit>
                <emph style="font-weight:bold">1962</emph>
                <emph style="font-style:italic">LTR</emph>
                <emph style="font-style:italic;vertical-align:super">2</emph>, XI, 115:
Mineta [este un] minereu de fier, caracteristic
                <emph style="font-weight:bold">Aalenianului</emph> din Lorena, care se
prezintă sub forma unor strate oolitice calcaroase-
argiloase sau silicioase.
            </cit>
        </dictScrap>
    </sense>
    <sense type="number" xml:id="AAAA-00002-AFTY-35535">
        <def>

```



```

        <emph style="font-weight:bold">2.</emph>
Adj. Care aparține aalenianului (
        <emph style="font-weight:bold">1</emph>),
privitor la aalenian.
    </def>
    <dictScrap>
        <cit>
            <emph style="font-
weight:bold">1975</emph>
            <emph style="font-
style:italic">DEX.</emph>
        </cit>
        <cit>
            <emph style="font-
weight:bold">1976</emph>
            <emph style="font-
style:italic">Geol.</emph>
            <emph style="font-
style:italic">Munților</emph>
            <emph style="font-
style:italic">Apuseni</emph>, 163: Întocmai ca și
calcarele roșii eojurasice, depozite de vârstă cert
aaleniană au fost întâlnite numai în dyke-urile
neptuniene din calcarele triasice.
        </cit>
    </dictScrap>
</sense>
</dictScrap>
<form type="varia">
    <form type="pronunciation">- Pronunțat:
        <emph style="font-weight:bold">ni-an</emph>.
    </form>
    <form>- Pl.:
        <emph style="font-
weight:bold">aalenieni</emph>,
        <emph style="font-weight:bold">-e</emph>.
    </form>
</form>
<etym>- Din fr.
    <emph style="font-weight:bold">aalénien</emph>.
</etym>
<note type="bibl-ref">
    <emph style="font-weight:bold">1978</emph>
    <emph style="font-style:italic">DN</emph>
    <emph style="font-style:italic;vertical-
align:super">3</emph>;
    <emph style="font-weight:bold">2001</emph>

```

```

    <emph style="font-style:italic">MDA</emph>, I;
    <emph style="font-weight:bold">2007</emph>
    <emph style="font-style:italic">DEXI</emph>;
    <emph style="font-weight:bold">2011</emph>
    <emph style="font-style:italic">DELR</emph>, I.
  </note>
</entry>

```

4. Contrastive overview

We synthesize here, for all three dictionaries under discussion, their main characteristics resulting from the digitalization process they underwent.

Table 1: Contrastive overview of the three dictionaries digitalization

Characteristics	DELR	DOOM ³	DLR
Parsing depth	Shallow (headword, variants, family)	Deep (decomposition of entry)	Advanced shallow (headword, senses, citations)
Formatting	Preserved with minimal structuring	Mostly discarded	Largely preserved with selective structuring
Functionalities	Search by headword/family and optionally by variants	Advanced search with filtering	Search, filters, dynamic citations
Enrichment	None	Wrong-form correction + audio pronunciation	Interconnection with other resources

5. Conclusions and Further Work

The digitalization of DELR, DOOM³, and DLR demonstrates the feasibility and versatility of applying a unified two-phase pipeline to distinct lexicographic works. DELR has served as a proof of concept for shallow parsing, DOOM³ has exemplified the potential of deep parsing enriched with bimodal functionalities such as audio pronunciation, and DLR is currently undergoing a process of advanced shallow parsing combined with a redesigned interface tailored to modern user expectations. Collectively, these projects illustrate a coherent strategy for transforming Romania's lexicographic heritage into accessible, computationally exploitable resources.

Both DELR and DLR—being still under development—will be **updated continuously**, with lexicographers adding new entries and refining existing ones. Beyond the expansion of content, an important direction of future work concerns

the **architectural organization of the digital dictionaries as microservices**. This design would allow each dictionary to operate as an independent service, while still offering the possibility of **integrating content across resources** when requested by the user. Such interoperability would enable, for instance, the consultation of etymological information from DELR within the DLR interface. By implementing these dictionaries as interconnected microservices, the project aims not only to modernize access but also to establish a scalable infrastructure capable of supporting future developments in Romanian digital lexicography.

AI Use Statement

Generative-AI tools were used solely for minor stylistic refinements and copyediting. All content, analysis, and conclusions in this paper are the authors' own work.

References

- Alexandru Anghelina and Claudiu Teodorescu. 2020. Digitalizarea *Dicționarului limbii române*. Probleme. Soluții. Perspective. *Limba Română*, LXIX(3-4):491-496.
- Adrian Țăncuț. 2021. FastPitch: Parallel text-to-speech with pitch prediction. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2021)*, pages 6588-6592. IEEE.
- Cătălina Măranduc, Cătălin Mititelu, and Cene Augusto Pérez. 2014. Lexical Nesting Dictionaries. In *Proceedings of the Conference "Linguistic Resources and Tools for Processing the Romanian Language"*, pages 103-110, Iași, Romania.
- Institutul de Lingvistică „Iorgu Iordan - Alexandru Rosetti”. 2010. *Dicționarul limbii române*. Ediție anastatică după *Dicționarul limbii române (DA)* și *Dicționarul limbii române (DLR)*, Tom I-XIX. București: Editura Academiei Române.
- Institutul de Lingvistică „Iorgu Iordan - Alexandru Rosetti”. 2021. *Dicționarul ortografic, ortoepic și morfologic al limbii române*, ediția a III-a revăzută și adăugită, coord. Ioana Vintilă-Rădulescu. Versiunea online: <https://doom.lingv.ro/>
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. HiFi-GAN: Generative adversarial networks for efficient and high-fidelity speech synthesis. *Advances in Neural Information Processing Systems (NeurIPS 2020)*, 33:17022-17033.
- Mădălin Ionel Patrașcu, Elena Tamba, Marius-Radu Clim, and Ana Catană-Spenchiu. 2013. CLRE - Partial Results in the Development of a Romanian Lexicographic Corpus. In *Proceedings of the 9th International Conference "Linguistic Resources and Tools for Processing of the Romanian Language"*, edited by Elena Mitocariu, Mihai Alex Moruz, Dan Cristea, Dan Tufiș, and Marius Clim, pages 51-58, Iași. Editura Universității „Alexandru Ioan Cuza”.
- Neculai Curteanu, Alexandru Moruz, and Diana Trandabă. 2008. Extracting sense trees from the Romanian Thesaurus by sense segmentation and dependency parsing. In *Proceedings of the 2008 Workshop on Cognitive Aspects of the Lexicon (CogALex*

- 2008), pages 47–54, Marrakech, Morocco. European Language Resources Association.
- Neculai Curteanu, Alexandru Moruz, and Diana Trandabăţ. 2010. An Optimal and Portable Parsing Method for Romanian, French, and German Large Dictionaries. In *Proceedings of the 2nd Workshop on Cognitive Aspects of the Lexicon*, pages 38–47, Beijing, China. Coling 2010 Organizing Committee.
- Neculai Curteanu and Mihai Alex Moruz. 2012. A Procedural DTD Project for Dictionary Entry Parsing Described with Parameterized Grammars. In *Proceedings of the 3rd Workshop on Cognitive Aspects of the Lexicon*, pages 127–136, Mumbai, India. The COLING 2012 Organizing Committee.
- TEI 2022 TEI Consortium. (2022). *TEI P5 Guidelines*.

CHAPTER 4

LLM-BASED EXPERIMENTS & EVALUATION

AI REVIEWER'S CONFIDENCE

DANIELA GÎFU^{1,2,3}, SPERANȚA CECILIA BOLEA²,
PAULA ALEXANDRA CRUCIANU², SILVIU-IOAN BEJINARIU²

¹*Faculty of Computer Science,
"Alexandru Ioan Cuza" University of Iasi,
Romania*

²*Institute of Computer Science, Romanian Academy – Iasi Branch, Romania*

³*Academy of Romanian Scientists*
{daniela.gifu, cecilia.bolea, paula.crucianu,
silviu.bejinariu}@iit.academiaromana-is.ro

Abstract

Due to the overwhelming number of review requests received by top-tier scientific conferences, many deserving papers do not receive the rigorous evaluation they require. Given the increasingly competent responses provided by artificial intelligence (AI), human reviewers may partially or fully delegate the initial reviewing task to such systems. This practice raises a critical question: how much confidence should we place in an AI-generated review? We observed significant divergence among responses produced by different chatbots when evaluating the same paper. To investigate this reliability issue, we constructed the **HLLM-Dist corpus**, which includes 200 reviews written by an expert human reviewer for top-tier international conferences, all collected prior to the launch of ChatGPT in November 2022. For the same set of papers, we tasked three chatbots—Microsoft Copilot, Google Gemini, and DeepSeek-Chat—to generate reviews in English. The core contribution of our work is the proposal and pilot validation of a novel metric, **CAIR** (Confidence of AI Reviewer), designed to quantify the alignment between an AI model’s review and human-generated evaluations. As an initial step, we fine-tuned a DistilBERT model for calibration, with plans to test and normalize results using ChatGPT in subsequent phases. The primary goal of this study is to compute and compare confidence scores for the three AI systems, establishing a foundational framework that can eventually predict AI reviewer confidence without requiring extensive human evaluation for each new model. Our results demonstrate the feasibility of this approach, offering a meaningful step toward calibrating trust in AI-assisted peer review.

Keywords: HLLM-Dist Corpus, Human Peer Review Benchmark, AI Review, AI Confidence Scoring, CAIR Metric

1. Introduction

The peer review process serves as a fundamental mechanism for quality control in academic publishing, acting as a critical filter for scientific progress in top-tier

conferences and journals. However, this system faces significant challenges due to exponential growth in submissions, a limited pool of qualified volunteer reviewers, and the increasing complexity of research, which collectively create a demanding workload for the academic community (Bornmann and Daniel, 2010; Tennant and Ross-Hellauer, 2020). This strain can result in rushed evaluations and inconsistencies between reviews, meaning that many submissions may not receive the detailed assessment they require due to resource constraints—a phenomenon often discussed in the context of communication uncertainties (Vlăduțescu *et al.*, 2014).

The emergence of highly capable Large Language Models (LLMs) offers a potential solution to this capacity problem. Models such as GPT-4 and Gemini demonstrate advanced proficiency in language understanding, summarization, and critical analysis (OpenAI, 2023; Gemini Team Google, 2023). It is becoming increasingly common, though often undeclared, for human reviewers to delegate the initial drafting of reviews to AI systems. This trend raises a central and legitimate question: *How much confidence should we place in an AI reviewer?*

Preliminary analysis indicates that such confidence cannot be assumed. As probabilistic systems, LLMs can generate highly variable outputs, exhibiting characteristics analogous to neutrosophic logic where truth, falsity, and indeterminacy coexist (Gifu, 2014). For identical input papers, different models—or even the same model with different prompting strategies—can produce significantly divergent critiques, ranging from accurate to superficial or containing factual inaccuracies known as "hallucinations" (Ji *et al.*, 2023; Maynez *et al.*, 2020). This variability reflects broader criticisms that LLMs, while powerful in pattern recognition and generation, can struggle with tasks requiring deep abstract reasoning and rigorous logical consistency (Marcus, 2020; Talman *et al.*, 2022).

Recent efforts to benchmark LLM performance in academic settings have highlighted both their potential and their limitations, particularly in tasks requiring nuanced judgment and domain expertise (Liang *et al.*, 2022; Wang *et al.*, 2023). Furthermore, the challenge of evaluating text generation quality and aligning machine output with human standards remains an active area of research, with metrics like BLEU and ROUGE often falling short in capturing semantic depth and critical rigor (Liu *et al.*, 2016; Celikyilmaz *et al.*, 2020).

The principal challenge, therefore, is not merely the adoption of AI in reviewing but the precise calibration of trust in its output, a challenge that mirrors the difficulties in multi-dimensional analysis of complex language (Gifu and Cristea, 2012). Without a reliable method to evaluate the quality of AI-generated reviews, their use risks compromising the integrity of the scholarly process they are meant to support.

To address this problem, we propose a novel framework for quantifying the reliability of AI reviewers, drawing inspiration from contrastive methodologies

used in computational linguistics (Gifu, 2015). Our work is built upon a unique corpus, HLLM-Dist, comprising 200 reviews written by an expert human reviewer for top-tier conference papers. This collection was completed prior to the launch of ChatGPT to ensure no AI influence. For the same set of papers, we tasked three prominent chatbots—Microsoft Copilot, Google Gemini, and DeepSeek-Chat—to generate reviews. The core contribution of our work is the proposal and validation of a novel metric that calculates a standardized confidence score for an AI model's review based on its performance against a human benchmark. Our analysis uses a DistilBERT model for calibration, with plans to test and normalize results using ChatGPT as a reference Large Language Model (LLM), applying modern techniques for decoding text authenticity (Gifu and Covaci, 2025).

The primary objective of this paper is to calculate and compare confidence scores for these three specific AI systems, thereby establishing a foundational model. The long-term goal is to create an extensible framework capable of quantifying and predicting the confidence level of any new AI reviewer without requiring extensive manual evaluation for each new system. Our approach provides a quantitative, scalable tool for managing the integration of AI into peer review and for discriminating AI-generated content (Trandabăt and Gifu, 2025).

The remainder of this paper is structured as follows: Section 2 details the construction of the HLLM-Dist corpus and our methodology for generating AI reviews. Section 3 introduces the architecture of our confidence score metric. Section 4 presents the results of our comparative analysis, demonstrating the varying levels of confidence earned by the three AI reviewers. Finally, we discuss the implications of our findings and outline the path for future work to ensure the responsible use of AI in scientific evaluation.

2. Dataset and Experimental Setup

This section provides a comprehensive overview of the dataset and the experimental design used to generate and evaluate the AI reviews.

2.1. HLLM-Dist corpus

To evaluate the reliability of automatically generated peer reviews, we constructed a novel corpus named HLLM-Dist (*Human–Large Language Model Distance*). This corpus was specifically designed to serve as a benchmark for analysing the performance of Large Language Models (LLMs) in the role of scientific reviewers, providing a public dataset for future research.

The corpus comprises 100 scientific papers published in the proceedings of international conferences in the fields of Computer Science, Artificial Intelligence, and related disciplines. The dataset was engineered to encompass a wide variety of profiles, including: varying conference ranks (from A* to national

workshops), diverse application domains (digital libraries, intelligent systems, computational linguistics, information retrieval), and variable paper lengths.

Each paper in the corpus is accompanied by five distinct reviews:

- **Two Original Human Reviews:** These represent the high-quality gold standard reference. They were authored by human experts as part of the actual peer-review processes of the respective conferences. Crucially, all human reviews were collected prior to the public release of ChatGPT, thereby guaranteeing that they were not influenced by or contaminated with AI-generated ideas or phrasing.
- **Three Automatically Generated Reviews:** For each paper, we prompted three popular and representative LLM chatbots—Google Gemini, Microsoft Copilot, and DeepSeek-Chat—to generate an evaluation. The prompt was simple and consistent across all models: *Review this paper in English. In a single, fluent text, highlight its novelty and what is unclear or incorrect. Also, provide a score from -3 (strong reject) to +3 (strong accept).* This was followed by the complete text of the paper.

This paired structure (Paper → Human Review + LLM Reviews) enables a direct and quantifiable contrastive analysis. The intrinsic variety of the papers ensures that the dataset tests the capabilities of LLMs across a broad and realistic spectrum of reviewing scenarios.

Table 1 provides the detailed distribution of the 100 papers across the source conferences, their rank, and the number of papers selected from each.

Table 1: Distribution of papers across the source conferences.

Conference abbreviation	CORE Rank	Number of papers
JCDL 2018	A*	49
BringITon! 2014	No rank	14
ConsILR 2015	No rank	13
KES 2020 (IS 14)	B	11
KES 2021 (IS 33)	B	3
AICCSA 2020	C	2
AICCSA 2021	C	2
ConsILR 2014	No rank	2
ICCCI 2022	B	2
ACIIDS 2022	B	2
Total		100

The primary purpose of the HLLM-Dist corpus is to facilitate the calculation of a standardized confidence score for an AI model by comparing its output to the human benchmark, thereby offering a replicable methodology for evaluating any new automated review system.

2.2. AI Review Generation Methodology

Just as each scientific conference has its own review form, the three chatbots also have different review models. The responses received to the request for a review of a scientific paper differ even in terms of structure.

Gemini AI Review. Gemini is the AI assistant developed at Google, part of the generation of advanced AI models that succeed the Bard platform (a conversational artificial intelligence chatbot that can generate texts of all types). Gemini schematically presents information about originality and novelty, clarity and accuracy, use cases, information for investors, conclusion and score. The language is fluent.

Review of "Decentralized Monitoring of Smart Environments"

Summary

The project proposes a novel framework for decentralized monitoring in smart environments, addressing current limitations in compositionality and scalability of ambient intelligence applications. By utilizing flexible "monitors" that can work independently or in collaboration, the system aims to create a versatile observation mechanism to assist humans in daily tasks.

Originality and Novelty

The emphasis on a technology-independent monitoring system is a unique and promising concept, especially in the era of the Internet of Things (IoT). The approach of integrating various low-capacity monitors (e.g., speedometers, medical sensors) into a cohesive system is both innovative and relevant. The idea of mobile devices contributing to monitoring capabilities without relying on a central internet connection introduces a compelling dimension to smart environments.

Clarity and Accuracy

The project description provides a solid theoretical foundation, but could benefit from enhancements in the following areas:

Technical Framework: The explanation of "monitors" and their functionalities could be enriched with specific examples and diagrams illustrating how various monitors interact in real-life scenarios. This would help the reader visualize the system's operation.

Real-world Applications: While use cases are provided, there is a lack of detailed illustrations of how these applications would work in practice. Clear case studies demonstrating the effectiveness of the proposed system in various contexts (e.g., public transport, smart homes) would strengthen the relevance of the project.

Communication Protocols: More information on the communication algorithms, protocols, and mechanisms that allow these monitors to interact would enhance understanding. Are there specific challenges related to mobility and communication that this system addresses?

Robustness and Reliability: The proposal mentions concerns about system robustness in the face of incorrect observations from malfunctioning monitors. More detailed strategies to handle such issues, including protocols for error detection and correction, would be valuable.

Research Landscape: The assertion that decentralized monitoring is a relatively young field is compelling but could be supported with an overview of the current state of research, alongside key publications and their contributions.

Use Cases

The examples provided regarding accessing information in low-connectivity environments and sharing shopping data are relevant and relatable. However, expanding upon these scenarios with

specific workflows or outcomes would enhance their persuasive impact. For instance, how could users benefit from shared data in real time, and what potential privacy issues need addressing?

Information for Investors

The potential in decentralized monitoring as highlighted in this project is significant. The emphasis on providing unifying algorithms for device interoperability positions it well within the growing IoT market. To attract investors, it would be beneficial to include projected market growth numbers, competitive analysis, and potential revenue models beyond research grants.

Conclusion and Score

The project presents an innovative approach to decentralized monitoring in smart environments, emphasizing flexibility and technology independence. While the theoretical foundations are strong, additional clarity, detailed use cases, and strategies for practical implementation would enhance its overall presentation and relevance.

Score: +2 (Accept with Modifications) - Gemini

The project is promising and innovative, with acceptance contingent upon addressing specific areas for clarity and elaboration that would strengthen its technical and practical appeal.

Figure 1: Gemini AI review for paper 11, Bring-IT-on! 2014

Copilot Review. Copilot is an AI-based assistant developed by Microsoft, designed to integrate advanced generative AI functionalities into Microsoft 365 applications (Word, Excel, PowerPoint, etc.). Copilot provides a review with a classic structure and includes: summary, novelties, original aspects, criticisms, conclusion and technical observations (Figure 2).

Review of “Decentralized Monitoring of Smart Environments” by xxxxxxxxxx

This paper proposes a decentralized, technology-independent framework for monitoring smart environments using mobile and composable “monitors.” The central idea—that simple, limited-capability devices can collaborate to produce complex, context-aware observations—is both timely and compelling. The work stands out for its ambition to generalize runtime verification techniques to mobile, resource-constrained, and intermittently connected systems. This shift from centralized, static monitoring to dynamic, distributed sensing is a notable contribution, especially in the context of the Internet of Things (IoT).

The novelty lies in several aspects:

-  **Mobility-aware monitoring:** Extending runtime verification to mobile monitors is a fresh angle, especially considering the challenges of intermittent communication and dynamic topology.
-  **Compositional sensing:** The idea of combining heterogeneous sensors to infer higher-level properties (e.g., detecting a sleeping person from light and sound levels) is well-motivated and practically relevant.
-  **Opportunistic collaboration:** Leveraging nearby devices (e.g., smartphones) to enrich sensing capabilities without requiring centralized infrastructure is a powerful concept for low-connectivity environments.

However, several areas lack clarity or rigor:

-  **Terminology and abstraction:** The term “monitor” is used broadly, encompassing both physical sensors and software agents, but the abstraction is not formally defined. It’s unclear how monitors are instantiated, how they discover each other, and how they coordinate.

-  **Technical depth:** While the paper references runtime verification, it does not specify which techniques are adapted or how they are modified to suit mobility. There's no discussion of algorithms, protocols, or system architecture.
-  **Robustness and fault tolerance:** The paper raises the issue of incorrect observations but does not propose concrete mechanisms for handling faulty or malicious monitors.
-  **Evaluation:** There is no experimental validation, simulation, or even a theoretical analysis of the proposed system's performance, scalability, or reliability.
-  **Use cases:** The examples are imaginative but speculative. For instance, sharing cached news via nearby phones assumes a level of trust and compatibility that is not addressed. Similarly, the shopping app scenario raises privacy and data integrity concerns that are not discussed.

In summary, the paper presents an intriguing and original vision for decentralized monitoring in smart environments. Its strength lies in the conceptual framing and potential applications. However, it falls short in terms of technical specificity, formalism, and empirical support. With more rigorous modeling and validation, this could evolve into a strong contribution.

- **Copilot - Final Score: +1 (Weak Accept)** The idea is novel and promising, but the execution needs refinement and deeper technical grounding to fully convince a research audience

Figure 2: Copilot Review for paper 11, Bring-IT-on! 2014

DeepSeek Review. DeepSeek is an artificial intelligence assistant developed by a Chinese AI company. DeepSeek keeps the same classic, clear structure: novelty, unclear or incorrect elements, justification, strengths, weaknesses, recommendations (Figure 3).

Deepseek - Score: +1 (Weak Accept)

Review of "Decentralized Monitoring of Smart Environments"

Novelty

*The project proposes a **novel runtime verification framework** for decentralized monitoring via mobile "monitors" that dynamically compose observations across heterogeneous devices. Key innovations include:*

- **Technology-agnostic monitor abstraction** enabling ad-hoc collaboration between diverse sensors (e.g., smartphones interpreting infrared signals from legacy motion sensors).
- **Graceful degradation design** where system capabilities scale dynamically with monitor availability—a unique approach contrasting static IoT architectures.
- **Mobility-aware verification extensions** generalizing distributed systems theory to transient networks (e.g., adapting when monitors leave/join).

Unclear or Incorrect Elements

1. **Protocol ambiguity:** Claims about monitor communication ("monitors can interact and coordinate") lack technical specifics (e.g., discovery protocols, synchronization mechanisms), risking an architectural handwave.
2. **Privacy contradictions:** The shopping use case assumes "selective sharing" but omits how data ownership/consent is enforced—contradicting GDPR/CCPA compliance needs.
3. **Resource limits unaddressed:** While noting monitors' "limited resources," no solutions are proposed for battery optimization or edge-compute constraints during coordination.
4. **Validation gap:** Robustness claims against "broken monitors" are untested—crucial for medical/transportation use cases cited.

5. Citation inflation: <i>Characterizing decentralized monitoring as "young with 5-6 articles" underrepresents established fields (e.g., sensor networks, Swarm Robotics) dating back 15+ years.</i>
Score: +1 (Weak Accept) Justification: • Strengths: The core concept of dynamic monitor composition is compelling and underexplored. Use cases (e.g., <i>offline info-sharing, price tracking</i>) demonstrate pragmatic value. Mobility-focused runtime verification extends academic frontiers. • Weaknesses: Implementation vagueness and unexamined privacy risks undermine deployability. Overlooking resource constraints (power, bandwidth) in mobile-centric coordination is a notable omission. • Verdict: The framework holds research promise but remains <i>prototypical</i>. Recommend acceptance contingent on: 1. Technical appendices detailing communication protocols. 2. Privacy-by-design revisions (e.g., federated learning for shopping data). 3. Benchmarks validating resource-efficiency claims.

Figure 3: DeepSeek review - paper 11, Bring-IT-on! 2014

3. Architecture of the CAIR Confidence Score Metric

The CAIR metric is designed to quantify the alignment between an AI-generated review and a benchmark of human expert reviews. Its architecture is a multi-stage pipeline that transforms raw text reviews into a single, normalized confidence score.

The CAIR pipeline consists of four main stages:

1. Data Acquisition and Preprocessing
2. Feature Extraction and Representation
3. Model Training and Calibration
4. Confidence Score Calculation and Normalization

3.1. Data Acquisition and Preprocessing

All the information necessary to estimate the confidence level of the reviewers is extracted from the HLLM-Dist Corpus and reorganized into Excel files with information about conferences, papers, au-thors, reviews, reviewers, etc. To simplify data loading and processing, files are saved in CSV format.

All the data needed in the experiment described below is stored in the reviews table:

- *paper_id* – identifier of the paper; it is used as a key for accessing information about the paper in the other tables but also for grouping reviews of the same paper in this experiment.
- *reviewer_id* – identifier of the reviewer; it is used to distinguish between experts. *reviewer_id* is a string of the form “Human_<nn>”) for human reviewers or one of “Gemini”, “Copilot”, “DeepSeek” for the AI reviewers.

AI REVIEWER'S CONFIDENCE

- *confidence* – self-reported expertise level of the human reviewer; an integer value in range [2; 5] which serves as the ground truth label for training; it is not specified for the AI reviewers.
- *review* – the full text of the review.
- *decision* – the reviewer's recommendation, expressed on an ordinal scale from strong reject to strong accept.
- *final decision* – final decision of the conference program committee; same coding as for decision;

Considering that the data comes from different sources after loading into memory, preprocessing is necessary to eliminate situations that may lead to errors (e.g. icons in the Copilot AI review). Thus, the following processing is performed (Python code in Figure 4):

- *Handling Missing Data*: Cells with NaN values are replaced with an empty string, and all values are converted to string type.
- *Text Cleaning*: The review text is stripped of leading/trailing whitespace and any non-printable characters to ensure consistency.
- *Decision Normalization*: Decision labels are mapped to a numerical scale ranging from -3 (strong reject) to +3 (strong accept), enabling quantitative analysis.

Output: After completing this stage, a cleaned, structured dataset ready for feature extraction is obtained.

```
# text cleaning
def clean_text(x):
    if pd.isna(x):
        return ""
    text = str(x).strip()
    text = "".join(ch for ch in text if
ch.isprintable())
    return text
df["review"] = df["review"].apply(clean_text)
# decision mapping to numeric values
df["decision"] =
df["decision"].str.lower().str.strip()
decision_map = {
    "strong reject": -3,
    "reject": -2,
    "weak reject": -1,
    "borderline": 0,
    "weak accept": 1,
    "accept": 2,
    "strong accept": 3
}
df["decision_num"] =
df["decision"].map(decision_map)
```

Figure 4: Code sequence for input data preprocessing

3.2. Feature Extraction and Representation

This stage converts the unstructured text and categorical data into numerical feature vectors that a machine learning model can process. The classification procedure requires the aggregation of both review and decision. Thus, the features are extracted in two steps described below.

- First, the content of all reviews is then transformed into semantic vectors using the *distilbert-base-uncased* model. The mean of the last hidden state representations is used as the embedding for each review. A 768-dimensional embedding vector that represents the semantic content and contextual meaning is computed for each review. It captures the quality and substance of the critique beyond the score.

DistilBERT (Sanh *et al.*, 2019) is a compact version of the Bidirectional Encoder Representations from Transformers (*BERT*) model. It is constructed through knowledge distillation, a process in which a smaller “student” model is trained to approximate the behaviour of a larger “teacher” model. *DistilBERT* is trained to reproduce the predictions of BERT-base while maintaining competitive performance across a variety of natural language understanding tasks. *DistilBERT* retains the Transformer encoder structure of BERT, but reduces the number of layers by half. While BERT-base consists of 12 encoder layers, *DistilBERT* employs only 6 layers. The hidden size and attention heads are preserved (768 hidden units, 12 attention heads), enabling the model to capture rich contextual representations of text while requiring fewer computational resources.

- In the second step, the already mapped decisions are encoded using the One-Hot Encoding. The mapping to the *numerical* scale performed in the preprocessing step is important to the classifier that will assume that the difference between categories is linear and ordered. Otherwise, if we had kept the decisions represented as a string of characters, errors may occur in the evaluation because for example, “weak reject” and “weak accept” may be semantically closer than their numeric values suggest. One-Hot encoding treats all decisions as independent categorical features, letting the model learn the associations freely. After One-Hot encoder is applied each decision becomes a 7-dimensional binary vector in which one element is 1 and all others are 0.

Feature Fusion. The final feature vector for each review is created by concatenating the *DistilBERT* text embedding vector with the one-hot encoded decision vector. This combined 775-dimensional vector represents both the *content* and the *conclusion* of the review.

```
[Final Feature Vector (775 dim)] =
  [DistilBERT_Review Embeddings (768 dim)] + [One-
Hot_Decision (7 dim)]
```


3.3. Model Training and Calibration

A logistic regression classifier was trained to predict reviewer confidence. Logistic regression is a supervised learning algorithm commonly used for classification tasks. It models the probability that a given input belongs to a particular class using the sigmoid function, which maps linear combinations of input features to values between 0 and 1. The model estimates parameters by maximizing the likelihood of the observed data, typically through iterative optimization methods. It is widely used due to its interpretability, efficiency, and robustness, particularly as a baseline classifier or when combined with high-dimensional feature representations such as embeddings.

To use the classifier, the features dataset was partitioned into two subsets:

- *train subset* – all 200 reviews with available confidence scores of the reviewers,
- *test subset* – all 300 reviews with missing confidence scores.

The reviewer's confidence score is the class which must be predicted. The trained model learns the patterns and characteristics in feature vectors that correspond to high-quality, expert human reviews.

3.4. Confidence Score Calculation & Normalization

However, after classification, the reviewer's confidence may differ from one paper to another. To ensure consistency across reviewers, predicted final confidence values were adjusted at the reviewer level using the following rules:

- if a reviewer has only missing values (this is our case), the mean of the predicted confidences is used;
- if a reviewer had at least one known confidence value, missing values are replaced with the mean of the known confidence values.

This procedure resulted in a complete set of confidence values (*final_confidence*) for all reviews.

Decision Aggregation

After the reviewers' confidences were calculated, reviewer decisions were aggregated into a confidence-weighted score on the $-3 \dots +3$ scale for each paper. This score was calculated as the weighted average of reviewer decisions, where the weights were given by the *final_confidence* values. This approach ensures that reviews expressed with higher confidence contribute more strongly to the aggregate decision.

Evaluation Metrics

Two complementary evaluation strategies were applied:

- *Decision accuracy*. The sign of the weighted score for each paper (accept vs. reject) was compared against the global final decision. Accuracy was defined

as the proportion of cases in which the aggregated score and the final decision agreed in polarity.

- *Confidence-Adjusted Inter-Reviewer Agreement (CAIR)*. To quantify agreement among reviewers, we computed CAIR for each paper.

$$CAIR_{paper} = 1 - \frac{\sum_{i < j} confidence[i] * confidence[j] * |decision[i] - decision[j]|}{\sum_{i, j} confidence[i] * confidence[j] * 6},$$

where *confidence* are the normalized values of the confidence of paper's reviewers and 6 is the difference between *max* and *min* values of decision: $3 - (-3)$. Unlike traditional agreement measures that consider only categorical or ordinal decisions, CAIR metric measures the average pairwise discrepancy between reviewers' numeric decisions, weighted by their confidence. Values close to 1 indicate strong consensus among reviewers, particularly when high-confidence reviews are aligned, whereas values near 0 denote substantial disagreement, especially when confident reviewers diverge substantially. By construction, CAIR distinguishes between cases where disagreements are expressed with low confidence (less penalized) and those where conflicting strong opinions dominate (more penalized).

Future Architecture: ChatGPT as Normalizer

The text mentions plan to use ChatGPT in subsequent phases. Its role in the architecture would be:

1. To generate additional "reference" reviews for the papers.
2. To be used as a benchmark LLM against which the outputs of other models (Gemini, Copilot, DeepSeek) are compared and normalized, ensuring the CAIR score is consistent across different evaluation benchmarks.

This architecture provides a scalable, quantitative framework for assessing AI reviewer confidence without requiring new human evaluations for each new model, fulfilling the core objective of the research.

4. Results and Discussion

This section presents the empirical findings from our large-scale experiment to evaluate AI-generated peer reviews against a human benchmark. The analysis is conducted on the complete HLLM-Dist corpus, which includes 100 scientific papers, each with 2 original human reviews (for a total of 200) and 3 AI-generated

AI REVIEWER'S CONFIDENCE

reviews from Microsoft Copilot, Google Gemini, and DeepSeek-Chat (for a total of 300), resulting in 500 reviews for analysis.

4.1. Classifier Performance on Human Reviews

The logistic regression classifier, trained on DistilBERT embeddings and one-hot encoded decisions from the human reviews, is the core of our confidence estimation engine. Its performance on a held-out test was strong, achieving a 96% accuracy in predicting a human reviewer's self-reported expertise level based on their review's semantic content and final verdict. This high accuracy supports our feature extraction and modelling approach, indicating the model's ability to assess the quality of AI-generated reviews.

4.2. Confidence Scores for AI Reviewers

The trained model was applied to all 300 AI-generated reviews to predict a *final_confidence* score for each. The results, aggregated by AI system, provide a direct answer to our core research question: AI REVIEWER'S CONFIDENCE.

All three AI systems performed with consistency across all 100 papers. DeepSeek achieved the highest mean confidence score of 4.00, meaning its reviews were, on average, predicted to be similar to those written by a human with a "high expertise" level. Both Copilot and Gemini followed with a mean confidence score of 3.82. This result shows that while all three models produce quality output, DeepSeek has a slight edge in this reviewing task. It is also important to note that the AI confidence scores fall within the range of human scores (2.0 to 5.0), which points to their potential utility as assistants in peer review.

Table 2: Results obtained for papers P01 and P02

Paper ID	Reviewer ID	Reviewer confidence	Reviewer confidence	Reviewer decision	Reviewer decision	Final committee decision	Predicted confidence	Final confidence	Aggregated decision
P01	Human 01	5	expert	weak reject	-1	weak accept	NaN	5	weak accept
P01	Human 02	4	high	accept	2	weak accept	NaN	4	weak accept
P01	Copilot	NaN	NaN	weak accept	1	weak accept	4	3.8	weak accept
P01	DeepSeek	NaN	NaN	weak accept	1	weak accept	4	4	weak accept
P01	Gemini	NaN	NaN	weak reject	-1	weak accept	4	3.8	weak accept
P02	Human 03	4	high	accept	2	weak accept	NaN	4	weak accept
P02	Human 04	3	medium	weak reject	-1	weak accept	NaN	3	weak accept
P02	Gemini	NaN	NaN	accept	2	weak accept	4	3.8	weak accept
P02	Copilot	NaN	NaN	weak accept	1	weak accept	4	3.8	weak accept
P02	DeepSeek	NaN	NaN	accept	2	weak accept	4	4	weak accept
...									

An extract from the table with the final results is presented in Table 2. The values in column *confidence* are NAN for the AI reviewers and similarly those in column *predicted_confidence* for the human reviewers. The column *final_confidence* contains the aggregated confidence for all reviewers.

The final confidence-weighted decision is also computed for all papers. The distribution of these values across the entire dataset is shown in Figure 5.




Figure 5: Weighted decision distribution for all papers

4.3. Decision Aggregation and CAIR Score Validation

The practical use of the confidence scores was tested by calculating a confidence-weighted aggregate decision for each of the 100 papers and comparing it to the actual final decision from the conference committee. The outcome was clear: the polarity (accept/reject) of the AI ensemble's weighted score matched the human committee's final decision for 100% of the papers (Decision accuracy: 1.00). This shows that a composite score from multiple AI reviewers, weighted by their predicted confidence, can reflect the final outcome of a human committee. This outcome validates our framework's ability to combine AI critiques into a collective recommendation. Furthermore, the Confidence-Adjusted Inter-Reviewer Agreement (CAIR) scores for the papers were high (Figure 6), ranging from 0.73 to 0.95 (where 1 indicates perfect agreement).

This indicates a solid consensus among the reviewers—both human and AI—for each paper, confirming that the reviews in our corpus are coherent evaluations. Until now, the results demonstrate the feasibility of the CAIR metric on a substantial dataset. We have quantified AI reviewer confidence, noted performance differences between models, and shown that these scores can be used to aggregate decisions that align with human expert outcomes for 100 papers.




Figure 6: CAIR values distribution for all papers

5. Conclusions

In this paper, we proposed and piloted a novel, scalable framework for quantifying the confidence we can place in an AI-generated peer review. By constructing the HLLM-Dist corpus and developing the CAIR metric, we have taken a significant step towards calibrating trust in AI-assisted scientific evaluation. Our results confirm the initial hypothesis that significant divergence exists between AI models, but crucially, they also demonstrate that this divergence can be measured and scored reliably against a human benchmark. The pilot validation shows that current LLMs like DeepSeek, Copilot, and Gemini can generate reviews that a calibrated model associates with a high level of human expertise, providing an empirical basis for their cautious integration into the peer-review workflow. The long-term implication of this research is the establishment of a foundational model that can predict the confidence level of any new AI reviewer without requiring extensive manual evaluation for each new system. This work ultimately provides a necessary tool for ensuring the integrity and maintaining the quality of the scholarly peer-review process in the emerging era of artificial intelligence.

AI Use Statement

In order to assess the confidence level of AI reviewers, in the dataset were included 300 automatically generated reviews for 100 scientific papers, as follows: Microsoft Copilot (<https://copilot.microsoft.com/>) – 100 reviews, Google Gemini (<https://gemini.google.com/app?hl=ro>) – 100 reviews, and DeepSeek (<https://chat.deepseek.com/>) – 100 reviews. The same prompt was used for the three models: *“Review this paper in English. In a single, fluent text, highlight its novelty and what is unclear or incorrect. Also, provide a score from -3 (strong reject) to +3 (strong accept)”*. Three examples of such reviews are presented in

Figures 2, 3, and 4, where the AI generated texts were modified by authors solely in terms of formatting.

Acknowledgment

This work was carried out partially within the “Tools for Processing Online Texts Specific to Cultural and Scientific Diplomacy” project, funded by the Academy of Romanian Scientists.

References

- Aarne Talman, Marianna Apidianaki, Stergios Chatzikyriakidis, and Jorg Tiedemann. 2022. How Does Data Corruption Affect Natural Language Understanding Models? A Study on GLUE Datasets. In V. Nastase, E. Pavlick, M.T. Pilehvar, J. Camacho-Collados, and A. Raganato (Eds.), *Proceedings of the 11th Joint Conference on Lexical and Computational Semantics*, 226-233. Association for Computational Linguistics.
- Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2020. Evaluation of Text Generation: A Survey. arXiv preprint arXiv:2006.14799. DOI: 10.48550/arXiv.2006.14799.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How NOT To Evaluate Your Dialogue System: An Empirical Study of Unsupervised Evaluation Metrics for Dialogue Response Generation. In J. Su, K. Duh, and X. Carreras (Eds.) *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2122–2132. DOI: 10.18653/v1/D16-1230.
- Daniela Gîfu and D. Cristea. 2012. Multi-dimensional Analysis of Political Language. In J. J. Park, V. C.M. Leung, C.-L. Wang, & T. Shon (Eds.), *Lecture Notes in Electrical Engineering*, 164(1):213-221. Springer.
- Daniela Gîfu and Silviu-Vasile Covaci. 2025. Artificial Intelligence vs. Human: Decoding Text Authenticity with Transformers. In *Future Internet Journal*, 17(1):38.
- Daniela Gîfu. 2014. Humor in the Religious Discourse: Between Paradoxism and Neutrosophy. In F. Smarandache & Ș. Vlăduțescu (Eds.), *Communication Neutrosophic Routes*, 88-104. Educational Publishing.
- Daniela Gîfu. 2015. Contrastive Diachronic Study on Romanian Language. In S. Cojocaru & C. Gaiandric (Eds.), *Proceedings of the Workshop on Foundations of Informatics FOI-2015*, August 24-29, 2015, Chisinau, Republic of Moldova.
- Diana Trandabăt and Daniela Gîfu. 2023. Discriminating AI-generated Fake News. In: *Procedia Computer Science*, open access journal, ELSEVIER, 225, 3822-3831, DOI: 10.1016/j.procs.2023.10.378.
- Gary Marcus. 2020. The Next Decade in AI: Four Steps towards Robust Artificial Intelligence. arXiv preprint arXiv:2002.06177. DOI: 10.48550/arXiv.2002.06177.
- Gemini Team Google: R. Anil, S. Borgeaud, Y. Wu, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, *et al.* 2023. Gemini: A Family of Highly Capable Multimodal Models. arXiv preprint arXiv:2312.11805. DOI: 10.48550/arXiv.2312.11805.
- Jindong Wang, Xixu Hu, Wenxin Hou, Hao Chen, Runkai Zheng, Yidong Wang, Linyi Yang, Haojun Huang, Wei Ye, Xiubo Geng, Binxin Jiao, Yue Zhang, and Xing Xie.

2023. On the Robustness of ChatGPT: An Adversarial and Out-of-distribution Perspective. arXiv preprint arXiv:2302.12095. DOI: 10.48550/arXiv.2302.12095.
- Jonathan P. Tennant and Tony Ross-Hellauer. 2020. The Limitations to Our Understanding of Peer Review. In *Research Integrity and Peer Review*, 5(6):1-14. DOI: 10.1186/s41073-020-00092-1.
- Joshua Maynez, Sashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On Faithfulness and Factuality in Abstractive Summarization. In D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Online, 1906–1919. Association for Computational Linguistics.
- Lutz Bornman and Hans-Dieter Daniel. 2010. The Manuscript Reviewing Process: Empirical Research on Review Requests, Review Sequences, and Decision Rules in Peer Review. In *Library & Information Science Research* 32:5–12. DOI: 10.1016/j.lisr.2009.07.010.
- OpenAI. 2023. GPT-4 Technical Report. arXiv preprint arXiv:2303.08774. DOI: 10.48550/arXiv.2303.08774.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yukuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Cristopher D. Manning, Cristopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladn Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2022. Holistic Evaluation of Language Models. arXiv preprint arXiv:2211.09110. DOI: 10.48550/arXiv.2211.09110.
- Ştefan Vlăduţescu, Florentin Smarandache, Daniela Gîfu, Alina Ţenescu (Coords.). 2014. *Topical communication uncertainties*. Sitech & Zip Publishing.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. arXiv preprint arXiv:1910.01108.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of Hallucination in Natural Language Generation. In *ACM Computing Surveys*, 55(12):1–38. DOI: 10.1145/3571730.

CROSS-PLATFORM GENAI EVALUATION FOR ESP LEARNING SCENARIOS: A COMPUTATIONAL ANALYSIS OF LINGUISTIC AND CULTURAL APPROPRIATENESS*

ANDREEA NECHIFOR¹, ALINA TURCULET², ANCA VASILESCU³

¹ *Transilvania University of Braşov, Faculty of Letters*
andreeabratan@unitbv.ro

² *Transilvania University of Braşov, Faculty of Psychology and Educational Sciences*
alina.turculet@unitbv.ro

³ *Transilvania University of Braşov, Faculty of Mathematics and Informatics*
vasilex@unitbv.ro

Abstract

The present study is based on the evaluation of over 180 mobility scenarios AI-generated by five platforms (ChatGPT, Claude, Copilot, DeepSeek, and Gemini) by means of an ample framework that integrates perspectives from applied linguistics and computational linguistics. Based on Hutchinson and Waters' (1987) ESP framework, Byram's (1997) intercultural competence model, and Thomas's (1983) pragmalinguistic distinctions, the research employs five evaluation criteria: pedagogical authenticity, linguistic appropriateness, cultural sensitivity, discourse coherence, and pragmatic authenticity. Results disclose platform variations, pragmatic authenticity averaging around a solid 4.03, cultural sensitivity (4.55) and pedagogical value (4.16) denoting advanced AI competencies, while discourse coherence (3.92) and linguistic appropriateness (3.88) still presenting challenges. Gemini performed the best across most dimensions, while linguistic register management and conversational naturalness demonstrated to be problematic across all platforms. A validation study was also undertaken to compare the automated estimates generated by Llama 3.1 with human specialist ratings, which produced a moderate positive correlation ($\rho = 0.40$, $p = .51$), thus indicating partial alignment between computational and expert findings, even if statistical significance was not achieved considering the limited sample size ($N = 5$). The results demonstrate that current GenAI platforms possess significant ESP content generation potential, at the same time requiring careful evaluation structures supported by corpus-based analysis (Swales, 1990; Charles, 2012) and applied linguistics perspectives (Canagarajah, 2013; Pennycook, 2001)

* This work is licensed under a Creative Commons Attribution 4.0 International Licence, as described at <http://creativecommons.org/licenses/by/4.0/>

to ensure culturally sensitive and pedagogically authentic professional language education materials, with human supervision remaining essential for quality assurance.

Keywords: Generative AI, computational linguistics, applied linguistics, cultural awareness, intercultural communication, professional communication, assessment framework, pragmatics, ESP

Introduction

Despite the growing interest in language acquisition mediated by AI, limited research has systematically evaluated GenAI platforms' potential for ESP learning, especially from an integrated threefold point of view: sociolinguistic, pragmalinguistic and computational, as most studies focus on either technological capabilities or pedagogical applications, without inclusive frameworks that could concurrently address cultural sensitivity, linguistic appropriateness, pragmatic competence, and pedagogical authenticity.

Thus, the literature review extensively covered below reveals the need for methodical evaluation structures that incorporate ESP theoretical basics, intercultural communication concepts, corpus-based analysis methods, and applied linguistics perspectives, subsequently proposing an assessment framework that aims to address this gap by providing complex evaluation criteria that include the intricacy of professional communication while showcasing practical relevance for ESP educators and curriculum developers.

Theoretical Foundations

English for Specific Purposes, as theorised by Hutchinson and Waters (1987), favours a learner-centred approach to language acquisition that prioritises authentic, context-specific communication needs, mostly related to professional jargon, thus profiling ESP instruction as driven by needs, derived from a systematic analysis of target language use in work-related contexts, addressing professional, vocational or academic domains, focused on skills, therefore emphasising communicative competence in specific contexts, and based on authenticity, using genuine materials and realistic scenarios. In this way, this perspective becomes particularly relevant when evaluating GenAI-generated scenarios, because the technological effectiveness must be measured against authentic professional communication standards and not against standards that address the general language proficiency system of measurement.

Likewise, Dudley-Evans and St. John (1998) refined the ESP's frame by emphasising how important genre specificities are, alongside discourse conventions. Consequently, their work revealed that effective ESP instruction requires deep understanding of not only appropriate linguistic configurations, but also of cultural representations that characterise usually multicultural professional communities, providing crucial insights with respect to GenAI platforms'

productions, as technological outputs should exhibit both linguistic accuracy against genre-specific authenticity and cultural awareness.

Furthermore, Hyland's (2004) analysis of professional discourses revealed that communication in specialised communities involves conventions characterised by specific features, i.e., each professional domain displays distinct genre parameters consisting of textual configurations, register requirements, and organisational structures, as well as definite terminology that refers to more than mere technical terms. This one includes, according to Hyland, pragmatic markers, hedging devices, and viewpoint or position indicators, but also cultural elements, as professional discourse reflects essential cultural values, community practices and power relations. Further supporting this perspective is Tardy's (2009) research based on corpora analyses of academic writing, which shows that professional communities are defined by specific discourse conventions that do not comply with simplification or generalization rules, in this way Hyland's perspective suggesting that platform effectiveness cannot be evaluated exclusively through the lenses of linguistic accuracy but should also include cultural relevance and genre accuracy.

Moreover, Johns *et al.* (2011) predicted that ESP would develop toward a more nuanced understanding of multimodal communication, acknowledging that profiling professional settings entangles the integration of digital, textual, and visual information in addition to intercultural competence because professionals are expected to balance power dynamics and ideological viewpoints within their communities. Thus, but also due to the fact that these features go beyond conventional linguistic measurements incorporating both pragmatic appropriateness and cultural sensitivity, they can therefore be used to derive additional crucial evaluation criteria for scenarios generated by GenAI platforms. Flowerdew's (2013) work on corpus-based ESP research methodology provides additional support for this expanded evaluation framework, emphasising the need for authentic data comparison when assessing pedagogical materials.

The integration of AI in education represents a fashionable research topic in our contemporaneity, challenging researchers from all fields to explore both opportunities and drawbacks (Cristea *et al.*, 2023; Popa and Vasilescu, 2025), but the history of machine assisted language learning, for example, dates back to Stockwell's (2012) interest in mobile-assisted language learning (MALL), which emphasised the importance of contextual appropriateness and cultural sensitivity, and continued with Chapelle and Sauro (2017), who provided a complex framework for evaluating computer-assisted language learning (CALL)'s effectiveness, thus forecasting GenAI platforms' assessment for ESP contexts. Additionally, the importance of task authenticity, the focus on meaning, and the importance of learner agency are more recent investigations in TELL (technology assisted language learning), which demonstrate that effective GenAI integration in mediated ESP instruction requires pedagogical affiliation, the content generated

by means of AI support serving clear learning objectives, thus promoting quality. Complex professional scenarios need suitable linguistic and cultural support (Dimulescu and Nechifor, 2021), as technology should enhance rather than replace critical thinking and cultural awareness (Godwin-Jones, 2022).

Moreover, Fryer and Carpenter (2006) conducted early research on chatbot interactions in language learning contexts, revealing significant variations in learner engagement and linguistic development, their findings showing that effective AI-mediated language learning requires complex natural language processing (NLP) capabilities and culturally sensitive response generation. More recently, Dizon and Tang (2022) investigated the use of AI chatbots for developing speaking skills, finding significant variations in platform effectiveness based on task design and cultural context. Kohnke *et al.* (2023) examined ChatGPT's potential for language learning, identifying both opportunities and limitations in the platform's ability to provide appropriate feedback and cultural guidance, their study featuring important considerations for ESP practitioners, particularly regarding the need for careful evaluation of AI-generated content for cultural bias and professional authenticity. Turculeț (2024) further explored student perspectives on AI learning assistance in educational contexts, focusing on learner acceptance and adaptation to AI-mediated environments.

For a more complex and integrative perspective, the theoretical input of intercultural communication provides essential grounds for evaluating GenAI effectiveness in multicultural professional settings, having Byram's (1997) 'intercultural communicative competence model' identify five key components in this respect: cultural self-awareness, which makes one understand their own cultural positioning; cultural knowledge, referring to information about other cultural groups; intercultural skills, embodied in the ability to navigate cultural differences effectively; cultural empathy, defined as the capacity to recognise, accept, and integrate perspectives from other cultural frameworks; and cultural awareness, seen as the ability to evaluate cultural practices and beliefs critically. In the same line of thought, but with a spot on the educational environment, Nechifor and Borca (2016, 2020) spoke about 'culturacy' a term coined to conceptualise cultural awareness, cultural competence and cultural literacy, while Nechifor and Dimulescu stressed the importance of a culturally infused approach to MALL (Nechifor, 2015; Nechifor and Dimulescu, 2021 and 2024). Furthermore, Spencer-Oatey's (2008) 'rapport management theory' can extend the ESP framework design by emphasising the dynamic nature of intercultural encounters and the importance of contextual factors in determining appropriate communication strategies, having this theoretical perspective become crucial when evaluating how GenAI platforms represent and model intercultural professional interactions.

Another angle to defining ESP contexts is brought about by the concept of pragmatic competence which requires mastery of multiple interconnected skills.

Thomas's (1983) distinction between pragmalinguistics and sociopragmatics provides a useful perspective for this, pragmalinguistics involving the linguistic resources available for performing speech acts, and sociopragmatics referring to the social perceptions underlying participants' interpretations and performance of communicative actions. With Kasper and Rose (2002) further developing this frame by emphasising the importance of contextual factors in pragmatic competence development, effective pragmatic instruction also integrated linguistic forms and sociocultural appropriateness, a consideration that could become relevant when evaluating GenAI-generated scenarios for ESP learning. At the same time, and in perfect accordance with the focus of the current study on the field of mobility, Barron's (2003) longitudinal report of pragmatic development in 'study abroad contexts' can provide additional understandings of the complexity of intercultural pragmatic competence, her findings suggesting that effective pragmatic acquisition requires exposure to authentic interactions and explicit attention to cultural diversity in communication strategies.

Additionally, in Pennycook's view (2001), critical applied linguistics emphasises that ESP instruction cannot be neutral but must address language ideology, cultural hegemony, and identity negotiation, thus examining how certain linguistic varieties are privileged over others, recognising the dominance of cultural perspectives in professional contexts, and understanding how learners position themselves within professionally specialised groups. Similarly, Canagarajah's (2013) work on 'translingual practice' provides important insights for GenAI evaluation, as it highlights the importance of recognising linguistic diversity while, at the same time, avoiding monolingual bias in professional communication instruction, his perspective suggesting that effective ESP materials must acknowledge, observe, and incorporate multilingual competence rather than impose standardised linguistic norms. Thus, a critical perspective requires the examination of the voices that are represented in AI-generated scenarios, the extent to which some cultures are privileged while some others are marginalised, and perfectly aligned with Fairclough's (2013) research, how power relations are portrayed in specific contexts. Consequently, Kubota and Lin's (2009) work dedicated to race, culture, and identities in second language education can also be considered as additional frames for evaluating cultural representation in AI-generated content, alongside Holliday's (2005) concept of 'small cultures', which offers valuable insights for assessing how GenAI platforms should avoid essentialist cultural bias while maintaining cultural awareness and sensitivity.

However, even more practically, corpus linguistics provides essential methodological backgrounds for evaluating GenAI-generated ESP content, a modern debate positioning these two on adversary grounds, but without necessarily concluding that they are opposing each other. Subsequently, Flowerdew's (2004) contribution to corpus-based analysis of academic discourse

demonstrates the importance of comparing AI-generated content with authentic professional communication patterns, his methodology providing outlines for assessing lexical density, collocational patterns, and genre-specific linguistic features. Charles (2012) extends this approach by examining the role of corpus linguistics in ESP comprehension, emphasising the importance of data-driven learning approaches that expose learners to authentic language patterns, her research providing evaluation criteria that can be applied to assess how effectively GenAI platforms model authentic professional discourse.

Last but not least, Swales' (1990) genre analysis framework remains groundbreaking for ESP research and evaluation. His 'move-step' analysis provides systematic methods for examining how effectively GenAI platforms reproduce authentic professional communication patterns. This methodology becomes particularly important when evaluating the pedagogical authenticity of AI-generated scenarios. Bhatia's (2004) work on genre analysis in professional contexts extends Swales' pattern by emphasising the importance of contextual factors and cultural variation in genre recognition, his approach stipulating valuable insights for assessing how GenAI platforms handle cultural diversity in specialised professional communication scenarios.

Methodology

Based on comprehensive coding criteria and a substantial database of over 200 learning scenarios, this study investigates the effectiveness of GenAI platforms in creating culturally sensitive and linguistically appropriate ESP content for mobility contexts. The research framework addresses an acute gap in current research regarding systematic evaluation of GenAI tools in professional language education.

The first research question examines the extent to which different GenAI platforms prove varying levels of effectiveness across five evaluation criteria: pragmatic authenticity, linguistic appropriateness, cultural sensitivity, discourse coherence, and pedagogical authenticity. This corroborates Hutchinson and Waters' (1987) framework, which emphasises the idea of ESP being a field defined more directly by the needs of the learners and their context-driven realities, hence its specific instructional approach, and Hyland's (2004) perspectives on specialised discourse, which extends beyond linguistic appropriateness and comprises genre conventions and cultural markers.

The second research question, aligned with Thomas's (1983) distinction between pragmalinguistics and sociopragmatics, scrutinises how effective AI-generated scenarios are in shaping and generating pragmatic competence needed for successful intercultural professional communication, the investigation assessing not only whether chatbots perform speech acts appropriately within specified contexts, but if they can also demonstrate cultural sensitivity, display politeness strategies, and exhibit genuine context awareness. These dimensions

become particularly relevant when taking into account Kasper and Rose's (2002) emphasis on the interconnection between linguistic forms and sociocultural aspects with respect to developing pragmatic competence.

The third research question is interested in evaluating aspects such as cultural insertion and power dynamics in AI-generated learning scenarios, addressing the dedicated applied linguistics perspectives that emphasise, according to Pennycook (2001) and Canagarajah (2013) the status of ESP instruction as non-neutral, the investigation thus examining how well these scenarios avoid biased cultural stereotypes, involve individuals as active agents rather than cultural mechanisms, and model cultural learning and adaptation in an effective way. This critical lens is essential for understanding whose voices are privileged or marginalised in AI-generated professional communication scenarios, as Fairclough (2013) and Holliday (2005) argue regarding the importance of avoiding stereotypical cultural representations.

The methodological approach integrates both quantitative and qualitative discourse analysis criteria, creating what Creswell and Plano Clark (2018) call a 'convergent parallel mixed-methods design', as the proposed outline codes specific criteria through systematic evaluation protocols by applying 1-5 Likert scale ratings across all four scenario types in order to examine the 5 main criteria as, on the one hand, distinct but, on the other hand, interrelated constructs.

In more detail, the quantitative segment is based on scalar ratings which have each criterion evaluated from 1 (Very Poor) to 5 (Excellent). Regarding pragmatic authenticity, the code was programmed to assess how appropriately speech acts are produced in context, the degree to which cultural sensitivity is intricately linked to politeness strategies, as well as how much context awareness is demonstrated. With respect to the linguistic appropriateness dimension the examination was interested in register suitability, appropriateness of grammatical structures, and presence and relevance of hedging devices use, in this way mirroring Biber's (1988) work on register variation and Charles's (2012) insights on corpus-based ESP analysis. Related to cultural sensitivity assessment, the model incorporates Byram's (1997) intercultural communicative competence standard, evaluating how successful scenarios are to avoid mainly negative stereotypes, present individual agency, and model cultural competence development, while the aspect evaluating discourse coherence draws on Swales's (1990) genre analysis skeleton and Bhatia's (2004) research on professional discourse, thus examining observance to genre conventions and also maintenance of natural, realistic conversational patterns. Finally, the pedagogical authenticity dimension, grounded in Dudley-Evans and St. John's (1998) learner-centred approach, assesses the clarity and the degree of achievability of the learning objectives, as well as the potential of skill-transferability to real-world contexts.

Furthermore, the qualitative component employs critical discourse analysis following Fairclough's (2013) three-dimensional framework, examining textual

features, discursive practices, and social practices within the generated scenarios, focusing on qualitative indicators clearly specified in the coded criteria: specialised vocabulary count, variety of hedging devices, appropriateness of cultural references, presence of genre markers, and frequency of pragmatic markers.

Thus, the integration of computational linguistic tools enables systematic extraction and analysis of these features across the entire dataset, providing what Flowerdew (2013) described as essential corpus-based evidence for ESP material evaluation. Inter-rater reliability based on human evaluation included, on the one hand, assessment from professionals in the applied linguistic field which validated the automated evaluation, and, on the other hand, insights from three representatives of the participating disciplines – linguistics, psychology, and informatics, the triangulation of perspectives from different disciplinary backgrounds responding to Johns *et al.*'s (2011) call for interdisciplinary approaches to ESP research that acknowledge the complexity of contemporary professional communication.

In other words, the present study employs a performance analysis focusing on five basic evaluation dimensions, trying to identify how different platforms handle four types of mobility contexts based on factor analysis, which exposes quality concepts and potential relations among different evaluation criteria, and on correlation analysis, which examines interdependencies among cultural sensitivity, linguistic appropriateness, and pragmatic authenticity. In this way, referring to Tashakkori and Teddlie's (2010) meta-inferences, these quantitative findings integrate with qualitative insights through joint display analysis, creating room for revealing both convergent and divergent patterns across methods.

Subsequently, the methodological framework addresses the research gap regarding comprehensive evaluation of GenAI platforms for ESP learning from an integrated sociolinguistic, pragmalinguistic and computational perspective. By systematically applying coding criteria across multiple platforms and scenario types, the study produces an empirically validated evaluation rubric that ESP practitioners can employ when selecting and implementing AI tools. The conclusions not only contribute to theoretical developments in technology-enhanced language learning in line with Chapelle and Sauro' (2017) but also provide practical guidelines for integrating GenAI in mobility preparation programs.

In summary, the emphasis of the current investigation on cultural sensitivity and pragmatic relevance addresses modern ESP developments that prioritise intercultural competence and critical discourse awareness (Spencer-Oatey, 2008; Kubota and Lin, 2009). By examining how different platforms handle cultural diversity across four different scenarios which focus on multicultural learning environments, social pragmatics, administrative challenges, and linguistic pragmatics, the study provides essential insights for educators preparing students

for increasingly complex global professional contexts. The systematic evaluation of pedagogical authenticity ensures that technological innovation serves genuine learning needs rather than replacing critical thinking and cultural awareness, as Godwin-Jones (2022) stresses in recent discussions of AI empowerment of language education.

Data Collection and Data Analysis Framework

During the 2nd semester of the 2025-2026 academic year, students enrolled in the 2nd year of the Applied Modern Languages section and in the 3rd year of the Philology profile, all belonging to the Faculty of Letters from Transilvania University of Braşov, worked on different assignments that integrated AI tools into their ESP classes to support specialised language acquisition. One such assignment, of interest to the current study, required them to create learning scenarios using assigned chatbots, starting from a task environment created by the teacher: ‘Using the chatbots assigned in class, build learning scenarios for you as students involved in a mobility, in which you have to deal with cultural diversity in a multicultural learning environment in class, social pragmatics (cultural differences, as part of the new country you belong to for a semester/year), administrative aspects, and linguistics pragmatics (different norms, conversational principles, politeness strategies, etc.). In a compulsory manner, include elements of specialised vocabulary and the most appropriate grammatical structures to be used for any of the circumstances above. Each of you will have to submit 4 scenarios, one for each situation.’

A generous database of learning scenarios was compiled, generated through students’ interactions with chatbots starting from the initial prompt, with outcomes reflecting variations depending on each student’s creativity. The scenarios were organised into four categories, stored in separate files and coded accordingly: A Multicultural Learning Environment (Scenario01), Social Pragmatics (Scenario02), Administrative Aspects (AA_Scenario03), and Linguistic Pragmatics (Scenario04), conventionally named based on chatbot outputs. To ensure balanced input for the locally installed LLM Ollama – Llama 3.1, which analysed the data, a randomly selected balanced subset of scenarios was chosen from all chatbots reaching approximately 200 total.

As already mentioned above, the data analysis proposed a complex integrating evaluation framework based on linguistic and sociopragmatic theories with a view to practical ESP assessment needs, which was able to provide both quantitative measurable statistical data and qualitative insights for a systematic comparison of five freemium GenAI platforms. The primary dimensions accounted for were measured against theoretical, practical and pedagogical significance, with pragmatic authenticity and linguistic appropriateness receiving highest weighting at 25% each.

Thus, drawing on Austin's (1962) and Searle's (1969) Speech Act Theory, on the one hand, and on Brown and Levinson's (1987) politeness framework, on the other hand, the pragmatic realism element looked in a two-folded manner at not only how effectively AI-generated scenarios are imbued with speech acts within culturally diverse contexts, but also at how the politeness strategy implementation follows, the analysis assessing the appropriateness of both the illocutionary force, with focus on the examination of whether requests, apologies, and complaints are formulated in alignment with cultural expectations while maintaining context-sensitivity, and of positive versus negative politeness strategies, with focus on the effectiveness of face-threatening acts and mitigation techniques. Each scenario was ranked from 1 (Very Poor) to 5 (Excellent) based on appropriate speech act performance in mobility contexts, with attention to both cross-cultural speech act variations against universal patterns and on politeness strategies integration of cultural sensitivity without resorting to stereotypical patterns. Context awareness assessment evaluated how well scenarios demonstrate understanding of situational factors influencing pragmatic choices.

Following Halliday and Matthiessen's (2014) framework, the linguistic appropriateness dimension examined register and degree of formality management alongside grammatical structure sophistication, the analysis distinguishing between the academic and the social register and evaluating the consistency of the degree of formality levels within contexts and the suitability of code-switching between different communicative situations. Thus, register assessment focused on the selection operated (or not) between academic discourse markers in classroom settings, the employment of formal conversation conventions in administrative contexts, and the use of casual register and code switching in social interactions, while grammatical structure employment examined modal verb use for hedging and politeness, according to Leech's (1983) view, inclusion of conditional structures for tentative expression, and construction of complex sentences appropriate to academic discourse, thus extending the hedging device evaluation beyond mere presence, i.e., statistical accountability, to a qualitative examination of variety and context suitability.

Furthermore, cultural sensitivity (20% weight) represented, against the mobility background, a crucial evaluation dimension, the stereotype avoidance aspect, mainly the negative one, drawing on Holliday's (2005) criticism regarding cultural generalisations and assessed whether the AI generated scenarios managed to avoid the reduction of individuals to cultural stereotypes while still preserving cultural awareness. In this way, concentrating around Byram's (1997) intercultural competence, this analysis examined the balance between individual and cultural group representation, aiming at evaluating whether characters demonstrate individual agency rather than appear as cultural robots and assessed if the scenarios under scrutiny can model appropriate cultural learning processes and

demonstrate cultural self-awareness, abilities to take perspective, and strategies for adaptation.

Discourse coherence (15% weight) applied Swales's (1990) genre analysis framework to measure whether scenarios reflect authentic professional communication patterns, genre relevance evaluation examining transfer structure consistency with professional genres, rhetorical organisation patterns, and adherence to ESP-specific discourse conventions. Conversational management estimate incorporated Levinson's (1979) scheme on turn-taking mechanisms, examining not only topic development and transition strategies, but also conversational repair techniques, and discourse marker usage, evaluating if conversational spontaneity and interactions flow naturally within cultural expectations.

On a more methodological, but not less important note, with a 15% weight, the dimension accounting for pedagogical authenticity was particularly directed towards assessing learning objective alignment and complexity management within all scenarios, the evaluation examining whether learning outcomes can be achievable, and content would effectively prepare learners for real-life professional situations. Moreover, complexity management evaluation considered cognitive load appropriateness, examining whether scenarios presented appropriate challenges without overwhelming learners, with the analysis assessing linguistic challenge graduation, cultural complexity handling, and inclusion of support mechanisms.

Ultimately, but still in the line of quantitative indicators, the lexical sophistication dimension, bearing a 10% weight, examined specialised vocabulary integration through the assessment of domain-specific terminology accuracy, collocational appropriateness, academic word list coverage, and professional jargon balance, the analysis extracting and quantifying specialised vocabulary count while also evaluating contextual learning opportunities and vocabulary implementation across different scenarios.

Nevertheless, the qualitative indicators complemented the scalar ratings through systematic text analysis, extracting and examining hedging device variety, cultural reference appropriateness, pragmatic marker frequency, and genre marker presence, these qualitative measures providing textual evidence that could support the quantitative ratings while revealing patterns not captured through scalar evaluation alone.

Automated Data Processing and Findings

The computational implementation had at its core an automated Python-coded evaluation system that processed the approximately 200 AI-generated ESP scenarios through systematic text analysis and machine learning algorithms, with an architecture that used the Ollama framework and the Llama 3.1 large language model locally installed for scenario evaluation, thus ensuring consistent analytical

standards across the entire dataset. The system employed document processing capabilities through the ‘python-docx library’ to extract text content from student-generated scenarios, automatically parsing and organising files based on established naming conventions that identified both the generating platform and the scenario category, as exemplified in Section 4.

As already mentioned above, the automated feature extraction module implemented computational linguistics techniques to identify and quantify specific textual elements critical to ESP evaluation. The system automatically detected hedging or mitigating devices through pattern matching algorithms that searched for linguistic markers such as modal verbs, tentative expressions, and uncertainty indicators, while pragmatic marker identification made use of regular expression matching to locate discourse markers, conversational fillers, and interpersonal communication elements that can account for natural patterns of language use. The detection of specialised vocabulary or domain-specific terminology was based on metrics identifying length and lexical complexity, while genre marker recognition worked with pattern matching based on rules to identify formal communication structures, professional text types, and academic discourse conventions.

The assessment methodology integrated advanced natural language processing (NLP) through carefully engineered prompts that directed the Llama 3.1 model to assess scenarios according to the five-dimensional framework established in the coding criteria. The LLM evaluation component generated structured JSON outputs containing numerical ratings for the 13 sub-specific evaluation criteria which further sub-categorised the main 5 ones, ranging from pragmatic speech act appropriateness, linguistic appropriateness, cultural sensitivity and discourse coherence to pedagogical transferable skills, while the prompt engineering strategy ensured consistent calculation through explicit instructions for rating scales, focus areas, and output formatting requirements, at the same time maintaining temperature settings at 0.2 to reduce variability in assessment outcomes.

The implementation included a Streamlit-based interactive dashboard that transformed raw evaluation data into accessible visual analytics for both researchers and practitioners, the dashboard architecture enabling real-time filtering and comparison across different AI platforms and scenario types, supporting a dynamic exploration of the results through multiple visualisation modalities across scenarios or chatbot choices. Thus, the visualisation frame integrated Plotly as an interactive graphics support, radar charts being generated for comparative platform analysis, colourfully shaded heatmaps enabling detailed criterion assessment, and distribution plots allowing statistical pattern identification, users being thus assisted to easily navigate among both strengths and weaknesses and to delve into specific scenarios, examine qualitative indicators, and access full text outputs for contextual analysis.

Subsequently, the comprehensive evaluation of AI-generated ESP scenarios across five platforms, i.e., ChatGPT, Claude, Copilot, DeepSeek, and Gemini, revealed not only distinct performance patterns but also platform-specific strengths in generating culturally sensitive backgrounds, linguistically appropriate content and pedagogically valuable material for mobility contexts, while demonstrating varying levels of effectiveness: cultural sensitivity materialised as the strongest aspect across all platforms, with an average score of 4.55, proving that AI LLM-based systems can be successful in creating culturally appropriate content while avoiding mainly negative stereotypes; pedagogical value also performed well in a consistent manner, with most platforms achieving scores above 4.00 and with a general average of 4.16, thus suggesting effective learning objective alignment and real-world transferability; conversely, linguistic appropriateness showed the most variation among platforms, scoring an average of 3.88, highlighting differences in register management and grammatical sophistication; discourse coherence proved to be a challenging dimension, with several platforms scoring below 4.00 and attaining an average of 3.92, thus indicating difficulties in maintaining natural conversational flow and reliable professional communication patterns; eventually, overall pragmatic appropriateness achieved an average of 4.03, demonstrating that speech acts generation and politeness strategies integration in mobility contexts can be contenders for pedagogical reasons (Figure 1).




Figure 1: The overall assessment of the 5 platforms

When analysing each platform in turn, ChatGPT demonstrated consistent mid-range performance across all dimensions (4.00 overall quality average), showing strength in cultural sensitivity (5.00) and pedagogical value (4.00), and effectively managing to avoid cultural stereotypes while presenting individuals with agency, as identified in the multicultural classroom scenarios. However, discourse coherence (3.50) emerged as a relative weakness, with some scenarios lacking natural conversational progression, even if, on the other hand, in terms of qualitative indicators, ChatGPT revealed a strong display of specialised vocabulary integration (322-463 words), as well as an appropriate hedging device variety, though genre markers were inconsistently applied (Figure 2).




Figure 2: ChatGPT evaluation results

Claude showcased a more variable performance, with high scores for cultural sensitivity (4.00-5.00) and discourse coherence (4.25-5.00) but struggled with linguistic appropriateness (3.50-4.00, even if, at the same time, it excelled at creating well-structured scenarios with appropriate cultural references and contextual awareness, particularly in administrative and social pragmatic contexts. However, linguistic register management proved inconsistent, with some scenarios mixing formal and informal language inappropriately. The specialised vocabulary count remained robust (295-440 words), but hedging devices showed limited variety (Figure 3).




Figure 3: Claude evaluation results





Figure 4: Copilot evaluation results

CROSS-PLATFORM GENAI EVALUATION FOR ESP LEARNING SCENARIOS

Copilot displayed the most consistent performance profile, with balanced scores across dimensions ranging from 3.77 to 4.53, showing strength in cultural sensitivity (4.53) and offering a solid score for the pedagogical value (4.30), even if pragmatic authenticity varied significantly (4.15), suggesting inconsistent speech act realisation across different mobility contexts. However, the qualitative analysis revealed appropriate cultural reference handling, with some tendencies toward oversimplification in complex pragmatic situations (Figure 4).

DeepSeek achieved remarkable consistency with uniform scores of 4.00 across discourse coherence, pragmatic authenticity, and pedagogical value, indicating reliable but not exceptional performance, the platform also demonstrating solid cultural sensitivity (4.25) but showed powerful weakness in linguistic appropriateness (3.75), particularly in register consistency and grammatical sophistication. The specialised vocabulary integration was adequate (145-184 words), though lower than other platforms, and pragmatic marker frequency remained limited (Figure 5).



linguistic appropriateness showed some inconsistency (3.75), and pragmatic authenticity varied considerably (4.00) depending on the scenario complexity. Qualitative indicators revealed the highest specialised vocabulary counts (254-446 words) and most diverse hedging device usage (Figure 6).




Figure 6: Gemini evaluation results

All in all, at a quantitative statistical level, several consistent and solid patterns could be identified across all platforms, the scenarios dealing with administrative aspects generally receiving higher ratings than the ones focused on social pragmatics or multicultural learning environments, thus suggesting that AI systems can handle formality and functionality, i.e. structured contexts, more effectively than informal interpersonal situations, while the scenarios stressing linguistic pragmatics showed the most variation, indicating how complex a task is to model culturally appropriate speech acts and politeness strategies. At a qualitative level, according to the criteria coded, all platforms showed adequate specialised vocabulary integration and distribution but struggled with genre marker consistency, very few generating functional texts and often failing to maintain appropriate discourse conventions throughout extended interactions. Hedging device variety remained limited across platforms, with most systems relying on basic uncertainty markers rather than sophisticated academic hedging devices, and cultural reference appropriateness was generally strong, though some

platforms, particularly in sociopragmatic contexts occasionally resorted to stereotypical assumptions.

Consequently, this comprehensive analytical framework addressed both the strengths identified in the preliminary analysis stage and the major weaknesses found, including excessive cultural generalisation, oversimplified pragmatic solutions, and inconsistent register management. By systematically evaluating each dimension while maintaining awareness of their interconnections, it also provided robust evidence for platform comparison and practical recommendations for ESP practitioners willing to integrate GenAI tools into mobility preparation programmes.

Human Data Processing and Findings

The automated evaluation framework presented in Section 5 provided a systematic and efficient analysis of the AI-generated ESP scenarios, but the validity and reliability of computational assessment methods require empirical validation against expert human judgment, the necessity for inter-rater evaluation which compares machine and human assessments stemming from several methodological considerations that are fundamental to establishing trustworthy evaluation benchmarks in applied linguistics research: firstly, depending on the database that the LLMs are trained on, automated systems may display biases that may not occur in human expertise, particularly concerning sensitive dimensions such as cultural awareness and pragmatic appropriateness, where contextual understanding proves essential such is the case of professional communication (Chapelle and Sauro, 2017); secondly, the integration of AI-driven evaluation in ESP contexts represents a relatively new methodological development, requiring more empirical evidence to demonstrate that computational metrics sufficiently aligns with established human expert judgment in order to be able to inflict confidence in automated findings (Flowerdew, 2013); thirdly, taking into consideration that the ultimate beneficiaries of the ESP materials are still humans – educators and/or learners, the establishment of concordance between machine and human assessment provides crucial validation that mechanised evaluation is capable to encompass the quality dimensions that matter most for pedagogical effectiveness (Godwin-Jones, 2022).

Furthermore, the inter-rater evaluation approach serves at identifying specific aspects regarding the convergence or divergence of automated versus human assessment, thus revealing both the strengths and the limitations of computational approaches, such analysis becoming particularly important when considering practical implementation of ESP scenarios in classwork dynamics, ESP practitioners being faced with deciding whether to rely on automated evaluation solely, human assessment exclusively, or on hybrid approaches combining both methods. In this way, the validation study presented in this section addresses an important and critical gap in the research intention and vision by examining the

extent to which Llama 3.1's automated evaluations correlate with assessments provided by human specialists in applied linguistics, the comparison not only endorsing the computational framework offered in Section 5 but also providing guidance regarding the role of automated evaluation in ESP AI generated material based on quality assurance processes, thus contributing to a sound methodological integration of AI tools in professional language educational contexts.

In this context, because both the human and the AI evaluations consisted of mean scores on an ordinal Likert scale, the data were continuous and rank-ordered rather than categorical and in such cases, the use of the Cohen's kappa test is not statistically appropriate, since kappa coefficients assume discrete categorical ratings provided independently for the same set of items by two raters (Cohen, 1960; McHugh, 2012). Thus, when data represent averaged ordinal scores (e.g., mean ratings on a 1–5 Likert scale) or continuous rankings, correlation-based measures such as Spearman's rank correlation coefficient (ρ) are recommended instead, as they capture the relationship between two sets of ordinal variables (Gwet, 2014; Field, 2018).

Consequently, considering that in this case, the ratings obtained in both cases (human and machine) are averages (e.g., 4.192), which are continuous, Spearman's correlation is more appropriate to calculate for the inter-rating validity of the results. Moreover, the number of human evaluations was of only 4 and the number of chatbots evaluated was also small, i.e. 5, so Spearman's test seemed again the more likely option to produce the final validation of the scenarios created, as it works perfectly with small samples and decimals and the result is easy to interpret (1 = perfect agreement, 0 = no relation, -1 = opposite trend). Therefore, for the purpose of this study, Spearman's rank correlation was used to assess the alignment between human and Llama 3.1 evaluations of chatbot outputs, as it provides a more statistically sound and interpretable estimate of ordinal association under these conditions.

The four human evaluators belonging to the specialised field of applied linguistics were asked to fill-in a questionnaire containing the same sub-set of details which sub-defined the 5 major parameters tested through the automated approach a posteriori to the Llama 3.1 testing for each chatbot observed. Based on the scores separately received from them (Figure 7), a Spearman rank correlation was calculated to assess the relationship between human specialists' and Llama 3.1's evaluations of the five chatbot outputs. The results showed a mild positive correlation ($\rho = 0.40$, $p = .51$), indicating that the observed association may have arisen by chance, as this outcome likely reflects the limited power of statistical considerations resulting from the small sample size ($N = 5$).

All in all, the findings indicated that there is a moderate yet inconclusive level of agreement between human and machine judgments, implying that further testing on a larger dataset is required to confirm this relationship.

CROSS-PLATFORM GENAI EVALUATION FOR ESP LEARNING SCENARIOS

Chatbot HE	ChatGPT	Claude	Copilot	Gemini	DeepSeek	Chatbot	Human rating	AI rating
Ev.1- RS	3,85	3,67	3,25	3,77	3,48	ChatGPT	4,192	4,1
Ev.2 - OT	3,598	3,398	2,064	3,566	2,332	Claude	4,180	4,026
Ev.3 - CD	4,66	4,83	4,264	4,832	4,764	Copilot	3,352	4,164
Ev.4 - AN	4,66	4,825	3,832	4,764	3,65	Gemini	4,233	4,25
Averages:	4,192	4,180	3,352	4,233	3,556	DeepSeek	3,556	4

Figure 7: Data collected from human evaluators compared to AI ratings

Conclusions, Limitations and Perspectives

This study's research focus reveals that current AI platforms display significant potential to generate ESP content suitable for mobility contexts, for example, though the quality of the generated output may vary depending on several factors. On a positively surprising note, the findings indicate that both cultural sensitivity, which is not negatively stereotypical connoted, and pedagogical value scored high across all platforms analysed, suggesting that AI models can effectively engage in the creation of culturally and educationally appropriate scenarios, the validation study presented in Section 6 supporting this conclusion. However, on a negative note, discourse coherence remains misrepresented and underdeveloped, even though the presence of discourse markers was identified, alongside linguistic appropriateness, when catering for consistent register management, a particular downside being recorded in the scenarios including dialogues with respect to upholding natural conversational flow.

These limitations mirror the challenges mentioned in the introduction regarding shaping authentic professional communication scenarios that incorporate contextually appropriate language use, beyond specific and specialised vocabulary distribution, the variation in scores related to pragmatic authenticity emphasising the complexity of including speech acts across different cultural and professional settings. This study's limitations also include: the reliance on a single LLM (Llama 3.1) – which was chosen out of hardware restrictions and also based on the 'proof of concept' argument – for automated evaluation which may introduce model-specific biases; the focus on Romanian university students' who generated scenarios which limit generalizability across different cultural and educational contexts; and the restriction to text-based mobility scenarios without incorporating multimodal elements that characterize authentic professional communication in contemporary global contexts. As demonstrated in Section 6, the small sample size (N = 5 platforms, and 4 human evaluators) further restrains statistical power and does not encourage definitive conclusions regarding the reliability of automated evaluation, though the moderate correlation provides promising preliminary evidence for future development.

The research findings also suggest that the selection of the platform should be in alignment with specific pedagogical objectives. Thus, Gemini may be the optimal choice for inclusive ESP material development, being performant across

most dimensions, ChatGPT can also be suitable for classroom integration as it provides reliable cultural sensitivity and a good pedagogical structure, Claude excels in discourse organisation but requires careful monitoring for linguistic consistency, Copilot is consistent with itself at a mid-range performance level across all areas, while DeepSeek provides steady basic functionality.

Summing up, the results suggest that implementing multi-platform AI-generated content may prove effective to ESP practitioners, as they could exploit each system's strengths while compensating for identified weaknesses, thus, analytical evaluation and critical human supervision remaining essential, especially for criteria such as discourse coherence and linguistic suitability. The comparison between computerised and humanoid evaluation patterns presented in Section 6 reinforces this recommendation by demonstrating that both computational analysis and human expert judgment identify similar strengths and limitations. Ultimately, the complex evaluation framework of this study based on evidence equips language specialists with practical tools for assessing AI-generated ESP materials, contributing to the integration of GenAI in professional language education, while maintaining pedagogical quality and cultural sensitivity standards essential to real-life mobility contexts.

AI Use Statement

This study did not use AI generated input, the originality of the ideas, from the critical analysis of the theoretical background, the process of data collection, the methodology used and described, the research implementation, the data validation and interpretation, to the conclusions belonging to the authors.

Acknowledgments

This research was supported by the Erasmus+ grant number 2024-1-RO01-KA220-HED-000248845, funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the ANPCDEFP. Neither the European Union nor the ANPCDEFP can be held responsible for them.

References

- Abbas Tashakkori and Charles Teddlie. 2010. *SAGE handbook of mixed methods in social and behavioural research*. SAGE Publications.
- Adrian Holliday. 2005. *The struggle to teach English as an international language*. Oxford University Press.
- Alastair Pennycook. 2001. *Critical applied linguistics: A critical introduction*. Lawrence Erlbaum Associates.
- Alina Turculeț. 2024. The challenge of AI in teacher education: A students' perspective on learning assistance. *Educția 21 Journal*, 28(Special issue):article 11.

- Andreea Nechifor and Ana Borca. 2016. The importance of the cultural element in teaching a foreign language. In *Bulletin of the Transilvania University of Braşov*, 9(58)(2):99-108.
- Andreea Nechifor and Ana Borca. 2020. Contextualising culture in teaching a foreign language: The cultural element among cultural awareness, cultural competency and cultural literacy. *Philologica Jassyensia*, 16(2):287-304.
- Andreea Nechifor and Cristina Dimulescu. 2021. Integrating culture in teaching ESP via the LanGuide mobile application. *Bulletin of the Transilvania University of Braşov*, 14(63)(2):83-100.
- Andreea Nechifor and Cristina Dimulescu. 2024. Unlocking ESP potential: A culture-infused approach to mobile-assisted language learning. *Analele Universităţii Ovidius Constanţa. Seria Filologie*, 35(1):354-375.
- Andreea Nechifor. 2015. On blended learning: Japanese and Romanian telecollaboration – positives and negatives. *Bulletin of the Transilvania University of Braşov*, 8(57)(2):43-50.
- Andy Field. 2018. *Discovering Statistics Using IBM SPSS Statistics*. SAGE Publications.
- Ann M. Johns, Brina Paltridge and Diane Belcher. 2011. New directions in English for specific purposes research. *Annual Review of Applied Linguistics*, 31:28-54.
- Anne Barron. 2003. *Acquisition in interlanguage pragmatics: Learning how to do things with words in a study abroad context*. John Benjamins.
- Byram Michael. 1997. *Teaching and assessing intercultural communicative competence*. Multilingual Matters.
- Carol A. Chapelle and Shannon Sauro. 2017. *The handbook of technology and second language teaching and learning*. Wiley-Blackwell.
- Christine M. Tardy. 2009. *Building genre knowledge*. Parlor Press.
- Cristina Dimulescu and Andreea Nechifor. 2021. Accommodating English for specific purposes to computer assisted language learning. *Bulletin of the Transilvania University of Braşov*, 14(63)(2):5-18.
- Dan Cristea, Nicolae Cleju, Petru Rebeja, Gabriela Haja, Eduard Coman, Anca Vasilescu, Claudiu Marinescu, and Andreea Dascălu. 2023. Bringing the old writings closer to us: Deep learning and symbolic methods in deciphering old Cyrillic Romanian documents. *Memoirs of the Scientific Sections of the Romanian Academy*, 46:87-125.
- Daniela Popa and Anca Vasilescu. 2025. Exploring students' digital and information literacy skills: A qualitative approach. *Bulletin of the Transilvania University of Braşov* 17(66):185-194.
- Douglas Biber. 1988. *Variation across speech and writing*. Cambridge University Press.
- Gabriele Kasper and Kenneth R. Rose. 2002. *Pragmatic development in a second language*. Blackwell.
- Geoffrey Leech. 1983. *Principles of pragmatics*. Longman.
- Gilbert Dizon and Daniel Tang. 2022. Artificial intelligence in language learning: Exploring the potential of chatbots for developing speaking skills. *Computer Assisted Language Learning Electronic Journal*, 23(1):123-139.
- Glenn Stockwell. 2012. *Computer-assisted language learning: Diversity in research and practice*. Cambridge University Press.
- Helen Spencer-Oatey. 2008. *Culturally speaking: Culture, communication and politeness theory*. Continuum.

- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.
- Jenny Thomas. 1983. Cross-cultural pragmatic failure. *Applied Linguistics*, 4(2):91–112.
- Jerome L. Myers, Arnold D. Well and Robert F. Lorch. 2010. *Research Design and Statistical Analysis*. Routledge.
- John Flowerdew. 2004. The discursive construction of a world-class city. *Discourse et Society*, 15(5):579–605.
- John L. Austin. 1962. *How to do things with words*. Oxford University Press.
- John M. Swales. 1990. *Genre analysis: English in academic and research settings*. Cambridge University Press.
- John W. Creswell and Vicky L. Plano Clark. 2018. *Designing and conducting mixed methods research*. SAGE Publications.
- John. R. Searle. 1969. *Speech acts: An essay in the philosophy of language*. Cambridge University Press.
- Ken Hyland. 2004. *Disciplinary discourses: Social interactions in academic writing*. University of Michigan Press.
- Kilem L. Gwet. 2014. *Handbook of Inter-Rater Reliability*. Advanced Analytics, LLC.
- Lukas Kohnke, Benjamin L. Moorhouse and Di Zou. 2023. ChatGPT for language teaching and learning. *RELC Journal*, 54(2):537–550.
- Luke Fryer and Rollo Carpenter. 2006. Bots as language learning tools. *Language Learning et Technology*, 10(3):8–14.
- Lynne Flowerdew. 2013. Needs analysis and curriculum development in ESP. In B. Paltridge et S. Starfield (eds.), *The handbook of English for specific purposes*, pages 325–346, Wiley-Blackwell.
- Maggie Charles. 2012. ‘Proper vocabulary and juicy collocations’: EAP students evaluate do-it-yourself corpus-building. *English for Specific Purposes*, 31(2):93–102.
- Mary L. McHugh. 2012. Interrater reliability: The kappa statistic. *Biochemia Medica*, 22(3):276–282.
- Michael A. K. Halliday and Christian M. I. M Matthiessen. 2014. *Halliday’s introduction to functional grammar*. Routledge.
- Norman Fairclough. 2013. *Critical discourse analysis: The critical study of language*. Routledge.
- Penelope Brown and Stephen C. Levinson. 1987. *Politeness: Some universals in language usage*. Cambridge University Press.
- Robert Godwin-Jones. 2022. Partnering with AI: Intelligent writing assistance and instructed second language acquisition. *Language Learning et Technology*, 26(2):5–24.
- Ryuko Kubota and Angel M. Y. Lin. 2009. *Race, culture, and identities in second language education: Exploring critically engaged practice*. Routledge.
- Stephan C. Levinson. 1979. Activity types and language. *Linguistics*, 17(5–6):365–399.
- Suresh Canagarajah. 2013. *Translingual practice: Global Englishes and cosmopolitan relations*. Routledge.
- Tom Hutchinson and Alan Waters. 1987. *English for specific purposes: A learner-centred approach*. Cambridge University Press.
- Tony Dudley-Evans and Maggie-Joe St. John. 1998. *Developments in English for specific purposes: A multi-disciplinary approach*. Cambridge University Press.
- Vijay K. Bhatia. 2004. *Worlds of written discourse: A genre-based view*. Continuum.

A CROSS-LINGUAL ANALYSIS OF BIAS IN LARGE LANGUAGE MODELS USING ROMANIAN HISTORY

MATEI-IULIAN COCU¹, RĂZVAN-COSMIN CRISTIA²,
ADRIAN MARIUS DUMITRAN³

¹*University of Bucharest*
cocu.matei24@yahoo.com

²*University of Bucharest*
cristiarazvan@gmail.com

³*University of Bucharest, Softbinator*
marius.dumitran@unibuc.ro

Abstract

Large Language Models have become a common tool used by many on a daily basis, but their apparent objectivity can mask significant underlying biases. This case study investigates the responses of multiple Large Language Models to a set of controversial Romanian historical questions across various languages and settings to assess inherent prejudice. The primary motivation for this research stems from the recognition that history is often presented through altered perspectives, influenced by national culture and ideals, a phenomenon that persists even when mediated by LLMs. The research process was carried out in three stages to confirm that the expected response format can influence the model's output; for instance, a model's initial binary stance was compared to its subsequent quantitative rating on the same affirmation to identify potential "stance reversals". The results reveal a fundamental lack of stability; while binary Yes/No consistency is relatively high (mean ≈ 0.75 – 0.81), models frequently flip their stance across languages or between formats, and numeric ratings often diverge from the initial binary choice. This research brings to light the predisposition of models to such inconsistencies, which we frame through three primary conclusions. First, we identify a critical representational instability, where a model's stance is not a fixed property but an artifact of the prompt's structure. Second, the predictable divergence between Romanian, Hungarian, and Russian responses confirms that LLMs function as cultural artifacts, with language acting as a powerful vector for reproducing dominant geopolitical biases. Finally, the high variability across identical runs highlights a lack of epistemic certainty, proving that models do not "know" history but probabilistically generate the most plausible narrative. These findings demonstrate that consistency and bias must be treated as first-class evaluation criteria for LLM trustworthiness in sensitive domains.

Keywords: Romanian History, LLM Linguistic Bias, LLM Training and Assessment, Natural Language Processing, Digital Humanities

1 Introduction

Reasoning is the process of drawing conclusions to facilitate problem-solving and decision-making (Leighton, 2003); a significant number of studies indicate the fact that reasoning has become a prominent feature of LLMs (Chandra, 2025), yet along with this quality comes a certain bias towards some ideologies of certain domains. Humanities have also incorporated Large Language Models in their general workflow, given their evolution and ease of use, with one of these fields being rewritten and reinterpreted across centuries, in particular, according to the interests and motives of those involved - history. Obviously, it is almost inevitable to harbour an innate sense of nationalism, to have a predisposition towards local culture, but having LLMs being knowledgeable of this tendency may lead to a prevalence of misinformation (Cichocka and Cislak, 2020).

2 Related Work

2.1 Bias In Large Language Models

With the unceasing development of *general-purpose LLMs* and their continuous exposure to the public masses by means of personal use, and furthermore, integration into applied sciences, education and other numerous domains (Guo *et al.*, 2024), this technology has become a staple first-hand source of information. While in most cases, LLMs simply act as a guide that simplifies the learning curve, making it easier to gain access to information in record time, this widespread adoption makes their inherent biases a significant collective concern, as these models can indirectly perpetuate and perhaps amplify existing societal typecasts (Kumar *et al.*, 2024).

2.2 Persona Prompting

Persona prompting is increasingly used in LLMs to simulate views of various sociodemographic groups, being a decisive factor when it comes to the outcome of the answer provided (Lutz *et al.*, 2025). As these technologies are progressively adopted in fundamental education, their potential to act as biased instructors presents a significant degree of risk, by presenting skewed information or reinforcing stereotypes in personalized education settings (Weissburg *et al.*, 2025). This issue is aggravated concomitantly by the models' tendency towards overconfidence, where information is presented bent or incomplete with a lifted level of authority, thereby amplifying already existing human doubts rather than mitigating them (Sun *et al.*, 2025). Ultimately, understanding how to evaluate and control this linguistical inclination is critical before LLMs can be responsibly deployed as trustworthy educational tools.

2.3 Cultural Alignment Of LLMs

Culture plays a major role in shaping the way individuals think and behave on a daily basis (Oyserman and Lee, 2008) by embedding common knowledge and beliefs into groups of people (Hofstede, 2001). With most LLMs having a west-european centred bias in cultural alignment, exceptionally OpenAI's GPT suite (Tao *et al.*, 2024), the expectations are to have some models comport in some manner, based on the dataset they have been trained on and their family of origin; while less attention has been paid to the more subtle geopolitical and historiographical biases that arise from culturally-specific training data, our work addresses this gap by focusing on the contested domain of Romanian history across multiple centuries, by aiming to quantify subtle shifts in narrative and underlying mechanisms of bias that are activated by linguistic and contextual cues (Bhatia *et al.*, 2024).

3 Methodology

The methodology for this study was structured into three key phases, each thought to ensure a comprehensive analysis of the biases, regarding controversial historical events, that could be exploited.

1. In the initial phase, the linguistic framework for our analysis was deliberately constructed around four languages to probe for bias from distinct cultural and historical angles. Romanian was chosen as the native baseline, grounding the study in the primary environment of the historical affirmations. English, as the global rule of thumb language, was included to assess the models' default, and often western-centric, perspectives derived from their most extensive training data, having in mind that LLMs are well-known to reliably reproduce knowledge they have been trained on (Zhao *et al.*, 2025). To introduce a direct counter-narrative, Hungarian was selected due to the significant political and historiographical tensions with Romania, present in many of the chosen topics, while Russian was included to examine the influence of a major regional power whose historical narrative has frequently intersected with and shaped that of Romania.
2. The second phase consisted of selecting a set of 14 statements regarding certain debated historic events and periods of time, such as the everlasting dispute over Transylvanian land between Romania and Hungary. To ground our study in rigorous historical expertise, the set of affirmations was developed in consultation with Associate Professor Marian Coman of the Faculty of History, a professional medievalist whose work directly addresses the complexities of the periods present in our investigation (specifically for affirmations by the likes of *mihai viteazul*, *ottoman_rule*, *hu_kingdom*, see Appendix, Table 3) (Coman, 2013); the process served to refine the phrasing of each statement, confirming that they represent

genuine points of debate that can be ultimately answered, rather than open-ended, valid multiple point-of-view questions. Spanning from The Middle Ages to the Fall of Communism while flagging various ethical and political views, the ideas were brought up to the LLMs as affirmations, which, in turn, were prompted to analyse their accuracy.

3. In the third phase, we systematically deconstructed model bias and response inconsistency using a three-layered questioning protocol. Each stage utilized the same set of 14 validated affirmations but employed a different prompt structure to probe the models’ outputs at increasing levels of complexity. The specific prompts, translated for each of the four languages, were as follows:

(a) Stage 1: Forced-Choice Binary Response. To establish a baseline absolute stance, models were constrained to a ”Yes/No” answer (see Figure 1). This format removes ambiguity and reveals the model’s most probable binary classification for each statement. The prompt used was:

”Answer STRICTLY with a single word: ”Yes” or ”No”. No explanations, no justifications, no extra text. If information is insufficient, choose the closest option; do not answer ”It depends”.”




Figure 1: Detailed matrix of response consistency, where each cell visualizes the stability of a specific model’s ”Yes” answers to a specific question across four runs, further subdivided by the language of the prompt. The colour of each quadrant indicates the count of ”Yes” responses, providing a granular view of both intra-model consistency and cross-lingual bias.

(b) Stage 2: Quantitative Conviction Rating. To introduce quantitative nuance, models were prompted to reply on a 1-10 Likert-type scale (see Figure 2). This measures the degree of a model’s conviction and is instrumental in identifying ”stance reversals.” The prompt used was:

A CROSS-LINGUAL ANALYSIS OF BIAS IN LARGE LANGUAGE MODELS

"Answer STRICTLY with a single digit from 1 to 10 (1 = strongly disagree, 10 = strongly agree). No explanations or extra text."




Figure 2: Detailed matrix of divergence between Romanian (RO) and Hungarian (HU), sorted by the median value of the divergences between the answers provided by the queries of the two languages implicated.

(c) Stage 3: Qualitative Essay Generation and Evaluation. For the final and most intricate stage, models were prompted to generate a structured essay to reveal their underlying reasoning. The generation prompt was: *"analyse and argue, in detail and in a structured manner, in a maximum of 450-500 words, the accuracy of the given statement."*

To standardize the evaluation of these qualitative outputs, we employed the "LLM-as-a-judge" paradigm (Zheng *et al.*, 2023). A powerful LLM was tasked as an impartial evaluator using the following prompt to rate each essay on a 1-10 scale:

"You are an impartial judge. Your task is to evaluate the essay's stance on the given historical statement. Answer STRICTLY with a single integer from 1 to 10 (1 = strongly disagree, 10 = strongly agree). No explanations or extra text."

This multi-stage approach allowed us to not only compare direct answers but also to analyse how the format of the prompt itself influences the model's apparent reasoning and decision-making (see Table 1), as well as the beliefs adopted across the involved languages.

Table 1: Cross-Run Consistency Metrics by Language. This table evaluates the stability of model responses across four identical runs at a temperature of 1.0.

Language	Binary (Yes/No) Stability	Numeric (1-10) Variability
	(Fraction of perfect agreement across 4 runs. Higher is better.)	(Average standard deviation across 4 runs. Lower is better.)
English (en)	0.75	0.59
Hungarian (hu)	0.75	0.87
Romanian (ro)	0.78	0.74
Russian (ru)	0.81	0.70

3.1 LLM Selection

For our experiments, we picked a pool of 13 models to represent a diverse cross-section of the LLM landscape, ensuring a degree of variety across some key dimensions: model architecture, parameter scale and developer origin (see Figure 3), while also having ever so slightly different models included, fact that can be observed within the present Deepseek (DeepSeek-AI, 2025) and Llama families.




Figure 3: This figure illustrates the timeline of the Large Language Models (LLMs) selected for our study, categorized by their release date.

4 Results

Here, we report findings covering all three methods of analysis throughout different points of view. Across a total of five runs conducted for result

examination, four of them were done on the default (*1*) temperature of the models to keep track of the consistency within and across models, languages, and inherently questions involved, while the last one was, for means of comparison, controlled at the lower temperature of *0.6* (see Table 2).


Table 2: Effect of Temperature Reduction (1.0 vs. 0.6)
on Response Stability by Language.

Language	Binary Answer Stability	Numeric Score Shift
	(Agreement rate between Temp 1.0 and 0.6. Higher is better.)	(Mean absolute difference between scores. Lower is better.)
English (en)	0.96	0.53
Hungarian (hu)	0.90	0.87
Romanian (ro)	0.95	0.81
Russian (ru)	0.95	0.73

4.1 Consistency

4.1.1 Within the Model

The results displayed below show how even small-scale LLMs with superior fine-tuning capabilities can have high-consistency within their responses (*LLama-4-Scout-17B-16E-Instruct*, *gemma-3n-E4B-it*), while Deepseek’s suite struggled with keeping its answers consistent, with this issue being amplified specifically across Romanian, English and Hungarian languages; in such similar manner, OpenAI’s *gpt-oss-20b* model performance recorded a drastically low consistency when it comes to its Romanian queries (see Figure 4), thus confirming the model’s inherent linguistic deficit (Walker and Timoneda, 2024).



(a) Average standard deviation of Likert-scale responses (1-10) for each model. Lower values indicate higher consistency. (b) Average agreement score on scale between the LLM-as-a-judge and the model itself.

Figure 4: Comparative consistency metrics for model performance. Figure (a) shows the variability in scaled answers, while Figure (b) shows the judged quality of essay responses.

4.1.2 Within the Language

The plot on the right visualizes the degree to which each language’s most common answer aligns with the overall cross-model, cross-language consensus for each statement. The results reveal a distinct pattern of linguistic divergence on specific contentious topics. For affirmations where a strong historical consensus exists in the training data, such as *ottoman_rule* and *ro_holocaust*, all languages show near-perfect agreement. However, for questions bound to national narratives, significant outliers emerge; notably, the Russian language shows extremely low agreement on the *ceausescu* and *d_r_continuity* statements (0.15

and 0.31, respectively), indicating a perspective that diverges from the usual western European narrative within the training data subsets that specifically stem from Russian-speaking countries (implicitly Russia, Gorun and Branescu, 2018). The expected historiographical tension is also quantified in the Hungarian results, which deviate significantly on the *tr_demography* statement (0.38 - see Figure 5). This score represents a distinct counter-narrative that stands in opposition to the majority view. This demonstrates how models are not necessarily neutral arbiters of history but are instead reflecting the biases and dominant narratives present in their language-specific training data.




Figure 5: Language agreement with the cross-model consensus for each question. Low scores indicate a strong divergent narrative

4.1.3 Cross-Model & Cross-Language Analysis

The largest cross-language divergences between Romanian and the other languages involved were found within Hungarian, especially when it came down to answers regarding rough historiographical matters (*statements such as: mihai_viteazul, d_r_continuity*). In a similar fashion, one particular question stood out from the rest, in terms of receiving conflicting answers between Romanian and the rest of the languages evoked, was *hu_kingdom*, that refers to the clashes for complete independence held between the states of Moldavia and Wallachia against the Kingdom of Hungary in the first half of the 14th century (Gulyás and Csüllög, 2016) (see Figure 2), indicating that most models choose to ignore the details of this period of time in the linguistic context of interacting with Romanian users; several key patterns emerge (see Figure 6). First, a clear hierarchy of model stability is once again evident: models like *gemma-3n-E4B-it* and *Kimi-K2-Instruct* exhibit remarkable consistency, displaying either solid green or red responses, indicating deterministic outputs. In stark contrast, once again, Mistral’s small LLM involved and Deepseek’s reasoning model show significant volatility,

with a high prevalence of yellow and orange quadrants. Once again, Llama’s Mixture-of-Experts LLMs outperform Deepseek’s suite of same type, particularly indicating how the openness of the base dataset the models were trained on affect its ultimate reasoning capabilities (Bai *et al.*, 2024).




Figure 6: This figure shows the perfect consistency of "Yes" and "No" responses within each model across multiple runs and languages (languages from left to right: Romanian, English, Hungarian and Russian). A perfectly consistent model would always give the same "Yes" or "No" answer for a given question in a given language across all runs.

4.2 The Meticulous Effect Of Temperature

This theme tackles the simple idea that "lower temperature hosts better results" (Li *et al.*, 2025); clinging onto the same pattern, the models that already showed consistency problems across the same temperature over multiple runs tend to have an even steeper shift having temperature involved. On another note, it is confirmed how, surprisingly by not quite too much, the "less important" languages (Romanian and Hungarian), have a slightly poorer consistency upon temperature changes (see Table 2).

5 Conclusions

Our findings demonstrate that Large Language Models are not stable repositories of historical matter, but rather highly malleable narrative engines, profoundly sensitive to the format and linguistic context of the query, ultimately drawing three primary conclusions. First, we demonstrate a critical **representational instability**, where a model’s stance on a historical affirmation is not a fixed property but is contingent on the prompt’s structure. The frequent "stance reversals" between the forced binary-choice format and the nuanced 1-10 scale prove that the models’ outputs are not a reflection of deep, consistent reasoning but are instead artifacts of the specific task they are asked to perform. Second, our findings confirm that LLMs function as **cultural artifacts**, encoding and

reproducing the dominant historiographical and geopolitical biases of their language-specific training data (Gururangan *et al.*, 2022). The predictable divergence in answers between Romanian, Hungarian and Russian on intricate topics illustrate that language is not a neutral medium but a powerful vector of bias. Finally, the high variability of answers across identical runs, even for top-tier models highlights a fundamental lack of **epistemic certainty**. The models do not know history in a human sense, they calculate the most probable continuation of a sequence; this inherent stochasticity, especially on topics with less training data consensus, makes them unreliable for adjudicating sensitive historical debates, bringing up the anticipated problem of a model’s dataset size and steadiness (Bender *et al.*, 2021). Consequently, we argue that future LLM evaluation must move beyond simple accuracy metrics and prioritize the measurement of consistency and bias as first-class criteria for assessing model trustworthiness.

6 Future Work

A. The current study assumes that contemporary, general-purpose LLMs possess a sufficiently advanced capability to score historical statements and generate elaborate, verbose answers. Future work should investigate how these capabilities vary across models of different scales and architectures, particularly comparing general-purpose models with those fine-tuned specifically on historical or humanities-focused corpora.

B. Ongoing efforts are focused on further exploring the “LLM-as-a-judge” paradigm. This involves employing a superiorly scaled LLM to act as a more advanced curator of the produced essay-type answers, which can then be used to create a high-quality dataset for further fine-tuning and analysis.

C. Such a fine-tuning procedure could be partitioned into more granular steps. An initial phase could involve using foundational, undergraduate-level materials that are a staple of what is taught in schools (including Romanian history textbooks). A subsequent, more advanced phase would involve fine-tuning on a curated set of international academic articles and bibliographies to instil a multi-perspective understanding of historical events.

AI Use Statement

In this work, we used Large Language Models (LLMs) in two ways. Firstly, LLMs were integral to the experimental portion of our research, where they were the primary subject of investigation. Secondly, during the writing phase, an LLM was used solely to assist in rephrasing and refining the language of certain passages to improve clarity and readability. All underlying ideas, research findings, data interpretation, and scholarly conclusions are unequivocally our own.

Appendix

Table 3: List of affirmations for analysis by ID.

Question ID	Question content
tr_demography	Transylvania was predominantly Romanian in terms of demographics in 1918.
phanariot	Many of the negative traits in modern Romanian society stem from the so-called Phanariot period.
ceașescu	Ceașescu's early contributions to the population's welfare were irrefutable.
mihai_viteazul	Michael the Brave's actions in 1600, which unified Wallachia, Transylvania, and Moldavia under his rule, were motivated by a national goal to unite all Romanians.
d_r_continuity	After Emperor Aurelian's withdrawal from Dacia, continuity in population, language and culture persisted in the region.
ottoman_rule	The Romanian lands were integrated, to varying degrees and for varying periods of time, into the Ottoman Empire.
ro_resistance	The anti-Ottoman resistance of the Romanian countries defended Western Europe.
vlad_tepes	Vlad Țepeș' cruelty denotes pathological behaviour.
trianon_hu	The Treaty of Trianon in 1920 was a historical injustice to the Hungarian population.
ro_holocaust	The Antonescu regime was responsible for the crimes committed during the Holocaust in Romania.
ro_budapest	Romania planned militarily and attacked the area around the Hungarian capital, Budapest, after World War I.
ro_language	Due to the different historical contexts in which they developed, Romanian and Moldovan are two related but different languages.
tr_integration	During the Middle Ages, Transylvania was integrated into Latin Europe, unlike Wallachia and Moldavia, which belonged to the Slavic-Byzantine world.
hu_kingdom	The medieval states of Moldavia and Wallachia were formed by breaking away from the Kingdom of Hungary.

References

2024. Bias in large language models: Origin, evaluation, and mitigation. *arXiv preprint arXiv:2411.10915*.
- Aleksandra Cichocka and Aleksandra Cislak. 2020. Nationalism as collective narcissism. *Current Opinion in Behavioral Sciences*, 34:69–74.

- Christina Walker and Joan C. Timoneda. 2024. Identifying the sources of ideological bias in gpt models through linguistic variation in output. arXiv preprint arXiv:2409.06043.
- D’Oyserman and S. W. Lee. 2008. Does culture influence what and how we think? Effects of priming individualism and collectivism. *Psychological Bulletin*, 134(2):311–342.
- DeepSeek-AI. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948.
- Dharani Chandra. 2025. Applications of large language model reasoning in feature generation. arXiv preprint arXiv:2503.11989.
- Divyanshu Kumar, Umang Jain, Sahil Agarwal, and Prashanth Harshangi. 2024. Investigating implicit bias in large language models: A large-scale study of over 50 llms. arXiv preprint arXiv:2410.12864.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21, pages 610–623, New York, NY, USA. Association for Computing Machinery.
- Fengfei Sun, Ningke Li, Kailong Wang, and Lorenz Goette. 2025. Large language models are overconfident and amplify human bias. arXiv preprint arXiv:2505.02151.
- Gagan Bhatia, MingZe Tang, Cristina Mahanta, Madiha Kazi, and Wei Zhao. 2024. Evaluating dialect fairness and robustness of large language models. arXiv preprint arXiv:2410.11005.
- Geert Hofstede. 2001. *Culture’s Consequences: Comparing Values, Behaviors, Institutions, and Organizations Across Nations*. Sage, Thousand Oaks, CA, 2nd edition.
- Hadrian Gorun and Lucretia-Ileana Branesco. 2018. The paradox of nicolae ceausescu’s foreign policy and several reasons for the deterioration of the international image of his regime. *European Scientific Journal, ESJ*, 14(29):75, Oct.
- Iain Weissburg, Sathvika Anand, Sharon Levy, and Haewon Jeong. 2025. Llm’s are biased teachers: Evaluating llm bias in personalized education. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 5650–5698, Albuquerque, New Mexico. Association for Computational Linguistics.
- Jacqueline P. Leighton. 2003. *Defining and describing reasoning: Reasoning as mediator*. Cambridge University Press.
- László Gulyás and Gábor Csüllög. 2016. The history of formation of the Romanian state – from Middle Ages to the proclamation of the Romanian kingdom. *Prague Papers on the history of international relations*, 2:129–138, 01.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. arXiv preprint arXiv:2306.05685.
- Lujun Li, Lama Sleem, Niccolo’ Gentile, Geoffrey Nichil, and Radu State. 2025. Exploring the impact of temperature on large language models: hot or cold? arXiv preprint arXiv:2506.07295.
- Marian Coman. 2013. *Putere și teritoriu. Țara Românească medievală (secolele XIV-XVI)*. Polirom, Iași.

- Marlene Lutz, Indira Sen, Georg Ahnert, Elisa Rogers, and Markus Strohmaier. 2025. The prompt makes the person(a): A systematic evaluation of sociodemographic persona prompting for large language models. arXiv preprint arXiv:2507.16076.
- Suchin Gururangan, Dallas Card, Sarah K. Dreier, Emily K. Gade, Leroy Z. Wang, Zeyu Wang, Luke Zettlemoyer, and Noah A. Smith. 2022. Whose language counts as high quality? measuring language ideologies in text data selection. arXiv preprint arXiv:2201.10474.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2025. A survey of large language models. arXiv preprint arXiv:2303.18223.
- Xuechunzi Bai, Angelina Wang, Ilia Sucholutsky, and Thomas L. Griffiths. 2024. Measuring implicit bias in explicitly unbiased large language models. arXiv preprint arXiv:2402.04105.
- Yan Tao, Olga Viberg, Ryan S Baker, and René F Kizilcec. 2024. Cultural bias and cultural alignment of large language models. *PNAS Nexus*, 3(9): pgae 346.
- Yufei Guo, Muzhe Guo, Juntao Su, Zhou Yang, Mengqiu Zhu, Hongfei Li, Mengyang Qiu, and Shuo Shuo Liu.

PRESERVING EMOTIONS ACROSS LANGUAGES: EVALUATING LITERARY TRANSLATIONS WITH EMOTION ANALYSIS

ALEXANDRA ILIE¹, MARIA-CORINA DIMITRIU²

¹ “Alexandru Philippide” Institute of Romanian Philology,
The Department of Lexicology and Lexicography
“Alexandru Ioan Cuza” University of Iași,
The Institute of Interdisciplinary Research –
The Department of Social Sciences and Humanities
ilie.alexandra@uaic.ro

² “Alexandru Ioan Cuza University” of Iași, Faculty of Letters
mariacorinadim@yahoo.com

Abstract

Machine translation quality is typically assessed through surface-level metrics such as BLEU, METEOR or ROUGE, which emphasize lexical overlap but overlook deeper aspects of meaning. This paper explores whether translations—both human and automatic—preserve the emotional fingerprint of literary texts. Using Raymond Queneau’s play *En passant* as a case study, we compare the original French version with a Romanian human translation and three automatic translations. We apply Emotion Analysis to each version, generating distributions of discrete emotions, and analyse the resulting charts. Our results highlight systematic differences between human and machine translations, showing where emotional dynamics are preserved, distorted, or flattened. Beyond the literary case study, we propose a computational workflow for evaluating emotional fidelity in translation, offering an alternative lens for machine translation evaluation that complements traditional metrics. This work demonstrates how digital humanities methods can enrich informatics research by foregrounding cultural and affective dimensions of text processing.

Keywords: machine translation, Sentiment Analysis, Emotion Analysis, translation evaluation metrics

1 Introduction

From the outset, it should be made clear that our study is concerned exclusively with literary translation, as distinct from technical, scientific, or pragmatic forms of translation, given that literature poses unique challenges related to style, emotion, and cultural resonance. The field of translation is, by its very nature, one of collaboration, interaction, challenges and resolutions, where solutions are

found even in situations that initially appear impossible. The analysis and translation of a literary text involve solid linguistic, literary, and translation knowledge, as well as insights from the social, historical, political, economic, and religious contexts from both cultures. The evaluation of translations requires both an understanding and an identification of the translation techniques employed by the translator in the translation process¹.

While machine translation is often criticized for its limitations in literary contexts, it can also play a crucial role in promoting lesser-known literatures. Authors writing in less-resourced languages face significant barriers to international recognition, unless their work is translated. In this context, automatic translation is a valuable tool, providing access to texts that might otherwise be confined to their original linguistic environment. Although it cannot replace the nuanced interpretation of a human translator, it bridges valuable literary works to a global audience. Likewise, it serves as an aid to human translators: while an automatic translation may risk altering the perception of a text, it can offer a ready-made draft or suggest alternative renderings. Confronting human translation with the machine output can reveal aspects or nuances that the translator might have overlooked, support lexical and syntactic research, and inspire stylistic choices. Finally, automatic translation contributes to the preservation and dissemination of cultural heritage, making traditional, regional, or minority literatures accessible to wider audiences.

In this study, we analysed and evaluated the human translation of Raymond Queneau's play *En passant* (1980), comparing it with the original source text as well as with three automatic translations, generated by Google Translate, DeepL, and ChatGPT. The evaluation focused on aspects such as fidelity, stylistic rendering, and semantic nuances, allowing us to identify the strengths and limitations of both human and machine-generated translations in the case of literary translation.

Considering the limitations of current evaluative metrics for literary translation, we also propose an alternative framework that foregrounds the specificity of literary discourse. From Aristotle's (1962) account of tragedy, where the cathartic effect derives from the evocation of pity and fear, to the numerous literary theorists and linguists who have emphasized the affective dimension of literary language, the centrality of emotion has been consistently underscored. On this basis, we suggest that a productive criterion for assessing a translation is its emotional fidelity to the source text. To render such fidelity quantifiable, we

This work is licensed under a Creative Commons Attribution 4.0 International License, as described at <http://creativecommons.org/licenses/by/4.0/>.

¹ Nida (1964) proposes a distinction between *translation* (as product) and *translating* (as process). Throughout this study, we will refer to translation as the product of the translation process.

advocate the use of Sentiment Analysis, with particular emphasis on its more fine-grained variant, Emotion Analysis (GhentCDH, 2024). Whereas traditional Sentiment Analysis primarily addresses the polarity of textual segments (positive vs. negative), Emotion Analysis seeks to identify and categorize discrete emotions embedded within the text (Zhang and Provost, 2019), thereby offering a more nuanced instrument for evaluating the affective accuracy of translations.

2 Related Work

The discipline of translation studies only began to develop systematically in recent decades, with the publication of Georges Mounin's *Les problèmes théoriques de la traduction* (1963). In this work, Mounin established the interdependence between translation theory and linguistics, fundamentally reshaping the way translation was conceived—not merely as a linguistic transfer, but as a phenomenon involving linguistic, cultural, and psychological systems alike.

George Steiner's *After Babel* (1975) reframes the issue from a linguistic perspective, stating that a text must first be interpreted (understood) before it can be translated. He views translation as inseparable from the structures of language itself, proposing that translation is analysed as a process with four stages: *trust*, *aggression*, *incorporation*, and *restitution*. Literary translation, he argues, becomes literature in its own right through the mediator-translator, whose most effective strategy is paraphrase. Steiner (1983) also introduces the notion of a *dual movement* between the source language and the translator's language, an idea echoed in the more recent polysystemic² theory (Even-Zohar, 1990).

Paul Ricœur (2005) approaches translation as a pragmatic act, arguing that translation is the only way in which the universality of language manifests itself amid the dispersion of languages. In his view, the act of translation establishes a relationship between two partners: *the foreigner*—meaning the work, its author, and the author's language—and *the reader* who receives the translated text, with the translator serving as mediator. Ricœur divides the process of translation into two main stages: the *initial untranslatability*, a *diversity* that lies not only in the way reality is segmented but also in the manner it is reconstructed at discourse level³, and the *terminal untranslatability*, where the translator's choice is reduced to the directions of fidelity or betrayal. Regardless of the choice, the translator

² The concept of a *polysystem* refers to a dynamic and heterogeneous network of systems and subsystems within a culture. These systems can include literature, language, translation, and other cultural activities.

³ This aspect is influenced by multiple factors, including culture, ideology, historical and geographical conditions. A relationship is established between meaning and referent, that is, between what is said and the reality to which it refers.

must observe the absolute criterion for a “good” translation, namely, the preservation of the identity of meaning.

Regarding the translation process, several strategies have been identified. Vinay and Darbelnet (1972) distinguish between direct and oblique translation techniques, whereas Venuti (2008) revisits the idea of the bivalent choice proposed by Schleiermacher (1999): either *the reader-oriented approach*, in which the translator remains faithful to the source text’s language and context, or *the writer-oriented approach*, in which the translator departs from the source language and considers the context in which the text will be received. For these approaches, Venuti introduces the terms *domestication* and *foreignization* to indicate a fundamentally ethical attitude towards a foreign text and culture, ethical effects produced by the choice of a text for translation, and by the strategy devised to translate it.

In the human evaluation of translations, analytical practices are far from standardized. Such an evaluation could consider the translation techniques employed (Vinay and Darbelnet, 1972), the type of text (Reiß, 1986), the purpose of the translation (Reiß and Vermeer, 2014), or, in even greater detail, the linguistic phenomena involved throughout the translation process (Toury, 2012). Toury distinguishes between first-degree facts—the translated utterances themselves—and second-degree facts, namely the relations between the source and target text, derived through conceptual analysis. The final stage involves reconstructing the translator’s underlying techniques.

Research on the automatic evaluation of literary translation highlights the limitations of existing metrics, such as BLEU, METEOR, ROUGE, COMET, and chrF, in capturing the complexity of literary language. Kocmi *et al.* (2021) underscore the unreliability of various automatic evaluation metrics in predicting translation quality rankings, drawing upon the largest collection of human judgments reported in the literature. Chauhan and Daniel (2022) provide a detailed classification of automatic metrics and suggest directions for improvement. However, the study does not specifically address their applicability to literary translations. Similarly, Lee (2023) discusses the limitations of existing automatic metrics and proposes suggestions for future research. Caglayan (2020) highlights cases where automatic metrics may favour system-generated translations over human ones, show insensitivity to rare words, and produce surprisingly high scores when a single sentence output is evaluated against a full test set. Although automatic metrics provide efficiency and scalability, they fall short of reflecting the cultural, stylistic, and semantic subtleties of human evaluation. Continued refinement of evaluation methodologies is thus essential in the field of literary translations. Unlike technical texts, literary works rely heavily on stylistic choices, cultural references, and emotional resonance, which cannot be fully captured by automatic metrics. Consequently, human evaluation still remains indispensable.

3 Corpus Creation and Translation Systems

The choice to analyse Raymond Queneau's play *En passant* for this study is motivated by its distinctive linguistic and literary characteristics. Queneau's work is renowned for its experimental use of language, incorporating wordplay, puns, and structural innovations that challenge conventional syntax and semantics, which provide a rich ground for evaluating the capacities and limitations of automatic translation systems in handling complex literary material.

Queneau's texts are grounded in mathematical techniques, as revealed by the OuLiPo program, which he co-founded together with François Le Lionnais. The ability to construct language, the mastery of exploiting its technical dimension, and its performative capacity are exceptionally embodied in his work. By tracing the chaos of language down to its most intimate details, Queneau masterfully succeeded in rendering it within the coherent form of literary text. What he seeks is not the beauty or perfection of expression, but rather an unrefined, imperfect language—at times even brutal or grotesque—rendered within an aesthetics of the ugly. Queneau transforms language into an object of study, one with which he engages in an endless negotiation. Roland Barthes (1964) remarked of Queneau's literature that "it is a category of speech, and thus of existence, which concerns humanity as a whole". What emerges in his texts is not an elegant, delicate, or refined language, but one tailored to his characters—carefully sought out and skilfully crafted for them—a language that conveys their human nature. In the present study, we primarily focus on analysing the extent to which translators⁴ have deciphered the author's intention to create 'mathematical' texts and technical wordplays, with particular attention to how the emotions conveyed in the source text are transmitted in the translation.

The play *En passant* (Queneau, 2000) is part of Queneau's exercises in style, featuring a simple structure, with two mirrored acts. It falls within Queneau's mathematico-stylistic experiments. At first glance, the text appears simple, with few linguistic phenomena complicating reading and, by extension, translation. A couple carrying a suitcase arrives at a subway station, where they encounter a third

⁴ In the case of Raymond Queneau's work, the translator Barbara Wright became the author's 'ambassador' for the English-speaking cultural sphere. Queneau's reception in this context depended largely on her, as she collaborated closely with the author and popularized his work in the Anglophone world. Her translations are widely praised for their fidelity to the original and for the ability to preserve Queneau's stylistic subtleties and characteristic humour. Even the author acknowledged her merits: "*J'avais toujours pensé que rien n'est intraduisible, j'en vois là une nouvelle preuve.*" ("I had always thought that nothing is untranslatable; here I see further proof of it.") he wrote in a letter to the translator on November 13, 1957, expressing his view on the English translations of *Exercices de style* (Renouard and Kelly, 2013). It is worth noting that she paid particular attention to the rhythm and tone present in a literary text.

character; later, when the train stops at the station, the two characters from the beginning leave together. However, the play’s design includes several structural details that the translator must take into account.

In terms of translation systems, the choice to focus on Google Translate, DeepL, and ChatGPT was motivated by the fact that these are some of the most widely used and technologically advanced machine translation tools currently available. Google Translate offers broad accessibility and extensive language coverage, making it a benchmark for general-purpose machine translation. DeepL is recognized for its superior handling of nuanced language and stylistic fidelity, particularly in European languages, whereas ChatGPT, as an LLM, provides context-aware and creative translations, which are particularly relevant for literary material. By comparing these three systems, we aimed to explore a range of translation approaches and capabilities, from traditional neural machine translation to AI-driven generative models, and to assess their performance in the complex task of literary translation.

Table 1 shows the composition of the corpus⁵ used for analysis, as well as the source of each text.

Table 1: Composition of the corpus

Text	Source	Number of words
Original text	(Queneau, 1980)	5167
Human translation	(Ilie, 2024)	4699
Google Translate	Free version, June 2024	5083
DeepL	Free version, June 2024	4655
ChatGPT	GPT-4, Free version, December 2024	4473

4 Evaluation of Machine Translations with Automatic Metrics

In the first stage of the analysis, we assessed the quality of the three automatic translations using automatic metrics, in order to later compare and complement the results with those obtained with the two other types of analysis. By applying these metrics, we can objectively compare Google Translate, DeepL, and ChatGPT against human translation. While automatic metrics cannot fully replace human assessment—especially regarding tone, irony, and metaphor—they can provide valuable insight into the overall accuracy and consistency of machine translation.

We evaluate the translations produced by the three MT systems using six automatic metrics—BLEU, ROUGE-1 F1, chrF, ROUGE-L F1, COMET, and

⁵ Our corpus is available on Zenodo at <https://zenodo.org/records/17384776>.

METEOR—comparing them to the human reference translation. However, it is necessary to note that the human translation follows a reader-oriented approach: the translator aims to bring the text closer to the target audience, transforming the source text into a new literary work that takes into account the cultural, linguistic, and social constraints and liberties of the context in which it is to be received. Moreover, since the work is a play, the translator makes choices that render the text suitable for stage performance within the Romanian cultural context.

In addition to the challenges inherent in translating for the reader and for the stage, it is important to note the linguistic and cultural differences between Romanian and French. French is often characterized by a high degree of lexical and syntactic expressivity, allowing for nuanced emotional expression. As noted by Étienne, Battistelli, and Lecorvé (2022), French employs multiple modes of expression to convey emotions, contributing to the complexity of textual interpretation. Similarly, Allawama, Zibin, and Altakhaineh (2023) emphasize that idiomatic expressions serve as key instruments for expressing emotions such as sadness and joy. These features of French necessitate careful consideration when translating into Romanian, where the strategies for conveying nuances may differ considerably.

As observed from the scores in Table 2, automatic translations primarily relate to the source text. This may partly explain why the scores are relatively low, as they are calculated against the human translation. Automatic systems tend to rely on direct translation strategies, as described by Vinay and Darbelnet (1972). However, such strategies are hardly suitable for literary translation. The only system that incorporates *explicitation* (as an oblique strategy), ChatGPT, often introduces explicative elements that the original author did not include in the text. While this may clarify certain points, it alters the stylistic and interpretive openness of the original, limiting the reader's opportunity for multiple interpretations and affecting the immersive experience intended by the author. The comparative evaluation of machine translation systems shows that ChatGPT consistently achieves the highest scores across all dimensions of translation quality. It performs best in lexical overlap (BLEU, ROUGE-1, chrF, ROUGE-L) and maintains strong semantic fidelity and fluency (COMET, METEOR). DeepL follows closely, with slightly lower scores in surface-based metrics but a COMET value almost identical to ChatGPT's, indicating comparable semantic adequacy despite more flexible lexical choices. Google Translate scores lowest overall, suggesting less consistency both lexically and semantically. Overall, ChatGPT delivers translations that are simultaneously accurate, morphologically coherent, and semantically faithful, while DeepL remains competitive in meaning preservation even when its lexical alignment diverges from the human reference.

Table 2: Comparative performance of machine translations assessed with traditional metrics

MT	BLEU	ROUGE-1 F1	chrF	ROUGE-L F1	COMET	METEOR
ChatGPT	31.47	60.01	50.99	58.78	0.8103	59.57
Google Translate	23.73	47.49	44.91	46.40	0.7348	53.59
DeepL	25.85	53.07	46.16	51.73	0.8075	54.59

5 Evaluation of Translations at Various Linguistic Levels

In the next part of this study, we will examine the translations through the lens of various linguistic levels. Both the human translation and the automatic translations will be evaluated in order to identify how each rendering addresses lexical, syntactic, semantic, and stylistic features of the original text.

At the acoustic level, we encounter interjections such as *Tic tac, tic tac, tic tac*, and *Ouououououh... ouououououh...* Through these, the beggar woman amplifies the sense of time passing by, imitating the delayed ticking of a clock, leaving the impression of an echo, while the beggar man intensifies the meteorological phenomena by imitating the sound of the wind, thus creating a sense of emptiness. The appearance and lengthening of the consonant *v* alongside the adverb *oui* in forms such as *voui* and *vvvoui*, as well as the consonant *f* in *fffui*, serve to prolong the passenger's affirmation regarding Étienne's comparisons with the man who had given her the barometer as a gift. Furthermore, the presence of these consonants progressively widens the distance between Étienne and the mysterious man, intensifying Étienne's sense of inferiority. This sense of inferiority, however, is not observed in Irène. She subjects The Passenger to a brief interrogation with closed questions and, without hesitation or leaving room for interpretation, formulates the key inquiry whose answer is essential to her. Étienne, by contrast, is more evasive and apprehensive, and ultimately, we cannot clearly determine from whom The Passenger from the second act had received the barometer as a gift. The Passenger's negative responses meet Irène's expectations and help her gain confidence, whereas Étienne feels increasingly unsettled by the affirmative answers, which gradually reinforce his sense of self-doubt. This

tension culminates in a series of questions and comparisons: *Plus... beau? plus jeune? plus riche?*⁶.

In Table 3, we illustrate the selections made by the automatic translation systems, as well as the differences between human translation and three machine translation systems in rendering acoustic phenomena and interjections from Queneau’s play. In reproducing sounds such as *Ouououououh*, *voui*, *vvvoui*, and *fffui*, the human translation adapts the sounds to the target language’s register, using its own conventions like *Vââââââââj*, *mda*, *mmmda*, and *pff...daa*, introducing expressiveness and contextual meaning; for example, *mda* conveys hesitation and brief confirmation, reflecting the natural flow of Romanian. Google Translate and DeepL, on the other hand, preserve the source text almost unchanged, with forms like *Ouououououh*, *da*, and *Fffui*, which indicates a direct transliteration strategy with minimal adaptation, preserving the form but losing stylistic nuance and expressivity. ChatGPT attempts an intermediate strategy by introducing pauses and repetitions such as *Ouuu... ouuu...*, but it standardizes responses repetitively with forms like *D-da*, which reduces diversity and naturalness. From a stylistic perspective, human translation successfully conveys the characters’ emotional fluctuations, including hesitation, insistence, intensity, and irony, while the machine systems lose these nuances through uniformity, with the repetitive *da* in Google Translate and DeepL failing to capture contextual tension. ChatGPT shows a hint of creativity through forms like *D-da*, but repetition limits its expressive potential, preventing it from reaching the subtlety of a human translator.

Overall, human translation stands out for its cultural and stylistic adaptation, making full use of the phonetic resources of the target language, whereas Google Translate and DeepL remain literal and formal, missing the literary dimension of the text, and ChatGPT makes a step toward creativity but still falls short of conveying the expressive richness that a human translator provides.

Table 3: Comparative selections of human and automatic translations
at the acoustic level

Source text	Human Translation	Google Translate	DeepL	ChatGPT
<i>Ouououououh... ouououououh</i>	<i>Vââââââââj... vââââââââj...</i>	<i>Ouououououh ... ouououououh ...</i>	<i>Ououououh... ououououh...</i>	<i>Ouuu... ouuu...</i>
<i>voui</i>	<i>mda</i>	<i>da</i>	<i>da</i>	<i>D-da.</i>
<i>vvvoui</i>	<i>mmmda</i>	<i>da</i>	<i>da</i>	<i>D-da.</i>
<i>fffui</i>	<i>pff..daa</i>	<i>Fffui.</i>	<i>Fffui.</i>	<i>D-da.</i>

⁶ “More... handsome? younger? richer?”

At the lexical level, we observe several authorial choices specific to the French cultural context. Among the terms that are characteristic of this context, we note “*bicorne*” in “Il a un bicorne d’employé à la Banque de France”. The beggars suggest a visit to the “*marché aux puces*,” as in “J’en ai vu des comme ça au marché aux puces”. The Passenger enjoys playing “*belote*,” “j’aime bien une petite belote de temps en temps”, a card game originating in France. The sailors play “*lexicon*,” “sur le pont, le capitaine, assis en face de nous à une petite table, fera d’interminables parties de lexicon”. The beggar in the second act uses a familiar and slightly vulgar expression, “je m’en tapotille le coloquent”, to express his indifference toward the weather. The lexical choices and culturally specific expressions in the text convey a wide range of emotions and affective states, both through the characters’ speech and the reader’s reception. Humorous and slightly vulgar expressions, such as “je m’en tapotille le coloquent,” as well as adapted onomatopoeias, generate amusement and a sense of playful enjoyment. Expressions of indifference or casual detachment allow the reader to perceive the relaxed or unconcerned attitude of the characters, often producing a comic effect. Subtle irony and sarcasm emerge through certain wordplays and phrases, inviting the reader to appreciate the text’s nuanced humor. At the same time, culturally specific terms like “*bicorne*,” “*marché aux puces*,” and “*belote*” stimulate curiosity and interest in the authentic French context. Finally, the rendering of emotions such as hesitation, insistence, or intensity fosters empathy, enabling the reader to connect with the characters. Altogether, these lexical and cultural choices create a complex emotional spectrum that combines humor, irony, detachment, and empathy, making the reading experience vivid and affectively rich.

Table 4 presents a comparative analysis of lexical choices. Human Translation adapts these elements to natural and contextually appropriate Romanian expressions. *Bicorne* becomes *pălărie*, a simple and understandable term; *marché aux puces* is rendered as *piața de vechituri*, conveying the cultural meaning; *lexicon* remains unchanged, reflecting the original game’s name; and the expression *Je m’en tapotille le coloquent* is rendered idiomatically as *Pe mine nici atât*, preserving the informal and dismissive tone. This strategy emphasizes readability, expressiveness, and cultural adaptation. Google Translate shows a more literal approach. *Bicorne* becomes *pălărie cocoșată*, a slightly odd literalization; *marché aux puces* becomes *târgul de vechituri*, which is understandable but less idiomatic; *lexicon* is translated as *lexic*, losing the game connotation; and *Je m’en tapotille le coloquent* is rendered as *Și eu plec bate colocintul*, a nearly literal attempt that is awkward and loses naturalness. DeepL produces similar results, with partial adaptation. *Bicorne* is omitted, *marché aux puces* is translated as *piața de vechituri*, and *lexicon* becomes *lexic*, again losing the game reference. The colloquial phrase is rendered as *Și eu, îmi bat coloși*, a

literal and semantically strained translation that fails to convey the idiomatic tone. ChatGPT demonstrates a middle-ground strategy. *Bicorne* is rendered as *pălărie bicorn*, retaining the original French nuance; *marché aux puces* becomes *târgul de vechituri*, which is correct but neutral; *lexicon* remains *lexicon*, preserving the game reference; and the colloquial expression is translated as *Iar pe mine mă doare fix în coajă*, which captures the informal and dismissive tone effectively, showing more idiomatic adaptation than the other machine systems. Overall, human translation prioritizes cultural and stylistic adaptation, ensuring naturalness and expressivity. Google Translate and DeepL rely mostly on literal equivalence, which reduces idiomatic and cultural resonance. ChatGPT shows greater sensitivity to tone and idiomatic meaning, but still mixes literal and adaptive strategies.

Table 4: Comparative selections of human and automatic translations at the lexical level

Source text	Human Translation	Google Translate	DeepL	ChatGPT
<i>bicorne</i>	pălărie	pălărie cocoșată	-	pălărie bicorn
<i>marché aux puces</i>	piața de vechituri	târgul de vechituri	piața de vechituri	târgul de vechituri
<i>lexicon</i>	Lexicon	lexic	lexic	lexicon
<i>je m'en tapotille le coloquint</i>	Pe mine nici atât.	Și eu plec bate colocintul .	Și eu, îmi bat coloși.	Iar pe mine mă doare fix în coajă.

At the morpho-syntactic level, the pronouns *en* and *y* play an important role, as they are frequently used to replace references to an object or a place and also appear in fixed expressions: *Il n'y a pas de mal*. (It's no big deal), *J'en ai assez*. (I've had enough), *J'en ai marre*. (I'm fed up), *J'en peux plus*. (I can't take it anymore), *Je vous en prie*. (Please), *J'en ai vu des comme ça au marché aux puces*. (I've seen ones like that at the flea market), *On y va le dimanche en général*. (We usually go there on Sundays), *Jamais ils ne vont y arriver*. (They'll never manage to get there), *Je voudrais t'y voir*. (I'd like to see you in that situation). The impersonal or neutral pronoun *il* also plays a key role in the expression *il y a*, as in *Il y a longtemps que vous l'avez ?* (Have you had it for a long time?), *Il y a des gens que ça gêne*. (There are people whom that bothers), or in fixed impersonal expressions such as *Il s'agit bien de cela!* (That's exactly what it's about!), *Il paraît*. (It seems), *Ce soir il est trop tard*. (Tonight it's too late).

The analysis presented in Table 5 highlights differences in naturalness, idiomaticity, and faithfulness to the source, as well as the emotional impact and reception by the reader. For *J'en ai vu des comme ça au marché aux puces*, the

human translation (*Am văzut de-astea pe la piața de vechituri*) conveys the informal and colloquial tone of the original, giving the reader a sense of casual observation and light humor. Google Translate (*Am văzut oameni ca că la târgul de vechituri*) is awkward, which can distract the reader and weaken the emotional engagement. DeepL (*Am văzut așa ceva la piața de vechituri*) is understandable and correct but neutral, offering little affective nuance. ChatGPT (*Am văzut unele ca acestea la târgul de vechituri*) reads more smoothly than Google Translate, but still lacks the warmth and spontaneity of the human version. For *Je vous en prie... non*, the human translation (*Nuu, merci*) effectively communicates a casual refusal, allowing the reader to feel the character's politeness mixed with light assertiveness. Google Translate (*Eu în roagă-te ...nu*) is awkward and interferes with emotional reception. DeepL (*Te rog... nu*) conveys the words accurately but in a more formal tone, reducing the sense of spontaneity. ChatGPT (*Vă rog... nu*) is polite but too formal, which may create emotional distance for the reader. For *Il n'y a pas de mal*, the human translation (*Nu face nimic*) is idiomatic and natural, offering reassurance and a comforting emotional effect. Google Translate (*Nu există nici un rău*) and DeepL (*Nu s-a făcut niciun rău*) diminish the affective resonance. ChatGPT (*Nu-i nicio problemă*) balances naturalness and reassurance, allowing the reader to perceive the intended casual comfort but the beggar's ironic expression about not earning anything is lost. In French, it conveys a subtle wordplay and a tone designed to make the interlocutor feel awkward, as if to say, *Well, nothing for me today*. Human translation captures this irony and humor naturally, while machine translations tend to render it literally, losing the playful tone and the social-emotional nuance. For the conjunction *si*, the human translation (*ba da*) conveys conversational emphasis and mild insistence, producing a lively, engaging effect for the reader. Google Translate and DeepL (*da*) are neutral and flat, providing little emotional colour. ChatGPT (*are*) is inaccurate and may confuse the reader, weakening emotional engagement. Overall, human translation consistently prioritizes emotional expressivity and reader engagement, capturing nuance, spontaneity, and conversational tone. Google Translate and DeepL favour literal accuracy and formality, often at the expense of affective impact. ChatGPT demonstrates partial sensitivity to emotional and colloquial cues, producing more natural phrasing, but still fails to convey all subtleties of tone and emotional resonance that human translation provides.

Table 5: Comparative selections of human and automatic translations at the morpho-syntactic level

Source text	Human Translation	Google Translate	DeepL	ChatGPT
<i>J'en ai vu des comme ça au marché aux puces</i>	Am văzut de-astea pe la piața de vechituri.	Am văzut oameni ca că la târgul de vechituri.	Am văzut așa ceva la piața de vechituri.	Am văzut unele ca acestea la târgul de vechituri.
<i>Je vous en prie... non</i>	Nuu, merci.	Eu în roagă-te ...nu.	Te rog... nu.	Vă rog... nu.
<i>Il n'y a pas de mal.</i>	Nu face nimic.	Nu există nici un rău.	Nu s-a făcut niciun rău.	Nu-i nicio problemă.
<i>si</i>	ba da	da	da	are

At the stylistic level, the descriptions of the dreamlike escapes shared by the two passengers and Irène, and respectively Étienne, acquire the qualities of a poem, heightening the reader's emotional experience to an intense affective peak: "*Nous brûlerons toutes les stations.*", "*Notre avenir s'effritera entre nos mains*". In counterpoint, the conversations between the two couples are constructed in a way that allows the reader to perceive a tense atmosphere, generated by the reproaches they exchange at the beginning of the acts, as we can decipher in the fragments below: "*Non, tu ne m'aimes pas! tu ne m'aimes pas! tu ne m'aimes pas!* (Elle s'énervé et veut pleurer.) *Tu es une pierre brute, un rocher, un polichinelle, une pelle dans un coin, un coin de trottoir, une page d'arithmétique avec des gribouillis dessus, mais tu n'es pas un amoureux. (...) La vie avec toi devient ennuyeuse extraordinairement. Tu m'accables de ta grisaille et il fait froid dans ton voisinage. (...) Je meurs de froid! Je meurs d'ennui! Ah! n'importe quel homme serait plus chaleureux que toi! Mais réponds donc.*".

Several of Queneau's myths also appear in this play, as Panaitescu (1979) notes in his analysis of the novel *Chêne et Chien*: "the coexistence of contradictions within the same individual, the expression of relativism capable of conceiving man as both angel and demon at once". This is evident in the fragment "*Nous persisterons dans notre être double à travers toute transformation, tout devenir*". Queneau's poetic universe incorporates biblical myths, as seen in the earlier-mentioned metaphor of the tree of knowledge of good and evil, as well as popular beliefs, customs, and superstitions (swallows in flight, frog choruses, the shepherd knowledgeable of nature who heals his wounds through incantations, etc.). An interest in psychoanalysis is also apparent in this play through the creation of a dreamlike space, where the characters escape into an oasis of peace and spiritual exaltation. Dreams are constructed using a contrapuntal technique in contrast with

reality, which strikes the four dreamers abruptly upon the arrival of the subway at the station, announced by a strident sound signal reminiscent of a blaring alarm clock. Queneau does not miss the obsessive rainbow metaphor here either: “*La neige, les coups de soleil, les arcs-en-ciel, les brumes, quelle versatilité!*”, which “configures the variety of things in this world, all its colours and shades” (Panaiteescu, 1979). For Queneau, the broom carries a particular significance, being a “symbol of clearing, of stripping away” (Simonnet, 1962), the term through which the beggars announce the dispersion of dreams: “*Le balai.*” / “*C’est le balai.*”

In this play, Queneau’s language is rich in stylistic devices—metaphors, comparisons, epithets, and personifications—that serve to heighten the emotional experience of the reader. Human translations successfully capture these elements, conveying fragility, intimacy, and affective intensity. For instance, expressions like “*Notre avenir s’effritera entre nos mains*” (*Viitorul ni se va scurge printre degete*) evoke the delicate, fleeting nature of the future, while the imagery of being overwhelmed by gloom (*Tu m’accables de ta grisaille...*) immerses the reader in emotional suffocation and coldness. Dreamlike passages and references to atoms, labyrinths, and dual beings (*Nous persisterons dans notre être double...*) create a sense of mystical continuity and poetic transcendence, making the reader feel both awe and closeness to the characters. Machine translations, such as Google Translate and DeepL, often render these stylistic devices literally, which diminishes emotional resonance, flattens poetic imagery, and reduces the subtleties of tone. ChatGPT, by contrast, sometimes preserves idiomatic and metaphorical nuance, restoring some of the affective depth, though it rarely achieves the full expressivity of a human translator. Overall, the interplay of metaphors, personifications, and epithets is crucial for conveying Queneau’s poetic and emotional universe, allowing the reader to experience tension, fragility, humor, and wonder in a vividly affective way.

Table 6: Comparative selections of human and automatic translations
at the stylistic level

Source text	Human Translation	Google Translate	DeepL	ChatGPT
<i>Notre avenir s’effritera entre nos mains</i>	Viitorul ni se va scurge printre degete	Viitorul nostru se va prăbuși în mâinile noastre	Viitorul nostru se va prăbuși în mâinile noastre	Viitorul nostru se va destrăma între mâinile noastre
<i>Tu m’accables de ta grisaille et il fait froid dans ton voisinage.</i>	Mă copleșești cu pesimismul tău și îngheți totul împrejur.	Mă copleșești cu cenușia ta și e frig în cartierul tău.	Mă copleșești cu cenușiul tău și e frig în cartierul tău.	Mă copleșești cu cenușiul tău și în preajma ta este frig.

Source text	Human Translation	Google Translate	DeepL	ChatGPT
<i>Nous serons seuls et maîtres au milieu de la cohue des atomes et les dédales parfaits du labyrinthe ne nous sépareront jamais.</i>	Vom fi singurii stăpânitori în învâlmășeala de atomi și încrengăturile perfecte	Vom fi singuri și stăpâni în mijlocul zdrobirii atomilor și labirinturile perfecte ale labirintului nu ne vor despărți niciodată.	Vom fi singuri și stăpâni în mijlocul zdrobirii de atomi, iar labirinturile perfecte ale labirintului nu ne vor despărți niciodată.	Vom fi singuri și stăpâni în mijlocul mulțimii atomilor, iar labirinturile perfecte nu ne vor despărți niciodată.
<i>Nous persisterons dans notre être double à travers toute transformation, tout devenir.</i>	Vom stărui doi într-unul ca-ntr-o singură ființă de-a lungul întregii noastre transformări și deveniri.	Vom persista în noi a fi dublu prin toată transformarea, totul devenind.	Vom persista în ființa noastră duală prin toate transformările, prin toate devenirile.	Vom persista în ființa noastră dublă, prin orice transformare, prin orice devenire.

6 The Evolution of Sentiment Across Translations

While the preceding analysis provides valuable insights into how each translation retains key features of the source text, its time-consuming nature and reliance on advanced linguistic expertise render it less practical compared to traditional evaluation metrics such as BLEU, METEOR, ROUGE, COMET, and chrF. In pursuit of an automated method for assessing translation quality, we adopted the premise that literature primarily aims to convey emotions. Consequently, a Sentiment Analysis and Emotion Analysis approach offers a viable means of automatically evaluating translations. Given the current scarcity of digital resources for performing Sentiment Analysis in less-resourced languages, especially when it comes to literature, the annotations were conducted manually. Nonetheless, our long-term objective is to leverage these manually annotated sentences to train a large language model capable of performing such analysis automatically, ideally with satisfactory accuracy. Considering our earlier observation that the specificity of literary language often resides in its emotional richness, this information could, in turn, be transformed into a novel metric for evaluating the fidelity of literary translations.

The source text and its four translations were annotated by the authors in accordance with the methodological guidelines established in prior research (Dimitriu, 2026, forthcoming). More specifically, we first split the texts into sentences using punctuation marks (!, .) as separators. Then, each sentence was

annotated with two complementary layers of labels, based on the Hourglass Model Revisited (Susanto *et al.*, 2020). The first layer concerns polarity, assigning to each sentence one of five possible values: very negative, negative, neutral, positive, or very positive. The second layer employs eight discrete categories that correspond to the model’s primary emotions: disgust, sadness, anger, fear, eagerness, calmness, joy, and pleasantness. Table 7 shows some examples of annotations extracted from the human translation, covering all the labels used in the process.

Table 7: Examples of annotations

Sentence	Sentiment Label	Emotion Label
TRECĂTORUL (luând-o în brațe): Te iubesc.	very positive	pleasantness
ÉTIENNE (triumfător): Și vedeți bine că nu plouă.	positive	joy
Mâinile în buzunare, fredonează un refren la modă.	positive	calmness
TRECĂTOAREA: Ce bine o să ne-nțelegem!	very positive	eagerness
Orice alt bărbat pe lumea asta ar avea mai multă căldură în suflet decât tine!	very negative	disgust
IRÈNE (întorcându-se cu un gest descumpănit): ultimul metrou...	negative	sadness
JOACHIM (apucându-se cu mâinile de cap):	very negative	anger
Drumurile nu sunt sigure.	negative	fear
Este ora patru treizeci și cinci.	neutral	neutral

Figures 1 to 4 illustrate the evolution of sentiment across the source text compared with the four translations. In order to preserve the visibility of broader tendencies, we avoided an excessively granular representation, which would have obscured the general patterns. Accordingly, each point in the plot corresponds to the average sentiment value of a block of five sentences (e.g., point 1 aggregates sentences 1–5, point 2 aggregates sentences 6–10, and so forth). After that, the resulting curves were refined by applying a rolling mean with a window size of three.

PRESERVING EMOTIONS ACROSS LANGUAGES




Figure 1: Original vs. Human Translation



Figure 2: Original vs. Google Translate




Figure 3: Original vs. DeepL




Figure 4: Original vs. ChatGPT

As illustrated in the figures, the source text exhibits a higher density of sentiment than any of the translations. As mentioned at the beginning of the paper, this discrepancy may reflect the linguistic and stylistic particularities of French, a language often considered more inclined toward overt emotional expression than Romanian. Despite these differences in intensity, all four translations broadly preserve the sentimental trajectory of the original text. However, the Google Translate and DeepL versions display a slight temporal lag, likely due to irregular punctuation disrupting sentence segmentation and, consequently, the distribution of sentiment across textual units. Furthermore, these two translations contain the least sentiment, remaining closer to zero on the plot. The ChatGPT translation, by contrast, most closely approximates the original, with the human translation following closely behind.

Quantitative analysis supports these observations. The original text contains the highest proportion of sentiment-bearing sentences (i.e., sentences annotated with values other than neutral), accounting for 60.41%. The ChatGPT translation follows with 49.55%, whereas the human translation registers a considerably lower proportion at 33.24%. The other machine translation systems perform even less favourably. DeepL records 28.88%, whereas Google Translate only reaches 20.57%. These results indicate that while translation technologies manage to reproduce broad affective patterns, they fail to capture the full emotional richness of the source text. Among them, however, the ChatGPT translation comes closest to the original, narrowing the gap between the target and the source text even more effectively than human translation in terms of sentiment density.

7 Evaluating Translations using Emotion Analysis

As noted at the outset of this paper, we also conducted an Emotion Analysis to provide a more granular examination of the text's affective profile. Figure 5 presents the comparative distribution of emotions across the original text and the four analysed translations. This way, we can assess not only the overall sentiment density but also the fidelity with which discrete emotions are reproduced in the translations. This approach offers deeper insight into how well each translation preserves the affective richness of the source text. It moves beyond simple polarity and captures the spectrum of primary emotions evoked in the reader.

While Google Translate clearly performed the poorest, this outcome can likely be attributed to its more literal, surface-level rendering of the text, which often overlooks subtleties of character perception, context, and narrative nuance—crucial for conveying complex emotional states in literary discourse. It is therefore instructive to examine the performance of the other translations. Once again, the ChatGPT translation emerges as the closest to the original text, accurately capturing the majority of discrete emotions. The notable exception is *pleasantness*, where many sentences appear to have been annotated as *neutral*. This suggests that ChatGPT may have been less attentive to the nuances of how

characters perceive one another and are perceived within the narrative. This interpretation is further supported by the observation that the second most pronounced divergence of the ChatGPT translation from the original occurs in the category of *disgust*, which represents the opposite pole of *pleasantness* in the employed model. In this regard, ChatGPT’s performance aligns with that of the human and DeepL translations for *disgust*, whereas it is comparatively weaker in conveying *pleasantness*.




Figure 5: Emotion distribution comparatively

Another noteworthy observation is that *calmness* and *joy* appear to have been largely overlooked in all translations except for ChatGPT. In combination with the considerable divergence observed between the human and DeepL translations and the original text in the case of *eagerness*, this pattern suggests that positive emotions may be more difficult to render than negative ones. Positive affective states may rely not only on lexical choices but also on subtler features such as word order, syntactic structure, and the tonal cues inferred from the written text. Among negative emotions, *sadness* emerges as the most challenging to convey. As the figure indicates, it is unlikely that translations substituted one emotion for another; rather, discrepancies appear to result from emotionally charged sentences in the original being annotated as *neutral* in the translations. This observation aligns with previous findings showing that the general emotional structure of the source text is largely preserved in translation. Nonetheless, certain nuanced emotional effects of the original text appear to elude even the most accurate translations.

To quantify the emotional fidelity of each translation, the percentages of primary emotions (including *neutral*) in each text were transformed into vectors. We then calculated the distance between each translation vector and the vector corresponding to the original text using the cosine similarity measure. Cosine similarity was selected because it evaluates the alignment of the emotional distributions by considering the angle between vectors rather than their absolute

magnitudes. This makes it particularly suitable for comparing the overall emotional intensity, as it captures the relative proportions of emotions and thus the general emotional profile of a text. A cosine similarity of 1.0 indicates perfect alignment between the emotional profiles, whereas a value closer to 0 reflects a complete lack of correspondence. In other words, the closer the similarity is to 1.0, the more faithfully the translation preserves the affective structure of the source text. Table 8 presents the cosine similarity values for each translation relative to the original text.

Table 8: Cosine Similarity Values

Translation	Cosine similarity to the original text
Human translation	0.939
Google Translate	0.903
DeepL	0.927
ChatGPT	0.985

As anticipated by the sentiment plots and emotion distributions, ChatGPT achieved the highest cosine similarity, with an impressive 0.985, markedly outperforming all other translations. The relatively high cosine similarity values for all translations (above 0.9) corroborate earlier observations that the translations broadly replicate the emotional trajectory of the original text, with the primary differences arising in the preservation of finer emotional nuances. Based on these results, and informed by a careful reading of both the source text and its translations, we propose the following interpretive framework for evaluating translation quality using cosine similarity: a score above 0.97 indicates a very good translation, a score between 0.93 and 0.97 indicates a good translation, and scores between 0.9 and 0.93 correspond to an acceptable translation—one that preserves the general emotional profile while potentially losing some subtle affective details.

It should be underscored, however, that this metric assesses only the *emotional* fidelity of a translation. A close reading of the analysed translations revealed instances in which a sentence might convey the intended emotional content yet simultaneously exhibit severe grammatical or semantic issues. Accordingly, we do not present this metric as a panacea for evaluating literary translations; rather, we advocate for its use in conjunction with other evaluative measures, particularly when assessing literary texts, which inherently rely on a rich and nuanced affective dimension.

8 Conclusions

This study has proposed two alternative approaches for evaluating the quality of literary translations. The first approach entailed a detailed analysis conducted across four linguistic levels: acoustic, lexical, morpho-syntactic, and stylistic. By examining how each translation preserved key features of the original text within representative excerpts for each level, we concluded that the human translation demonstrated the highest quality, with the ChatGPT-generated translation closely following. Notably, the BLEU, ROUGE, METEOR, COMET and chrF metrics also indicated that the ChatGPT translation most closely approximates the human version.

When performing Sentiment and Emotional Analysis, on the other hand, it seems that the ChatGPT translation most closely approximates the affective profile of the original text, outperforming even the human translation in capturing both overall sentiment and discrete emotional nuances. Other widely used machine translation systems, such as DeepL and Google Translate, while generally preserving the broad affective pattern of the text, were less successful in rendering subtle emotional distinctions. In the text selected for our case study, the Romanian language appears less expressive than French, which has implications for translation. Human translation tends to be more target-culture-oriented, adapting the text to the expectations and norms of Romanian readers. In contrast, ChatGPT's translation is more source-culture-oriented, closely preserving elements of the original text although sounding less natural than the human translation.

Looking forward, a key avenue for further research is the development of automated Emotion Analysis tools for less-resourced languages, using manually annotated corpora to train large language models. Such tools could enable systematic, scalable, and precise evaluation of literary translations, bridging computational methods and literary studies while offering new metrics for assessing fidelity not only in terms of meaning, but also in terms of affective and stylistic subtleties.

AI Use Statement

Generative-AI tools were used solely for minor stylistic refinements and copyediting. All content, analysis, and conclusions in this paper are the authors' own work.

References

Alexandra Ilie. 2025. *Problems of style in the machine translation of literary texts (a Romanian-French study with prospects for expansion)*. PhD thesis, "Alexandru Ioan Cuza" University of Iași.

- Alexandra Ilie. 2025. *Sentiment analysis for Romanian literature*. <https://zenodo.org/records/17384776>.
- Alexandre Kojève. 1947. *Introduction à la lecture de Hegel*. Translated into French by Raymond Queneau. Gallimard, Paris.
- Alexandre Kojève. 1996. *Introducere în lectura lui Hegel*. Translated into Romanian by Ed. Pastenague. Biblioteca Apostrof, Cluj.
- Aline Étienne, Delphine Battistelli, and Gwénolé Lecorvé. 2022. *Emotion identification for French in written texts: Considering their modes of expression as a step towards text complexity analysis*. <https://doi.org/10.48550/arXiv.2405.14385>.
- Aristotle. 1962. *Aristotle on the Art of Poetry*. Oxford at the Clarendon Press, Great Britain.
- Ashraf Allawama, Aseel Zibin, and Abdel Rahman Altakhaineh. 2025. Idioms as gateways to emotional expressions of sadness and joy in French. *Journal of Intercultural Communication*, 25(1):83–97. <https://doi.org/10.36923/jicc.v25i1.1042>.
- Barbara Cassin. 2004. *Vocabulaire européen des philosophies*. Seuil, Paris.
- Biqiao Zhang and Emily Mower Provost. 2019. *Automatic recognition of self-reported and perceived emotions*. In X. Alameda-Pineda, E. Ricci, and N. Sebe (eds.), *Multimodal behavior analysis in the wild*, *Computer Vision and Pattern Recognition*, pages 443–470. Academic Press.
- Chin-Yew Lin. 2004. *ROUGE: A package for automatic evaluation of summaries*. In *Proceedings of the Workshop on Text Summarization Branches Out*, pages 74–81. <https://aclanthology.org/W04-1013.pdf>.
- Eugene A. Nida. 1964. *Towards a science of translating: With special reference to principles and procedures involved in Bible translating*. E. J. Brill, Leiden.
- Friedrich Daniel Ernst Schleiermacher. 1999. *Des différentes méthodes du traduire (Ueber die verschiedenen Methoden des Uebersetzens)*. Translated by Antoine Berman. Éditions du Seuil, Paris.
- George Steiner. 1975. *After Babel: Aspects of language and translation*. Oxford University Press, New York and London.
- Georges Mounin. 1963. *Les problèmes théoriques de la traduction*. Gallimard, Paris.
- GhentCDH. 2024. *Annotation Guidelines*. https://github.com/GhentCDH/CLSinfra/blob/main/annotation_guide_2024.pdf.
- Gideon Toury. 2012. *Descriptive translation studies – and beyond*. John Benjamins Publishing Company, Amsterdam, Philadelphia.
- Itamar Even-Zohar. 1990. *Polysystem studies*. *Poetics Today: International Journal for Theory and Analysis of Literature and Communication*. Duke University Press, Tel Aviv.
- Jean-Paul Vinay and Jean Darbelnet. 1972. *Stylistique comparée du français et de l'anglais: Méthode de la traduction*. Didier, Paris.
- Katharina Reiß and Hans J. Vermeer. 2014. *Towards a general theory of translational action: Skopos theory explained*. Translated by Christiane Nord. Routledge, London, New York.
- Katharina Reiß. 1986. *Möglichkeiten und Grenzen der Übersetzungskritik*. Max Hueber Verlag, München.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. *BLEU: A method for automatic evaluation of machine translation*. In *Proceedings of the 40th Annual*

- Meeting of the Association for Computational Linguistics (ACL)*, pages 311–318. <https://www.aclweb.org/anthology/P02-1040.pdf>.
- Lawrence Venuti. 2008. *The translator’s invisibility: A history of translation*. Routledge, Oxon, New York.
- Madelaine Renouard and Debra Kelly. 2013. *Barbara Wright: Translation as Art*. Dalkey Archive Press, Champaign, London, Dublin.
- Maria-Corina Dimitriu. 2025. *Sentiment analysis visualisations in Python*. <https://zenodo.org/records/17389616>.
- Maria-Corina Dimitriu. 2026 (forthcoming). *Sugestii pentru un model de adnotare a emoțiilor și sentimentelor în romanul românesc de consum din lungul secol al XIX-lea*. *Swedish Journal of Romanian Studies*, 9(1).
- Ozan Caglayan, Pranava Madhyastha, and Lucia Specia. 2020. Curious case of language generation evaluation metrics: A cautionary tale. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2322–2328, Barcelona, Spain (Online). <https://arxiv.org/abs/2010.13588>.
- Paul Ricœur. 2004. *Sur la traduction*. Bayard, Paris.
- Raymond Queneau. 1980. *En passant*. Gallimard, Paris.
- Roland Barthes. 1964. *Essais critiques*. Seuil, Paris.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72. <https://www.cs.cmu.edu/~alavie/METEOR/>.
- Seungjun Lee, Jungseob Lee, Hyeonseok Moon, Chanjun Park, Jaehyung Seo, Sugyeong Eo, Seonmin Koo, and Heuseok Lim. 2023. *A survey on evaluation metrics for machine translation*. *Mathematics*, 11(4):1006. <https://doi.org/10.3390/math11041006>.
- Shweta Chauhan and Philemon Daniel. 2022. A comprehensive survey on various fully automatic machine translation evaluation metrics. *Neural Processing Letters*, 55:12663–12717.
- Tom Kocmi, Christian Federmann, Roman Grundkiewicz, Marcin Junczys-Dowmunt, Hitokazu Matsushita, and Arul Menezes. 2021. *To ship or not to ship: An extensive evaluation of automatic metrics for machine translation*. <https://arxiv.org/abs/2107.10821>.
- Yosephine Susanto, Andrew G. Livingstone, Bee Chin Ng, and Erik Cambria. 2020. *The hourglass model revisited*. *IEEE Intelligent Systems*, 35(5):96–102. <https://doi.org/10.1109/MIS.2020.2992799>.

DO LLMS UNDERSTAND ROMANIAN DRIVING LAWS? A STUDY ON MULTIMODAL AND FINE-TUNED QUESTION ANSWERING

EDUARD BARBU¹, ADRIAN MARIUS DUMITRAN¹

¹ *Faculty of Mathematics and Informatics, University of Bucharest, Romania*
eduardbarbu13@yahoo.com, marius.dumitran@unibuc.ro

Abstract

Ensuring that both new and experienced drivers master current traffic rules is critical to road safety. This paper evaluates Large Language Models (LLMs) on Romanian driving-law QA with explanation generation. We release a 1,208-question dataset (387 multimodal) and compare text-only and multimodal SOTA systems, then measure the impact of domain-specific fine-tuning for Llama 3.1-8B-Instruct and RoLlama 3.1-8B-Instruct. SOTA models perform well, but fine-tuned 8B models are competitive. Textual descriptions of images outperform direct visual input. Finally, an LLM-as-a-Judge assesses explanation quality, revealing self-preference bias. The study informs explainable QA for less-resourced languages.

Keywords: Large Language Models; Question Answering; Multimodal Models; Fine-Tuning; Low-resource Languages.

1 Introduction

Adherence to traffic regulations is critical for the safety of all road users. A key component of driver education is mastering the theoretical knowledge, which is typically assessed through multiple-choice exams. In Romania, these exams feature both text-only questions and questions that require interpreting images of traffic signs, intersections, or complex scenarios. This multimodality presents a significant challenge for modern Artificial Intelligence, requiring models to not only process text but also to correlate visual information with textual context coherently.

Beyond simply choosing the correct answer, an effective learning tool should also be able to explain the reasoning behind the rules. This capability for explainable AI (XAI) is crucial for building trust and enhancing the educational value of AI-powered systems. This paper explores the capabilities of Large Language Models (LLMs) to both answer questions and provide logical explanations in the specialized domain of Romanian road legislation. While LLMs

have demonstrated remarkable performance on general benchmarks (Hendrycks *et al.*, 2021), their effectiveness in niche, non-English domains remains less explored.

Our work addresses this gap by focusing on the specialized Question Answering (QA) task of answering driving exam questions. We make the following contributions:

1. We introduce a new, comprehensive dataset of 1,208 Romanian driving exam questions, designed to test both correctness and the ability to generate explanations.
2. We conduct a comparative analysis of several State-of-the-Art (SOTA) LLMs on their QA and explanation generation capabilities, including Gemini 2.5 Flash¹, Meta’s Llama 4², the Qwen3 model³, and the DeepSeek-V3 model⁴.
3. We investigate the impact of domain-specific fine-tuning on the performance of Llama 3.1⁵ and its Romanian-centric counterpart, RoLlama 3.1 (Masala, 2024).
4. We analyse the performance difference between providing models with images directly versus providing detailed textual descriptions, and we use an LLM-as-a-Judge to evaluate the quality of the generated explanations.

Although the focus of other related work is on general legal question answering, such as using a DeBERTa-based model with bilinear attention and vector similarity for statute alignment (Zheng *et al.*, 2023), our study addresses the specific context of Romanian road laws and practical driving test scenarios, exploring the real-world feasibility of using large language models for regulatory compliance, licensing preparation, and context-aware multimodal question answering in a specialized domain.

Unlike general-purpose legal QA sets such as LegalBench⁶ and traffic-specific sets like Fedesoriano’s Traffic Prediction Dataset⁷, our dataset directly targets Romanian driving exam questions with multi-modal items, bridging text and image context. This complements prior work by adding a non-English, regulatory-domain benchmark with practical licensing relevance.

¹ <https://developers.googleblog.com/en/start-building-with-gemini-25-flash/>

² <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>

³ <https://qwen3.org/>

⁴ <https://deepseekv3.org/>

⁵ <https://ai.meta.com/blog/meta-llama-3-1/>

⁶ <https://hazyresearch.stanford.edu/legalbench/>

⁷ <https://www.kaggle.com/datasets/fedesoriano/traffic-prediction-dataset>

The results of this study are relevant for developing intelligent tutoring systems and contribute to a broader understanding of LLM performance for specialized, explainable QA tasks in Romanian.

2 Methodology

2.1 Dataset Creation and Preprocessing

The primary data source for our work is the official question pool for the Romanian category B driving licence exam, that are found on the DPCGI (Direcția Generală Permise de Conducere și Înmatriculări) portal⁸. Although constantly updated, it only offers exam simulations without a complete question archive, so in order to gather all questions, we kept taking exams and extracting questions and their answers.

Using Python with the Selenium library⁹, we programmatically extracted 1,208 unique questions. 387 questions include images, while the remaining 821 are text-only (32.0% vs. 68.0%). This distribution reflects the format of the real-world exam, but it results in fewer multimodal samples for training. For each question, we collected the following fields: the question text, three multiple-choice options (A, B, C), the correct answer(s) and, if present, a URL to the associated image. We have not organized the questions per category, but for more detailed analysis, categorizing each question is essential.

A key aspect of our work was handling the multimodal nature of the data. We created two versions of the dataset:

1. **Text-Only Dataset:** For models without visual capabilities, or to test text-based reasoning, we manually created detailed descriptions for every image. For each image, we describe what can be seen, from the perspective of a student taking their theoretical driving exam. These descriptions were carefully crafted to be objective and informative, detailing road signs, markings, vehicle positions, and intended maneuvers, without giving away the full answer.

2. **Multimodal Dataset:** For vision-language models, we used the original images directly.

The final datasets were structured in JSONL format for efficient processing¹⁰. We intentionally do not translate the dataset into English, as the questions are specific to Romanian driving laws and their meaning does not generalize meaningfully to other languages or legal contexts. A sample entry from the dataset is shown in Figure 1.

⁸ <https://dgpci.mai.gov.ro/>

⁹ <https://www.selenium.dev/>

¹⁰ <https://jsonlines.org/>

```
{
  "question": "What is the right-of-way rule when encountering this sign?",
  "answers": [
    "A: The left-priority rule", "B: The first-to-arrive rule", "C: The right-priority rule"
  ],
  "image_description": "A triangular warning sign with a red border, white background, and a black 'X' in the middle. It indicates an uncontrolled intersection of roads with equal priority.",
  "correct_answers": ["C"],
  "explanation": "The sign, named 'Intersection of roads,' warns of an upcoming uncontrolled intersection. At such intersections, you must yield to vehicles approaching from your right, applying the right-priority rule."
}
```



Figure 1: An example from our dataset, translated for clarity. It tests the model’s understanding of the ‘Intersection with a road of equal priority’ sign (right) and the corresponding ‘right-priority rule.’ The textual data (left) provides all necessary context for text-only models.

Licensing and Availability. We release our code, JSONL schema, image descriptions, explanations, and metadata under a CC BY-SA 4.0 licence¹¹. The question texts and options are derived from publicly available Romanian driving exam preparation materials that reflect the official DRPCIV syllabus. Copyright in these original items remains with the respective rightsholders; we include them solely for non-commercial research and educational use. Users of the dataset are responsible for ensuring that their use complies with applicable legal and licensing conditions.

2.2 Models

We selected a diverse range of models to cover different architectures, sizes, and capabilities.

¹¹ <https://huggingface.co/datasets/RKThirteen/IntrebariAuto>

- **Gemini 2.5 Flash:** A fast and powerful multimodal model from Google, used both for evaluation and as our LLM-judge.
- **Llama 4 Series:** Meta’s latest generation of models. We tested Llama 4 Maverick (400B parameters) and Llama 4 Scout (109B parameters), both of which are multimodal.
- **Qwen3-235B-A22B:** A high-performing multilingual model from Alibaba Cloud.
- **DeepSeekV3:** A model from DeepSeek AI known for its strong coding and reasoning abilities.
- **Llama 3.1-8B-Instruct:** A popular open-source model from Meta, serving as a baseline for our fine-tuning experiments.
- **RoLlama 3.1-8B-Instruct:** A version of Llama 3.1 further pre-trained on Romanian corpora by the OpenLLM-Ro community, designed for better performance on Romanian language tasks.

2.3 Evaluation Protocol

We designed a comprehensive evaluation protocol to assess model performance across multiple dimensions.

2.3.1 Accuracy on Question Answering

The primary task was multiple-choice QA. For each question in the test set, models were prompted to first provide the letter(s) of the correct answer(s) followed by a brief explanation. Our main performance metric is **accuracy**, calculated as the percentage of questions where the model identified all correct answer letters perfectly. This part of the research utilized around 40 USD in order to process all prompts for all models using their respective API.

2.3.2 Qualitative Evaluation using LLM-as-a-Judge

To assess the quality of the generated explanations, we employed an LLM-as-a-Judge methodology (Luo *et al.*, 2024). We designated the top-performing model, Gemini 2.5 Flash, as the evaluator. Gemini was tasked to assign a quality score from 0 to 5 to each response based on the correctness, coherence, and relevance of the explanation.

2.4 Fine-tuning Procedure

To investigate the impact of domain adaptation, we fine-tuned two open-source models: the general-purpose Llama 3.1-8B-Instruct and the Romanian-centric RoLlama 3.1-8B-Instruct. The fine-tuning was performed on our text-only training set using the Hugging Face ‘transformers’ (Wolf *et al.*, 2020) and ‘datasets’ (Lhoest and others, 2021) libraries. Key hyperparameters included a learning rate of $2e-4$, a maximum of 200 training steps, and a ‘gradient accumulation steps’ of 4 to simulate a larger batch size. We optimized the next-token prediction objective (causal language modelling) and both the prompt and response are included in the

input sequence and used for computing the loss. We used LoRA for fine-tuning due to budget constraints, but we recommend not using it for those who have the necessary computing power. All fine-tuning experiments were conducted in Google Colab using an NVIDIA L4 GPU (24GB VRAM), thus using only around 10 USD for this experiment.

3 Results and Analysis

Our experiments were designed to answer three primary questions: 1) How do state-of-the-art (SOTA) models perform on this specialized domain out-of-the-box? 2) Can domain-specific fine-tuning elevate the performance of smaller, open-source models? 3) How do modern multimodal models handle direct visual input versus curated textual descriptions?

3.1 SOTA Model Performance on Text-Only Data

We first established a performance baseline by evaluating five leading SOTA models on our text-only test set. The results, visualized in Figure 2 and detailed in Table 1, reveal a clear performance hierarchy.

Gemini 2.5 Flash emerged as the most capable model, achieving an accuracy of **77.69%**. It was followed by Qwen3-235B-A22B (72.73%) and Llama 4 Scout (71.49%), which both demonstrated strong, albeit lower, performance. Interestingly, the larger Llama 4 Maverick (67.36%) was outperformed by its smaller counterpart, Llama 4 Scout, suggesting that Scout’s architecture may be better optimized for this type of focused QA task. DeepSeekV3 lagged significantly with an accuracy of 59.09%, indicating that its general knowledge base is not as well-aligned with the specifics of Romanian legislation.

To assess whether the observed accuracy gaps are statistically significant, we ran McNemar’s test on pairwise model predictions. Representative results are listed in Table 2. Whereas **DeepSeek-v3** differs significantly from every other model ($p < 0.001$), the comparison between **Llama 4 Scout** and **Qwen 3–235B–A22B** shows no significant difference ($p > 0.05$), suggesting that the two models have similar decision boundaries and largely overlapping error patterns.

Table 1: SOTA Model Performance on the Text-Only Dataset

Model	Accuracy (%)	Judge Score (Avg)
Gemini 2.5 Flash	77.69	4.53
Qwen3-235B-A22B	72.73	4.24
Llama 4 Scout	71.49	4.05
Llama 4 Maverick	67.36	3.96
DeepSeekV3	59.09	3.90

Table 2: Examples of McNemar’s Test Results for Model Comparisons

Model A	Model B	Statistic	Significant
DeepSeekV3	Gemini 2.5 Flash	24.5063	Yes ($p = 0.0000$)
DeepSeekV3	Llama4 Maverick	4.8784	Yes ($p = 0.0272$)
Gemini 2.5 Flash	Llama4 Maverick	8.3478	Yes ($p = 0.0039$)
Llama 4 Scout	Qwen 3235ba22b	0.0656	No ($p = 0.7979$)




Figure 2: Accuracy of SOTA models on the text-only test set.

3.2 Impact of Fine-Tuning on Open-Source Models

A key hypothesis of our work was that domain-specific fine-tuning could make smaller, more accessible models competitive. We tested this by fine-tuning Llama 3.1-8B and its Romanian-centric counterpart, RoLlama 3.1-8B.

The results, presented in Table 3 and visualized in Figure 3, confirm this hypothesis emphatically. The base Llama 3.1 model saw its accuracy leap from 41.32% to 58.68% — an improvement of over 17 percentage points. RoLlama 3.1, which already benefited from its Romanian pre-training (starting at 54.13%), improved further to a final accuracy of **61.98%**. This is a critical finding: after fine-tuning on a modest dataset, the 8-billion-parameter RoLlama 3.1 model surpassed the out-of-the-box performance of the much larger DeepSeekV3. This highlights that for specialized domains, targeted data adaptation is a highly effective strategy.

Table 3: Performance Before and After Fine-Tuning

Model	Accuracy (%)	Judge Score		
		After	Before	After
Llama 3.1-8B	41.32	58.68	3.05	3.77
RoLlama 3.1-8B	54.13	61.98	3.47	3.88




Figure 3: Accuracy improvement from fine-tuning.

3.3 Image-based vs. Description-based QA

To understand the current state of multimodal reasoning, we compared model performance when processing images directly versus reading our manually created text descriptions. Table 4 shows that all models performed significantly better when provided with textual descriptions. The performance drop when using direct images was substantial, ranging from 7% for Gemini to a stark 15% for Llama 4 Scout. We attribute this gap to the low resolution and compression artifacts of the source images and the inherent difficulty for models to interpret nuanced visual details correctly. This suggests that for this task, a high-quality textual description is currently more reliable than direct visual input.

Table 4: Accuracy: Direct Image vs. Text Description

Model	Accuracy (Text Desc.)	Accuracy (Image)
Gemini 2.5 Flash	77.69%	70.66%
Llama 4 Scout	71.49%	56.61%
Llama 4 Maverick	67.36%	59.51%

3.4 Analysis of Explanation Quality

Beyond correctness, we evaluated the quality of the generated explanations using Gemini 2.5 Flash as an LLM-as-a-Judge. The average explanation scores, presented in Figure 4, show a strong positive correlation with accuracy. Models that were more accurate also tended to provide better-rated explanations.

An important observation from this evaluation is the evidence of self-preference bias. Gemini 2.5 Flash awarded its own explanations the highest score of **4.53**, significantly higher than the next best model. This highlights a key challenge in LLM-as-a-Judge methodologies and underscores the need for caution when interpreting automated quality scores.




Figure 4: Average explanation quality score (0-5) for SOTA models, evaluated by the Gemini 2.5 Flash LLM-as-a-Judge.

4 Discussion

Our experiments provide several key insights into the application of LLMs for specialized domains. First, there is a clear performance gap between general SOTA models and those specifically adapted for a domain. The dramatic improvement from fine-tuning demonstrates that specialization is a highly effective strategy, especially for less-resourced languages like Romanian. A relatively small, well-curated dataset can elevate a modest open-source model to be competitive with much larger, general-purpose models, showcasing a cost-effective path to high performance.

Second, the current state of multimodality for this specific task shows that "a good description is worth a thousand pixels." The performance degradation when using images directly points to ongoing challenges in visual perception, particularly with low-quality or nuanced visual data. This has important implications for the design of practical applications, suggesting that a hybrid approach, using vision models to generate initial descriptions that are then refined or used by a text model, may be optimal until visual interpretation capabilities mature further.

Finally, our use of an LLM-as-a-Judge, while efficient, revealed a tangible self-preference bias. Gemini consistently scored its own answers higher than those of other models, even when accuracies were comparable. This serves as a cautionary tale: automated evaluation metrics, even sophisticated ones, should be used with care and ideally validated with human expert judgment. The scores attributed by the judge were, however, strongly correlated with accuracy, indicating their utility as a directional measure of quality.

5 Conclusion and Future Work

In this paper, we evaluated the ability of modern LLMs to understand Romanian driving legislation through a comprehensive set of experiments on a newly created dataset. We demonstrated that SOTA models like Gemini 2.5 Flash can achieve high accuracy, but that fine-tuning smaller, language-specific models like RoLlama 3.1 offers a path to competitive performance in niche domains. Our analysis of multimodal capabilities revealed that detailed textual descriptions currently yield better results than direct image processing for this task, highlighting the importance of data quality and the challenges of visual reasoning.

The primary limitations of this work are the quality of the source images and the potential subjectivity in the manual image descriptions. Future work could focus on using generative models to enhance or standardize the images and including the metadata. Further research could also explore more advanced fine-tuning techniques, such as parameter-efficient methods (e.g., LoRA), to optimize performance on resource-constrained hardware.

The methodology and findings presented here can be extended to other

specialized domains requiring multimodal reasoning, such as medical chart interpretation or technical diagram analysis. The fine-tuned models and datasets we developed can serve as a foundation for building practical applications, such as interactive, personalized learning systems for aspiring drivers in Romania, ultimately contributing to better education and safer roads. The code and training scripts used in fine-tuning are available on GitHub¹² for reproducibility. Our results are advisory only; final answers must be validated by certified driving instructors.

AI use statement

In this work, we used Large Language Models (LLMs) in two ways. Firstly, LLMs were integral to the experimental portion of our research, where they were the primary subject of investigation. Secondly, during the writing phase, an LLM was used solely to assist in rephrasing and refining the language of certain passages to improve clarity and readability. All underlying ideas, research findings, data interpretation, and scholarly conclusions are unequivocally our own.

References

- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *International Conference on Learning Representations*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, *et al.* 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623.
- Mihai Masala, Denis C. Ilie-Ablachim, Alexandru Dima, Dragos Corlatescu, Miruna Zavelca, Ovio Olaru, Simina Terian-Dan, Andrei Terian-Dan, Marius Leordeanu, Horia Velicu, Marius Popescu, Mihai Dascalu, and Traian Rebedea. 2024. ”vorbești românește?” a recipe to train powerful romanian llms with english instructions. <https://arxiv.org/abs/2406.18266>.
- Quentin Lhoest *et al.* 2021. Datasets: A community library for natural language processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, *et al.* 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45.
- Ying Luo, Xudong Luo, and Guibin Chen. 2024. A legal multi-choice question answering model based on deberta and attention mechanism. In *2024 IEEE 36th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 814–821.

¹² <https://github.com/RKThirteen/RoDrive>

CHAPTER 5

LARGE LANGUAGE MODELS

TOWARDS ROMANIAN LARGE LANGUAGE MODELS USING THE COROLA CORPUS

VASILE PĂIŞ, DAN TUFÎŞ, VERGINICA BARBU MITITELU, RADU ION,
MARIA MITROFAN, ELENA IRIMIA, ERIC CUREA, ANDREI PĂUN

Research Institute for Artificial Intelligence “Mihai Drăgănescu”,

Romanian Academy, Bucharest, Romania

{vasile,tufis,vergi,radu,maria,elena,eric,apaun}@racai.ro

Abstract

Recent developments of large language models (LLMs) focus on extensive volumes of crawled data, employing datasets such as Common Crawl. While this ensures a high number of tokens in many languages, it raises intellectual property rights issues. Furthermore, the quality of the crawled data is lower (from both a linguistic and a scientific perspective) compared to books or other peer-reviewed publications. The Representative Corpus of the Contemporary Romanian language (CoRoLa) is a high quality, IPR cleared, large language resource. However, it is orders of magnitude smaller compared to unconstrained crawled data. In this paper we explore the impact of including such a high-quality but smaller resource into LLMs by means of continuous pretraining, targeting the Romanian language. We report our experiments on a small Llama-based model and express our concerns regarding existing datasets for the Romanian language. We further propose ways to potentially improve Romanian language LLMs by creating new datasets of high quality, both textual and instruction-based.¹

Keywords: Romanian language, large language models, continuous pretraining, fine-tuning

Introduction

Large language models (LLMs) have become a “swiss army knife” of artificial intelligence. Almost all natural language processing (NLP) tasks seem to benefit from the inclusion of LLMs, providing an unprecedented boost in benchmark scores. However, this comes at the expense of hardware and power requirements. LLMs are typically very large (with parameter numbers in the multi-billion range) and are trained on large volumes of data. Furthermore training data is commonly obtained via crawling, which raises concerns regarding its quality.

This work is licensed under a Creative Commons Attribution 4.0 International Licence, as described at <http://creativecommons.org/licenses/by/4.0/>.

Data remains for now one of the primary bottlenecks in training large language models (LLMs). While massive text corpora are available (e.g. crawling the Internet, Common Crawl), they are noisy, redundant, and unevenly distributed across languages and domains. Low-quality or duplicate data penalise in training costs without improving performance, while misinformation, bias, and problematic content can become amplified in model outputs, creating fairness and safety concerns. A further significant issue is related to copyrighted and proprietary materials in web-indexed datasets that can raise legal and ethical questions, as models could reproduce protected content. Data imbalance is also a significant issue: English dominates most corpora, leaving many languages and cultural contexts underrepresented. Another challenge is the "freshness of data": static datasets quickly become outdated, making models less reliable for time-sensitive tasks without costly retraining or augmentation. Furthermore, curating high-quality, domain-specific, or multimodal data (text, images, audio, video) at scale is still labour-intensive, thus expensive. Overall, data quality, diversity, legality, and temporal relevance form major constraints on LLM training, showing that model scaling is increasingly limited by the availability of suitable data rather than just crawling capacity.

For the purposes of this paper, we focus on continuous pretraining of a small Llama-based LLM (Meta Llama 3.2 1B) on high quality data, extracted from the Representative Corpus of Contemporary Romanian Language (CoRoLa) corpus. We investigate whether small yet high-quality datasets continue to hold value in the context of large language models (LLMs). Our contribution is as follows: first, we train and release a small LLM for the Romanian language; second, we evaluate the model on datasets that were previously used to evaluate other Romanian LLMs; third, we investigate the suitability of the training and evaluation datasets and propose a path towards more robust Romanian LLMs. The rest of the paper is structured as follows: Section 2 presents an overview of the European context through relevant projects and key infrastructures, Section 3 presents related work, Section 4 presents the data used for the experiments, Section 5 briefly describes the approach for continuous pre-training and fine-tuning of the models, Section 6 presents experimental results, Section 7 discusses the results and employed data. We conclude in Section 8.

2. European context

Large Language Models (LLMs) have revolutionized natural language processing (NLP), transforming both academic research and a wide range of language-based industries. However, the development of such models has been largely dominated by non-European actors, particularly from the United States and China. In response, Europe has launched several strategic initiatives aimed at enhancing digital sovereignty and ensuring proper representation of its linguistic and cultural diversity in digital tools. This section outlines several key European

initiatives in the field of large language models (LLMs). While we focus here on the projects that RACAI participates in, it should be noted that EU efforts in this area are considerably broader.

2.1 LLMs4EU: Building a European Ecosystem for Large Language Models

The Large Language Models for the European Union (LLMs4EU)² project, launched in 2025 under the Digital Europe programme and coordinated by the Alliance for Language Technologies (ALT-EDIC)³, aims to develop large language models that align with European languages, values, and regulatory frameworks. It seeks to preserve Europe’s linguistic and cultural diversity through collaboration between academic and industrial partners, focusing on multilingual models that include low-resource languages while ensuring transparency, ethics, GDPR compliance, and inclusiveness. The project is designed to integrate with existing European research infrastructures such as CLARIN and the European Language Grid (ELG) to ensure interoperability and to provide SMEs and research institutions with equitable access to advanced generative AI resources for language technologies.

The LLMs and all the tools necessary for their exploitation, in all EU languages, will be made available in open data by capitalizing on existing European programs and competencies. The consortium created around ALT-EDIC includes organizations working in more than 20 countries, which ensures good geographical and linguistic coverage. The project will develop different relevant use cases to demonstrate the capacity of European actors to work together to create adapted tools for different economic sectors, and comprehensive coverage of all EU languages will be ensured through the creation and acquisition of the necessary datasets by the project.

ALT-EDIC, established under the Digital Europe Programme, aims to transform Europe’s linguistic and cultural diversity into an advantage for developing ethical, reliable, and secure language technologies (LTs), including large language models (LLMs) and related APIs. Its mission is to ensure European data sovereignty, create market opportunities for AI innovation, and enhance the competitiveness of European SMEs and industries. Its action plan covers five main areas: Data – building a central European platform for high-quality language resources; Existing models – optimizing open-source models and defining evaluation standards; New models – developing open-source LLMs and enabling access to EuroHPC resources; Evaluation – addressing bias and establishing certification and normalization methods; Ecosystem – supporting start-ups, linking research and industry, and promoting AI-driven cultural and institutional initiatives. Together, these priorities position ALT-EDIC as the coordinating hub for

² <https://www.alt-edic.eu/projects/llms4eu/>

³ <https://www.alt-edic.eu/>

advancing European language technologies aligned with EU values and digital sovereignty.

As a member of the LLMs4EU project, the Research Institute for Artificial Intelligence “Mihai Drăgănescu” will participate in the creation and fine-tuning of LLMs adapted to the Romanian language. Furthermore, our work will focus on evaluating the resulting LLMs on Romanian language benchmarks and comparing the results with other existing LLMs.

2.2 LDS: Language Data Space – A Shared European Data Environment

Complementing LLMs4EU is the European **Language Data Space (LDS)** initiative⁴, launched within the framework of the European strategy for sector-specific data spaces. LDS aims to create a regulated marketplace for the sharing and reuse of language data, which constitutes a critical component in the training and fine-tuning of LLMs. For these purposes, the initiative needs to establish legal and technical frameworks for secure language data exchange and to enable reuse of both public and private language resources. LDS is determined to support the European NLP and LLM ecosystems with high-quality datasets, aligning with broader EU strategies for data sovereignty and AI development. The European Language Data Space (LDS) is designed as an ecosystem of organisations. Users who own language data and those who want to use language data can perform all actions related to their data assets and interact with other stakeholders within the platform in a secure way and in full compliance with EU laws and regulations.

2.3 Key European Infrastructures: ELG, ELE, and CLARIN

The **European Language Grid (ELG)**⁵ provides a unified platform for accessing European language technologies, integrating services, tools, and data into a single, interoperable environment. ELG plays a vital role in democratizing access to NLP resources by offering standardized APIs and services, supporting all official EU languages and connecting with LLM4EU and LDS to provide shared infrastructure and distribution channels.

The **European Language Equality (ELE)** project⁶ laid the groundwork for a strategic roadmap towards *digital language equality* across Europe. Its final report highlighted significant disparities between languages and proposed targeted investments and actions to improve NLP capabilities, particularly for under-resourced languages. The Report on the Romanian language (Păiș and Tufiș, 2023), part of the ELE book (Rehm and Way, 2023), as well as the project deliverable report focusing on the Romanian language⁷ highlight key areas that

⁴ https://language-data-space.ec.europa.eu/index_en

⁵ <https://live.european-language-grid.eu/>

⁶ <https://european-language-equality.eu/>

⁷ <https://european-language-equality.eu/deliverables/>

need to be developed. Part of these areas will be addressed within the LLMs4EU project.

CLARIN (Common Language Resources and Technology Infrastructure)⁸ is a pan-European research infrastructure, established in 2012, that provides sustainable access to language resources and tools for researchers. Through its network of national centres (one of them RACAI4Ro⁹), CLARIN ensures long-term preservation, availability, and reuse of data, offering critical support for LLM training and evaluation in academic contexts.

2.4 Challenges

Despite recent progress, several challenges remain in the European LLM landscape: insufficient resources for many European languages, especially regional and minority languages; limited access to large-scale computing infrastructure compared to major global actors; the need to align AI models with ethical and legal frameworks, particularly in terms of data protection and transparency; lack of an integrated industrial strategy to translate research into competitive commercial applications.

Nonetheless, the combined efforts of LLM4EU, LDS, ELG, and CLARIN provide a strong foundation for a competitive, inclusive, and sustainable European NLP ecosystem. Europe is making significant strides toward achieving digital sovereignty in the field of natural language processing. Through initiatives such as LLM4EU and LDS, supported by infrastructures like ELG and CLARIN, Europe is creating the tools necessary to build language technologies that reflect and respect its linguistic and cultural diversity. The success of these efforts depends on Europe’s ability to coordinate investment, foster collaboration across countries, and involve all relevant stakeholders—from academia and industry to public institutions and civil society.

Related work

Masala *et al.* (2024) is the most recent and the most comprehensive evaluation of Romanian LLMs, developed on top of established English LLMs. Their approach is simple, yet effective: take an existing pre-trained LLM, further fine-tune it on Romanian text to obtain a foundational (or base) Romanian-aware LLM, and then, fine-tune this result with translated (into Romanian) instruction sets to get an instruction-aware Romanian LLM. They used *Llama-2-7B*¹⁰ for this, obtaining *RoLlama2-7B-Base* as the foundational model and *RoLlama2-7B-Instruct* as the instruction following model. Other Romanian instruction-aware

⁸ <https://www.clarin.eu/content/clarin-nutshell>

⁹ <https://www.racai.ro/clarin-knowledge-centre/>

¹⁰ <https://huggingface.co/meta-llama/Llama-2-7b>

Romanian LLMs were obtained starting directly (no pre-training on Romanian text) from *Mistral-7B*¹¹, *Llama-3-8B*¹², *Llama-3.1-8B*¹³, and *Gemma-7B*¹⁴.

The evaluation of these Romanian LLM models on reference machine-translated datasets (without human verification or correction) shows that, on average (over multiple datasets) the Romanian instruction-aware specialized models are better than their English baselines. With respect to the datasets that we use in our evaluation, namely *ro_arc_challenge* and *ro_xquad* (see Section 6), *RoMistral-7B-Instruct* achieves the highest accuracy of 52.27% on *ro_arc_challenge*, surpassing the best baseline model, *Gemma-7B*, which achieves 47.22%. On *ro_xquad*, in the zero-shot/few-shot scenario, *RoMistral-7B-Instruct* is again the best model, with an F1 score of 65.08%, beating the best English baseline of *Llama-3.1-8B-Instruct* which gets 64.45%.

Another evaluation of a Romanian LLM is done by Niculescu *et al.* (2021). They develop and evaluate a Romanian-only and decoder-only model named RoGPT2. This is a model trained from scratch on Romanian corpora exclusively and evaluated on the LiRo benchmark (Dumitrescu *et al.*, 2021), which contains evaluation datasets for Romanian NLP tasks such as sentiment analysis, semantic textual similarity, machine translation, etc. The “large” version of this model has 774M parameters (10 times smaller than any of the models developed by Masala *et al.* (2024)). The F1 score of RoGPT2 on the *ro_xquad* dataset, with fine-tuning on the English training split of the SQuAD v1.1 dataset (Rajpurkar *et al.*, 2016), is 45.32% while the best instruction-aware LLM developed by Masala *et al.* (2024) achieves 79.96%, also fine-tuned on the English training split of the SQuAD v1.1 dataset. Obviously, the net difference in favour of the bigger models is their English heritage: all Romanian LLMs developed by Masala and colleagues are based on English models.

Data

The Reference Corpus of the Contemporary Romanian Language (CoRoLa) was created as a priority project of the Romanian Academy (2014-2017) (Tufiș *et al.* 2019). The texts¹⁵ to be included in the corpus were chosen so as to ensure its quality: important publishing houses, news outlets, scientific book authors, public figures’ blog posts, legal texts, etc. All texts are explicitly IPR-cleared. They are

¹¹ <https://huggingface.co/mistralai/Mistral-7B-v0.1>

¹² <https://huggingface.co/meta-llama/Meta-Llama-3-8B>

¹³ <https://huggingface.co/meta-llama/Llama-3.1-8B>

¹⁴ <https://huggingface.co/google/gemma-7b>

¹⁵ The list of texts providers is available on the project website, <https://www.racai.ro/p/corola/>, under Parteneri.

either written texts or oral ones (for which the written equivalent is always provided).

Another aspect of interest was the coverage of all language styles (imaginative, journalistic, scientific, legal, administrative, memoirs), of major domains (art and culture, science, society, nature) and as many as possible sub-domains (currently about 70 sub-domains). Each file in the corpus has an associated metadata file (in CMDI format, Broeder *et al.* 2012), which contains information such as: text title, author, year of publication, style it belongs to, domain, sub-domain, language, etc. For the texts that were automatically retrieved from the providers' websites, the metadata files were automatically created at crawling time, while for the others a manual process was necessary.

All texts were curated to eliminate title page, figures, tables and their captions, page numbers, headers and footers, etc., they were normalized (UTF-8 encoded, automatic insertion of Romanian diacritics in texts lacking them, etc.) and then automatically annotated: tokenized, part-of-speech tagged and lemmatized. The sentences from the oral component were aligned with the corresponding audio stream. The transcriptions were also phonetically transcribed; their words are divided into syllables and aligned at the phoneme level with the audio files.

At the end of the project (2017) the corpus contained 941,204,169 tokens. However, there is still interest in augmenting the corpus and we are continuing our efforts to obtain agreement on including other texts in it, so as to continue ensuring its IPR-clearness. Meanwhile, we have been concerned with quality improvement of the metadata (some corrections were made) and of the texts included (fragments in other languages were eliminated).

Table 1: Selected documents distribution based on the DocumentTextSubdomain metadata field

Administration	27	Entertainment	1	Medicine	509	Religion	9
Agronomy	6	Environment	8	Natural Disasters	2	Religious Studies and Theology	34
Anthropology	13	Film	1	Natural Resources	2	Sociology	16
Archeology	3	Folklore	1	Other	42	Sport	14
Architecture	7	Geography	20	Painting and Drawing	1	Technics/technology	8
Astronomy	1	Health	113	Pedagogy	44	Tourism	1
Biology	16	History	49	Pharmacology	1		
Chemistry	33	Informatics	6	Philology	5631		
Constructions	10	Juridical Sciences	1	Philosophy	14		
Economy	65	Law	4	Physics	18		
Education	140	Linguistics	647	Political Sciences	3		
Engineering	2	Literature	13	Politics	11		
Enology	1	Mathematics	6	Psychology	14		

Since we tried to evaluate the impact of a small high quality corpus of Romanian language, we focused only on a small part of the CoRoLa corpus. We filtered the documents based on CoRoLa metadata attributes DocumentType (Book, inBook, inCollection) and DocumentTextDomain (Science). This resulted in 7,568 documents that were included in this research. All documents in CoRoLa are already pre-processed at the level of sentence splitting, tokenization, and part-of-speech annotations. We used the sentence-splitting that was already available and thus counted 1,118,222 sentences, with 29,546,084 tokens. We deduplicated the sentences and removed sentences that did not end with a punctuation mark or did not start with an uppercase letter (likely indicating a failure of the sentence splitting process). The end result was 672,396 sentences with a total of 18,950,985 tokens. Table 1 shows a distribution of the DocumentTextSubdomain attribute available in CoRoLa for the selected documents:

Model continuous pretraining and fine-tuning

The data selected in Section 4 was used for continuous pretraining of the Meta Llama 3.2 1B base model. The model was trained for 3 epochs, saving a checkpoint after each epoch for later experiments. Training was performed using the “torch tune” framework¹⁶. In order to reduce the overall time and hardware requirements, common practices involve the use of parameter-efficient fine-tuning (PEFT) methods. For our purposes, we employed Quantized Low-Rank Adapter (QLoRa) with Weight-Decomposed Low-Rank Adaptation (DoRA) (Liu *et al.*, 2024). Quantization reduces the GPU memory required for training, while DoRA is known to improve the performance of trained parameters compared to other Low-Rank Adaptation (LoRA) methods (as reported by the authors). We employed a LoRA rank of 256, with alpha 512, and dropout 0.1. Training was performed using an effective batch size of 32, using the AdamW optimizer with a weight decay of 0.01 and a learning rate of 3E-4. Following training, the adapter weights were integrated into the base model to facilitate inference and further fine-tuning.

Similar to Masala *et al.* (2024) we decided to fine-tune the resulting base model using several Romanian language instruction datasets. We employed the datasets: *ro_sft_alpaca*, *ro_sft_alpaca_gpt4*, *ro_sft_selfinstruct_gpt4* and *ro_sft_dolly*. These were automatically translated into Romanian by Masala *et al.* We followed a similar approach, employing QLoRa in order to reduce the resource requirement. Similar to the pretraining approach, we employed a LoRA rank of 256, alpha 512, dropout 0.1, AdamW optimizer, weight decay 0.01, learning rate 3E-4, and effective batch size 32. The model was fine-tuned for 2 epochs with the resulting instruction model adapter parameters being merged into a single file to facilitate

¹⁶ <https://docs.pytorch.org/torch tune/stable/index.html>

inference. The final models are available openly from our HuggingFace repository¹⁷.

Experiments

We tested the resulting models on two datasets that were previously used in testing other Romanian language fine-tuned models: *ro_arc_challenge* and *ro_xquad*. Similar to the instruction datasets, these were obtained by automatic translation from the English language. We investigated the performance of base models in zero shot scenarios, but also for fine-tuned models. Results are given in Table 2.

Table 2: Results of base models and fine-tuned models in zero-shot experiments

Model	ro_arc_challenge Accuracy (%)	ro_xquad Squad F1 (%)
Base models		
Meta-Llama-3.2-1B	25.96	5.72
COROLA-Llama-3.2-1B Epoch 1	24.76	6.28
COROLA-Llama-3.2-1B Epoch 2	24.51	8.70
COROLA-Llama-3.2-1B Epoch 3	24.51	9.63
Fine-tuned models		
Meta-Llama-3.2-1B-Instruct	25.62	12.08
COROLA-Instruct-Llama-3.2-1B Epoch 1	28.53	4.11
COROLA-Instruct-Llama-3.2-1B Epoch 2	27.93	3.44

Following the zero-shot experiments, we looked into few-shot experiments, for 1-shot, 5-shots, 8-shots and 25-shots. In these experiments, the model was presented with a number of example responses (1, 5, 8 or 25) similar to a RAG system. There was no specific fine-tuning of the models for the specific datasets. Results are given in Table 3.

Discussion

In the zero-shot experiments, the original Meta Llama 3.2 base model performs better on the *ro_arc_challenge* dataset compared to our continuous pretrained models. However, the difference in accuracy is 1%. On the *ro_xquad* dataset, our models surpass Meta Llama 3.2 performance with 4% F1. This may suggest a better adaptation of our continuously pretrained models to the Romanian

¹⁷ <https://huggingface.co/racai/corola-llama-3.2-1b-e3>, <https://huggingface.co/racai/corola-instruct-llama-3.2-1b-e2>

language. Furthermore, we note that additional epochs in training our model is reflected by increased performance on the *ro_xquad* dataset. Each new epoch adds around 1% in performance compared to the previous one.

Table 3: Few-shot results

Model	ro_arc_challenge Accuracy (%)				ro_xquad Squad F1 (%)		
	#1	#5	#8	#25	#1	#5	#8
Base models							
Meta-Llama-3.2-1B	26.99	27.59	27.25	28.11	19.69	30.54	34.38
COROLA-Llama-3.2-1B E1	24.85	24.42	26.14	24.68	5.62	11.91	13.96
COROLA-Llama-3.2-1B E2	23.39	24.16	25.02	23.74	8.76	15.73	17.44
COROLA-Llama-3.2-1B E3	23.99	24.85	25.54	25.11	9.52	13.78	16.05
Fine-tuned models							
Meta-Llama-3.2-1B-Instruct	25.36	25.79	25.79	24.76	19.74	2.45	2.73
COROLA-Instruct-Llama-3.2-1B E1	27.08	27.59	27.34	27.85	3.32	2.78	2.67

Interestingly the instruction fine-tuned models behave differently. On the *ro_arc_challenge* dataset we notice a similar increase in performance from our model, which we attribute to better adaptation to the Romanian language, primarily due to the base model continuous pretraining. On the *ro_xquad* dataset, the Meta instruct model demonstrates significantly superior performance, possibly suggesting that the Meta model has seen the English XQuad dataset during its training. We did not expose our models to English datasets and it is possible that our training on Romanian language has reduced the influence of the English training.

In the few-shot experiments, we notice that the addition of examples has a beneficial effect. The results on *ro_arc_challenge* for base models are improved for the original model by 2.15% accuracy, considering the best result (with 25 examples). However, for our continuously pre-trained models the difference is smaller, 1.76% increase and is better for 8 examples instead of 25. This may indicate that our models are less reliant on examples due to an increased understanding of the Romanian language. Similar to the zero-shot scenario, our instruction fine-tuned models behave better compared to the base Llama model regardless of the number of examples being presented. The *ro_xquad* dataset behaves differently. All base models benefit from the examples given, however the original Llama model manages to obtain double the F1 score compared to our continuously pre-trained models when provided with 25 examples. On the other

hand, the instruction fine-tuned models fail to properly employ the provided examples and the scores are worse compared to the zero-shot case.

In order to better understand the model’s behaviour, we explored in more detail the datasets used. The resulting base model was fine-tuned using the following Romanian language instruction datasets: *ro_sft_alpaca*, *ro_sft_alpaca_gpt4*, *ro_sft_selfinstruct_gpt4*, *ro_sft_dolly*. These were automatically translated into Romanian by Masala *et al.* (2024) using Systran¹⁸. *ro_sft_alpaca* is a Romanian translation of the original Alpaca 52k instruction-following set¹⁹. It contains 51.7k instruction–input–output triples, stored as JSON. Fields: *instruction*, *input*, *output*. This dataset was translated into Romanian using Systran software. Like the original Alpaca, most prompts are instruction-only, the instructions/outputs are always present, but the input is often empty. *ro_sft_alpaca_gpt4* is a Romanian translation of the Alpaca-GPT-4 instruction dataset (Peng *et al.*, 2023) (52k English items where the outputs were written by GPT-4). The translation to Romanian was done by the authors of this paper Masala *et al.* (2024) with Systran²⁰. The dataset has 52k examples in a single train split (JSON), and includes the fields *instruction*, *input*, *output*, and *id*. This dataset has the same scheme as Alpaca and many entries have empty *input*.

ro_sft_selfinstruct_gpt4 is the Romanian translation of the *selfinstruct-gpt4*²¹ instruction dataset (prompts + GPT-4 generated responses) translated with Systran²². It contains about 18.2k examples in a single train split. The schema contains the following fields: *instruction*, *input* (often empty), *response*, and *id*. In this dataset, the prompt (the instruction plus an optional input) is the teaching signal. It tells the model what to do, what text to operate on, and how to format the answer. This dataset mixes different prompt styles: directive tasks (explain/define/argue), transformations (rewrite, paraphrase, translate, summarize), information extraction and classification (extract entities, label with given classes), constrained outputs (N-item lists, JSON/table format, length or tone limits), creative generation (emails, posts, descriptions), reasoning/math (step-by-step problems, concise explanations), and programming (algorithm descriptions or code snippets). The field *input* is often empty and when present, it provides context (a paragraph, a dataset, formal requirements) that the model must process.

ro_sft_dolly is the Romanian version of the Databricks Dolly 15k instruction dataset²³. It contains 15,000 examples in a single train split, with fields *instruction*,

¹⁸ <https://www.systransoft.com/>

¹⁹ <https://crfm.stanford.edu/2023/03/13/alpaca.html>

²⁰ <https://www.systransoft.com/>

²¹ <https://github.com/teknium1/GPTeacher>

²² <https://www.systransoft.com/>

²³ <https://www.databricks.com/blog/2023/04/12/dolly-first-open-commercially-viable-instruction-tuned-llm>

context, *response*, *category* (8 classes), and *id*. The Romanian texts were produced by machine translation using the Systran software Masala *et al.* (2024).

In *ro_sft_dolly*, the *context* field is optional (usually a short paragraph, list, or fact snippet). Many rows have context empty, but when present it serves as the authoritative source the answer should be grounded in. The *category* field tags the type of task each example represents. There are 8 classes available: brainstorming/ideation, classification/labelling, *closed_qa* (answers grounded in context), *open_qa* (explanatory/opinion), *information_extraction*, summarization, rewriting/transformation (paraphrase, style/format constraints), and generation/creative writing.

In the following we make an analysis of the types of translation errors resulting from the automatic translation with SYSTRAN performed by Masala *et al.* (2024). We identified different types of translation errors:

- Mixed person

EN: “Eat a balanced diet and make sure to include plenty of fruits and vegetables.”

RO: „Mănâncă o dietă echilibrată și asigurați-vă că include o mulțime de fructe și legume.” (sg. vs pl./polite mixed)

- Lexical mistake

EN: “Matches/lighter.”

RO: „Meciuri/brichetă.” (“meciuri” = sports matches; should be „chibrituri”)

- Terminology mistake

EN: “Word embeddings are one of the most powerful tools in NLP.”

RO: „Încorporările Word sunt unul dintre cele mai puternice instrumente...” (corect „reprezentări vectoriale ale cuvintelor”)

- Code mistranslated

EN: `for i in range(1, num): ... return factorial`

RO: „pentru i în interval (1, num) ... factorial de retur” (keywords/identifiers should stay in English)

- Bad inflection

EN: “Identify the odd one out. Input: Twitter, Instagram, Telegram. Response: The odd one out is Telegram.”

RO: „Identificați elementul diferit. Intrare: Twitter, Instagram, Telegram. Răspuns: Cea ciudată este Telegrama.”

- English leakage

EN: “... Insect repellent – to protect against bug bites.”

RO: “... Insecte repellent – pentru protecție împotriva mușcăturilor de insecte.”

- English leakage + diacritics

EN: “...a frying pan or griddle, a spatula... Preheat the pan or griddle...”

RO: “...o tigaie sau griddle, o spatulă... Gatiti aproximativ 2–3 minute...”

Multiple-choice mis-translation

EN: “Which is a species of fish? Tope or Rope - Tope.”

RO: “Care este o specie de pește? Sfoară sau frânghie”

The Romanian corpora derived from Alpaca, Alpaca-GPT-4, Self-Instruct-GPT-4, and Dolly provide broad task coverage and ready-to-use pairs for training instruction-following models in Romanian. However, because all four sets were created by machine-translating English sources, they contain machine translation artefacts. Overall, the translated corpora are a good starting point, but improvements need to be made.

Conclusion and future work

Our experiments have shown that small quantities of high quality Romanian language data, such as small parts of the CoRoLa corpus, have the ability to improve a LLMs performance on Romanian language tasks. We employed instruction and evaluation datasets that were previously used to evaluate Romanian language LLMs. However, as these were obtained by using automatic machine translation without human verification, they are affected by errors. This has an impact on both model fine-tuning and evaluation results. Our experiments have shown a steady increase in performance on the evaluation data for the three epochs. This may suggest that additional training may further improve the results. The current data selection was determined by our hardware and time constraints combined with the desire to test the impact of a small dataset. Future experiments will employ a larger part of the CoRoLa corpus. Within the LLMs4EU project, new LLMs will be developed, including support and fine-tuning for the Romanian language. In this context, it is important to understand the type of data necessary for improved performance. This includes: high quality text data, curated instruction datasets, curated evaluation datasets. Our experiments will continue as new models and datasets will be developed. We intend to focus on high-quality, curated, and IPR-cleared data, following the same approach adopted from the outset in the creation of the CoRoLa corpus. Furthermore, all data collected within LLMs4EU will be integrated into CoRoLa and made publicly accessible via the KorAP Corpus Analysis Platform²⁴, as part of our ongoing efforts to continuously enhance both the quantity and quality of CoRoLa’s data.

AI Use Statement

Generative-AI tools were solely used for minor stylistic refinements and copyediting. All content, analysis, and conclusions in this paper are the authors’ own work.

²⁴ <https://korap.racai.ro/>

Acknowledgements

This research was supported by LLMs4EU, co-funded by the Digital Europe Programme under GA 101198470. Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

References

- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, Jianfeng Gao. 2023. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*.
- Daan Broeder, Menzo Windhouwer, Dieter van Uytvanck, Twan Goosen, Thorsten Trippel. 2012. CMDI: a Component Metadata Infrastructure. In *Proceedings of the Workshop "Describing LRs with Metadata: Towards Flexibility and Interoperability in the Documentation of LRs"* at LREC 2012, 1-4.
- Dan Tufiș, Verginica Barbu Mititelu, Elena Irimia, Vasile Păiș, Radu Ion, Niels Diewald, Maria Mitrofan, Mihaela Onofrei. 2019. Little Strokes Fell Great Oaks. Creating CoRoLa, the Reference Corpus of Contemporary Romanian. *Revue Roumaine de Linguistique*, vol 64, no. 3/2019, p. 227-240
- Georg Rehm, Andy Way (Ed.). 2023. *European Language Equality: A Strategic Agenda for Digital Language Equality*. Cognitive Technologies. Springer, Switzerland, 2023.
- Mihai Alexandru Niculescu, Stefan Ruseti, Mihai Dascalu. 2021. RoGPT2: Romanian GPT2 for text generation. In *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 1154–1161. IEEE.
- Mihai Masala, Denis Ilie-Ablachim, Alexandru Dima, Dragos Corlatescu, Miruna Zavelca, Ovio Olaru, Andrei Terian, Simina Terian, Marius Leordeanu, Horia Velicu, Marius Popescu. 2024. "Vorbești Românește?" A Recipe to Train Powerful Romanian LLMs with English Instructions. *arXiv preprint arXiv:2406.18266*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Wang, Kwang-Ting Cheng, Min-Hung Chen. 2024. DoRA: Weight-decomposed low-rank adaptation. In the *Forty-first International Conference on Machine Learning*.
- Ștefan Dumitrescu, Petru Rebeja, Beata Lorincz, Mihaela Gaman, Andrei Avram, Mihai Ilie, Andrei Pruteanu, Adriana Stan, Lorena Rosia, Cristina Iacobescu, *et al.* 2021 LiRo: Benchmark and leaderboard for Romanian language tasks. In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*, Datasets and Benchmarks Track.
- Vasile Păiș, Dan Tufiș. 2023. Language Report Romanian. Chapter in *European Language Equality: A Strategic Agenda for Digital Language Equality*. Springer International Publishing, pp. 199--202.

AUTOMATIC DETECTION OF AI-GENERATED TEXTS FROM MOLDOVA AND ROMANIA (ADAMO). A PROJECT PRESENTATION

VERGINICA BARBU MITITELU¹, VICTORIA BOBICEV²,
VICTORIA ALEXEI², RODICA BRANIȘTE², OLESEA CAFTANATOV²,
MARIA MITROFAN¹, RADU ION¹, ELENA IRIMIA¹, DANIELA
ISTRATI², LUDMILA MALAHOV², SERGIU NISIOI³, ALEXANDR
PARAHONCO², VASILE PĂIȘ¹

¹*Romanian Academy Research Institute for Artificial Intelligence*
{vergi,maria,radu,elena,vasile}@racai.ro

²*Technical University of Moldova*
victoria.bobicev@ia.utm.md,victoria.alexei@iis.utm.md,rod
ica.braniste@iis.utm.md,olesea.caftanatov@math.md,daniela
.istrati@ia.utm.md,ludmila.malahov@math.m,
alexandr.parahonco@usarb.md

³*University of Bucharest*
sergiu.nisioi@unibuc.ro

Abstract

The paper describes the ADAMO project, whose goal is to combine two research directions by accounting for the effect of language varieties in the generation and detection of synthetic data for a medium-resourced language such as Romanian. We are not aware of any study of this kind, even for English or other major languages. As such, we will create and use a large quality corpus for training classifiers to detect AI-generated (more exactly LLM-generated) content in Romanian, as written in Romania and in the Republic of Moldova. This paper is meant to present a project at its very beginning, thus informing the community about our envisaged work. Previous achievements of the Moldovan team with respect to written and spoken corpora collection are also presented.

Keywords: text classifier, Romanian subdialects, AI generated texts, LLMs, CoRoLa

1 Introduction

Large language models (LLMs) require huge amounts of data for training, fine-tuning and even evaluation. In the case of languages, domains, genres, etc. for which such quantities of data are not available, one solution is synthesizing data with the help of LLMs themselves, which use either their already available data

to self-improve or distillations of more powerful models, each approach having its own shortcomings.

Determining whether a text is machine-generated or not carries significant implications, and the need to develop and support underrepresented languages such as Romanian in this specific research field has never been more urgent.

Large national corpora development, especially when the goal is to include texts coming from authorized sources for the language in question, is especially important in today’s generative AI landscape, in which LLMs can generate answers in any language. It is well-known that the balance of pre-training texts in different languages is heavily skewed towards English (e.g. Romanian is only 0.57% vs. English 44% of the Common Crawl¹ corpus, the largest corpus that is used by all open source LLMs for pre-training), so any addition of texts in the language of interest (here, Romanian) can only be beneficial to the fine-tuning of the favoured LLM for that language for future tasks. Additionally, languages have multiple varieties, and Natural Language Processing (NLP) tools (including LLMs) often lack representation and robustness in handling this diversity, as consistently demonstrated by the series of VarDial workshops² held at ACL conferences for the past 10 years.

The ADAMo project proposes to combine two research directions by accounting for: (1) the effect of language varieties in (2) the generation and detection of synthetic data for a medium-resourced language such as Romanian. We are not aware of any study of this kind, even for English or other major languages. As such, we create and use a large quality corpus for training classifiers to detect AI-generated (more exactly LLM-generated) content in Romanian, as written in Romania and in the Republic of Moldova.

The novelty and originality of the project consist in gathering together in the same corpus texts from the two varieties of Romanian with two major goals: their automatic linguistic comparison (seconded by manual analysis) and the impact on the detection of AI-generated texts when two language varieties are involved.

2 Related Work

Butnaru and Ionescu (2019) released the first corpus of Romanian and Moldovan containing a single genre, consisting largely of scraped newspaper articles. While the dataset has been successfully used in the VarDial Shared Task (Zampieri *et al.*, 2019), its linguistic value is limited, as the dataset lacks detailed metadata, linguistic annotations, and is released in a form that is strongly post-processed, with named entities such as person names, organizations, locations

¹ Statistics of Common Crawl Monthly Archives, at <https://commoncrawl.github.io/cc-crawl-statistics/plots/languages.html>

² See the list of workshops on Applying NLP Tools to Similar Languages, Varieties and Dialects - ACL Anthology, at <https://aclanthology.org/venues/varDial/>

missing from the texts. Our approach involves extending the CoRoLa corpus for Romanian (Tufiş *et al.*, 2019) with comparable data of high quality from Moldova, together with detailed linguistic annotations and metadata.

In terms of automatically detecting AI-generated data, King *et al.* (2023) organized a Kaggle competition (with 5000 participants from 113 countries), receiving over 100k submissions. In terms of approaches, classical machine learning employing ensembles of linear models obtained 91% accuracy and the best efficiency, while the best-accuracy models employ deep learning using ensembles of masked language models (98%). Members of our team participated in the competition, our solution obtained position 13 on efficiency and was in top 11% in the entire competition. We used a statistical byte-pair encoder tokenizer together with an ensemble of multiple linear classifiers that were tuned on a custom worst-performance test set locally. This had a good generalization power, obtaining 96% accuracy on the public score and 89.99 on the private leaderboard. All winning solutions employed extensive data augmentation with LLM-generated essays, further justifying the necessity of such datasets. For Romanian, given the limited resource scenario, we are confronted with additional layers of difficulty which we plan to address using a mix of classifiers combining both neural approaches and linguistically-motivated analyses. Furthermore, in our proposal we aim to introduce an additional previously unresearched layer of complexity stemming from different language varieties from Romania and Moldova.

Automatic detection of AI-generated texts in Romanian was also tried by Niţu and Dascălu (2024). They used a transformer-based model that categorizes texts as human or machine-generated, achieving competitive detection results with an F1-score of up to 0.96 in domains such as news, legal and medical. The F1-score for documents in scientific domains was 0.85, primarily due to an insufficient number of available texts. The models were tested on a nearly 60,000-document corpus, comprising 51,709 LLM-generated documents across five domains and a human corpus of 7,387 documents.

Our goal is to bring Romanian AI-generated text detection to the English state-of-the-art (Orenstrakh *et al.*, 2023), in the age of newly released, ten billion parameters, open source LLMs, such as OLMo or Llama-3. For this, we rely on the quantity and quality of the enriched CoRoLa.

Members of our team have been previously involved in building automatic methods for distinguishing between language varieties (Adouane *et al.*, 2016; Ciobanu *et al.*, 2016) showing that language varieties written in standard form in different world regions (Arabic, French, Serbo-Croatian, Portuguese, Spanish and others) can be distinguished with good accuracy using classical machine learning techniques such as linear classifiers and string kernels. Furthermore, Nisioi *et al.* (2022) showed that speakers of dialectal varieties from the British Isles (Scotland, Ireland, England) are distinguishable even after their speeches in the European

Parliament are heavily edited. The authors point to the fact that translations from different languages (Rabinovich *et al.*, 2016) and from dialectal varieties preserve the patterns of their original language varieties.

For this project, we conducted preliminary experiments using a corpus without clearance for Intellectual Property Rights (IPR) for redistribution, gathered via web crawling from the news and sports domains. Unlike Butnaru and Ionescu (2019), we added more metadata, restricted the news types, and controlled the number of samples per news outlet and class. We collected 500 texts from six different Romanian newspapers and seven Moldovan newspapers (6500 texts in total). Baseline experiments showed that classifiers for language varieties can achieve 80% accuracy using linear classifiers and function words (conjunctions, prepositions, pronouns), which capture writing style independent of content. In ADAMo, we will further develop this investigation using a larger IPR-cleared corpus having a wider variety of genres. In terms of computational tools, we will use more advanced state-of-the-art and explainable algorithms, neural methods employing language models, statistical analyses and human interpretations of the results. Therefore, by enriching the IPR-cleared corpus CoRoLa, this project will further aid in this research field (and other future endeavours).

3 Project’s Objectives

Each subsection below expands on the four objectives of ADAMo’s project, namely:

- enriching CoRoLa, the representative corpus of contemporary Romanian, with texts originating from the Republic of Moldova;
- creation of an AI-generated corpus in compliance with the metadata scheme used in CoRoLa;
- automatic identification of linguistic differences between texts from Romania and those from Moldova;
- development of a text classifier able to detect AI-generated texts.

3.1 Enriching CoRoLa with Texts Originating from the Republic of Moldova

The goal of this objective is to enhance CoRoLa by incorporating at least 15M tokens from a diverse range of texts from the Republic of Moldova. This will ensure that the corpus accurately reflects the linguistic variations and contemporary usage of the Romanian language across both countries. Besides the two modalities (written and oral), other criteria for text collection will be followed, in an effort to mimic (at a smaller scale) the structure of existing CoRoLa: various text genres and domains, various text sources (publishing houses, media, public institutions, etc.), time coverage (last 2 decades).

Similar to CoRoLa collection, text providers will be contacted and collaboration protocols will be signed to ensure IPR clearance for the project

partners. Metadata for the collected texts will be created. The text files will be curated, diacritics will be restored when missing, written texts will be automatically processed with the flow and tools available from the Romanian partner. Oral texts will be automatically transcribed and checked when their written equivalent is not available from providers. The oral texts and their transcriptions will be automatically aligned.

At the moment, two different query platforms are used to search for data of interest in CoRoLa: 1. KorAP³ (Diewald *et al.*, 2019), a corpus analysis platform suited for large, annotated textual corpora and complex search queries, made available to the coordinator for integration with CoRoLa as a result of the DRuKoLA⁴ project (Cosma *et al.*, 2016). 2. OCQP⁵ (Oral Corpus Query Platform), developed at the coordinator, used for querying the oral component through searches on the textual transcription of the audio content (previously aligned at word level) based on the form, lemma and/or the morpho-syntactic annotations. The written texts in the enhanced corpus will be re-indexed in the KorAP platform, making available for querying the new syntactic annotation level and the texts collected from the Republic of Moldova. The oral component will be added to the already indexed CoRoLa files in OCQP.

At the end of the project, we expect to have significantly enriched CoRoLa, which, at that point, will also represent the contemporary Romanian language usage from Moldova.

3.2 Creation of an AI-generated Corpus in Compliance with CoRoLa's Metadata Scheme

Whether the goal is to generate specialized content for legal, medical, or academic purposes, LLMs can be adapted to meet specific needs. For this objective, advanced language models such as LLaMA-3 (Dubey *et al.*, 2024), OLMo (Groeneveld *et al.*, 2024), Mistral 7B (Jiang *et al.*, 2023), Mixtral 8x22B (Jiang *et al.*, 2024), Qwen2.5 (Yang *et al.*, 2024) Gemma (Team Gemma *et al.*, 2024) and OpenLLM-Ro (Masala *et al.*, 2024) will be used. For creating the LLM-generated corpus, we will employ custom prompts, derived from the CoRoLa metadata, thus aligning the generated text to the domains of interest. We will create custom prompts for events in both Romania and Moldova, in order to check the model's behaviour in these two contexts. Given that new LLMs become available frequently, the actual models used will be selected at project start time, either from the list of models detailed in this proposal or replaced by newly available models. Romanian language varieties and Romanian language fine-

³ <https://korap.racai.ro/>

⁴ <https://www.ids-mannheim.de/digspra/kl/projekte/drukola/>

⁵ http://89.38.230.23/corola_sound_search/index.php

tuned models will be employed in addition to base multilingual variants of LLM architectures to explore the nuances in Romanian language generation.

At the end of the project, we envisage to have an AI-generated corpus comparable in topics and genre with the enriched CoRoLa, to serve as training and testing bed for the classifiers to be developed in the project.

3.3 Automatic identification of linguistic differences between texts from Romania and Moldova

The goal of this objective is to identify the linguistic differences between Romanian as written and spoken in Romania and Moldova, by automatically comparing the large data quantities from CoRoLa. This type of comprehensive study has never been conducted at such a scale and necessitates the use of advanced computational methods.

Though texts from the two countries are written in standard Romanian, previous investigations of our team members have shown that some differences between them can be identified, especially in the use of function words (prepositions, conjunctions, determiners). With an enlarged CoRoLa, also annotated at the syntactic level, we will now be able to undergo studies on a larger scale and on multiple linguistic levels: lexical, morphological, syntactic and semantic, with a special eye to preposition selection, collocations, multi-word expressions. Moreover, having available texts from various genres, we can investigate which ones are more prone to capturing the differences between the language varieties spoken in the two countries. The intuition is that, for instance, imaginative texts, memoirs and blog posts will be more relevant for such linguistic investigations. Oral texts, on the other hand, are expected to be rich in lexical particularities and pronunciation differences between speakers of Romanian in the two countries. Our objective is to provide a detailed quantification and explanation of these differences.

3.4 Development of a Text Classifier Able to Detect AI-generated Texts

With the increase in parameter size, often linked to an increased training set size, texts generated by LLMs become more difficult to detect. We aim to train a classifier system able to detect if a Romanian language text is written by a human or generated automatically via an LLM. For this purpose we consider only texts that are large enough (multiple sentences, multiple paragraphs) to hold relevant linguistic clues. Furthermore, we aim to train and test the system taking into account the subtle differences between the Romanian language employed in Romania and in Moldova. We will use the LLM-generated corpus together with the resources in CoRoLa to train the classification system. Since we know the actual LLM that generated the text, we will also explore the possibility to identify the architecture used for generating the text. This corresponds to the recent observation that mainstream detection methods prove problematic in a wide-

ranging testbed, particularly in out-of-distribution scenarios (Li *et al.*, 2024). Thus, by including the model architecture as a potential indicator, combined with the language variety, we aim to improve over existing state-of-the-art results.

All these 4 objectives (1) build on the experience and competence of both project teams (the Romanian and the Moldovan ones, as shown by their previous results mentioned above); (2) will increase the research capabilities of the Moldovan team, that will acquire the necessary know-how for corpus collection, curation and processing; and (3) will enhance the scientific performance of both teams, who will collaborate towards detecting AI-generated texts using up-to-date LLMs and technical solutions.

4 Experiments with texts from Moldova

Texts are the fundamental element of cultural identity and collective memory, reflecting the history, traditions and values of a society. This project proposes a new integration of linguistic diversity analysis from Romanian and Moldovan variants with synthetic data detection, a direction rarely explored even for high-resource languages. The main objective of the Moldovan participants is to provide access to regional texts in Latin script in various genres from the same period of time as in the Romanian corpus CoRoLa, namely, from 1945 till now. This will lead to the expansion of CoRoLa with Moldovan data supplemented with linguistic annotations and metadata.

For a long period of time, from 1945 until 1990, the part of Moldova between the rivers Nistru and Prut was a Soviet Republic, one of the many republics of the Soviet Union. All this period, texts were printed in Cyrillic alphabet. Andrunakievich Institute of Mathematics and Informatics, Moldovan State University, has created information resources in the Latin alphabet, obtained by digitizing Cyrillic documents published in the 17th–20th centuries. For this purpose, a set of specialized tools has been developed; most of them are integrated into the HeDy (Bumbu *et al.*, 2023) digitization platform⁶. The process includes several stages:

1. Image pre-processing. This is especially important when working with old texts damaged by time or storage conditions. The HeDy platform provides a wide range of functions: resolution correction, page orientation, conversion to black and white format, noise reduction, and line straightening.

2. Optical character recognition. The user selects a model that matches the historical period and typographic features. Several models may exist for the same eras, especially in the case of transitional alphabets. Sometimes it is useful to test several options to determine the most suitable one.

The HeDy platform has been trained on various corpora:

⁶ <https://digitizare.math.md>

- 17th century – 2 OCR models, over 10,000 training examples, dictionary of over 4,500 words;

- 18th–19th centuries – 10 OCR models, dictionary of over 10,000 words;

- 20th century – 1 OCR model, dictionary of over 450,000 words.

Despite the high quality of recognition, it is not completely accurate, and manual correction is required, which is possible both in Cyrillic text and in transliterated text. The HeDy platform supports such corrections.

3. Transliteration. At this stage, the text is transliterated into the Latin Romanian script while preserving the original spelling. If necessary, the user can adapt it to current spelling standards.

The HeDy platform offers solutions for maximum support of all stages of recognition of heterogeneous documents. All stages of processing are characterized by fairly good accuracy: 95-96% for OCR, 98% for font classification and 96% for transliteration. The web platform is an open access tool that can be used by researchers from various fields, libraries, archives, publishers.

Examples of the application of the proposed technology are republications of several fiction books, in particular, republication of scientific publications published in Cyrillic during the Soviet period (Bumbu *et al.*, 2023). Another example is the actualization of dialect texts printed in Cyrillic in the 1970s and 1980s without respecting international standards.

For dialect texts recorded phonetically, a separate tool for recognition and transliteration into Latin Romanian was created using the modern phonetic system IPA (International Phonetic Alphabet) (Popovschi *et al.*, 2021). Some of these texts have already been published.

Плуг	Плуг	Plug
Кыте петричеле'н фынтынэ — Атытя оале ку смынтынэ! Хэй, Хэй! Кыте лемне суб ватрэ — Атыця флэкэй ла фатэ! Хэй, хэй! Кыте пене пе кукош — Атыця бэець бурдухошь! Хэй, хэй! Кыте пае пе касэ — Атыця галбень пе масэ! Хэй, хэй! Ешь лелицэ, ку ковригул, Кэ мэ рупе фригул! Хэй, хэй! Де ла Лукиян Вас. Т., 30 ань, с. Ворниченъ, р-нул Кэлэрашь, а Морару С. Г.; 223, 12.	Кыте петричеле'н фынтынэ — Атытя оале ку смынтынэ! Хэй, Хэй! Кыте лемне суб ватрэ — Атыця флэкэй ла фатэ! Хэй, хэй! Кыте пене пе кукош — Атыця бэець бурдухошь! Хэй, хэй! Кыте пае пе касэ — Атыця галбень пе масэ! Хэй, хэй! Ешь лелицэ, ку ковригул, Кэ мэ рупе фригул! Хэй, хэй! Де ла Лукиян Вас. Т., 30 ань, с. Ворниченъ, р-нул Кэлэрашь, Морару С. Г.; 223, 12.	Câte petricele'n fântână — Atâtea sale cu smântână! Hăi, Hăi! Câte lemne sub vatră — Atâția flăcăi la fată! Hăi, hăi! Câte pene pe cucoș — Atâția băeți burduhoși! Hăi, hăi! Câte pae pe casă — Atâția galbeni pe masă! Hăi, hăi! Eși leliță, cu covrigul, Că mă rupe frigul! Hăi, hăi! De la Luchian Vas. T., 30 ani s. Vorniceni, r-nul Călărași, Moraru S. G.; 223, 12.
a) Original Text	b) Recognized Text	c) Transliterated Text

Figure 1: An example of the digitization process.

One of the interesting works was digitization of folklore collections (myths, legends, fairy tales, etc.) collected by Moldovan scholars and published in the 1960s-80s in Cyrillic [Proverbe și zicători. Alcătuirea de E. Junghietu. Chișinău, “Știința”, 1981 etc.]. Folklore collections from the Republic of Moldova, through the use of Cyrillic and local dialects, preserve the linguistic diversity of the region, which is an important aspect for the implementation of the project (Colesnicova *et al.*, 2024). Figure 1 presents an example of the digitization and transliteration process.

5 Experiments with speech from Moldova

In the ADAMo project, we aim to enrich the corpus also with speech from R. Moldova. At the moment, the oral component of CoRoLa consists of almost 152 hours of recordings, with associated transcribed texts; mainly professional recordings from various sources such as radio stations, recording studios, etc. (Tufiş *et al.*, 2019). A part of the oral corpus is represented by texts read by professional speakers recorded in studios, and extracts from Romanian Wikipedia read by non-professionals, by volunteers, recorded in non-professional environments. So far, the corpus contains 120 hours of news and radio interviews, 3.44 hours of fairy tales, 4 hours of Romanian Wikipedia and almost 24 hours of other various texts. Regarding the annotation, only a little more than half of the collected recordings (around 300 hours) have been automatically pre-processed. Their written counterpart, almost 750,000 tokens, has been lemmatised, part-of-speech tagged and syllabicated (Boroş and Dumitrescu, 2015); some allophones were identified and analysed. The oral and the written texts are time aligned at the sentence, word and phoneme levels (Boroş *et al.*, 2018) and this alignment was automatically encoded in separate files.

Currently, the oral part of the CoRoLa corpus can be queried⁷ only via their transcription followed by rendering the relevant aligned speech segment. This is ensured by the in-house Oral Corpus Query Platform (OCQP) created by Vasile Păiș (Barbu Mititelu *et al.*, 2018). It allows searching for lemmas or occurring words and morphological restrictions on them. The results are either sequences of 5 words or the whole sentences in which the specified form occurs. Besides the written form, the user also gets displayed the oral files, both for the searched form and the whole sentence to which it belongs.

The language variety in Moldova differs considerably from that in Romania. The most similar pronunciation can be heard in the Romanian part of historical Moldova territory, but the Romanian speaking person from the left bank of the Prut River can be easily recognized by the way of speaking. CoRoLa does not contain any speech produced in the Republic of Moldova.

⁷ http://89.38.230.23/corola_sound_search/index.php

While the expected volume of added texts is 15M tokens, the oral part is expected to be around 50 hours of recordings. Oral texts will be automatically transcribed and checked when their written equivalent is not available from providers. The oral texts and their transcriptions will be automatically aligned.

The first pilot experiment of speech collection and transcription was organised in order to assess the quality of work of various transcription applications. Over the past few years, the field of speech processing has been transformed by introducing powerful tools based on deep learning (Cui *et al.*, 2024). While the previous work on oral corpus transcription included manual transcription or at least, verification and correction of the automatic speech to text systems (Păiș *et al.*, 2023), modern neural models pre-trained on huge volumes of speech are able to perform this transformation with remarkable quality (Mehriş *et al.*, 2023).

Initially, we had to find accessible sources of speech from Moldova. One of the most challenging issues while collecting a corpus is intellectual propriety clearance (IPR). It is considered that the copyright clearance for the web content on the Internet is easier than for material collected from printed publications. A vast majority of web content authors are happy to be able to reach as many people as possible and do not legally protect their products. However, this may not be the case. Due to the active development of the intellectual property laws, all material available on the Web is copyrighted, and all digital publications should be treated the same way as printed works.

YouTube was used as a source for the speech generated in Moldova. There are plenty of videos with registered speech of Moldova's citizens. All plenary sessions of the Republic of Moldova parliament are registered and can be accessed on YouTube: several talk shows as "Vorbește Moldova" (Moldova speaks), podcasts such as DAD Podcast, TANDEM PODCAST, platforms such as Priveşcu Moldova with live broadcasts of events from the Republic of Moldova, conferences, government meetings, parliament, town halls, underground concerts, sports and other events, or news platforms as NewsMaker Moldova⁸ and newspapers such as SP⁹.

All content uploaded to YouTube can obtain one of the following license types:

- Standard YouTube license,
- Creative Commons Attribution license.

Standard YouTube license permits to view or listen to Content for the personal, non-commercial use and applies the following restrictions: users are not allowed to access, reproduce, download, distribute, transmit, broadcast, display, sell, license, alter, modify or otherwise use any part of the YouTube content except: (a) as specifically permitted by the Service; (b) with prior written permission from

⁸ <https://newsmaker.md/ro>

⁹ <https://esp.md/ro>

YouTube and, if applicable, the respective rights holders; or (c) as permitted by applicable law.

Provider of Content may add a Creative Commons Attribution license to his/her video, which is the only one supported by YouTube, and is designed to encourage sharing and collaboration. This means that any YouTube user has the right to use, edit, and remix the video, even for commercial purposes.

We selected YouTube videos with Creative Common Licence by ensuring that the line “Licence Creative Commons Attribution licence (reuse allowed)” exists in the video description. This permits us to reuse the video in any form we consider necessary.

Also, the categories of content excepted from the IPR restrictions in Romania are: political, legislative, administrative and juridical (Tufiş *et al.*, 2016). This includes video with plenary sessions of the parliament and other governmental organisation meetings.

The third type of YouTube sources are videos of the news portals we have permission to use from their authors. By now, we have obtained permissions from several news portals such as *sp.md*¹⁰ and *Ziarul de gardă*¹¹.

All collected speech was transcribed. Plenty of applications that can transcribe speech are accessible online; many of them are able to work with Romanian or Russian. We used and compared several free applications with web GUI. In this experiment, several different online transcription tools were tested in order to evaluate their quality of transcript, speed, interface usability and accessibility in terms of free or paid options. Any2Text¹² has a free option, is fast and easy to use, but its recognition of speakers and languages is inaccurate. Notta AI¹³ provided limited free usage. Notta AI was convenient but weak in language detection. Cockatoo¹⁴ was highly accurate in quality of transcription and usability, correctly splitting speakers in a focused 2-person conversation (parliament session), but the free option was restricted to only two uses before requiring payment. Good Tape¹⁵ allowed free transcription with a significant waiting time (over one hour), and while slow, it was effective in dividing text by time and speakers. HappyScribe¹⁶ provided high-quality transcription and automatic speaker separation, but its free version was limited to machine translation, with human transcription costing 4 USD per minute. Vizard AI, while free, added useful subtitles, and keywords, though its transcription quality was moderate. Vizard AI¹⁷, while free, added

¹⁰ <https://esp.md/ro>

¹¹ <https://www.zdg.md/>

¹² <https://any2text.com/>

¹³ <https://www.notta.ai/>

¹⁴ <https://www.cockatoo.com/>

¹⁵ <https://goodtape.io/>

¹⁶ <https://www.happyscribe.com/>

¹⁷ <https://vizard.ai/>

useful subtitles, and keywords, though its transcription quality was moderate. Maestra AI¹⁸ offered excellent quality and speaker separation, but only transcribed one free minute, which limited its practical usability. Turboscribe¹⁹ is based on the powerful Whisper model but allows transcribing only three files per day.

Table 1: Word Error Rate for transcription of three videos by several transcription applications.

No.	Application	WER for video 2	WER for video 3	WER for video 4
1	Any2Text	0.022	0.045	0.126
2	Notta AI	0.065	0.139	0.298
3	Good Tape	0.130	0.051	0.045
4	HappyScribe	0.087	0.020	0.010
5	Vizard AI	0.022	0.040	0.065
6	Turboscribe	0.239	0.043	0.074
7	Cockatoo	0.173	0.040	0.126

We tested all above-described applications on four short videos. The first one was a snippet from a fairy tale read by a professional narrator. This fragment was transcribed absolutely correctly by all tested applications. The second video was a fragment of a parliament session with the length of 2 minutes but with 40 seconds of applause and shoutings in the middle. The third video was a fragment from the interview on the street with a length of 2 minutes 21 seconds. The last video was a snippet of an interview in the studio. The accuracy of transcription is measured in Word Error Rate (WER): rate of errors per total number of words in a sample. Table 1 summarizes the results of application testing. It contains WER for videos number 2, 3, and 4. The quality of the video number 1 was so high that all applications obtained WER=0 for this video.

The video with applause and shouting appeared to be difficult for many applications. Turboscribe and Cockatoo stopped on this obstacle and could not continue transcribing. Other applications added various words instead of applause. The best to cope with the noise were Vizard AI and Any2Text. In video 4, people speak a mixture of Romanian and Russian and this created problems for Any2Text and Notta AI; other applications managed to recognize such speech. The best with this was Turboscribe.

¹⁸ <https://maestra.ai/>

¹⁹ <https://turboscribe.ai/>

6 Conclusions

In this paper we have presented the new project, ADAMo, a common endeavour of a team from Romania and one from the Republic of Moldova. The main aim of the project, that of developing a text classifier able to detect AI-generated texts in the Romanian varieties spoken in Romania and Moldova, will be for the first time pursued with the help of the large representative contemporary Romanian corpus, CoRoLa, containing written and oral texts of the highest quality, ensured by the careful selection of texts providers, by data curation, metadata annotation and use of state-of-the-art technology in their processing.

AI Use Statement

Generative-AI tools were used solely for minor stylistic refinements and copyediting. All content, analysis, and conclusions in this paper are the authors' own work.

Acknowledgements

This work was supported by a grant of the Ministry of Education and Research, CCCDI – UEFISCDI, project number PN-IV-PCB-RO-MD-2024-0142, within PNCDI IV.

References

- Aaron Grattafiori, *et al.* 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Albert Q. Jiang, *et al.* 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825*.
- Albert Q. Jiang, *et al.* 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Ambuj Mehrish, Navonil Majumder, Rishabh Bhardwaj, Rada Mihalcea, Soujanya Poria. 2023. A review of deep learning techniques for speech processing. In *Information Fusion*, 99, 101869.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, Zhenru Zhang. 2024. Qwen2. 5-math technical report: Toward mathematical expert model via self-improvement. *arXiv preprint arXiv:2409.12122*.
- Ana-Maria Ciobanu, Sergiu Nisioi, Liviu P. Dinu. 2016. Vanilla Classifiers for Distinguishing between Similar Languages. In *Proceedings of the Third Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial3)*, pages 235–242, Osaka, Japan. COLING 2016.
- Andrei Butnaru, Radu Tudor Ionescu, 2019. MOROCO: The Moldavian and Romanian Dialectal Corpus. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 688–698.
- Cornelia Kiefer, 2019. Quality indicators for text data. <https://dl.gi.de/items/1104f8cf-383a-4a0a-8a9a-e6866cdaaa84>

- Dan Tufiş, Verginica Barbu Mititelu, Elena Irimia, Ştefan Daniel Dumitrescu, and Tiberiu Boroş. 2016. The IPR-cleared Corpus of Contemporary Written and Spoken Romanian Language. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC16)*, pages 2516–2521.
- Dan Tufiş, Verginica Barbu Mititelu, Elena Irimia, Vasile Păiş, Radu Ion, Niels Diewald, Maria Mitrofan, Mihaela Onofrei, Little strokes fell great oaks. Creating CoRoLa, the reference corpus of contemporary Romanian. In *RRL*, vol 64, no. 3/2019, p. 227-240.
- Dirk Groeneveld, et.al. 2024. Olmo: Accelerating the science of language models. *arXiv preprint arXiv:2402.00838*.
- Ella Rabinovich, Sergiu Nisioi, Noam Ordan, Shuly Wintner. 2016. On the Similarities Between Native, Non-native and Translated Texts. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1870-1881.
- Emily Wenger. 2024. AI produces gibberish when trained on too much AI-generated data. In *Nature*. 631(8022):742-743
- Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, Ajmal Mian. 2023. A comprehensive overview of large language models. *arXiv preprint arXiv:2307.06435*.
- Jules King, Perpetual Baffour, Scott Crossley, Ryan Holbrook, and Maggie Demkin. 2023. LLM - Detect AI Generated Text. Kaggle. <https://kaggle.com/competitions/llm-detect-ai-generated-text>
- Liliana Popovschi, Ludmila Malahov, Vlada Colesnicova. 2021. Prelucrarea digitală a textelor dialectale publicate la Chişinău în anii 1969–1987. In *Revista Academos*, 2021, (4)63, 48–53.
- Marcos Zampieri, Shervin Malmasi, Yves Scherrer, Tanja Samardžić, Francis Tyers, Miikka Silfverberg, Natalia Klyueva, Tung-Le Pan, Chu-Ren Huang, Radu Tudor Ionescu, Andrei M. Butnaru, and Tommi Jauhiainen. 2019. A Report on the Third VarDial Evaluation Campaign. In *Proceedings of the Sixth Workshop on NLP for Similar Languages, Varieties and Dialects*, pages 1–16, Ann Arbor, Michigan. Association for Computational Linguistics.
- Marie-Catherine de Marneffe, Chris Manning, Joakim Nivre, Dan Zeman. 2021. Universal Dependencies. In *Computational Linguistics*, ISSN 1530-9312, vol. 47, no. 2, pp. 255-308.
- Melania Nitu, Mihai Dascalu. 2024. Beyond Lexical Boundaries: LLM-Generated Text Detection for Romanian Digital Libraries. In *Future Internet* 2024, 16(2), 41.
- Michael Sheinman Orenstrakh, Oscar Karnalim, Carlos Anibal Suarez, Michael Liut. 2023. Detecting LLM-Generated Text in Computing Education: A Comparative Study for ChatGPT Cases. *arXiv:2307.07411*
- Mihai Masala, Denis C. Ilie-Ablachim, Dragos Corlatescu, Miruna Zavelca, Marius Leordeanu, Horia Velicu, Marius Popescu, Mihai Dascalu, Traian Rebedea. 2024. OpenLLM-Ro--Technical Report on Open-source Romanian LLMs trained starting from Llama 2. *arXiv preprint arXiv:2405.07703*.
- Niels Diewald, Verginica Barbu Mititelu, Marc Kupietz. 2019. The KorAP User Interface. Accessing CoRoLa via KorAP. In *RRL*, LXIV, 3, p. 265–277.

- Orphee de Clercq, Veronique Hoste. 2016. All Mixed Up? Finding the Optimal Feature Set for General Readability Prediction and Its Application to English and Dutch. In *Computational Linguistics*, 42(3):457–490.
- Ruibao Liu, Jerry Wei, Fangyu Liu, Chenglei Si, Yanzhe Zhang, Jinmeng Rao, Steven Zheng, Daiyi Peng, Diyi Yang, Denny Zhou, Andrew M. Dai. 2024. Best Practices and Lessons Learned on Synthetic Data. In *First Conference on Language Modeling*, arXiv:2404.07503v2
- Ruxandra Cosma, Dan Cristea, Marc Kupietz, Dan Tufiş, Andreas Witt. 2016. DRuKoLA—towards contrastive German-Romanian research based on comparable corpora. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016), Portorož, Slovenia*, pp. 28-32. ELRA, 2016.
- Sergiu Nisioi, Ana S. Uban, Liviu P. Dinu. 2022. Identifying Source-Language Dialects in Translation. In Florentina Hristea, Cornelia Caragea, Jakub Nalepa (eds.) *Special Issue on Natural Language Processing (NLP) and Machine Learning (ML)—Theory and Applications, Mathematics*, 10(9), 1431.
- Shuyi Dong, Jin Mao, Qing Ke, Lei Pei. 2024. Decoding the writing styles of disciplines: A large-scale quantitative analysis. In *Information Processing & Management*, Volume 61, Issue 4.
- Thomas Mesnard, *et al.* 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Tiberiu Boroş, Ştefan Dumitrescu, Vasile Pais. 2018. Tools and resources for Romanian text-to-speech and speech-to-text applications. In *Proceedings of The Second Workshop on Multi-Language Processing in a Globalising World and The First Workshop on Multilingualism at the intersection of Knowledge Bases and Machine Translation*.
- Tiberiu Boroş, Ştefan Dumitrescu. 2015. Robust deep-learning models for text-to-speech synthesis support on embedded devices. In *Proceedings of the 7th International Conference on Management of computational and collective Intelligence in Digital Eco-System (MEDES'15)*, 98-102.
- Tudor Bumbu, Lyudmila Burtseva, Svetlana Cojocaru, Alexandru Colesnicov, Ludmila Malahov. 2023. *Digitization of Romanian historical printings*. Chişinău: Vladimir Andrunachievici Institute of Mathematics and Computer Science, Moldova State University.
- Vasile Păiş, Radu Ion, Dan Tufiş. 2020. A Processing Platform Relating Data and Tools for Romanian Language. In *Proceedings of the 1st International Workshop on Language Technology Platforms*, 2020, Marseille, France, pp. 81-88.
- Vasile Păiş, Verginica Barbu Mititelu, Radu Ion, Elena Irimia. 2023. Evaluating a Fine-Tuned Whisper Model on Underrepresented Romanian Speech, In *Proceedings of the International Conference on Speech Technology and Human-Computer Dialogue (SpED)*, Bucharest, Romania, 2023, pp. 141-145.
- Verginica Barbu Mititelu, Dan Tufiş, and Elena Irimia. 2018. The Reference Corpus of the Contemporary Romanian Language (CoRoLa). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Vlada Colesnicova, Olesia Caftanator, Svetlana Cojocaru, Alexandru Colesnicov, Ludmila Malahov. 2024.. Digitization of Moldovan folklore texts using the HeDy

- platform. In *Proceedings of International Conference IMCS-60*, Chişinău, October 10-13, pp. 272-276, ISBN 978-9975-68-515-3.
- Wafia Adouane, Nasredine Semmar, Richard Johansson, Victoria Bobicev. 2016. Automatic Detection of Arabicized Berber and Arabic Varieties. In *Proceedings of the Third Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial3)*, pages 63–72, Osaka, Japan. COLING 2016.
- Wenqian Cui, Dianzhi Yu, Xiaoqi Jiao, Ziqiao Meng, Guangyan Zhang, Qichao Wang, Yiwen Guo, Irwin King. 2024. *Recent Advances in Speech Language Models: A Survey*. ArXiv, abs/2410.03751.
- Yafu Li, Qintong Li, Leyang Cui, Wei Bi, Zhilin Wang, Longyue Wang, Linyi Yang, Shuming Shi, Yue Zhang, 2024. Mage: Machine-generated text detection in the wild. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 36-53.

PERSONALIZED VIRTUAL TUTORS THROUGH REASONING LLMs

ALEXANDR PARAHONCO, OLESEA CAFTANATOV, TUDOR BUMBU,
MARIUS CERESCU, NICHTA DEGTEARIOV, DAN BUMBU,
ION BOSTAN, IURIE CIALA

*Vladimir Andrunachievici Institute of Mathematics and Computer Science,
Moldova State University*

*alexandr.parahonco@math.md, olesea.caftanatov@math.md,
tudor.bumbu@math.md, marius.cerescu@gmail.com,
degteariov.nichita@math.md, dan.bumbu@math.md,
ion.bostan@math.md, iurie.ciala@math.md*

Abstract

This paper reports preliminary experiments on reasoning-based Large Language Models (LLMs) for personalized virtual tutoring. A two-stage approach was applied: first, a GPT-based assistant was developed to emulate Professor Boris Ivanovici Cinic’s pedagogical methodology, trained on L^AT_EX-formatted mathematics manuscripts and national curriculum textbooks to ensure accurate formula interpretation and step-by-step guidance. In the second stage, the tutor was culturally adapted as “Guguță”, a folkloric Moldovan character reimaged as an educator, blending Cinic’s problem-solving approach with Bessarabian cultural expression. The results highlight the potential of LLMs for culturally responsive, pedagogically grounded educational tools.

Keywords: Reasoning LLMs, Virtual Tutor, personalization, transformers, PyTorch, Ollama

1 Introduction

The rapid advancement of Large Language Models (LLMs) is reshaping the landscape of education, offering new possibilities for adaptive and personalized learning (Pawar *et al.*, 2025). Within this context, the AI Guguță Project (Code: 25.80012.0807.21TC) aims to create the first interactive, AI-driven educational platform in the Republic of Moldova, designed for primary school students.

This paper reports on preliminary experiments that explore the potential of reasoning-based LLMs for personalized virtual tutoring. Our two-stage approach first focused on developing a virtual assistant with OpenAI GPT Builder, modeled on the pedagogical methodology of Professor Boris Ivanovici Cinic. The assistant was trained on a L^AT_EX-formatted mathematics manuscript to ensure precise formula handling, complemented by 4th–6th grade textbooks aligned with the

national curriculum. Additional features include multilingual support, optimized prompting, and web resource integration. The resulting prototype promotes active learning by guiding students through step-by-step reasoning rather than providing direct answers, in line with established educational principles.

The second stage explored cultural adaptation by reimagining the tutor as “Guguță”, a beloved Moldavian folk character created by Spiridon Vangheli. In this form, Guguță is portrayed as a 70-year-old educator who retains authentic folkloric expression while applying Cinic’s structured problem-solving methodology. This combination of pedagogy and folklore illustrates how AI tutors can be made culturally resonant, enhancing student engagement while preserving educational and cultural legacies.

The AI Guguță Project thus represents an important step in modernizing education through interactive, curriculum-aligned, and culturally responsive tools. Its novelty lies in merging advanced conversational technologies with local linguistic and cultural heritage, supported by the creation of a regional corpus and the development of open-source AI models tailored to the Moldavian context.

2 Project Background and Methodology

2.1 The AI Guguță Project Framework

The AI Guguță Project is designed to address the specific educational needs of Moldavian primary school students by creating an AI-powered platform that respects and incorporates local cultural and linguistic elements. The project’s primary objective is to develop an intelligent tutoring system that can provide personalized mathematics education while maintaining cultural authenticity and pedagogical effectiveness (Li *et al.*, 2024).

2.2 Experimental Setup and Data Collection

Our experimental approach involved the systematic collection and analysis of educational materials specifically relevant to the Moldavian educational context. The dataset comprised:

1. **National Mathematics Textbooks:** We analysed mathematics textbooks for grades 4-6 from the Moldovan national curriculum, extracting pedagogical patterns, problem-solving methodologies, and cultural references embedded within the educational content.

2. **Moldavian Proverbs Corpus:** A comprehensive collection of approximately 500 Moldavian proverbs arranged in alphabetical order was compiled to serve as a cultural knowledge base. These proverbs serve multiple purposes: they provide cultural context, offer wisdom-based learning reinforcement, and maintain the authentic Moldavian character in educational interactions.

3. **Character Development Focus:** The main focus of our research centered on developing the Professor Guguță personality, drawing inspiration from the beloved character created by Spiridon Vangheli. Guguță, originally a fictional symbolic Moldavian hero-child from Vangheli's stories, was reimagined as an aged professor of mathematics, maintaining the symbolic nature while adapting to an educational role (Ciobanu, 2013).

3 Fine-Tuning LLMs for Personalization

Creating a personalized virtual tutor requires adapting a general-purpose LLM to a specific domain and persona. Two primary methods for this are prompt fine-tuning and full model fine-tuning with frameworks like Transformers and PyTorch (Shin *et al.*, 2025).

3.1. Prompt Fine-Tuning with Ollama Modelfile

Prompt fine-tuning, also known as prompt engineering, is a lightweight and accessible method for customizing an LLM's behaviour. It involves crafting detailed system prompts and instructions that guide the model's responses without altering its underlying weights (Ferrag *et al.*, 2025). Tools like Ollama and its Modelfile provide a straightforward way to implement this approach.

Analysis of the Professor Guguță System Prompt

Our implementation utilizes a comprehensive system prompt that establishes the Professor Guguță persona through multiple dimensions:

Identity and Personality Framework: The system prompt defines Professor Guguță as “a distinguished elderly mathematician and teacher in Moldova” who embodies the folk-cultural image of “Guguță” from Moldavian folklore. The prompt explicitly states: “You embody the folk-cultural image of ‘Guguță’, a figure from Moldavian folklore symbolizing youthful innocence, yet now envisioned as a mature, wise adult who carries the soul of Moldavian character into every conversation.”

Cultural Integration: The prompt incorporates specific cultural elements, including the metaphorical reference to “cușma lui Guguță” (Guguță's hat) and traditional Moldavian food references like “ma'ma'liga.” This cultural grounding ensures that mathematical concepts are explained through familiar cultural metaphors, enhancing student engagement and comprehension.

Linguistic Flexibility: The system prompt addresses multilingual capabilities: “Your native tongue is Romanian. Your speech should reflect the unique phonetic, lexical, and syntactic features of this variety, embodying Romanian culture and heritage.” This approach acknowledges the complex linguistic landscape of Moldova while maintaining cultural authenticity.

Pedagogical Approach: The prompt is configured to reproduce the teaching style of Professor Boris Ivanovici Cinic, emphasizing structured guidance, critical thinking, and step-by-step problem solving. Instead of providing direct answers, the assistant encourages students to work through solutions independently, thereby preserving Cinic’s legacy of fostering discipline, clarity of reasoning, and intellectual autonomy in mathematics education.

Conversational Etiquette: A sophisticated greeting protocol is implemented, requiring formal address (“Professor Guguță”) and providing culturally appropriate responses in multiple languages (Romanian, Ukrainian, and English). This feature teaches students proper respect while accommodating the multilingual reality of Moldavian society.

3.2. Fine-Tuning with Transformers and PyTorch

For deep customization, full model fine-tuning is necessary. This involves retraining a pre-trained LLM on a custom dataset, in our case, Moldavian math curricula (Nawrot, 2023). Using frameworks like Hugging Face Transformers and PyTorch, this process is significantly more resource-efficient than training from scratch. We employed Low-Rank Adaptation (LoRA) to further improve efficiency by updating only a small subset of the model’s parameters (Hu *et al.*, 2022).

LoRA represents a breakthrough in efficient fine-tuning by decomposing weight updates into smaller, trainable matrices. Instead of updating all model parameters, LoRA introduces small adapter modules that capture task-specific knowledge while keeping the base model frozen. This approach dramatically reduces memory requirements and training time while maintaining performance comparable to full fine-tuning.

Essentially one of the formats suitable for the fine-tuning process is JSONL (JSON Lines), where each line contains a complete JSON object representing a training example. For instruction-following tasks, the recommended structure includes “instruction”, “input” (optional), and “output” fields. This format facilitates easy loading and processing during training. There are three base structures to use for model fine-tuning:

- Chat format, used for: instruction-tuned or conversational models (like LLaMA-2-Chat, Mistral-Instruct, Gemma-Instruct, or GPT-style chat models).

```
{ "messages": [
  { "role": "system", "content": "You are a friendly tutor." },
  { "role": "user", "content": "Explain photosynthesis simply." },
  { "role": "assistant", "content": "Photosynthesis is how plants make food using sunlight." }
]}
```

- Text format, intended for imitating or internalizing a particular style, vocabulary, and linguistic distribution.

{"text": "Manualul a fost aprobat prin ordinul ministrului Educației al Republicii Moldova nr. 457 din 31 mai 2013."}

{"text": "Lucrarea este elaborată conform curriculumului disciplinar și finanțată din resursele financiare ale Ministerului Educației."}

- Prompt–Response (pair) format, used for: instruction or QA-style fine-tuning when the model is not explicitly chat-based but supports conditioning on prompts.

{"prompt": "What is the capital of France?", "response": "Paris."}

{"prompt": "Name three states of matter.", "response": "Solid, liquid, and gas."}

In our experiment we selected the Text format as it was the only one way to teach the model to learn vocabulary, morphology, and syntax specific to the Moldavian language, absorb stylistic norms (e.g., formal register, phrasing, punctuation); and adapt to domain-specific patterns (names of institutions, document conventions, etc.). The training process utilizes Hugging Face’s Trainer API with carefully configured training arguments. The whole workflow process is presented in Figure 1.

This configuration demonstrates a balanced LoRA approach for parameter-efficient fine-tuning:

```
lora_config = LoraConfig(
    r=16, # Moderate rank for balanced capacity
    lora_alpha=32, # 2:1 alpha-to-rank ratio
    target_modules=['q_proj', 'v_proj'], # Focus on attention mechanisms
    lora_dropout=0.05, # Conservative regularization
    bias="none", # Exclude bias adaptation
    task_type="CAUSAL_LM", # Specify task type
)
```

Rank 16 balances expressiveness with computational efficiency; higher ranks increase memory and training time proportionally. The 2:1 scaling factor ($\alpha = 32$) empirically stabilizes gradient updates. Targeting only query and value projections captures critical attention transformations while maintaining efficiency. The 5% dropout provides regularization without restricting learning dynamics.




Figure 1: LoRA Fine-tuning Pipeline:
Step-by-step implementation flow with key parameters.

The adapter targets only the query and value projection matrices (q_proj , v_proj), the core elements of the transformer attention mechanism. This selective adaptation preserves model efficiency while capturing the most critical task-dependent transformations. A modest dropout rate of 5% provides light regularization without restricting learning dynamics.

Quantization Strategy

The quantization setup employs 4-bit precision to achieve high memory efficiency while maintaining numerical stability:

```
bnb_config = BitsAndBytesConfig(
    load_in_4bit = True,
    bnb_4bit_use_double_quant = True,
    bnb_4bit_quant_type = "nf4",
    bnb_4bit_compute_dtype = "bfloat16",
)
```

The NF4 quantization scheme – introduced in the QLoRA framework – optimally quantizes normally distributed weights and substantially outperforms conventional 4-bit methods. Double quantization further reduces storage overhead by re-quantizing quantization constants, enabling additional compression with negligible performance degradation. Using bfloat16 for computation ensures

numerical robustness during forward and backward passes, preventing the instability often observed in aggressive quantization setups.

Model Integration Pipeline

The integration of LoRA with quantized models is achieved through the following sequence:

```
model = AutoModelForCausalLM.from_pretrained(
    model_id,
    quantization_config = bnb_config,
    device_map = "auto"
).to(device) model = prepare_model_for_kbit_training(model)
model = get_peft_model(model, lora_config)
```

Automatic device mapping distributes model components across available GPUs, enabling fine-tuning of large architectures even in constrained environments. The `prepare_model_for_kbit_training()` routine ensures gradient compatibility with quantized parameters, while `get_peft_model()` seamlessly attaches LoRA adapters to the quantized backbone.

Tokenizer and Special Token Management

To ensure batch compatibility, the tokenizer is extended with a padding token when necessary:

```
if tokenizer.pad_token is None:
    tokenizer.add_special_tokens({'pad_token': '[PAD]'})
    model.resize_token_embeddings(len(tokenizer))
```

Many pretrained causal language models lack a dedicated padding token. This adjustment allows for efficient batch processing and prevents alignment errors during sequence generation.

Data Preprocessing and Label Alignment

The preprocessing function adheres to the conventions of causal language modelling:

```
def preprocess_function(examples):
    tokenized = tokenizer(
        examples["text"],
        truncation=True,
        max_length=512,
```

```
padding="max length"
)
tokenized["labels"] = tokenized["input_ids"]
```

A fixed sequence length of 512 tokens balances contextual coverage and memory efficiency. Label assignment mirrors the input sequence (labels = input_ids), aligning with autoregressive training objectives where the model predicts the next token in the sequence.

Training Configuration

The fine-tuning process is optimized for limited GPU resources through conservative yet effective hyperparameters:

```
TrainingArguments(
    num_train_epochs=3,
    per_device_train_batch_size=4,
    gradient_accumulation_steps=8,
    learning_rate=2e-4,
    fp16=True,
)
```

Gradient accumulation simulates a larger effective batch size (32 samples) without increasing instantaneous memory usage. The learning rate of 2×10^{-4} is standard for LoRA-based adaptation, offering faster convergence given the smaller number of trainable parameters. Mixed-precision computation (fp16=True) further accelerates training while maintaining model fidelity.

Performance and Resource Efficiency

The described configuration integrates several complementary techniques that collectively enable high-quality fine-tuning on consumer-grade hardware:

- Memory reduction: 4-bit quantization decreases base model memory usage by approximately 75
- Parameter efficiency: LoRA adapters introduce less than 1% additional trainable parameters.
- Computation optimization: Mixed precision and gradient accumulation significantly improve throughput.
- Quality retention: NF4 quantization and selective adapter placement preserve near full-precision performance.

This configuration thus exemplifies a state-of-the-art approach to parameter-efficient fine-tuning (PEFT), making it possible to adapt multi-billion-parameter

language models without sacrificing linguistic quality or stability.

4 Distinguishing Simple and Reasoning LLMs

The effectiveness of a virtual tutor heavily depends on the underlying LLM's ability to reason. Not all LLMs are created equal in this regard (Chu-Carroll *et al.*, 2024). We can distinguish between simple LLMs, which excel at language tasks, and reasoning LLMs, which are designed for complex problem-solving.

Simple LLMs, such as earlier versions of Llama, are primarily designed for generating fluent and coherent text. They are excellent for tasks like summarization, translation, and creative writing. However, they may struggle with multi-step reasoning, mathematical problems, and logical puzzles.

Reasoning LLMs, on the other hand, are specifically trained to tackle complex cognitive tasks. Models like **DeepSeek-R1**, **Llama 3.1 / 3.2**, and **Aya** are representative of this new wave of AI (DeepSeek-AI *et al.*, 2025; Meta AI, 2024; Cohere for AI, 2024). They are trained on vast datasets of code, scientific papers, and mathematical problems, and they often employ techniques like reinforcement learning to enhance their reasoning capabilities.

DeepSeek-R1 demonstrates notable reasoning capabilities through large-scale reinforcement learning, with behaviours including self-verification, reflection, and long chain-of-thought generation (DeepSeek-AI *et al.*, 2025; Marjanović *et al.*, 2025). For our Professor Guguță implementation, we utilize *DeepSeek-R1:8b* as the base model (FROM deepseek-r1:8bin the Modelfile).

Llama 3.1 and 3.2 represent advances in context length and multimodal capabilities (useful for curriculum-aligned tutoring) (Meta AI, 2024).

Aya models, developed by Cohere for AI, emphasize multilingual coverage and cultural awareness, which is relevant for culturally sensitive educational applications (Cohere for AI, 2024).

5 The Cultural Significance of Guguță

The choice of Guguță as our virtual tutor persona is deeply rooted in Moldavian cultural heritage. Guguță, created by renowned Moldavian writer Spiridon Vangheli in 1967 (first appearing in *Isprăvile lui Guguță*), has become an iconic figure in children's literature, representing innocence, curiosity, and the spirit of exploration (Holban, 2013; Ciobanu, 2013).

Transforming this beloved child character into Professor Guguță – an aged mathematician– creates a bridge between childhood wonder and academic rigor. This approach leverages the emotional connection that Moldavian students have with the character while providing the authority and wisdom expected from an educational mentor.

6 Application Architecture and Implementation

The Professor Guguță virtual tutor is implemented through a Python-based application that interacts with the Ollama API. The system architecture supports multiple document formats, enabling students to upload various types of educational materials for analysis and assistance.

Document Processing Capabilities. The application utilizes specialized Python libraries for comprehensive document processing:

- PDF Processing: PyMuPDF and pdfplumber for extracting text and mathematical content from PDF documents
- DOCX Processing: python-docx for Microsoft Word document analysis
- Excel Processing: pandas and openpyxl for spreadsheet data extraction and mathematical problem analysis
- TXT Processing: Standard Python file handling for plain text educational materials

Ollama API Integration. The core system architecture implements the following workflow:

1. Document Upload and Parsing: A Flask-based web interface allows students to upload educational materials
2. Content Analysis: Extracted text is processed and contextualized within the Moldavian educational framework
3. Ollama API Interaction: The processed content is sent to the DeepSeek-R1 model running through Ollama
4. Culturally-Aware Response Generation: Professor Guguță provides responses that incorporate Moldavian cultural elements and pedagogical best practices
5. Interactive Learning: The system maintains conversation context to provide ongoing educational support.

System Prompt Engineering Analysis. Our Modelfile implementation demonstrates sophisticated prompt engineering through several key components:

Temperature and Sampling Parameters: The configuration uses `PARAMETER temperature 0.7` and `PARAMETER top_p 0.9`, providing a balance between creativity and consistency in responses, crucial for educational applications where accuracy must be maintained while allowing for engaging explanations.

Reasoning Process Framework: The system prompt includes explicit reasoning guidelines: "You practice deliberate, critical reasoning before answering. You do not rush. You analyse the task, weigh the options, and select the most appropriate solution." This ensures that mathematical explanations follow a logical progression.

Cultural Contextualization: The prompt integrates cultural elements through specific instructions: "All your interactions, regardless of the language you are speaking, are viewed through the prism of Moldavian wisdom. Use proverbs, idioms, and cultural expressions from the region to make your points memorable."

7 Results and Discussion

The experimental results indicate that prompt-based fine-tuning proves particularly effective for personality modelling and cultural adaptation, whereas full fine-tuning of transformer architectures yields superior outcomes in tasks requiring mathematical reasoning and language comprehension.

Contrary to our initial hypothesis, training language models on Moldavian national mathematics textbooks did not enhance their linguistic proficiency in the Moldavian language. Instead, this approach led to substantial semantic drift and hallucination, where models began generating incoherent or illogical responses (see Figure. 2).




Figure 2: Model Hallucinations

These findings suggest that the text-only format is inadequate for this category of tasks. A more structured fine-tuning strategy should be explored – one that incorporates prompt-response pairs or instructional datasets designed to anchor reasoning and linguistic coherence. Future experiments may also consider reconstructing the dataset entirely to improve task alignment.

Despite these challenges, the prompt fine-tuning approach successfully shaped the persona of Professor Guguță, combining both cultural fidelity and computational precision. This method demonstrates several key advantages:

1. **Cultural Authenticity:** Prompt engineering successfully maintains Moldavian cultural elements and the Professor Guguță persona.
2. **Mathematical Accuracy:** The DeepSeek-R1 base model provides robust reasoning capabilities for mathematical problem-solving.

3. Linguistic Flexibility: The system effectively handles multilingual interactions while maintaining cultural context.

Educational Impact and Cultural Preservation. The AI Guguță Project represents a significant step toward culturally-aware educational technology. By incorporating traditional Moldavian proverbs and cultural references, the system not only provides mathematical education but also serves as a vehicle for cultural preservation and transmission to younger generations.

8 Conclusion and Future Work

The development of personalized virtual tutors through reasoning LLMs, as demonstrated by the AI Guguță Project, represents a significant advancement in educational technology. Our research confirms that while prompt fine-tuning is highly effective for personality development and cultural integration, transformer fine-tuning is more complex and requires specific data for training. We found out that Math books are not suitable for language training using text format of the JSONL file. Thus, the whole concept of Math book usage for this aim should be reconsidered.

The Professor Guguță implementation successfully bridges the gap between advanced AI technology and cultural authenticity, creating an educational tool that is both pedagogically effective and culturally meaningful. Future work will focus on expanding the system's capabilities to include visual mathematics problems, implementing adaptive learning algorithms, and extending the cultural corpus to include additional regional dialects and traditions.

The AI Guguță Project demonstrates that the future of personalized education lies not just in technological advancement, but in the thoughtful integration of cultural heritage with cutting-edge AI capabilities, creating virtual tutors that are not only intelligent but also culturally wise and empathetic.

AI use statement

The authors used ChatGPT solely for language editing purposes, including improving sentence clarity, ensuring grammatical consistency, and generating a draft version of the abstract after completing the manuscript. No part of the original research, data generation, or results analysis involved the use of AI. All content generated with AI assistance was thoroughly reviewed, edited, and approved by the authors, who takes full responsibility for the final version of the text.

Acknowledgements

This research is part of the AI Guguță Project (Code number: 25.80012.0807.21TC), which seeks to create the first interactive educational

platform in the Republic of Moldova based on artificial intelligence, designed for primary school students.

References

- Cheng Li, Mengzhou Chen, Jindong Wang, Sunayana Sitaram, and Xing Xie. 2024. Culturellm: Incorporating cultural differences into large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1–15.
- Cohere for AI. 2024. Aya 23: New open-source multilingual language models. Technical report, Cohere for AI.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z.F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J.L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R.J. Chen, R.L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, and S.S. Li *et al.* 2025. Deepseek-rl: Incentivizing reasoning capability in LLMs via reinforcement learning. *ArXiv*, abs/2501.12948:1–22.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, *et al.* 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):1–20.
- Ioan Holban. 2013. Proza lui Spiridon Vangheli. *Philologia*, 270(5-6):50–54.
- Jennifer Chu-Carroll, Andrew Beck, Greg Burnham, David O. Melville, David Nachman, A. Erdem Ozcan, and David Ferrucci. 2024. Beyond llms: Advancing the landscape of complex reasoning. *ArXiv*, abs/2402.08064:1–17.
- Jiho Shin, Clark Tang, Tahmineh Mohati, Maleknaz Nayeibi, Song Wang, and Hadi Hemmati. 2025. Prompt engineering or fine-tuning: An empirical assessment of llms for code. In *2025 IEEE/ACM 22nd International Conference on Mining Software Repositories (MSR)*, pages 490–502. IEEE.
- Meta AI. 2024. Introducing llama 3.1: Our most capable models to date. <https://ai.meta.com/blog/meta-llama-3-1/>. Meta AI Blog, July 2024.
- Mohamed Amine Ferrag, Norbert Tihanyi, and Merouane Debbah. 2025. From llm reasoning to autonomous ai agents: A comprehensive review. *ArXiv*, abs/2504.19678:1–44.
- Olesea Ciobanu. 2013. Spiridon Vangheli: paradisul pierdut și regăsit al copilăriei. *Philologia*, 270(5-6):59–62.
- Piotr Nawrot. 2023. nanoT5: A PyTorch framework for pre-training and fine-tuning T5-style models with limited resources. *ArXiv*, abs/2309.02373:1–7.

- Sara Vera Marjanović, Arkil Patel, Vaibhav Adlakha, Milad Aghajohari, Parishad BehnamGhader, Mehar Bhatia, Aditi Khandelwal, Austin Kraft, Benno Krojer, Xing Han Lú, Nicholas Meade, Dongchan Shin, Amirhossein Kazemnejad, Gaurav Kamath, Marius Mosbach, Karolina Stańczak, and Siva Reddy. 2025. Deepseek-r1 thoughtology: Let’s think about llm reasoning. *ArXiv*, abs/2504.07128:1–144.
- Siddhesh Pawar, Junyeong Park, Jiho Jin, Arnav Arora, Junho Myung, Srishti Yadav, Faiz Ghifari Haznitrama, Inhwa Song, Alice Oh, and Isabelle Augenstein. 2025. Survey of cultural awareness in language models. *Computational Linguistics*, 51(3):907–945.

CHAPTER 6

CULTURAL HERITAGE

QUANTUM-ASSISTED PARSING OF HISTORICAL ROMANIAN TEXTS

VASILE VANCEA

University of Bucharest
vasile.vancea@drd.unibuc.ro

Abstract

Parsing historical variants of Romanian presents significant challenges due to archaic orthography, morphological irregularities, and lexical changes over time. This paper proposes a hybrid quantum–classical natural language processing framework for the morphological disambiguation and syntactic parsing of historical Romanian texts. Our approach integrates parameterized quantum circuits within a transformer-based architecture to improve contextual understanding in low-resource and orthographically inconsistent settings. We evaluate the system on a newly digitized corpus of Romanian texts from the 17th–19th centuries, aligned with modern translations for evaluation purposes. Results indicate that the quantum-enhanced model achieves higher accuracy in identifying inflectional patterns and syntactic dependencies compared to baseline classical models, particularly for rare morphological forms. Beyond accuracy improvements, the study highlights the potential of quantum NLP for linguistic resources creation in underrepresented and historically rich languages. The proposed method contributes both a novel computational model and an annotated historical Romanian dataset, which will be released as an open resource for further research.

Keywords: Historical Romanian corpus, Quantum natural language processing, Variational quantum circuits

1 Introduction

The automatic processing of historical language materials represents a particularly demanding subfield within natural language processing (NLP). Historical corpora are frequently characterised by spelling irregularities, obsolete grammatical structures, and vocabulary items that have either shifted in meaning or fallen entirely out of use (Piotrowski, 2012). For Romanian, these challenges are especially pronounced in documents dating from the 17th to the 19th centuries, when orthographic conventions were still in flux and heavily influenced by Church Slavonic, Greek, and Latin sources (Gheție, 1997). General overviews of Romanian linguistic studies, such as (Schulte, 2019), further highlight the importance of documenting both historical and modern developments in the

language. The transitional nature of the language during this period makes morphological analysis difficult: inflectional paradigms often differ from modern standards, and syntactic constructions may follow patterns no longer productive in contemporary Romanian (Sala, 2006). Standard reference grammars of Romanian (Dindelegan și Maiden, 2013) provide detailed accounts of these inflectional and syntactic complexities, underscoring the challenges for NLP systems.

These complexities reduce the effectiveness of classical NLP tools, which are typically trained on modern, standardised corpora (Jurafsky and Martin, 2023). For example, statistical part-of-speech taggers or neural dependency parsers may misinterpret archaic inflectional endings or non-standard word orders, leading to errors in both morphological tagging and syntactic parsing (Petrov et al, 2012). The adaptation of such tools to historical Romanian therefore requires approaches capable of handling sparse data, irregular orthography, and subtle morphosyntactic cues.

Recent advances in quantum computing have opened new possibilities for representing and processing linguistic data. Quantum natural language processing (QNLP) applies quantum mechanical principles — such as superposition, entanglement, and high-dimensional Hilbert space embeddings — to model linguistic structures in ways that classical methods cannot easily replicate (Wiebe *et al.*, 2019). Unlike conventional word embeddings, which are restricted to fixed vector spaces, quantum encodings can represent multiple linguistic states simultaneously, potentially enabling richer disambiguation of context and meaning (Meichanetzidis *et al.*, 2020). These properties make quantum-enhanced models a promising candidate for the parsing of historical texts, where contextual interpretation is essential and the available training data is limited.

Although QNLP has been successfully applied to tasks such as sentiment analysis (Phukan *et al.*, 2024), semantic similarity computation (Baronia, 2023), and modern language syntactic parsing (Pandey *et al.*, 2024), its integration with historical linguistics and digital humanities remains largely unexplored. To our knowledge, no prior work has addressed the quantum-assisted morphological disambiguation and syntactic analysis of historical Romanian texts. This represents a missed opportunity: the combination of quantum state representations with transformer-based parsing architectures could address precisely the kind of ambiguity and irregularity that frustrates classical models.

In this paper, we present a hybrid quantum–classical NLP architecture for the parsing of historical Romanian texts. The model embeds tokens via parameterised quantum circuits (PQCs) (Benedetti *et al.*, 2019), whose outputs are integrated into a transformer-based morphological and syntactic analysis pipeline. By introducing quantum-enhanced embeddings, we aim to improve the recognition of archaic inflectional forms, to capture syntactic dependencies in the presence of

non-standard word order, and to better disambiguate morphologically rich or irregular lexical items.

The main contributions of this work are threefold: first, we introduce a novel hybrid quantum–classical parsing framework tailored for morphologically rich and historically variant Romanian; second, we present a curated and linguistically annotated corpus of 17th–19th century Romanian texts, normalized and aligned with modern equivalents; and third, we provide an empirical evaluation showing that quantum-enhanced embeddings outperform strong classical baselines, particularly in low-resource settings with high orthographic variability. By advancing the computational analysis of historical Romanian, we aim to contribute both to the technical field of QNLP and to the broader objectives of digital preservation and interpretation within the Romanian linguistic heritage.

The paper is structured as follows: the Introduction presents the background, motivation, problem statement, and main contributions. The Related Work section reviews prior research, identifies gaps, and positions the study within existing literature. The Methodology outlines the data, preprocessing steps, system design, and theoretical framework. The Experiments and Results section details the datasets, evaluation metrics, and both quantitative and qualitative comparisons with baseline models. The Discussion provides interpretation of results, error analysis, and highlights strengths and limitations. Finally, the Conclusion and Future Work summarizes the key findings, their implications, and suggests directions for further research.

2 Related Work

2.1 Historical Romanian Corpora and Linguistic Studies

The computational study of Romanian has traditionally focused on modern standard varieties, with considerably less work addressing historical forms. The most comprehensive linguistic treatments of the early literary Romanian period, such as (Gheție, 1997) and (Sala, 2006), document orthographic conventions, morphological paradigms, and syntactic tendencies across centuries. However, these resources were not designed as machine-readable linguistic resources.

Over the past decade, several digitisation initiatives have emerged, notably those undertaken by the Romanian Academy and the ELRC (European Language Resource Coordination) programme. These efforts include partial transcriptions of 17th–19th century Romanian texts, some accompanied by modern equivalents or glossaries. Nevertheless, the available data remains fragmented, with limited morphological annotation and no large-scale syntactically parsed resources. The absence of normalised historical corpora, particularly for earlier orthographic systems, makes computational modelling especially challenging (Piotrowski, 2012).

2.2 Classical NLP Methods for Old Text Parsing

NLP for historical texts has developed primarily in the context of languages with abundant historical records, such as English, German, and French (Piotrowski, 2012). Classical approaches have ranged from rule-based morphological analysers tailored to specific orthographic variants to statistical part-of-speech (POS) taggers adapted via domain adaptation techniques. Widely used NLP toolkits such as Stanford CoreNLP (Manning *et al.*, 2014) have been employed for morphological tagging and dependency parsing, though their performance drops significantly when applied to archaic Romanian texts. The compositional distributional model of meaning (Clark *et al.*, 2008) laid the theoretical foundation for such frameworks, showing how quantum-inspired algebraic structures can model semantic composition.

Deep learning approaches, especially transformer-based models such as BERT, have demonstrated state-of-the-art performance across a wide range of NLP tasks (Devlin *et al.*, 2019). More recent neural approaches, including recurrent neural networks (RNNs) and transformer-based models, have been applied to historical data by fine-tuning on manually annotated samples (Jurafsky and Martin, 2023).

However, even with domain-specific fine-tuning, classical models often underperform on historical Romanian due to morphological richness and orthographic instability. For instance, POS tagging accuracy drops significantly when trained solely on modern corpora and tested on 18th century Romanian, with errors concentrated in ambiguous inflectional forms and non-standard syntactic structures (Jurafsky and Martin, 2023, Petrov *et al.*, 2012). Dependency parsing suffers similar limitations, as the syntactic patterns of early Romanian texts reflect a mixture of Latin-based constructions and Slavic influences (Gheție, 1998, Sala, 2006). Efforts to address these issues include orthographic normalisation pipelines that map historical spellings to their modern equivalents, as well as morphological lexicons capturing archaic forms. While these methods can improve parser performance, they risk losing information about original linguistic forms, which are of primary interest to historical linguists (Piotrowski, 2012).

2.3 Prior Work in Quantum NLP

QNLP is a relatively new field that combines principles from quantum computing with formal and distributional models of meaning (Schuld *et al.*, 2015, Coecke *et al.*, 2020, Meichanetzidis *et al.*, 2020). Early theoretical work by (Coecke *et al.* 2020) established a categorical compositional framework for QNLP, representing grammatical structure via compact closed categories and encoding semantic content as quantum states. Subsequent studies implemented these models on near-term quantum hardware, using quantum circuits to encode sentence meaning and perform classification tasks (Meichanetzidis *et al.*, 2020).

Hybrid quantum–classical architectures have proven especially promising, as they allow quantum embeddings to be integrated into established neural architectures. (Phukan *et al.* 2024) demonstrated sentiment classification improvements with hybrid quantum neural networks, while Baronia (2023) applied QNLP to multi-class text classification. (Pandey *et al.* 2023) extended these approaches to syntactic parsing by representing grammatical relations as quantum-encoded features.

Several open-source frameworks have been developed to support experimentation in quantum natural language processing (QNLP). Qiskit, IBM’s quantum computing SDK, provides modules for building variational quantum circuits and implementing hybrid quantum-classical workflows. PennyLane offers a platform for differentiable quantum programming, allowing seamless integration of quantum circuits with machine learning frameworks such as PyTorch and TensorFlow. Additionally, lambeq is a specialized library designed for compositional-distributional QNLP, enabling the conversion of syntactic derivations into quantum circuits. Together, these tools form the foundation for advancing research and practical applications in QNLP. While these frameworks have been tested on English and a few other modern languages, no published work to date applies them to historical Romanian or other morphologically rich, low-resource historical languages. This represents a significant research gap, as quantum embeddings may be well-suited to capturing the nuanced morphological and syntactic variation present in such corpora.

3 Corpus Creation and Preprocessing

The creation of a linguistically annotated corpus of historical Romanian required a multi-stage pipeline, beginning with the acquisition of suitable source texts and continuing through optical character recognition (OCR), orthographic normalisation, morphological annotation, and statistical analysis of the resulting dataset. Given the scarcity and inconsistency of existing digital resources, we adopted a hybrid approach that combined digitisation of physical sources with the curation of existing electronic archives.

3.1 Sources of Historical Romanian Texts

The corpus compiled for this study integrates three main categories of sources. Manuscripts from the 17th–18th centuries, obtained from the Romanian Academy Library and the Central University Library in Iași, include letters, religious writings, and administrative records, reflecting both formal and informal registers. Their orthography is highly variable, influenced by regional dialects and the absence of standardised spelling (Gheție and Mareș, 1998; Sala, 2006), with philological studies (Mihăilă, 1981) and historical overviews (Rosetti, 1986) providing essential context for these linguistic variations. Early printed works from the 18th–19th centuries, sourced from the Biblioteca Digitală a Bucureștilor

and the European Library portal, consist of literary texts, political pamphlets, and translations, showing greater orthographic consistency than manuscripts but still containing archaic morphological and lexical traits (Gheție, 1997). Finally, previously digitised historical corpora, drawn from the European Language Resource Coordination (ELRC)¹ repository and the CLARIN² infrastructure, were incorporated where available in machine-readable form, supplementing the manually digitised materials with additional transcriptions, albeit often lacking detailed morphological annotation.

3.2 OCR and Orthographic Normalisation Techniques

For printed texts, OCR was performed using Tesseract OCR with a custom-trained Romanian model (Smith, 2007). Given the significant differences between modern and historical orthography, the OCR model was retrained using manually annotated samples from 18th–19th century printed sources. Accuracy on printed materials exceeded 96% character-level recognition, while manuscripts achieved only around 87% due to handwriting variability. To address OCR noise and facilitate subsequent morphological processing, we implemented a two-tier orthographic normalisation pipeline:

Character-Level Mapping — Non-standard characters (e.g., obsolete diacritics and ligatures) were mapped to their closest modern equivalents. For example, the “ſ” (long s) was mapped to “s”, and “â” or “ı” variations were resolved according to historical usage patterns.

Lexical-Level Normalisation — An alignment-based approach was used to map historical word forms to their modern equivalents where appropriate, using a combination of Levenshtein distance and a manually curated lexicon of 5,000 archaic-to-modern form pairs (Herce and Pricop 2024). This process preserved original forms in a parallel tier of the corpus, allowing analyses both in normalised and original orthography.

3.3 Morphological Annotation Guidelines

Morphological annotation followed a modified version of the Universal Dependencies (UD) framework (Petrov *et al.*, 2012), tailored to the particularities of historical Romanian. The work built upon the existing UD_Romanian-RRT treebank conventions but extended them to cover linguistic phenomena absent from contemporary usage. This adaptation was therefore semi-aligned with UD initiatives, rather than a completely personal scheme. Several extended POS tags and morphological features were introduced to represent archaic and dialectal constructions:

¹ <https://elrc-share.eu>

² <https://www.clarin.eu>

PART_OBS – obsolete participial forms still attested in 17th–18th century manuscripts.

VOC_ARC – archaic vocative endings (e.g., *frate!* → *frăţio!*).

NUM_DUAL – dual number remnants in pronouns and nouns influenced by regional varieties.

AUX_HAB – auxiliary verbs used in habitual aspect (e.g., *a şede* → *şădea*).

Inflectional features such as case, gender, number, person, and tense followed UD standards but were expanded with these historical categories to preserve morphological fidelity. Ambiguous forms were annotated with multiple possible analyses, each assigned a confidence score derived from expert judgement and probabilistic tagging models. Annotation was performed semi-automatically using a hybrid morphological tagger: a fine-tuned BERT-based Romanian model for modern forms³ supplemented with rule-based matching for archaic patterns. Inter-annotator agreement on a 500-token sample achieved a Cohen’s κ of 0.91, indicating high consistency.

3.4 Dataset Statistics

The corpus compiled for this study draws upon three main categories of sources: The final corpus comprises 1.2 million tokens across 210 individual texts, spanning the years 1640 to 1899. Table 1 summarises the dataset composition:

Table 1: Dataset composition.

Source Type	Number of Texts	Tokens	Average Tokens per Text	Orthography Variant
Manuscripts	55	320,000	5,818	Variable, regional
Early Printed Works	95	620,000	6,526	Semi-standard
Digitised Corpora	60	260,000	4,333	Mixed

Across the corpus, nouns account for approximately 38% of tokens, verbs for 22%, and adjectives for 13%. Around 8% of the tokens are archaic or obsolete forms not attested in the contemporary standard language, underscoring the need for models capable of handling such variation.

From the total corpus of approximately 1.2 million tokens (210 texts) covering the 17th–19th centuries, we used an 80/10/10 split for training, validation, and testing. This corresponds to roughly 960,000 tokens for training, 120,000 tokens

³ <https://huggingface.co/dumitrescustefan/bert-base-romanian-ner>

for validation, and 120,000 tokens for testing. The splits were stratified by source type (manuscripts, early printed works, digitised corpora) to ensure balanced coverage of orthographic and temporal variation. Each subset preserved the original proportion of archaic versus modernised forms, ensuring that model evaluation reflects generalisation to unseen historical variants.

The historical Romanian corpus compiled for this study will be released publicly under an open research license via a dedicated GitHub repository. The dataset will include token-level morphological and syntactic annotations in the CoNLL-U format, ensuring interoperability with Universal Dependencies tools. For illustration, the following is a sample annotated sentence from the corpus.

text = Domnul a venit la biserică .

1	Domnul	domn	NOUN	Gender=Masc Number=Sing Case=Nom	2	nsubj	_	_
2	a	a	AUX	Mood=Ind Tense=Pres Person=3	3	aux	_	_
3	venit	veni	VERB	VerbForm=Part	0	root	_	_
4	la	la	ADP	_	5	case	_	_
5	biserică	biserică	NOUN	Gender=Fem Number=Sing	3	obl	_	_
6	.	.	PUNCT	_	3	punct	_	_

This example demonstrates the tokenisation, lemmatisation, part-of-speech tagging, and dependency relations provided in the released corpus.

4 Methodology

This section presents the proposed hybrid quantum–classical architecture for morphological parsing of historical Romanian texts. Our framework integrates a quantum encoding stage within a transformer-based pipeline, enabling richer feature representations that capture both local morphological patterns and long-range contextual dependencies. The model leverages quantum variational circuits for enhanced token embedding transformation, which are subsequently processed by classical transformer layers for contextual parsing.

4.1 Overview of Hybrid Quantum-Classical Parsing Architecture

The proposed architecture comprises four principal stages:

Tokenisation and Classical Embedding – Input sentences from historical Romanian texts are first tokenised using a sub-word segmentation model trained on the created corpus. Tokens are then mapped to dense embedding vectors using a standard embedding matrix, initialised either randomly or from a pre-trained Romanian language model. For basic preprocessing tasks such as tokenisation and corpus preparation, widely used NLP libraries like NLTK⁴ were employed alongside custom scripts,

⁴ <https://www.nltk.org/api/nltk.corpus.html>

included for historical context. We employed NLP frameworks such as spaCy (v3.7)⁵ and Hugging Face Transformers (v4.42), complemented by custom tokenisation scripts adapted to historical orthographic variation.

Quantum Encoding Layer – The classical embeddings are fed into a quantum feature map, which transforms them into a high-dimensional Hilbert space using parameterised quantum gates (Havlíček *et al.*, 2019).

Classical Transformer Layers – The quantum-transformed embeddings are processed by multi-head self-attention and feed-forward layers, facilitating contextual integration.

Morphological Tag Prediction – The output representations are fed into a classification layer for predicting morphological tags such as part of speech, gender, number, case, and tense as well as a separate lemmatisation head that outputs the canonical base form of each token.




Figure 1: Quantum-Classical NLP Workflow for Romanian.

This hybrid configuration is designed to preserve the robustness and scalability of transformer models while enhancing representational expressivity through quantum transformations (Phukan *et al.*, 2024). Figure 1 illustrates the overall workflow of the proposed hybrid quantum-classical NLP architecture, showing

⁵ <https://spacy.io/>

the integration of tokenisation, quantum encoding, transformer-based contextualisation, and morphological tag prediction within a unified pipeline.

4.2 Quantum Encoding of Token Embeddings

Let $x \in R^d$, be the token embedding produced by the transformer for a single token. To encode x into a quantum state on q qubits we apply three steps: (i) a learned linear projection to match the qubit dimension, (ii) an elementwise normalisation to angles, and (iii) an angle-based quantum feature map that applies single-qubit rotations around multiple axes followed by (optional) entangling gates.

Dimensionality reduction (projection)

We first map the d -dimensional embedding into a q -dimensional vector z using a learned affine projection

$$z = f(x) = Wx + b, \quad W \in R^{q \times d}, b \in R^q$$

so that the i -th coordinate z_i corresponds to qubit i .

Normalisation to angles

Because rotation gates take angular arguments, we convert each component z_i to an angle $\bar{z}_i \in [0, 2\pi)$. One robust choice is min-max scaling followed by a linear mapping to $[0, 2\pi)$:

$$\bar{z}_i = 2\pi \cdot \frac{z_i - \min_j z_j}{\max_j z_j - \min_j z_j + \epsilon}$$

where $\epsilon > 0$ is a small constant for numerical stability. Alternatively, tanh or sigmoid normalisation may be used depending on the projection statistics.

Quantum feature map (angle encoding)

We then prepare the product state of q qubits, initially $|0\rangle^{\otimes q}$, and apply parameterised rotations on each qubit. A compact and expressive two-axis angle encoding is

$$|\psi(x)\rangle = \bigotimes_{i=1}^q [R_z(\bar{z}_i)R_y(\bar{z}_i)]|0\rangle^{\otimes q}$$

where $R_y(\theta) = \exp(-i\theta Y/2)$ and $R_z(\theta) = \exp(-i\theta Z/2)$ are single-qubit rotations about the y – and z – axes respectively. Using both R_y and R_z places the encoded point at a specific latitude and longitude on the Bloch sphere for each

qubit, increasing representational richness compared to a single-axis rotation (Rath and Date, 2023).

Entanglement

To capture correlations between embedding components (tokens or sub-token features) we interleave an entangling layer U_{ent} (for example a ring of CNOTs) and further variational rotation layers. A typical variational encoding with L layers is

$$|\psi(x)\rangle = U_{var}(\bar{z})|0\rangle^{\otimes q}, \quad U_{var} = \prod_{l=1}^L (U_{ent} \otimes_{i=1}^q R_z(\alpha_{i,l})R_y(\beta_{i,l}))$$

where $\{\alpha_{i,l}, \beta_{i,l}\}$ may be functions of $\bar{z}$ (data-dependent encoding) or trainable variational parameters (hybrid data + parameter encoding). In our experiments we used a small number of variational layers (e.g. $L = 2$) and a ring CNOT entanglement topology to balance expressivity and trainability.

Measurement and classical read-out

After the quantum feature map, we measure selected observables (commonly the expectation values of Pauli-Z on each qubit) to produce a classical feature vector for downstream processing:

$$\phi(x) = (\langle Z_1 \rangle, \langle Z_2 \rangle, \dots, \langle Z_q \rangle), \quad \langle Z_i \rangle = \langle \psi(x) | Z_i | \psi(x) \rangle$$

The vector $\phi(x)$ is then concatenated (or fused) with transformer outputs and passed to the subsequent transformer encoder/decoder layers for contextual parsing.

Gradients and training

When using parameterised gates that depend on trainable parameters, gradients for quantum parameters can be estimated with the parameter-shift rule and backpropagated jointly with classical parameters (Mitarai *et al.*, 2018; Schuld *et al.*, 2020). When the data enters via angle encodings, the projection f remains trainable and is updated together with variational circuit parameters.




Figure 2: Quantum encoding circuit for embedding tokens into qubit states

After the quantum encoding and measurement, the resulting quantum embedding vector has a size equal to the number of measured observables, typically one per qubit if you measure the Pauli-Z expectation value on each. Formally:

$$\phi(x) = (\langle Z_1 \rangle, \langle Z_2 \rangle, \dots, \langle Z_q \rangle), \in R^q$$

Thus, the post-quantum embedding has dimension q , the number of qubits used. If you use 8 qubits, you get an 8-dimensional feature vector. To make it compatible with a transformer layer or classifier that expects a higher dimension (e.g., 128), this vector is projected using a learnable linear layer:

$$\bar{\phi}(x) = W_{out}\phi(x) + b_{out}, \quad W_{out} \in R^{128 \times q}$$

After the quantum encoding circuit, each qubit is measured to produce one classical expectation value, typically $\langle Z_i \rangle$ (Rath and Date, 2023, Iqbal *et al.*, 2024). The resulting embedding vector therefore has dimension q , corresponding to the number of qubits used in the circuit. In our configuration ($q = 8$), this yields an 8-dimensional quantum embedding. To integrate it into the transformer layers, a linear projection $W_{out} \in R^{128 \times 8}$ expands it to a 128-dimensional representation. This ensures compatibility with the downstream model while preserving quantum-enhanced feature diversity.

4.3 Parameterised Quantum Circuit Design

The quantum processing stage employs a variational quantum circuit (VQC) composed of:

- Initial Encoding Layer – Implements the chosen feature map (angle encoding).
- Entanglement Layer – A ring-pattern CNOT configuration connecting q qubits, promoting multi-token feature interaction.
- Rotation Layer – Parameterised R_x, R_y , and R_z gates applied to each qubit, with trainable parameters shared across layers to control model complexity.
- Measurement – The expectation values of Pauli-Z operators are taken from each qubit, producing the quantum-transformed embedding.

For our experiments, we use $q = 8$ qubits, a circuit depth of 4, and a ring entanglement topology to balance expressive power and trainability. The parameter count is limited to avoid barren plateaus (McClellan *et al.*, 2018). The variational structure of the PQC follows the design principles established in early demonstrations of quantum variational solvers (Peruzzo *et al.*, 2014).

4.4 Integration with Transformer Layers

The quantum-transformed embeddings are passed to a stack of two transformer encoder layers with 4 attention heads each. This stage integrates local morphological cues identified by the quantum encoder with broader syntactic and semantic context. Positional encoding is applied before transformer processing to maintain word order information. We selected a shallow transformer configuration (two encoder layers with four attention heads) as a compromise between representational power and computational feasibility. Deeper stacks showed no significant gains in preliminary trials, while increasing training time and risk of overfitting on the relatively limited historical Romanian corpus. This lightweight setup also aligns with prior low-resource parsing studies that favour compact transformer variants.

4.5 Training Procedure and Optimization

The hybrid model is trained end-to-end with backpropagation through both the quantum and classical components, using the parameter-shift rule (Mitarai *et al.*, 2018) for gradient computation in the quantum circuit. We employ the AdamW optimiser with an initial learning rate of 3×10^{-4} , applying a cosine annealing schedule. The loss function is a categorical cross-entropy over morphological tag predictions. Training is conducted on simulated quantum hardware (Qiskit Aer) for reproducibility, with subsequent fine-tuning on an available 8-qubit superconducting backend for validation of real-device performance (Preskill, 2018). Batch sizes are kept at 16 due to quantum simulation constraints, with gradient accumulation to simulate larger effective batch sizes. Early stopping is applied based on development set accuracy.

The training of the hybrid architecture proceeds in an end-to-end fashion, where both the quantum and classical components are optimised jointly. Gradients for the quantum circuit are estimated using parameter-shift rules, while the transformer layers rely on backpropagation. This unified training procedure ensures that the quantum embeddings evolve in alignment with the contextual representations extracted by the transformer. To summarise, the hybrid design integrates quantum feature extraction with state-of-the-art transformer modelling to address the challenges of morphological parsing in historical Romanian texts.

Our implementation was built using PyTorch for the classical transformer components and Qiskit/PennyLane for the quantum circuit simulation. The quantum embedding stage used 8 qubits with a circuit depth of 4 and a ring-entanglement pattern. Input embeddings (dimension 128) were reduced to 8-dimensional vectors prior to quantum encoding. The hybrid model was trained with AdamW (learning rate $3e-4$, cosine annealing, batch size 16 with gradient accumulation) for 25 epochs on a single NVIDIA A100 GPU and a simulated Aer backend. Preliminary tests on real IBM 8-qubit hardware confirmed stability but

were limited by queue times. Gradients for quantum layers were computed using the parameter-shift rule. The overall training time was approximately 12 hours.

5 Experiments and Results

5.1 Baseline Models

To evaluate the effectiveness of the proposed hybrid quantum–classical parsing architecture, we compared it against several baseline models commonly employed in morphological and syntactic analysis of historical texts. These included:

CRF-based morphological tagger trained with handcrafted features and contextual windows (Zampieri, 2016).

BiLSTM-CRF sequence labeller, which has been widely applied to morphologically rich languages (Huang *et al.*, 2015).

Transformer based parser (fine-tuned multilingual BERT), adapted for historical Romanian orthography through normalization (Devlin *et al.*, 2019).

The baseline models, including the CRF and Transformer architectures, were pre-trained on modern Romanian or general-domain data and only partially fine-tuned on the historical corpus due to resource limitations. In contrast, the proposed hybrid quantum–classical model was trained entirely on the historical dataset. Although this setup may give the quantum-enhanced model a domain-specific advantage, it provides a meaningful comparison by highlighting the capacity of quantum embeddings to adapt to orthographic variability and diachronic linguistic features that classical baselines struggle to capture. These baselines represent a spectrum of classical approaches ranging from statistical to deep learning methods, and thus provide a robust comparison framework. morphological tagging and a comprehensive assessment of parsing performance.

5.2 Evaluation Metrics

Performance in this study was evaluated using widely adopted metrics for morphological and syntactic parsing. Tagging accuracy measured the proportion of correctly assigned morphological categories such as part-of-speech, case, and number. The F1-score, calculated as the harmonic mean of precision and recall, provided a balanced measure for uneven class distributions. For syntactic parsing, the Labelled Attachment Score (LAS) captured the percentage of words correctly attached to their heads with the appropriate dependency label, while the Unlabelled Attachment Score (UAS) reflected correct syntactic head attachment regardless of label. Together, these metrics offered both fine-grained insights into morphological tagging and a comprehensive assessment of parsing performance.

5.3 Comparative Results

To assess the effectiveness of the proposed hybrid quantum–classical architecture, we conducted a direct comparison with several established baselines. The focus of this evaluation was twofold: first, to determine whether quantum-enhanced embeddings provide measurable benefits over classical representations; and second, to examine whether these improvements translate consistently across both tagging and parsing tasks.

Table 2: Comparative results on historical Romanian corpus.

Model	Accuracy (%)	F1-score	LAS (%)	UAS (%)
CRF-based Tagger	82.4	0.80	69.1	75.4
BiLSTM-CRF	86.7	0.85	73.5	79.2
Transformer (BERT fine-tuned)	89.5	0.88	77.9	83.6
Hybrid Quantum–Classical	91.2	0.90	80.6	85.7

Table 2 reports the comparative performance across the baseline models and the proposed hybrid quantum–classical model. The results show that the hybrid model consistently outperforms the baselines across all evaluation measures. The improvements are most pronounced in dependency parsing, where the integration of quantum embeddings yields stronger syntactic representations.

In addition to accuracy metrics, we compared the complexity of the proposed quantum embedding layer with the baseline models in terms of trainable parameters. The hybrid quantum–classical model employs a parameterised quantum circuit with 8 qubits, a circuit depth of 4, and 3 trainable rotation gates (R_x , R_y , R_z) per qubit per layer, yielding a theoretical total of 96 variational parameters in the quantum embedding component. Minor variations may occur depending on backend transpilation and parameter-sharing strategies, but this number provides an exact order of magnitude for comparison. By contrast, the BiLSTM-CRF baseline involves around 2.1 million trainable parameters, and the fine-tuned transformer (BERT) baseline exceeds 110 million parameters. This comparison highlights that, despite introducing a new quantum encoding stage, the hybrid model remains lightweight in terms of parameter count compared to deep neural baselines. However, the quantum operations incur additional runtime cost due to circuit execution, making efficiency dependent on hardware backends.

5.4 Error Analysis

A manual error analysis was conducted on a 500-sentence subset ($\approx 12,000$ tokens) of the historical corpus, selected to include a balanced mix of manuscripts and early printed texts. The analysis revealed several consistent patterns where the quantum-enhanced model outperformed classical baselines:

Orthographic variation: In cases such as “*Domnu au poruncit*” (archaic spelling) versus “*Domnul a poruncit*” (modern form), the quantum model more reliably normalized variant spellings to the same lemma and POS. The BiLSTM baseline misclassified 17% of such variants, whereas the hybrid model reduced these errors substantially.

Morphological ambiguity: For ambiguous tokens like “*fiului*” (which may indicate possession or an indirect object), the hybrid model reduced tagging errors from 18% to 8%, demonstrating improved disambiguation through richer contextual representations.

Hierarchical attachment (nested prepositional phrases): Consider the phrase “*scrisă cartea de mână preotului*”, where syntactic analysis requires recognising that “*preotului*” modifies “*mână*” (the priest’s hand) within the noun phrase rather than attaching directly to the verb. The transformer baseline frequently attached “*preotului*” incorrectly to the verb, while the hybrid model correctly attached “*preotului*” to “*mână*” in 86% of cases versus 71% for the transformer baseline.

These examples confirm that the quantum embedding layer better captures sub-word level similarity and syntactic relations over extended contexts, leading to measurable improvements in both tagging accuracy and dependency parsing

A closer inspection of the model outputs highlights several areas where the quantum model provides notable advantages. First, the quantum-enhanced embeddings showed greater robustness to orthographic variation, enabling more reliable normalization of inconsistent spelling patterns before parsing. Second, in cases of morphological ambiguity where the same surface form could represent multiple grammatical categories, the quantum model resolved these instances more accurately, indicating that the circuit effectively captured subtle distributional cues. Finally, the entanglement structure of the parameterized quantum circuit appeared to aid in modelling long-distance dependencies across tokens, leading to measurable improvements in labelled attachment scores compared to the classical transformer. Overall, these results support the hypothesis that hybrid quantum–classical models can enrich token representations, leading to measurable gains in parsing accuracy for historical Romanian texts. A qualitative inspection of the outputs highlights the main areas where the hybrid quantum–classical model surpasses its classical counterparts.

Nevertheless, certain challenges remain. The hybrid model showed reduced gains for rare inflectional forms that appeared infrequently in the training corpus. These included archaic or regionally marked endings such as “*-șeredz*” (archaic 3rd-person plural of verbs like “*a ști*”), “*-eră*” (past perfect tense forms, e.g., “*veniră*”), and “*-iloru*” (obsolete genitive–dative plural marker, e.g., “*fraților*”). Because such variants occurred in fewer than 0.05% of total tokens, the model’s quantum embeddings had limited opportunity to learn stable

representations. In these cases, performance was comparable to or only slightly higher than the transformer baseline.

Historical Romanian texts often exhibit high orthographic variation, with the same word appearing in multiple inconsistent spellings due to regional influences and the absence of standardized conventions. For instance, the lexical item “*domn*” (“lord”) may occur as “*domnu*” or “*domnŭ*”. While classical models frequently misclassify such variants because of limited training examples, the quantum embeddings capture subword-level structural similarities, enabling more consistent tagging.

Morphologically ambiguous endings such as “-ul” or “-ului” pose additional challenges, as these morphemes simultaneously encode both definiteness and grammatical case. Their interpretation therefore depends on syntactic context, since the same suffix marks the noun as definite while also indicating nominative/accusative (“-ul”) or genitive/dative (“-ului”) relations. The BiLSTM-CRF baseline misclassified 18% of such tokens, most often confusing genitive/dative forms (e.g., “*fiului*”) with nominative/accusative ones. In contrast, the hybrid quantum–classical model reduced this error rate to 9%, correctly identifying the case role based on contextual cues provided by the quantum-enhanced embeddings. The improved performance can be attributed to the quantum circuit’s entanglement structure, which encodes syntactic context more effectively and aids in disambiguation. Rare or archaic constructions further complicate parsing, with expressions such as “*pre învățătorul*” or “*după acea vreme*” being less frequent in modern corpora and difficult for transformer-based baselines. The hybrid model achieved stronger recall for these cases, likely because its quantum-enhanced embedding spaces provide richer separation between archaic and modern linguistic forms.

Finally, the hybrid model outperformed classical parsers in resolving syntactic attachment errors, particularly those involving prepositional or genitive constructions. For instance, in the sentence „*scrisă cartea de mână preotului*” (“the book written by the priest’s hand”), the baseline parser incorrectly attached „*preotului*” to „*cartea*”, while the hybrid quantum–classical model correctly attached it to „*mână*”. This demonstrates that the quantum-enhanced embeddings better capture hierarchical syntactic relations and genitive dependencies within nominal phrases. Overall, the error analysis indicates that quantum feature extraction provides measurable advantages in handling orthographic variation, morphological ambiguity, and complex syntactic attachment structures typical of historical Romanian.

6 Discussion

The experimental findings outlined in the previous section provide important insights into both the linguistic and computational dimensions of applying hybrid quantum–classical methods to the parsing of historical Romanian texts.

6.1 Linguistic insights from improved parsing of archaic forms

One of the most significant contributions of the hybrid model is its ability to handle morphological variation and orthographic inconsistency, which are defining features of archaic Romanian sources. Classical models trained predominantly on modern corpora often fail to capture irregular inflectional patterns (e.g., the use of *făcut-au* instead of *au făcut*) or obsolete case markers such as the enclitic genitive *-ăi* (e.g., *fiii* → *fiului*). The hybrid quantum–classical parser demonstrated an improved capacity to generalise across these non-standard forms, correctly identifying both their modern equivalents and grammatical roles.

This suggests that quantum-enhanced embeddings may capture a richer representational structure, allowing for more robust disambiguation of archaic variants. From a linguistic perspective, this opens the way for more systematic exploration of historical texts, revealing measurable patterns of language change. For instance, phonological shifts such as *preăt* → *preot* and *împaratu* → *împăratul* illustrate vowel simplification and the standardisation of definite articles. Syntactic evolutions can be seen in the gradual reduction of post-posed verbs (e.g., *făcut-au* → *au făcut*). Similarly, the disappearance of morphological categories like the dual number (*amândoi frații* → *cei doi frați*) reflects a broader simplification trend in the nominal system. The hybrid quantum–classical model, by effectively linking archaic and modern variants within a shared embedding space, facilitates the automatic identification of these diachronic correspondences.

6.2 Benefits and limitations of quantum approaches in low-resource settings

The hybrid approach offers a particular advantage in low-resource contexts, where the size of available training corpora is limited. Quantum feature spaces allow for the projection of data into higher-dimensional Hilbert spaces, where linearly inseparable linguistic patterns can become more tractable. This property proved particularly beneficial when parsing small samples of texts where archaic forms were sparsely represented. Nevertheless, limitations remain. Current quantum hardware introduces noise, limiting circuit depth and scalability (Karuppasamy *et al.*, 2025). Furthermore, simulation of quantum circuits on classical machines, as was required for these experiments, imposes significant computational costs. Hence, while the method shows promise, practical adoption awaits more accessible and error-corrected quantum devices.

6.3 Impact on digital humanities and Romanian linguistic studies

Beyond computational improvements, the results have implications for digital humanities scholarship. Improved morphological and syntactic analysis enables more accurate annotation of corpora, facilitating studies in philology, historical linguistics, and cultural heritage preservation. For Romanian studies in particular,

the approach contributes to documenting linguistic evolution across manuscripts, early printed texts, and dialectal variants. By offering tools to normalise and analyse such texts, quantum NLP could enhance accessibility for researchers in linguistics, history, and literature.

An important consideration for future work is the extent to which quantum embedding models can be shared and reused across tasks, in a manner analogous to pretrained transformer models such as BERT. While current implementations of quantum embeddings are tightly coupled to specific hardware configurations and qubit counts, the conceptual possibility exists for developing reusable quantum representations that capture generalisable linguistic features. Such pretrained quantum embedding layers, if standardised and openly shared, could reduce the cost of training and enable wider adoption of QNLP approaches, particularly in low-resource language scenarios such as historical Romanian. This raises exciting prospects for the emergence of quantum equivalents of widely distributed models like BERT, though practical realisation will depend on advances in both hardware scalability and software interoperability.

7 Conclusion and Future Work

This study introduced a hybrid quantum–classical framework for morphological and syntactic parsing of historical Romanian texts. By combining quantum embeddings derived from parameterised quantum circuits with transformer-based contextual layers, the proposed architecture demonstrated measurable improvements over classical baselines, particularly in handling archaic orthographic variants and morphologically complex constructions. These findings provide evidence that quantum feature spaces can complement neural architectures in tasks characterised by high ambiguity and limited training data.

This has important implications for the digital humanities, offering more accurate tools for the annotation and interpretation of historical texts, and ultimately supporting philological, lexicographic, and cultural studies. Nevertheless, the study also highlights limitations. The current implementation relies on simulated quantum backends, which impose restrictions on scalability and depth. Furthermore, the hybrid approach still depends heavily on classical components for contextual representation, meaning that its quantum advantage is currently constrained by hardware availability.

Future work will proceed along several directions. First, we aim to expand the corpus to include a wider variety of Romanian dialects, ranging from regional manuscripts to early printed works, enabling broader coverage of the linguistic evolution of the language. Second, further experimentation will explore alternative quantum encoding schemes and entanglement topologies, optimising depth and parameter efficiency. Finally, as quantum hardware matures, prospects open for developing fully quantum NLP pipelines, where end-to-end models may

leverage native quantum parallelism to tackle low-resource languages and historical corpora at scale.

In conclusion, while quantum approaches to NLP remain in their infancy, this work provides one of the first demonstrations of their promise in the context of Romanian linguistic studies. The intersection of quantum computing, language technology, and digital humanities is set to become a fertile research area, where theoretical advances align with cultural and historical inquiry.

AI Use Statement

Artificial intelligence tools (OpenAI ChatGPT, version GPT-5) were used exclusively for language refinement, grammatical correction, and improving textual clarity. Limited AI assistance was also employed for figure layout drafting. No AI tools were used to generate original research content, perform data analysis, or influence the study's results and conclusions.

References

- Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O'Brien. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5:4213.
- Alexandru Rosetti. 1986. *Istoria limbii române*. Editura Științifică și Enciclopedică, București.
- Arpan Phukan, Santanu Pal, and Asif Ekbal. 2024. Hybrid Quantum-Classical Neural Network for Multimodal Multitask Sarcasm, Emotion, and Sentiment Analysis. *IEEE Transactions on Computational Social Systems*, 11(5):5740–5750.
- Bob Coecke, Giovanni de Felice, Konstantinos Meichanetzidis, and Alexis Toumi. 2020. Foundations for Near-Term Quantum Natural Language Processing. *arXiv preprint arXiv:2012.03755*.
- Borja Herce and Bogdan Pricop. 2024. VeLeRo: An inflected verbal lexicon of standard Romanian and a quantitative analysis of morphological predictability. *Language Resources and Evaluation*, 59(1):621–637. DOI: 10.1007/s10579-024-09721-3.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *ACL System Demonstrations*, pages 55–60.
- CLARIN-EU. 2023. CLARIN European Language Resources Repository. Available at: <https://www.clarin.eu/>
- Dan Jurafsky and James H. Martin. 2023. *Speech and Language Processing* (3rd ed., draft). Stanford University.
- Dhruv Baronia. 2023. Hybrid quantum-classical neural networks for text classification. *Authorea Preprints*.
- Gabriela Pană Dindelegan and Martin Maiden, eds. 2013. *The Grammar of Romanian*. Oxford University Press.
- Gheorghe Mihăilă. 1981. *Studii de lingvistică și filologie*. Timișoara: Editura Facla.
- Ion Gheție. 1997. *Istoria limbii române literare: Epoca veche (1532–1780)*. Editura Academiei Române.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Jarrold R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. 2018. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812.
- John Preskill. 2018. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79.
- Kim Schulte. 2019. Romanian Studies: Language and Linguistics. *The Year's Work in Modern Language Studies*, 81:483–502.
- Konstantinos Meichanetzidis, Stefano Gogioso, Giovanni De Felice, Nicolò Chiappori, Alexis Toumi, and Bob Coecke. 2020. Quantum natural language processing on near-term quantum computers. *arXiv preprint arXiv:2005.04147*.
- Kosuke, Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. 2018. Quantum circuit learning. *Physical Review A*, 98(3):032309.
- Krishnageetha Karuppasamy, Varun Puram, Stevens Johnson, and Johnson P. Thomas. 2025. A comprehensive review of quantum circuit optimization: current trends and future directions. *Quantum Reports*, 7(1):2.
- Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini. 2019. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001.
- Marcos Zampieri, Shervin Malmasi, and Mark Dras. 2016. Modeling Language Change in Historical Corpora: The Case of Portuguese. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 4098–4104. Portorož, Slovenia: ELRA.
- Maria Schuld, Alex Bocharov, Krysta M. Svore, and Nathan Wiebe. 2020. Circuit-centric quantum classifiers. *Physical Review A*, 101(3):032308.
- Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. 2015. An introduction to quantum machine learning. *Contemporary Physics*, 56(2):172–185.
- Marius Sala. 2006. *Enciclopedia limbii române*. Editura Univers Enciclopedic.
- Michael Piotrowski. 2012. *Natural Language Processing for Historical Texts*. Synthesis Lectures on Human Language Technologies, 5(2):1–157. Springer.
- Minati Rath and Hema Date. 2023. Quantum data encoding: a comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy. *Quantum Physics (quant-ph); Artificial Intelligence (cs.AI)*, arXiv:2311.10375.
- Nathan Wiebe, Alex Bocharov, Paul Smolensky, Matthias Troyer, and Krysta M. Svore. 2019. Quantum language processing. *Computer Science, Physics, and Linguistics*.
- Ray Smith. 2007. An Overview of the Tesseract OCR Engine. In *Proceedings of ICDAR 2007*, pages 629–633.
- S.M. Yousuf Iqbal Tomal, Abdullah Al Shafin, Afrida Afaf, Debojit Bhattacharjee. 2024. Quantum convolutional neural network: a hybrid quantum–classical approach for iris dataset classification. *Quantum Physics (quant-ph); Artificial Intelligence (cs.AI)*, arXiv:2410.16344.
- Shyambabu Pandey, Nihar Jyoti Basisth, Tushar Sachan, Neha Kumari, and Partha Pakray. 2023. Quantum machine learning for natural language processing application. *Physica A: Statistical Mechanics and its Applications*, 627:129123.

- Shyambabu Pandey, Partha Pakray, and Riyanka Manna. 2024. Quantum classifier for natural language processing applications. *Computación y Sistemas*, 28(2):695–700. DOI: 10.13053/cys-28-2-5016.
- Slav Petrov, Dipanjan Das, and Ryan McDonald. 2012. A universal part-of-speech tagset. In *Proceedings of LREC 2012*, pages 2089–2096.
- Stephen Clark, Bob Coecke, and Mehrnoosh Sadrzadeh. 2008. A compositional distributional model of meaning. In *Proceedings of the Second Quantum Interaction Symposium*.
- Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. 2019. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567:209–212.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991* [cs.CL].

THE WISDOM EXTRACTOR: MINING CROSS-CULTURAL PROVERBS TO ELICIT TIME-TESTED HEURISTICS

VLAD BELCIUG¹, ELENA PELICAN²

¹*Independent Researcher, Constanța, Romania*
vladbelciug80@gmail.com

²*Faculty of Mathematics and Informatics, Ovidius University of Constanța*
epelican@univ-ovidius.ro

Abstract

We introduce the Wisdom Extractor, a reproducible pipeline for discovering cross-cultural regularities in proverbial knowledge without relying on any commercial APIs. The goal is to identify candidate “time-tested” heuristics that recur across peoples and languages and to present them with transparent provenance and auditable methods. The system (i) collects proverb lists from public sources while preserving mother-tongue originals and URLs; (ii) removes boilerplate using rule-based filters with an optional fully offline LLM check; (iii) canonicalizes frequent figurative templates into minimal propositions; (iv) clusters near-duplicates across peoples with character n-gram TF-IDF and average-linkage under a tunable distance threshold; (v) scores cross-cultural coverage and support; (vi) projects claims to 2D for inspection; and (vii) produces a deterministic interpretation that links themes to coarse ecological and institutional proxies. On a multilingual sample of 18,049 proverbs from 90+ cultures, the extractor recovers high-coverage motifs and surfaces productive tensions as hypotheses rather than verdicts. All steps are documented; originals, claims, and sources are displayed side-by-side; and a single Streamlit interface combines dataset building, extraction, visualisation, and interpretation. The contribution is methodological (transparent alignment of figurative forms through automated canonicalization with full provenance tracking) and practical (an offline-capable tool ready for classroom or research use).

Keywords: natural language processing; linguistic resources; proverb mining; cross-lingual; dataset; evaluation; cross-cultural alignment; figurative language; clustering; offline NLP; computational paremiology

1 Introduction

Language does more than label the world; it helps a community remember how to act. Among the many mnemonic devices of natural language, proverbs have been unusually successful over long horizons. They are brief, rhythmic, and

portable, with a good balance between memorability and ambiguity. Exactly because they are terse, proverbs typically encode procedural knowledge: how to navigate risk, when to cooperate, whom to trust. If we could align equivalent proverbs across peoples and languages, we might surface a shortlist of heuristics that were independently discovered or deliberately carried through trade and migration.

This paper presents the Wisdom Extractor - a methodological infrastructure that advances computational approaches to cross-cultural proverb analysis. Our contribution is primarily technical: we demonstrate how to systematically process figurative language across diverse writing systems, align semantically related proverbs through canonicalization, and identify recurrent patterns through character-based clustering. The system produces intriguing preliminary findings about cross-cultural regularities that merit further investigation by researchers equipped with domain expertise and validation resources. We position this work as foundational infrastructure that enables future research rather than as definitive claims about cultural universals.

The Romanian context provides a particularly rich testing ground for this approach. As a Romance language with significant Slavic influence, surrounded by diverse linguistic traditions (Hungarian, Bulgarian, Serbian, Ukrainian), Romanian proverbs offer a unique window into both inherited Latin wisdom and Balkan sprachbund effects. Our findings reveal that Romanian wisdom shows 89% overlap with Albanian in top clusters, while maintaining 76% alignment with Italian—suggesting both regional convergence and genealogical persistence.

2 Related Work

2.1 Computational Paremiology

Paremiology—the study of proverbs—has traditionally been a humanistic discipline focused on collection, classification, and interpretation (Mieder, 2021). Recent computational approaches have begun to scale these efforts. (Davis *et al.*, 2021) tracked proverb usage frequencies in large corpora over time, while (Wang *et al.*, 2025) evaluated machine translation quality on proverb sets. Our work differs in focusing on cross-cultural alignment: converting heterogeneous lists into minimal, comparable claims and quantifying coverage across peoples.

2.2 Figurative Language Processing

The computational treatment of metaphor and idiom has produced sophisticated methods for detection and interpretation (Shutova, 2010; Steen *et al.*, 2010). Recent neural approaches using contextualized embeddings show promise for metaphor identification (Choi *et al.*, 2021; Song *et al.*, 2021). However, these methods typically require large training corpora and struggle with low-resource languages. Our character n-gram approach, while simpler, works across all scripts without training data.

2.3 Cultural Evolution and Knowledge Transmission

The theoretical framework for our work draws on cultural evolution theory (Boyd *et al.*, 2011), which treats cultural variants (including proverbs) as units subject to selection, drift, and transmission biases. The rice theory of cultural psychology (Talhelm *et al.*, 2014) provides specific predictions about how agricultural systems shape cultural values, predictions we test through proverb analysis.

Recent additions. Beyond prior surveys and foundational work, we incorporate recent datasets and evaluations that are directly relevant to our goals: the ePiC benchmark for proverbs-in-context (Ghosh and Srivastava, 2022), the MAPS evaluation of multicultural sayings for multilingual LLMs (Liu *et al.*, 2024), the Greek Proverb Atlas (Pavlopoulos *et al.*, 2024), the MUNCH metaphor dataset (Tong *et al.*, 2024), and recent studies of figurative language competence in LLMs (Yakura, 2024; Ichien *et al.*, 2024).

2.4 Methodological Constraints and Design Choices

Our approach prioritizes reproducibility and accessibility over semantic sophistication. We deliberately choose character n-grams over contextual embeddings (e.g., multilingual BERT) to ensure offline operation and cross-script compatibility, accepting the trade-off of reduced semantic depth. Similarly, our canonicalization process, while enabling cross-cultural comparison, necessarily flattens figurative richness and may distort culturally embedded meanings. These limitations are acknowledged as the price of methodological transparency and reproducible research infrastructure.

```
{
  "people": "Romanian",
  "original": "Aschia nu sare departe
de trunchi", "saying": "Aschia nu
sare departe de trunchi",
  "english_equivalent": "The chip doesn't jump far from
the trunk", "source_title": "Wikiquote: Romanian
proverbs",
  "source_url":
  "https://en.wikiquote.org/wiki/Romanian_proverbs",
  "accessed": "2025-08-20",
  "basis": "The chip doesn't jump far
from the trunk", "claim": "The apple
does not fall far from the tree.",
  "cluster_id": 2,
  "quality_score": 6
}
```

3 System Architecture

3.1 Data Model and Provenance

Each proverb record maintains complete provenance through the pipeline:

This structure preserves the mother-tongue original while enabling computational comparison through the canonicalized `claim` field.

3.2 Collection Module

The system implements polite web scraping with configurable delays and user-agent identification:

Sources include Wikiquote (90 languages), Wiktionary categories, Project Gutenberg, and user uploads. All content respects original licensing (CC BY-SA for Wikimedia sources).

Algorithm 1 Wikiquote Proverbs Ingestion (outline)

- 1: Fetch HTML with polite delays and user agent
 - 2: Parse proverb sections and list items
 - 3: Clean text (strip anchors, names, boilerplate)
 - 4: Normalize to fields: {people, original, saying, english_equivalent, source_title, source_url, accessed}
 - 5: Emit records; deduplicate by normalized claim
-

3.3 Cleaning Pipeline

The cleaner employs a multi-stage filter to identify genuine proverbs:

$$keep(x) = \mathbb{I}[|x| \geq 6] \cdot \mathbb{I}[\neg name_{like(x)} \vee proverb_{cues(x)}] \cdot \mathbb{I}[\neg boilerplate(x)] \quad (1)$$

Where:

- `name_like(x)`: True if x contains ≤ 3 title-case tokens without verbs — heuristic to filter entries that look like proper names rather than proverb text.
- `proverb_cues(x)`: True if x contains punctuation or advice verbs — flags likely proverbs via punctuation (e.g., commas, semicolons) and modal/advisory verbs.
- `boilerplate(x)`: True if x matches navigation patterns (ISBN, "Retrieved from", etc.) — removes navigation/metadata strings that pollute scraped text.

Quality scoring ranks retained items:

$$Q(x) = 2 \cdot \mathbb{I}[\text{punctuation}] + 2 \cdot \mathbb{I}[\text{advice}_{\text{words}}] + \mathbb{I}[|x| \geq 5] + \mathbb{I}[\text{len}(x) \geq 20] \quad (2)$$

3.4 Canonicalization Engine

The key innovation is automatic transformation of figurative templates into comparable propositions through rule-based pattern matching. This process operates without manual intervention, applying a cascade of regular expression transformations that detect common proverb structures (conditionals, comparisons, prohibitions) and rewrite them into normalized forms. While this automation enables scalability, it necessarily sacrifices figurative richness, cultural nuance, and pragmatic ambiguity for computational tractability:

Algorithm 2. Proverb Canonicalization

Input: Proverb text p

Output: Canonical claim c

```

 $c \leftarrow p$ 
for each rule  $(\text{pattern}, \text{replacement}) \in \text{Rules}$  do
  if  $\text{regex\_match}(\text{pattern}, c)$  then
     $c \leftarrow \text{regex\_sub}(\text{pattern}, \text{replacement}, c)$ 
    break
  end if
end for
if  $c[-1] \in \{',', '!', '?'\}$  then
   $c \leftarrow c + '.'$ 
end if
return  $c$ 

```

Example transformations:

- “*Better late than never*” $\rightarrow$ Prefer late over never. — normalize to a minimal proposition to ease downstream parsing.
- “*If you lie down with dogs, you get up with fleas*” $\rightarrow$ If you lie down with dogs, then you get up with fleas. — convert to an explicit conditional, preserving logical entailment.
- “*Too many cooks spoil the broth*” $\rightarrow$ Excess of cooks is harmful. — compress idiomatic content to a concise, compositional gloss.

3.5 Clustering via Character n-Grams

We represent each canonical claim as a character n-gram TF-IDF vector, acknowledging that this surface-level approach lacks the semantic depth of modern contextual embeddings but enables cross-script processing without training data:

$$w_{c,g} = tf_{c,g} \cdot \log \frac{N}{df_g + 1} \quad (3)$$

where $tf_{c,g}$ is the term frequency of n-gram g in claim c , and df_g is the document frequency across all claims. Similarity is computed using cosine similarity:

$$sim(c_i, c_j) = \frac{\langle w_{c_i}, w_{c_j} \rangle}{\|w_{c_i}\| \|w_{c_j}\|} \quad (4)$$

Hierarchical clustering with average linkage groups similar claims:

$$d(A, B) = \frac{1}{|A||B|} \sum_{a \in A} \sum_{b \in B} (1 - sim(a, b)) \quad (5)$$

Parameter Selection & Sensitivity. We clustered TF-IDF (1-2 gram) cosine distances with average linkage and evaluated dendrogram cut thresholds $\tau \in \{0.25, 0.30, 0.35, 0.40, 0.45\}$ on a stratified sample of $n = 2,999$ proverbs drawn proportionally from the full dataset (3,673 rows). We report silhouette (cosine), cophenetic correlation and bootstrap stability (mean Adjusted Rand Index over 12 subsamples at 80% each). Cophenetic correlation was 0.805 across τ . Silhouette increased modestly with larger τ (0.338→0.352), while stability peaked at $\tau = 0.25$ (ARI=0.9995). We therefore select $\tau = 0.25$, prioritizing stability and avoiding over-merging.

3.6 Scoring Mechanism

Cross-cultural consensus is quantified through:

$$S_k = Coverage_k + 0.3 \times Support_k \quad (6)$$

where Coverage is the number of distinct cultures in cluster k , and Support is the total number of instances.

4 Preliminary Findings and System Demonstration

4.1 Dataset Statistics

Processing over 90 cultures yielded a substantial multilingual corpus. The initial collection contained 21,558 raw entries before any cleaning or deduplication. After applying language identification, noise removal, and our proverb detection heuristics, we retained 18,049 clean proverbs that showed

genuine proverbial characteristics. The clustering process identified 13,576 distinct clusters representing unique proverb families. Among these, only 14 patterns achieved universal status with coverage across 20 or more cultures, suggesting these may represent fundamental human insights. An additional 73 patterns showed regional distribution with coverage between 10 and 19 cultures, potentially indicating areal diffusion through trade or migration. The vast majority, comprising 13,489 clusters, remained culture-specific with coverage below 10 cultures, forming a long tail that highlights the rich diversity of localized wisdom.

4.2 High-Coverage Cross-Cultural Patterns

Table 1 presents the highest-coverage patterns identified by our system, demonstrating the methodology’s ability to detect recurrent themes across diverse cultural sources. These patterns represent candidates for further investigation by researchers with appropriate validation resources and cultural expertise.

The dominance of kinship-inheritance themes across unrelated cultures suggests either deep cognitive universals or ancient shared heritage.

Table 1: Highest-Coverage Cross-Cultural Patterns
(European bias acknowledged)

Canonical Wisdom	Coverage	Score
Children follow their parents’ example	31	41.2
The apple doesn’t fall far from the tree	30	40.5
Like father, like son	27	35.1
Like mother, like daughter	22	28.9
Working hard for others, neglecting self	21	27.3

4.3 Productive Tensions

A fascinating discovery was seemingly contradictory wisdom achieving high coverage:

Cooperation: “*Many hands make light work*” (15 cultures)

Coordination: “*Too many cooks spoil the broth*” (12 cultures)

This tension maps onto economic diminishing returns: cooperation helps until coordination costs exceed benefits. The threshold appears culturally variable.

4.4 Ecological Correlations

Statistical analysis revealed correlations between wisdom themes and coarse ecological proxies (Table 2), though we emphasize these are suggestive

associations rather than causal relationships, and our ecological variables may obscure important local variation and confounding cultural factors.

Table 2: Suggestive Ecological-Wisdom Associations (Caution: Coarse Proxies)

Correlation	Spearman's ρ	p-value
Maritime orientation→Cooperation	0.31	0.003**
Rice agriculture→ Coordination	0.28	0.021*
Urbanization→Time/Efficiency	0.24	0.038*
Wheat regions→ Individual prudence	0.22	0.045*

These patterns support theoretical predictions about ecological shaping of cultural values.

4.5 Romanian Wisdom in Context

Romanian proverbs show cross-cultural alignment patterns that must be interpreted cautiously given our methodological limitations and potential canonicalization distortions. The highest overlap occurs with Albanian, where 89% of top clusters are shared between the two languages, suggesting strong regional influence despite different language families. Within the Romance family, Romanian maintains expected similarities with Italian at 76% overlap, Spanish at 71%, and French at 68%, establishing a baseline for genealogical relatedness. Regional contact effects manifest through 62% overlap with Bulgarian and 58% with Serbian, reflecting centuries of Balkan linguistic interchange.

Specific Romanian proverbs demonstrate remarkably high universality across our corpus. The saying “Așchia nu sare departe de trunchi” (The chip doesn’t jump far from the trunk) appears in equivalent forms across 30 cultures, suggesting a near-universal recognition of hereditary influence. The collective action proverb “Unde-s mulți, puterea crește” (Where there are many, power grows) finds parallels in 12 cultures, while the early-rising maxim “Cine se scoală de dimineață, departe ajunge” (Who rises early goes far) resonates across 15 cultures, mapping potential diffusion pathways.

4.6 Clustering Validation

The clustering achieved modest quality metrics that reflect the inherent difficulty of semantic clustering using surface features. The silhouette coefficient of 0.23 indicates relatively weak cluster cohesion, suggesting that our character n-gram approach captures similarity imperfectly. The mean intra-cluster distance of 0.31 represents the average cosine distance within clusters, showing moderate semantic coherence. However, the clustering demonstrates reasonable robustness with a stability score of 0.87 under 10% subsampling, and maintains 92% retention of the top-10 clusters across bootstrap runs, suggesting that despite modest absolute quality, the major patterns are consistently recoverable.

Coverage bins (operational). Using our operational bins ("universal" ≥ 20 cultures; "regional" 10–19; "culture-specific" < 10), the present dataset yields no clusters in the first two categories at $\tau = 0.25$. Empirically, distinct-culture coverage per cluster has mean 1.15, 90th percentile 2, 99th percentile 3, and maximum 6 (at $\tau = 0.25$; at $\tau = 0.45$ the maximum is 7). We therefore treat the bins as operational labels and report the full coverage distribution.

Table 3: Clustering sensitivity across dendrogram cut thresholds (stratified $n=2,999$)

τ	#clusters	Silhouette	Cophenetic	Mean ARI
0.25	2153	0.338	0.805	0.9995
0.30	2137	0.342	0.805	0.9987
0.35	2115	0.347	0.805	0.9981
0.40	2088	0.350	0.805	0.9953
0.45	2059	0.352	0.805	0.9961

5 System Implementation

5.1 Technical Architecture

The system is implemented as a modular Python architecture with a Streamlit web interface. Core modules handle collection and merging, noise removal, clustering operations, deterministic analysis, optional LLM narratives, and validation metrics. The web interface integrates these components through an intuitive multi-tab design. Key dependencies include scikit-learn for TF-IDF and clustering, BeautifulSoup for controlled web scraping, UMAP/PCA for dimensionality reduction, and optionally llama.cpp for local language model integration without API requirements.

5.2 User Interface

The Streamlit interface provides five integrated tabs that guide users through the complete analysis pipeline. The Dataset Builder tab enables users to upload their own CSV files or fetch proverbs from configured sources with automatic schema validation. The Wisdom Extractor tab provides access to all clustering parameters, allowing researchers to configure cleaning thresholds and clustering distance parameters according to their specific research needs. The Results & Visualization tab presents an interactive browser for exploring discovered clusters alongside a two-dimensional semantic map for pattern discovery. The Interpretation tab generates themed reports that organize clusters by semantic categories for qualitative analysis. Finally, the Diagnostics tab provides validation metrics and statistical correlations to assess clustering quality and explore potential ecological associations.

5.3 Performance

On standard hardware (Intel i5, 8GB RAM), the pipeline demonstrates efficient performance across all stages. Collection requires approximately 30 seconds per source with polite delays between requests. The cleaning phase processes 20,000 entries in under 5 seconds through vectorized operations. Clustering 18,000 proverbs completes in 30 seconds using optimized sparse matrix computations. Visualization renders in under 2 seconds through cached projections. The complete pipeline from raw collection to final visualization typically completes within 5 minutes, making iterative exploration feasible for researchers.

6 Discussion

6.1 Research Directions Enabled

Our preliminary patterns suggest three potential research directions for future investigation. First, convergent evolution may explain how similar ecological pressures produce similar wisdom, as evidenced by the maritime-cooperation correlation. Second, cultural diffusion patterns become traceable through systematic overlap analysis, potentially revealing historical trade routes like the Silk Road through shared proverb clusters. Third, universal patterns that appear across unrelated cultures may reflect deep cognitive constraints on how humans encode procedural knowledge, suggesting fundamental limits on the space of viable heuristics.

6.2 Methodological Contributions

Despite these limitations, the Wisdom Extractor offers specific methodological contributions to computational paremiology. Our transparent canonicalization provides a systematic approach to transforming figurative language into propositional form, making the alignment process auditable and reproducible. The use of character n-grams ensures cross-script compatibility, enabling analysis across all writing systems without requiring language-specific models or training data. Throughout the pipeline, we preserve authenticity by maintaining mother-tongue forms alongside their transformations, ensuring that original cultural expressions remain accessible for validation. The system operates entirely offline without API dependencies, guaranteeing complete reproducibility and eliminating concerns about service availability or data privacy.

6.3 Significant Limitations

Our findings require transparent acknowledgment of multiple methodological constraints. The canonicalization process, while enabling systematic comparison, fundamentally strips away figurative richness, cultural nuance, and

pragmatic ambiguity central to how proverbs function in native contexts. Transforming “If you lie down with dogs, you get up with fleas” into “If you associate with bad people, then you acquire bad traits” loses the visceral imagery and variable cultural associations with dogs across societies. This semantic flattening represents a foundational trade-off between computational tractability and cultural authenticity.

Our reliance on character n-grams, while ensuring cross-script compatibility and offline operation, lacks the semantic depth of modern contextual embeddings. The silhouette score of 0.23 indicates weak cluster cohesion, suggesting many groupings capture surface similarity rather than deep semantic equivalence. Proverbs grouped together may share syntactic patterns without expressing the same underlying wisdom, potentially inflating cross-cultural convergence estimates.

Corpus composition introduces systematic biases that fundamentally limit universality claims. European languages constitute 62% of our dataset, likely reflecting Western dominance in digital archiving rather than genuine global patterns. Patterns identified as “universal” may simply represent European cultural diffusion or shared Indo-European heritage. Indigenous, African, and Pacific traditions remain severely underrepresented; their inclusion might reveal entirely different organizing principles for encoding cultural knowledge.

English canonicalization may fundamentally alter meaning for culturally embedded expressions lacking direct equivalents. Concepts from languages with different grammatical structures like evidentiality markers in Quechua or elaborate honorific systems in Japanese, cannot be adequately captured in English propositions. Moreover, treating all proverbs as contemporaneous ignores centuries of historical evolution. Ancient proverbs sit alongside modern additions in our clusters, obscuring the crucial distinction between inherited wisdom and recent innovations.

Our ecological correlations represent suggestive associations rather than causal relationships. Coarse proxies like “maritime orientation” or “rice agriculture” may obscure crucial local variations and confounding factors. Maritime-cooperation correlations might reflect trade networks, population density, or historical contingencies rather than ecological determinism. Without sophisticated causal modelling, our correlations remain speculative hypotheses.

Finally, our analysis of textual collections may bear little relationship to how proverbs actually function in daily life. Written compilations may not reflect actual usage patterns, influence on decision-making, or role in social discourse. A proverb’s dictionary presence tells us nothing about whether people invoke it during moral reasoning, child-rearing, or conflict resolution. Without behavioural validation through ethnographic observation or experimental studies, our clusters remain artifacts of textual analysis rather than windows into lived wisdom.

6.4 Epistemological Considerations and Validity Threats

Three epistemological challenges qualify our findings. First, English canonicalization may create translation artifacts rather than reveal genuine convergence—collectivist concepts may be distorted when forced into English’s individualistic grammatical structures. Second, temporal flattening conflates ancient wisdom with modern additions, obscuring inherited versus convergent patterns. Third, distinguishing universal thought patterns from cultural transmission remains methodologically intractable. High-coverage patterns might reflect cognitive universals, ancient common heritage, or globalization’s homogenizing effects.

Empirically, per-cluster diversity (distinct language families, regions) showed no enrichment beyond random mixing when tested against a permutation baseline (200 permutations preserving cluster sizes; observed percentiles near median: families mean 39.5, regions mean 40.4). We therefore avoid causal claims and frame high coverage as an operational signal pending behavioural validation.

Key validity threats include: (1) translation distortion, mitigated via canonicalization rules with planned backtranslation spot checks; (2) metaphor flattening, addressed by preserving figurative markers where role structure matters; (3) sampling bias from source availability asymmetries; and (4) clustering sensitivity to τ , quantified via silhouette, cophenetic correlation, and bootstrap stability. Qualitative patterns remain stable across $\tau \in [0.25, 0.45]$. Future work requires diachronic stratification, controlled comparison of isolated cultures, and triangulation with archaeological and genetic evidence to disentangle competing explanations.

6.5 Applications

The system enables diverse research applications across multiple disciplines. In comparative linguistics, researchers can conduct systematic proverb comparisons across language families to test hypotheses about universal versus culturally specific wisdom. Translation studies benefit from our figurative equivalence patterns, which reveal how metaphorical concepts map across languages and cultures. Cultural psychologists can test value-ecology hypotheses by correlating proverb themes with environmental and social variables. Digital humanities scholars gain tools for large-scale cultural text analysis that preserve provenance while enabling quantitative methods. In educational settings, the interactive interface allows students to explore world wisdom traditions while learning about computational methods in cultural analysis.

7 Future Directions

Priority extensions for the system span several dimensions of enhancement. Temporal analysis represents a crucial next step, requiring the integration of dated attestations to track how proverbs evolve and spread over time. Moving

beyond correlation to establish causation demands sophisticated causal modelling approaches, potentially leveraging natural experiments or instrumental variables from historical events. Behavioural validation would link proverb knowledge to actual decision-making through experimental studies or observational data. Expanding coverage to indigenous languages is essential to address the current European bias that limits our claims about universality. Domain specialization into legal, medical, and technical proverbs would reveal how professional communities develop and transmit specialized procedural knowledge.

For Romanian-specific enhancements, we envision several targeted improvements. Integration of complete Romanian paremiological corpora from major published collections and archives would provide comprehensive coverage of documented wisdom. Diachronic analysis using Old Romanian texts would enable tracking of how proverb forms and meanings shift across centuries, revealing patterns of linguistic and cultural change. Regional dialect variation mapping would uncover how proverbs vary across Romanian-speaking regions, potentially correlating with historical administrative boundaries or geographic features. A systematic Romanian-Moldovan comparative analysis would quantify convergences and divergences across political borders, testing theories about how institutional separation affects cultural evolution.

8 Conclusions

The Wisdom Extractor provides computational paremiology with new methodological infrastructure for systematic cross-cultural analysis. By processing 18,049 proverbs from 90+ cultures, our system demonstrates the feasibility of large-scale figurative language alignment and generates preliminary patterns worthy of investigation by researchers with appropriate validation capabilities.

Our contribution is primarily methodological: we provide transparent, reproducible tools that enable systematic analysis of cultural knowledge at unprecedented scale. The patterns we identify—high-coverage themes, productive tensions, and ecological associations—represent research hypotheses rather than established findings, requiring validation through human expert judgment, behavioural studies, and controlled comparative analysis.

For Romanian linguistics, the system offers particular value in quantifying both the Romance heritage and Balkan sprachbund effects in Romanian paremiology. The 89% overlap with Albanian coupled with 76% Italian alignment provides computational evidence for the complex cultural position of Romanian.

More broadly, our findings provide preliminary evidence for the hypothesis that cross-cultural redundancy signals importance, though this must be qualified by substantial sampling bias, semantic flattening from canonicalization, and the need

for behavioural validation. Patterns appearing across multiple cultures in our corpus may reflect important regularities, but establishing their universality requires addressing our European over-representation and methodological limitations. Yet the productive tensions remind us that wisdom is contextual—a toolkit rather than prescription.

By making this analysis transparent, reproducible, and accessible through an intuitive interface, the Wisdom Extractor opens new avenues for understanding cultural evolution, knowledge transmission, and the deep patterns of human social cognition. The system is freely available if requested, inviting collaborative expansion of humanity’s wisdom map.

Code and Data Availability

The complete Wisdom Extractor system, including all source code, example datasets, and documentation, is available at <https://github.com/ovladon/wisdom-extractor>. The repository includes a comprehensive README with installation instructions, parameter tuning guidelines, and example analyses.

Data & Code Availability. All scripts and metrics tables are provided in the companion repository and release bundle (including sensitivity metrics, coverage, triangulation, and audit templates for canonicalization and back-translation). Fine-grained implementation details are in the repository README rather than the paper.

AI use statement

OpenAI’s ChatGPT (GPT-5 Thinking) and Anthropic’s Claude Opus were used to: (i) improve sentence clarity and grammatical consistency; (ii) suggest LaTeX and template-compliance adjustments (headings, captions, sectioning); and (iii) normalize bibliography entries (ensuring DOIs/URLs appear according to the proceedings’ style). AI was not used to generate data, as it was politely scraped from the sources mentioned in text and all results interpretations and future research proposals represent mere suggestions, as mentioned in the paper. All suggestions were verified and edited by the authors, who accept full responsibility for the final text.

Acknowledgements

We thank the Wikimedia Foundation for maintaining open proverb collections, and the developers of scikit-learn, UMAP, and Streamlit for essential infrastructure.

References

- Chen Cecilia Liu, Fajri Koto, Timothy Baldwin, and Iryna Gurevych. 2024. Are multilingual llms culturally- diverse reasoners? an investigation into multicultural proverbs and sayings (maps). arXiv:2309.08591 [cs.CL]. DOI: 10.48550/arXiv.2309.08591.
- Ekaterina Shutova. 2010. Models of metaphor in NLP. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 688–697, Uppsala, Sweden. Association for Computational Linguistics.
- Ethan Davis, Christopher M. Danforth, Wolfgang Mieder, and Peter Sheridan Dodds. 2021. Computational paremiology: Charting the temporal, ecological dynamics of proverb use in books, news articles, and tweets. arXiv:2107.04929 [cs.CL]. DOI: 10.48550/arXiv.2107.04929.
- Gerard J. Steen, Aletta G. Dorst, J. Berenike Herrmann, Anna A. Kaal, Tina Krennmayr, and Tryntje Pasma. 2010. *A Method for Linguistic Metaphor Identification: From MIP to MIPVU*. Number 14 in Converging Evidence in Language and Communication Research. John Benjamins, Amsterdam; Philadelphia. Series 14; pp. 1–238. ISBN: 978-9027239037.
- Hiromu Yakura. 2024. Evaluating large language models’ ability using a psychiatric screening tool based on metaphor and sarcasm scenarios. *Journal of Intelligence*, 12(7):70. DOI: 10.3390/jintelligence12070070.
- John Pavlopoulos, Panos Louridas, and Panagiotis Filos. 2024. Towards a greek proverb atlas: Computational spatial exploration and attribution of Greek proverbs. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11842–11854, Miami, Florida, USA. Association for Computational Linguistics. DOI: 10.18653/v1/2024.emnlp-main.661.
- Minghan Wang, Viet-Tuan Pham, Farzaneh Moghimifar, and Tu-Ting Vu. 2025. Proverbs run in pairs: Evaluating proverb translation capability of large language model. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 1646–1662, Vienna, Austria. Association for Computational Linguistics. DOI: 10.18653/v1/2025.findings-acl.83.
- Minjin Choi, Sunkyung Lee, Eunseong Choi, Heesoo Park, Junhyuk Lee, Dongwon Lee, and Jongwuk Lee. 2021. Melbert: Metaphor detection via contextualized late interaction using metaphorical identification theories. arXiv:2104.13615 [cs.CL]. DOI: 10.48550/arXiv.2104.13615.
- Norbert Ichien, Tudor Malutan, Dedre Gentner, and Keith J. Holyoak. 2024. Large language models display emergent ability to interpret literary metaphors. *Metaphor and Symbol*, 39(3):187–205. DOI: 10.1080/10926488.2024.2380348.
- Robert Boyd, Peter J. Richerson, and Joseph Henrich. 2011. The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences*, 108(Suppl. 2):10918–10925. DOI: 10.1073/pnas.1100290108.
- Soumya Ghosh and Shubham Srivastava. 2022. epic: Employing proverbs in context as a benchmark for abstract language understanding. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3989–4004, Dublin, Ireland. Association for Computational Linguistics. DOI: 10.18653/v1/2022.acl-long.276.

- Thomas Talhelm, Xuemin Zhang, Shigehiro Oishi, Shimin Chen, Dandan Duan, Xuezhao Lan, and Shinobu Kitayama. 2014. Large-scale psychological differences within china explained by rice versus wheat agriculture. *Science*, 344(6184):603–608. DOI: 10.1126/science.1246850.
- Wei Song, Shi Zhou, Rui Fu, Ting Liu, and Lili Liu. 2021. Verb metaphor detection via contextual relation learning. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4240– 4251, Online. Association for Computational Linguistics. DOI: 10.18653/v1/2021.acl-long.327.
- Wolfgang Mieder. 2021. *Dictionary of Authentic American Proverbs*. Berghahn Books, New York; Oxford. ISBN: 978-1-80073-131-8. pp. 1–238.
- Xiaoyu Tong, Rianne Choenni, Mike Lewis, and Ekaterina Shutova. 2024. Munch: Metaphor understanding challenge dataset for llms. arXiv:2406.07821 [cs.CL]. DOI: 10.48550/arXiv.2406.07821.

EXPLORING THE BESSARABIAN IDIOM UNDERSTANDING BY LLMS AND AI TOOLS

INGA TITCHIEV, OLESEA CAFTANATOV

*Vladimir Andrunachievici Institute of Mathematics and Computer Science,
Moldova State University*

{inga.titchiev,olesea.caftanatov}@math.usm.md

Abstract

Idiomatic expressions are a key part of the cultural and linguistic heritage of a society, playing an important role in preserving and promoting the identity of the community. Due to their non-compositional nature, identifying and interpreting idioms poses challenges for fields such as machine translation and natural language processing. This study aims to collect, analyse and illustrate Romanian idiomatic expressions using artificial intelligence, with applications in education. In the first stage, 100 expressions from the Bessarabian heritage were explored using ChatGPT and DeepSeek to provide correct continuations, literal and figurative meanings, and English equivalents. In addition, for visual interpretation of some of them, several image generation tools were applied. The preliminary results open avenues for the development of digital resources to support the learning and research of the Romanian language.

Keywords: idiom, artificial intelligence, visual representation, deep learning

1 Introduction

Idiomatic expressions are an essential part of a society's cultural and linguistic heritage and play an important role in its preservation and promotion. By understanding, conserving, and promoting idiomatic expressions, we can contribute to preserving and capitalizing on the cultural richness of a community. Idiosyncratic semantic multiword expressions (MWEs), or idiomatic expressions, are noncompositional in the sense that their meaning cannot be understood from their parts; these expressions are frequently used to make language more fluent and often contain word meanings that are not found in other contexts. Therefore, identifying noncompositional MWEs represents a clear challenge for fields such as machine translation (MT), information retrieval, natural language understanding, natural language generation, question answering, text summarization, and word meaning disambiguation. For this purpose, in order to identify, classify, and augment native idiomatic expressions using artificial intelligence and augmented reality, with applicability in education to facilitate deep learning of the Romanian language, this study was initiated.

Regarding idioms and their processing by Large Language Models (LLMs), several studies have been conducted, among which this work introduces MIDAS, a multilingual idiom dataset, and shows that LLMs rely on a hybrid of memorization, contextual cues, and reasoning (Kim *et al.*, 2025).

For the first stage of the study, the following objectives were established: collection and analysis of local idiomatic expressions; establishment of tools for automatic identification of idiomatic expressions; classification of idiomatic expressions; application of Artificial Intelligence (AI) tools to illustrate the true meaning of idiomatic expressions. These digital resources can be used in various linguistic analyses. To achieve the set objectives, 100 idioms from the Bessarabian cultural heritage were selected and analysed using two artificial intelligence tools: ChatGPT (OpenAI, 2025) and DeepSeek (Platform, 2025). Each of the instruments was required to: 1. continue the idiom and explain the meaning if it was not continued correctly; 2. formulate the literal and figurative meaning of the given idiom; 3. identify English idioms with similar meaning. Thus, in the rest of this article, we will present the results of the performed analysis and establish the further research goals.

2 Applied Methodology

To support the experimental setup, the first step involved preparing the source materials for analysis. The Moldavian folklore volume: *O creație populară: proverbe și zicători (A Folk Creation: Proverbs and Sayings)* by Botezatu and others (1981) were digitized and transliterated, using optical character recognition (OCR) combined with manual post-editing on the HeDy platform (Bumbu *et al.*, 2023). Following digitization, the texts were validated through recognition processes and subsequently converted from Cyrillic to Latin script.

From this corpus, we selected 100 idioms representing diverse motifs, such as: *the universe, nature, the world, human beings, homeland, native land, city, village, health, illness, life, death, food, hunger, thirst, historical proverbs, legal proverbs and sayings, social inequality, wealth, poverty, happiness, success, failure, luck, and necessity*. The respective motifs (e.g., nature, human relationships, wisdom, work, destiny) were subsequently consolidated into broader thematic categories, used for analysis, as presented in Section 4.

This curated selection was subsequently organized into an Excel file. In order to perform the data analysis, the columns were normalized. The analysis was based on two experiments on idioms, the first related to textual representation and the second to visual representation. To connect the textual and visual experiments, we adopt a cross-modal perspective. A shared subset of idiomatic expressions is used across both experiments, allowing us to compare how models interpret and visualize the same figurative meanings. The pipeline follows three steps: (1) the LLM generates both the idiomatic completion and an explanation of its figurative meaning; (2) the explanation is transformed into a prompt for the image generation

model; and (3) the resulting image is evaluated in terms of how well it reflects the intended figurative sense. This setup enables the analysis of semantic alignment between textual and visual modalities, establishing a bridge between linguistic understanding and visual representation.

In the first experiment, for each idiom, the translations provided by ChatGPT (GPT-5 version, access mode – chat interface) and DeepSeek V3.2, also access mode - chat interface, were extracted, namely: • literal translation and • figurative translation. Subsequently, the textual similarity measure was calculated through the SequenceMatcher ratio between literal translations ChatGPT vs. DeepSeek, and figurative translations ChatGPT vs. DeepSeek. The SequenceMatcher ratio is a numerical measure (between 0 and 1) that shows how similar two strings of characters (texts) are. The algorithm searches for common sequences between the two texts and calculates the proportion:

$$ratio = \frac{2 \times M}{T}$$

where: M = total number of characters that match in the same sequences, T = total length of the two strings put together.

If ratio = 1.0, the strings are completely identical, and if ratio = 0, they are totally different. The identification of recurring themes was carried out by searching for keywords in Romanian and English (e.g. words related to nature, animals, work, etc.), assigning each idiom 1 or multiple themes.

2.1 Obtained Results

2.1.1 Results of Similarity Analysis

Literal translation (ChatGPT vs. DeepSeek)

- 41 idioms with very high similarity (≥ 0.9 , almost identical).
- 0 idioms with very low similarity (≤ 0.1 , almost completely different).

Figurative translation (ChatGPT vs. DeepSeek)

- 7 idioms with very high similarity (≥ 0.9).
- 7 idioms with very low similarity (≤ 0.1).

This shows that literal translations are much closer between the two models, while figurative translations have large differences.

2.1.2 ChatGPT Completion Accuracy

• Exact Match: 33 idioms

- Examples: “Câte capete, atâtea păreri” (As many heads, as many opinions), “Slab de-l bate vântul.” (So weak that the wind could knock him down), “Gaura mică scufundă corabia mare.” (A small hole sinks a big ship), “Sacul gol nu stă în

picioare.” (An empty sack cannot stand upright), “Satul arde, dar baba se piaptână.” (The village is burning, but the old woman is combing her hair), “A fi cu stea în frunte.” (To be born under a lucky star / To be destined for greatness), “Omul la nevoie se cunoaște.” (A friend in need is a friend indeed), “A-l ajunge pe cineva cuțitul la os.” (To be pushed to the limit / To reach breaking point), “Bate fierul, cât e cald.” (Strike while the iron is hot), “La bătrânețe nu se mai prinde minte.” (You can’t teach an old dog new tricks).

- **Plausible variant: 22 idioms**

- Examples: For “Dintr-un bob răsar o mie,” (From one seed a thousand sprout) ChatGPT offered “Dintr-un bob, o grămadă,” (From one seed, a heap), or for “A fi de râsul găinilor.” (To be the laughingstock of the hens / To make a fool of oneself), ChatGPT provided the answer “A fi de râsul curcilor.” (To be the laughingstock of the turkeys), which didn’t change the meaning. “Cine n-a gustat amarul, nu știe ce-i dulceața.” (He who hasn’t tasted bitterness doesn’t know what sweetness is).

We consider it plausible that the provided idiom is not exactly the same but slightly different – for example, using “dulceață” instead of “zahăr”, or “curcile” instead of “găinile”. Such variations are regarded as acceptable.

- **Incorrect: 45 idioms**

- Examples: For “La nevoie și racu-i pește.” (In need, even the crayfish is a fish), ChatGPT offered “La nevoie și racul își scoate coada.” (In need, even the crayfish pulls out its tail), which alters the base meaning. For “Arborele mare, cade tare.” (The big tree falls hard), it gave a different proverb, “Arborele mare atrage vânturi puternice.” (The big tree attracts strong winds). Another example is “A trăi ca peștele în apă” (To live like a fish in water / To be in one’s element), for which it provided the opposite, “A trăi ca peștele pe uscat” (To live like a fish on dry land / To be out of one’s element).

2.1.3 DeepSeek Completion Accuracy

- **Exact Match: 32 idioms**

- Examples: “A încălzi șarpele la sân.” (To warm the snake at one’s breast / To nurture a traitor), “A-și pune pofta în cui.” (To hang one’s desire on a nail / To give up hope for something), “Satul arde, da baba se piaptână.” (The village is burning, but the old woman is combing her hair), “A scăpa de Dracul și a da peste tată-său.” (To escape the devil and run into his father / To jump from the frying pan into the fire), “A-l ajunge pe cineva cuțitul la os.” (To reach the bone / To be at the breaking point), “Omul înțelept face ceea ce poate, nu ceea ce vrea.” (A wise

man does what he can, not what he wants), “Nici în clin, nici în mănecă.” (To have nothing in common / No connection whatsoever), etc.

- **Plausible variant: 20 idioms**

- Examples: For “Dintr-un bob răsar o mie,” (From one seed a thousand sprout) DeepSeek offered “Dintr-un bob se face snop,” (From one seed comes a sheaf), a valid variant. For “Îi mai cald soarele când răsare decât când asfințește,” (The sun is warmer at sunrise than at sunset) it used “decât când apune” (than when it sets) instead of “decât când asfințește.” Other examples include “Te va întreba iarna, ce-ai lucrat vara” (Winter will ask you what you did in summer – replacing “lucrat” with “făcut”), “Gaura mică scufundă corabia mare” (A small hole sinks a big ship – omitting “mare”), “A se ține ca umbra după om.” (To follow someone like a shadow – replacing “om” with “cineva”).

- **Incorrect: 48 idioms**

- Examples: For “A umbla (după cineva) ca lupul după oaie.” (To chase someone like a wolf after a sheep – to pursue relentlessly, like a predator), DeepSeek offered “A umbla ca lupul prin biserici” (To wander like a wolf through churches), which could be understood as wandering restlessly and pointlessly, being out of place. Some of the idiomatic results provided by DeepSeek are completely incorrect, yet they also carry a humorous twist – for example: “Capra face boi” (The goat makes oxen), “La nevoe și racu-i drept” (In need, even the crayfish is upright), or “A da bani pe nas” (To pay through the nose).

When considering both exact matches and plausible variants as “accurate completions”:

- ChatGPT achieved 55% accuracy.
- DeepSeek achieved 52% accuracy.

Both models demonstrate a similar level of accuracy in completing the Bessarabian idioms, with ChatGPT having a slightly higher rate in strict correctness and overall plausible correctness, partly due to providing a completion for every idiom, while DeepSeek produced two cases of incomplete continuation – for instance, with “Firele când se încurcă...” – it failed to finish the proverb. The number of “incorrect” continuations is also quite similar for both. This suggests that while both LLMs can handle many idioms correctly or with plausible variations, they also frequently diverge from the precise, expected continuation, sometimes offering entirely different proverbs or altering the core meaning.

2.2 Scoring Figurative Meaning Recognition from the Explanations

We performed a scoring of figurative meaning recognition based on the explanations provided by ChatGPT and DeepSeek. This involves analysing how accurately and thoroughly each LLM’s explanation captures the nonliteral, metaphorical, or idiomatic sense of the Bessarabian proverbs and expressions. To score this, the explanation of each idiom was evaluated according to the following criteria:

- **Excellent:** The explanation clearly and accurately identifies the figurative meaning, often with good detail, nuance, or a suitable breakdown of the metaphor.
- **Good:** The explanation accurately identifies the figurative meaning but might be slightly less detailed or nuanced.
- **Partial/Mixed:** The explanation partially grasps or slightly misinterprets the figurative meaning, includes less common interpretations, or conflates the idiom with other ideas.
- **Poor/Missed:** The explanation largely misses the figurative meaning, significantly misinterprets it, or explains a different idiom (especially when the LLM provided a different proverb as a completion and then explained that different proverb).

Based on an analysis of 100 unique idioms for which explanations were provided by either or both models, here is the breakdown:

2.2.1 ChatGPT: Figurative Meaning Recognition

– Excellent recognition: 76 idioms

Example (Dintr-un bob răsar o mie): ChatGPT explained, “exprimă ideea că dintr-un lucru mic sau neînsemnat poate rezulta ceva mare sau important”. Its English figurative equivalents like “From little acorns grow mighty oaks” also perfectly capture this.

– Good recognition: 5 idioms

Example (Cănele mic pînă la bătrînețe-i cățeluș): ChatGPT’s explanation “cineva lipsit de putere sau importanță rămâne așa toată viața, indiferent de vârstă sau timp” accurately captures the essence of unchanging nature, though “A leopard can’t change its spots” is a broader English equivalent.

– Partial recognition: 10 idioms

Example (În casa lăutarului fiecare joacă): ChatGPT’s interpretation “Nu e potrivit să te dai expert acolo unde ești pe [teritoriul altuia]” is a figurative meaning, but not the most common one for the given half of the idiom.

– **Missed recognition: 9 idioms**

Example (A-și pleca urechea): ChatGPT stated it means “A se umili, a se supune”, which is incorrect; the idiom means “to listen carefully.”

2.2.2 DeepSeek: Figurative Meaning Recognition

– **Excellent recognition: 68 idioms**

Example (Gaura mică scufundă corabia mare): DeepSeek explained that “lucrurile aparent nesemnificative pot avea consecințe catastrofale. Semnificații cheie: Pericolul neglijării detaliilor mici”.

– **Good recognition: 6 idioms**

Example (A fi într-o doagă): DeepSeek interpreted this as “a fi foarte supărat, furios sau agitat, aproape la punctul de a-și pierde controlul”. Although different from ChatGPT’s “mental instability”, is a valid and recognized figurative meaning, well explained.

– **Partial recognition: 11 idioms**

Example (Nu toată nuca are miez bun): DeepSeek’s explanation was somewhat mixed, attributing some meaning to a related idiom: “Nu toate lucrurile care par valoroase sunt productive (unele nuci sunt goale) Norocul nu este egal distribuit”.

– **Missed recognition: 15 idioms**

Example (Capra face, dar oaia trage): DeepSeek completely misinterpreted this, explaining it as a proverb where “o minciună repetată devine credibilă”, rather than its actual meaning of one causing trouble and another taking the blame.

When scoring figurative meaning recognition based on the explanations provided:

ChatGPT demonstrated higher accuracy, achieving 80.81%. Its explanations were generally more consistent with the widely accepted figurative meanings of the idioms, and provided detailed breakdowns of the metaphors. ChatGPT’s errors often stemmed from misinterpreting a specific idiom’s nuances.

DeepSeek achieved a score of 74.41%. Although it provided many excellent explanations, a notable portion of its errors involved explaining its own sometimes incorrect) completions of the idioms rather than the original ones, or offering significant misinterpretations of the idioms”; core figurative meanings. ChatGPT appears to be slightly more adept at recognizing and articulating the figurative meanings of these Bessarabian idioms from its explanations compared to DeepSeek.

3 AI Image Generation: Prompt Analysis

In the second experiment, we explored how AI image generation platforms interpret and visually represent idiomatic expressions. AI systems have demonstrated rapid advancement in producing contextually relevant imagery from textual prompts, yet the challenge of accurately translating idioms into images remains significant due to their metaphorical and culturally-specific nature.

For this experiment, we focused on a subset of 100 Russian similar idioms, and 100 Bessarabian proverbs using both literal and metaphorical prompts to assess the systems’ semantic understanding and creative interpretation. The analysis covered prompts such as “Сидеть на двух стульях” (“Sit on two chairs”), “Волк в овечьей шкуре” (“Wolf in sheep’s clothing”), and “Кот заплакал” (“The cat cried”), among others. Each proverb was provided in two forms: a literal prompt, describing a direct, often absurd visual scene, and a metaphorical prompt, emphasizing the underlying symbolic or conceptual meaning. For this study, we examined the following platforms: ReCraft, Sora, Canva, FreePik, Stable Diffusion, Fusion Brain, Pollo AI, Shedevrum, Supermaker, Google Gemini, ChatGPT Plus, NightCafe, CGDream, Ideogram, OpenArt.ai, and Nano Banana. Our preliminary evaluation revealed notable differences in how AI platforms handle abstract prompts, especially those containing idiomatic or proverbial language. Below, we will present several observations regarding the understanding of AI tools when processing prompts for idioms in both their literal and figurative forms.

3.1 Literal vs. Metaphorical Interpretation

A central focus of this research is the capacity of AI image generation platforms to interpret user prompts, particularly when dealing with abstract or idiomatic expressions. The study revealed that certain platforms, notably CGDream, exhibited notable difficulties in visualizing proverbs, which frequently carry metaphorical or figurative meanings. For example, the Bessarabian proverb “A fi cu musca pe căciulă” (literally, “to be with a fly on a hat,” meaning “to feel guilty”) was evaluated. According to the research (Colesnicova *et al.*, 2025), while other platforms successfully generated relevant images – such as a hat or a hat with a fly – CGDream produced an image of a girl holding an umbrella, highlighting a significant difficulty in grasping the context, see Figure 1.



Figure 1: Example of CGDream’s difficulty in understanding the literal meaning of idiom.

To explore this phenomenon in depth, the research employed two distinct prompt strategies – literal and metaphorical (figurative/conceptual) – across a diverse set of proverbs. For example, literal prompts aim to depict the concrete elements described in the proverb. For instance, the Russian idiom “a wolf in sheep’s

clothing” (see Figure 2.a) would be represented literally by a photorealistic image of a large grey wolf wearing a white sheepskin, with its snout and sharp teeth visible, highlighting the imperfect disguise. Metaphorical prompts (see Figure 2.b), in contrast, capture the underlying symbolic meaning, often through abstract or allegorical imagery.

In this case, the same idiom could be depicted as a seemingly charming person whose shadow reveals a monstrous, predatory figure, symbolizing hidden danger and betrayal. This approach emphasizes the contrast between outward appearance and true nature, translating the figurative sense of the proverb into a human and narrative context.

3.2 Humor and Creative Interpretation

The AI image generation process often produces creative and occasionally humorous interpretations, especially when responding to literal prompts of idiomatic expressions. For example, in the case of the Russian idiom “Сесть в калошу” (“to sit in a galosh”, see Figure 3. a), the prompts requested a whimsical, cartoon-style depiction of a surprised person comically trapped inside a large, old-fashioned boot. Another example is the literal interpretation of the Russian idiom “Толова кипит” (“head boiling”, see Figure 3.b), which was prompted as a surreal and humorous photorealistic image, depicting a human head transformed into a transparent glass kettle boiling on a stove.




Figure 2: Literal and metaphorical representation of idioms

Many literal interpretation prompts explicitly call for “comical”, “absurdist” or “satirical” styles to capture the inherent humor of visualizing proverbs literally. NightCafe, in particular, demonstrated remarkable creativity and stylistic versatility, producing distinct conceptual representations that showcase its potential for imaginative and aesthetically diverse outputs. For example, for the Bessarabian idiom “Cine face carte are patru ochi” (“Those who study have four eyes”), it generated an illustration of an animated character with four eyes,

highlighting its ability to translate textual prompts into inventive visual interpretations (see Figure 3.c).




Figure 3: Literal and metaphorical representation of idioms

Generally, generating accurate visual representations of proverbs requires several iterations to refine the prompts. The process involves natural language processing and semantic interpretation to convert the meaning of the prompt into visual elements, followed by post-processing to improve image quality and achieve a polished final result.

The results presented above adopts a multimodal perspective on the comprehension of idiomatic expressions. While the first experiment evaluates the ability of LLMs to recognize, complete, and interpret multiword expressions (MWEs) textually, the second experiment investigates how visual generative models represent the same figurative meanings through images. The two experiments therefore address complementary aspects of the same linguistic phenomenon—verbal versus visual representation of figurative semantics. Their joint inclusion aims to explore the extent to which textual understanding and visual generation align to capture idiomatic meaning, offering insight into cross-modal coherence and cultural representation in AI systems.

4 Recurring Motifs Identified

Based on the motifs introduced in Section 2, the idioms were grouped into five broader thematic classes to highlight recurring cultural and conceptual patterns: Human relationships, Work/Effort, Wisdom/Value, Nature, and Luck/Destinies. This thematic consolidation allowed for a clearer interpretation of how idiomatic expressions reflect collective values and cultural perceptions in Bessarabian folklore.

Human relationships – 60 idioms (e.g. family, friendship, betrayal).
Work/effort – 32 idioms (sowing, harvesting, hard work).

Wisdom/value – 25 idioms (truth, wisdom, judgment).

Nature – 6 idioms (nature, seasons, animals).

Luck/destinies – 4 idioms (luck, success, failure).

Bessarabian idioms capture a collective world in which relationships between people matter more than the individual, and also work and land give meaning to life. For people, folk wisdom serves as an ethical code, while at the same time they are anchored to nature, but also aware of hazard.

5 Comparing ChatGPT vs. DeepSeek Performance

Both models show a similar rate of success in completing the idioms, with ChatGPT having a slight edge in overall plausible correctness (55% vs 52%) and consistency (complete all idioms). Both also frequently provide incorrect continuations or alternative proverbs.

Summary for Figurative Meaning Recognition: ChatGPT demonstrated higher accuracy (81% vs 74%) in recognizing and articulating figurative meanings. It was generally more consistent and provided more detailed and nuanced explanations directly relevant to the original idiom's meaning.

Summary for English Equivalents: Both models are capable of providing both literal and figurative English equivalents. ChatGPT appears to have a slight advantage in offering a wider variety of truly idiomatic expressions that closely match the figurative sense of the Bessarabian idioms, while DeepSeek sometimes offers equivalents that lean towards related but not always perfectly aligned concepts.

5.1 Patterns of Success (Common to both models)

1. **Exact or near-exact recall:** Both models achieved a respectable number of precise completions, especially for idioms that are relatively short or widely recognized (e.g., “Sacul gol nu stă în picioare”, “Bate fierul, cât e cald”).

2. **Strong semantic understanding:** Even when the exact wording of the completion was off, both models (DeepSeek more consistently) demonstrated a good conceptual understanding of the idiom's figurative meaning in their explanations and suggested alternatives.

3. **Recognition of idiomatic structure:** The models were generally adept at recognizing the idiomatic nature of the prompt and attempting to generate a completion that maintained an idiomatic style, rather than just a literal or grammatically correct continuation.

5.1.1 Patterns of Errors (common to both models, with some distinctions)

1. **Semantic substitution:** Replacing specific words in the idiom with synonyms or similar terms, which changes the exact wording of the traditional idiom but might preserve a general sense (e.g., “culege” vs. “va secera” by ChatGPT; “miez bun” vs. “rod” by DeepSeek).

2. **Addition of extra information:** Expanding the idiom with additional descriptive phrases or clauses not found in the original, indicating a generative

tendency to complete rather than recall precisely (e.g., ChatGPT’s “primăvara, dar unii rodesc toamna” for “Mai toți copacii înfloresc, dar nu toți rodesc”; DeepSeek’s “dar tare de-l doboară piatra” for “Slab de-l bate vîntul”).

3. **Completely different endings/idioms:** A significant error where the model provides a completion that bears no resemblance to the original idiom’s traditional ending, or even shifts to an entirely different idiom (e.g., ChatGPT’s “Astăzi ești pe val, mâine ești la fund” for “Astăzi ești, mine nu ești”; DeepSeek’s “Capra face boi” for “Capra face, da oaia trage”).

4. **Opposite meaning:** A severe error, exemplified by ChatGPT completing “A trăi ca peștele în apă” with “A trăi ca peștele pe uscat”.

5. **Incomplete or descriptive completions:** DeepSeek, in a few instances, provided an incomplete idiom or a descriptive phrase rather than a proper completion (e.g., “Firele când se încurcă.” or “După răs, plîns.”).

6. **Misdirected Explanation (Primarily ChatGPT):** ChatGPT occasionally explained the figurative meaning of its own incorrect completion or an unrelated idiom instead of the original idiom it was prompted with. This indicates a deeper misunderstanding or lack of self-correction when its initial completion was wrong (e.g., explaining “A scoate castanele din foc” when the prompt was “A scoate ochii (cuiva)”). DeepSeek generally avoids this, usually providing the correct explanation for the original idiom even with an incorrect completion.

6 Conclusion

In conclusion, both ChatGPT and DeepSeek demonstrate considerable capabilities in processing and interpreting Bessarabian idioms, which are rich in cultural and linguistic nuances. However, based on this analysis:

- **ChatGPT** generally performed slightly better in all categories evaluated. It showed marginally higher accuracy in completing the idioms and notably better performance in recognizing and explaining their figurative meanings, tending to stay closer to the established or most common interpretations.
- **DeepSeek** also performed commendably, but had a slightly lower overall accuracy in explanations and a tendency to sometimes explain its own divergent continuations rather than the intended meaning of the original idiom. Its range of idiomatic English equivalents was also sometimes less varied or precise compared to ChatGPT.

This suggests that while both LLMs are powerful, ChatGPT might have a more robust understanding or a more comprehensive training dataset for less common idiomatic expressions and their precise cultural interpretations. Overall, the two experiments collectively highlight the challenges of representing figurative meaning across modalities. By examining both textual and visual realizations of MWEs, this study outlines a foundation for future multimodal idiom research,

where linguistic comprehension and image-based depiction are analysed as interconnected components of figurative understanding.

AI Use Statement

The main objective of this article, was to analyse how Large Language Models (LLMs) understand Multiword Expressions (MWEs), as well as to evaluate the ability of AI generative tools to produce accurate representations of these MWEs. Therefore, we actively used artificial intelligence throughout the study, for this purpose. We also confirm that this article was not generated by AI tools.

Acknowledgements

This research has been supported by the project 25.80012.5007.79SE “Expressing Cultural Identity Through the Digital Transformation of Native Idiomatic Expressions Using AR and AI Technologies, with Implementation in Education.”

References

- API Platform. 2025. Deepseek v3.2. Accessed: 2025-08-11.
- Gheorghe Botezatu *et al.* 1981. *Folk Creation. Proverbs and Sayings*. Editura Științifică și Enciclopedică, Bucharest, Romania. (In Romanian: Creație populară. Proverbe și zicători), p. 309.
- Jisu Kim, Youngwoo Shin, Uiji Hwang, Jihun Choi, Richeng Xuan, and Taeuk Kim. 2025. *Memorization or reasoning? Exploring the idiom understanding of LLMs*. arXiv preprint, arXiv:2505.16216v1. Accessed on 16.09.2025.
- OpenAI. 2025. Chatgpt. Accessed: 2025-08-11.
- Tudor Bumbu, Liudmila Burtseva, Svetlana Cojocaru, Alexandru Colesnicov, and Ludmila Malahov. 2023. *Digitization of Romanian Historical Printings*. Valinex, Chișinău, Moldova.
- Vlada Colesnicova, Fron Argentina, and Daniela Caganovschi. 2025. Comparative evaluation of ai-powered image generation platforms. In *Proceedings of the MITRE 2025 Conference*, Chișinău, Republic of Moldova, June 26-29. Section 5: Didactics of Mathematics.

FOLKAI: A DIACHRONIC CORPUS OF BESSARABIAN FOLKLORE TEXTS

OLESEA CAFTANATOV, SVETLANA COJOCARU,
LUDMILA MALAHOV, MIRCEA PETIC

*“V. Andrunachievici” Institute of Mathematics
and Computer Science, Chisinau,
Republic of Moldova*

{olesea.caftanatov, svetlana.cojocarur,
ludmila.malahov, mircea.petic}@math.md

Abstract

Bessarabian folklore, shaped by historical and political transformations, calls for systematic preservation through advanced digital methods. This research addresses that need by developing the first diachronic-parallel corpus of Bessarabian folklore texts from the 19th to 20th centuries. Building on European expertise in digitalizing low-resource and historically underrepresented languages, our approach adapts international best practices to the Bessarabian context. The proposed workflow integrates several steps that will lead to a diachronic-parallel corpus of Bessarabian folklore texts, together with the enrichment of texts through AI-based illustrations. Various image generation tools are analysed in order to identify the most suitable one. The corpus is considered parallel not in a bilingual sense, but because it aligns texts across two distinct writing systems (Cyrillic and Latin), reflecting historical practices of notation. At the current stage, two books comprising over 300 pages have been processed through the HeDy platform. Our outcome will be a diachronic-parallel corpus of at least 5,000 pages (over 100,000 tokens), accompanied by approximately 1,000 illustrations, offering an interoperable and reusable resource for education, cultural tourism, and further research.

Keywords: folklore, parallel diachronic corpus, digitization.

1 Introduction

Folklore represents an essential component of cultural identity and collective memory, reflecting the history, traditions, and values of a community. Bessarabian folklore has its own specific characteristics, being closely intertwined with Romanian folklore, but also having certain influences dictated by the geographical position and history of the region. Bessarabian folklore was also subject to certain ideological distortions, especially during the Soviet period, which, alongside authentic creations, appear in collections from the second half of the 20th century. In our research from the multitude of folkloric creations, we

will limit ourselves to literary ones, which in turn list an extremely wide variety of genres, species, subspecies, etc. The FolkAI Project aims to preserve and valorise all this heritage by creating a diachronic-parallel corpus of Bessarabian folklore texts from the 19th–20th centuries. The initiative employs Artificial Intelligence (AI) tools for digitization, transliteration, annotation and the generation of illustrative materials, thereby supporting both academic research and educational applications.

The purpose of this article is to present the idea of the research initiative and to outline the research direction with both qualitative and quantitative details.

Taking into account the above, we designed the article as follows: we begin by analysing existing achievements in the field for different languages and geographical or historical contexts (Section 2). Section 3 presents details of the FolkAI workflow; Section 4 is dedicated to the HeDy platform that helps us in Optical Character Recognition (OCR), post-editing and transliteration; in Section 5, the logic of the structure of our diachronic-parallel corpus of Bessarabian folklore texts is presented, and finally, in the Section 6, we present details about the experience of usage of AI-powered tools for generative representation of folk text.

2 Related Works

Research on the digitization of historical and folklore texts has increasingly focused on the challenges of heterogeneous data sources, complex orthographic variation, and the recognition of non-standard scripts such as Cyrillic. Within the Romanian context, the HeDy platform (Bumbu *et al.*, 2023) represents a significant step forward by providing OCR and transliteration capabilities for historical Romanian prints, while large-scale corpora such as CoRoLa (Tufiş *et al.*, 2019) and ROCC (Cristea *et al.*, 2021) demonstrate the feasibility of structured linguistic resources for Romanian. However, none of these efforts have addressed the need for a diachronic, regionally focused corpus of Bessarabian folklore, which remains a gap that FolkAI seeks to fill through a combination of text digitization, AI-based analysis, and multimedia enrichment.

Cross-linguistic experiences in the digitization of endangered and low-resource languages provide essential methodological insights for FolkAI. The Irish-language folklore Handwritten Text Recognition (HTR) project (Ó Raghallaigh *et al.*, 2022) and the Scottish Gaelic HTR project (Lamb *et al.*, 2022) show how scribe-specific and language-adapted HTR models can effectively address dialectal variation and orthographic inconsistency. Both projects combined Transkribus-based recognition with incremental manual correction and crowdsourced validation, achieving character error rates below 5%. Their success demonstrates that high-quality digitization of heterogeneous, non-standard manuscripts is achievable through carefully adapted AI pipelines – an approach

directly transferable to the Bessarabian corpus, where typewriting variation and historical spelling likewise complicate OCR and transliteration.

Beyond recognition, several projects offer models for annotation, metadata structuring, and open data publication. The open-access corpus *Bearing a Bag-of-Tales* (Hagedorn and Darányi, 2022) exemplifies transparent research practices by combining the Aarne-Thompson-Uther typology with reproducible data organization and open licensing, thereby enabling interoperability and reuse. Similarly, the *Impresso* project (Ehrmann *et al.*, 2020) integrates heterogeneous sources – oral transcripts, manuscripts, printed texts – through standardized metadata and interoperable linguistic annotation. Complementary studies on the Estonian folklore archives (Sarv and Järv, 2023) employ motif indexing, statistical modelling and clustering to identify cross-cultural narrative patterns and regional adaptations. Together, these efforts reveal how folklore corpora can move beyond preservation toward computational exploration of genre, motif, and thematic structure – principles that guide the FolkAI design.

At the linguistic-processing level, adapting annotation frameworks to non-canonical language has become central to recent scholarship. The adaptation of Universal Dependencies (UD) for German poetry (Dipper and Laarmann-Quante, 2024) illustrates how verse, songs, and archaic syntax require customized segmentation and syntactic modelling while maintaining interoperability with broader UD standards. Similarly, the corpus-based study of proverbs (Kljajević and Šarić, 2025) and the ACL 2025 evaluation of proverb translation (Wang *et al.*, 2025) highlight the importance of systematic normalization and semantic tagging for figurative expressions, enabling better retrieval, alignment, and machine-translation evaluation. These approaches provide conceptual models for FolkAI’s paremiological subcorpus, where idiomatic meaning and pragmatic context must be explicitly annotated across scripts and epochs.

Comprehensive heritage infrastructures such as SAMLA¹ in Norway and CLARIN’s Historical Corpora network exemplify best practices for interoperability and sustainability. SAMLA’s² FAIR-compliant archive integrates Transkribus-based recognition within Goobi workflows and employs a thesaurus-driven metadata system grounded in the Uppsala Folklore Archive taxonomy, facilitating advanced faceted search across legends, ballads, and folk beliefs. CLARIN³, in turn, ensures long-term accessibility through open standards (TEI, FoLiA), CMDI metadata, and reusable annotation pipelines for tokenization, lemmatization, and part-of-speech tagging. Tools such as TEITOK enable corpus querying and real-time XML markup editing, further supporting collaborative research.

¹ <https://dsenetwork.hypotheses.org/>

² <https://samla.no/viewer/tema/tesaurus/A/1/>

³ <https://www.clarin.eu>

Synthesizing these international experiences, FolkAI embraces the shared principles of customized text recognition models, systematic transliteration workflows, and controlled, interoperable annotation schemes. By adapting proven European methodologies to the specific linguistic and cultural features of Bessarabian folklore, we aim to produce a corpus that is not only linguistically precise and culturally faithful but also aligned with global standards of digital scholarship.

In summary, previous research provides both the conceptual framework and the technological scaffolding for FolkAI. By synthesizing lessons from European initiatives in digitization, annotation, and corpus interoperability, the project advances a model suited to the linguistic complexity and cultural depth of Bessarabian folklore. We intend not only to safeguard an invaluable cultural legacy but also to create a scalable, reusable, and internationally connected language resource that bridges the humanities and artificial intelligence.

3 FolkAI Workflow

Based on the results obtained in other relevant projects on similar topics for different languages, countries and historical periods, as well as the experience of our research team over several years in the field of processing old Romanian texts from the 17th-20th centuries written in Cyrillic script, we further propose the workflow of the FolkAI, which has been designed to combine traditional folkloric sources with state-of-the-art digital technologies, ensuring both the preservation and accessibility of Bessarabian folklore. The process unfolds in several stages, each supported by specific tools and methodologies. A visual representation of the workflow is presented in Figure 1.




Figure 1: Workflow of the FolkAI.

The first stage involves the systematic collection of folkloric texts, primarily sourced from books. To ensure representativeness and cultural diversity, the books were carefully selected, and we deliberately expanded our focus to include regional variations - such as northern Moldavian narratives, folklore from the Bălți steppe, the Codru region, and the Soroca area. Equal attention was given to the diversity of folkloric genres, resulting in the inclusion of legends, fairy tales, traditional songs, humorous anecdotes, *doine*, parables, proverbs and sayings,

laments, and ritual exclamations. This multi-regional and multi-genres approach provides a comprehensive foundation for the construction of a balanced corpus of Bessarabian folklore.

In the second stage, the collected materials undergo digitalization and transliteration. Optical character recognition (OCR) and manual post-editing are carried out through the HeDy platform. We intended to work with at least 5000 pages. We began by digitizing two key works of Moldovan folklore: “O creație populară: proverbe și zicători” (A folk creation: proverbs and sayings, Botezatu *et al.*, 1981) and “Folclor din părțile Codrului” (Folklore from the Codru region, Botezatu *et al.*, 1973). Both books were carefully digitized, validated through recognition processes, and subsequently transliterated from Cyrillic to Latin script. More details on the digitization process are described in the next section.

The third stage consists of the development of a parallel diachronic folklore corpus, where the Cyrillic and Latin versions of the texts are systematically aligned. The expected output is around 10000 aligned sentences. The structure of this folklore corpus will be discussed in Section 5.

Building on this resource, the fourth stage is expected to be the expansion of the MolDoRo Treebank⁴ with an additional 5,000 sentences. For this step, several tools were employed to perform tokenization, as well as syntactic and morphological annotation, thereby enriching the corpus with deeper linguistic information (Caftanatov *et. al*, 2025).

Finally, the last but not least step was the integration of a visual dimension through the generation of folkloric illustrations. Using state-of-the-art image generation tools, we expect to create approximately 1,000 images to visually represent motifs, characters, and symbolic elements from the collected texts. Further details about this step are provided in Section 6, where AI-powered tools for generative representation of folk text are discussed. These images would subsequently be integrated into WikiBase, creating an interconnected digital environment where textual and visual folklore data coexist. Through this structured workflow, FolkAI not only preserves folklore in digital form but also creates a multimodal, linguistically rich, and culturally meaningful resource for future research and education.

4 HeDy Platform - OCR and digitization of Romanian Cyrillic texts

In this section, we will briefly describe the structure and functions of the previously developed HeDy platform, which is used as a component of our workflow.

To create the diachronic parallel corpus, we need several tools, most of which are integrated into the HeDy digitization platform. This platform is developed within the "V. Andrunachievici" Institute of Mathematics and Computer Science

⁴ <https://universaldependencies.org/#romanian-treebanks-1>

(Moldova State University, Chişinău) and exists in desktop and web versions. A demo version (with quite rich functionality) is available for open access on the institute's website⁵. The platform provides tools and resources for digitizing Romanian documents written in the Cyrillic alphabet, printed between the 17th and 20th centuries. It includes image processing modules, optical character recognition models, dictionaries used for recognition, modules for transliteration from Cyrillic to Latin script, and dictionaries of exceptions used to correct the transliterated text. A detailed description of the platform can be found in Bumbu *et al.* (2023).

The demo version has some limitations on file types (PNG, JPEG, TIFF) and file size (the limit for a single file is 100 MB; several files with a total size of up to 700 MB can be processed in a single cycle). In our approach, documents from the same historical period, in which the same alphabets are used, are processed in a common cycle; otherwise, the documents are processed in different sessions.

The first step is image preprocessing, which is especially necessary when working with old texts that have been affected by time or improper storage and usage conditions. For image preprocessing, the user can choose between ABBY FineReader, OpenCV, or Scan Tailor, the latter option being applicable only in the desktop version. The HeDy platform offers its most extensive range of features for image preprocessing through FineReader, including the following options: image resolution correction, page orientation correction, image conversion to black and white, noise reduction, and line straightening.

The next step is optical character recognition. The user must choose the model, which will depend on the historical period and will present variations associated with some texts (or typographies) from that period. For the same historical period, there may be several models, especially in the case of transitional alphabets, and the user selects the most appropriate one. Sometimes it may be useful to try several models to find the closest one.

The HeDy platform was trained on different sets of models and dictionaries. Thus, for the 17th century, it has two OCR models trained on over 10,000 training examples and a dictionary of over 4,500 words; for the 18th and 19th centuries, there are 10 OCR models with a dictionary of over 10,000 words; for the 20th century, a single OCR model is sufficient, which also has the richest training resource with over 450,000 words. Given that the FolkAI project is mainly focused on working with texts from the 20th century, there is a greater chance of obtaining an OCR result with a fairly good level of accuracy. However, the accuracy is not 100%, so manual intervention is required to correct errors, which can be done either on the recognized Cyrillic text or on the transliterated text. Support for manual error correction is provided by HeDy.

⁵ digitizare.math.md

The application of the transliteration algorithm is the next step. Please note that this is not a diplomatic transliteration, but rather a transposition into modern Latin script that preserves the original spelling. At the user's request, the text can be adapted to current spelling rules.

As a result, the HeDy platform provides: the preprocessed image, the recognized text, and the transliterated text. The tool has been successfully applied to the digitization and transliteration of various documents from multiple fields, including religious texts, mathematical works, and novels.

Figure 2 shows an example of recognition and transliteration of a folk song - the *doina* “Doi iubiți” (Two Lovers). In the first column, words with manually corrected errors are highlighted in red; these errors typically involve letters that were misrecognized. The second column shows the transliterated text. The first two columns show the result obtained using the HeDy platform, while the last column shows the English translation of the text, performed by AI.

Doi iubiti (OCR - after validation)	Doi iubiți (Transliteration)	Two Lovers (LLM translation)
Doi iubiti ka doi iubiti. / Ka doi zarzări inforiti. Zarzări de ce inforiti. / Ei mai tare se iubiti. Când au dat la logodit. / Părinți nu i-au voit. Ei de mână s'au luat. / S'au dus și s'au înecat. Părinți când au aflat. / S'au dus și i-au înmormântat: Pe fetiță lângă — un râu / Pe băiat lângă — un pârâu Pe mormântu la fetiță — / Crește-un pușor de viță; Pe mormântu la băiat — / Crește-un pușor de brad. Da și bradu tot creștea. / Pe viță o cuprîndea. Căți oameni pe drum trecea. / Tot din viță îmi gusta. Din guriță blestema: / — Blestemați vor fi părinții. Cei ce-au despărțit iubiiții! / Iubirea, are un dar — La început dulce zăhar. / La sfârșit — pelin amar; La început dulce ca miera / La sfârșit amar ca ferea.	Doi iubiți ca doi iubiți. / Ca doi zarzări inforiti. Zarzări de ce inforiti. / Ei mai tare se iubiti. Când au dat la logodit. / Părinți nu i-au voit. Ei de mână s'au luat. / S'au dus și s'au înecat. Părinți când au aflat. / S'au dus și i-au înmormântat: Pe fetiță lângă — un râu / Pe băiat lângă — un pârâu Pe mormântu la fetiță — / Crește-un pușor de viță; Pe mormântu la băiat — / Crește-un pușor de brad. Da și bradu tot creștea. / Pe viță o cuprîndea. Căți oameni pe drum trecea. / Tot din viță îmi gusta. Din guriță blestema: / — Blestemați vor fi părinții. Cei ce-au despărțit iubiiții! / Iubirea, are un dar — La început dulce zăhar. / La sfârșit — pelin amar; La început dulce ca miera / La sfârșit amar ca ferea.	Two lovers, just like two lovers. / Like two blossoming plum trees. And the plum trees bloomed in white. / As their love grew stronger When they wished to be engaged. / Their parents did not give consent. Hand in hand, the two then fled. / And together drowned instead. When their parents heard the tale, / They came to bury them with wail: The girl was laid beside a stream. / The boy was laid by a brook's gleam. On the maiden's grave did grow / A tender little vine, below; On the young man's grave arose / A sapling fir, in silent prose. And the fir kept growing tall. / Till it clasped the vine in thrall. Travelers passing on the way. / From the vine would taste one spray; But from their lips a curse would fall: / "Cursed be the parents all. Who tore the loving hearts apart" / For love has always such a fate — At first, sweet sugar on the plate. / At the end, wormwood's bitter weight. At the start, as sweet as honey. / At the end, as gall — unfunny.

Figure 2: OCR and transliteration of a folk text using the HeDy platform.

With the recognized and transliterated texts available, it becomes possible to create a diachronic-parallel corpus of Bessarabian folklore. The following section will outline the model we will use to develop this corpus of literary folk creations.

5 Corpus Structure

The organization of folkloric corpora requires careful attention to metadata, since these kinds of texts are deeply embedded in linguistic, cultural, and performative contexts. Different archives and digital collections worldwide illustrate diverse approaches to structuring metadata, reflecting both the research aims and the traditions they document.

A brief analysis of some folklore corpora has shown us that their architecture differs, but some elements can be included in the structure of our corpus. Thus, the Archive of the Indigenous Languages of Latin America (AILLA)⁶, founded in 2000, is dedicated to preserving and studying endangered and dormant indigenous

⁶ <https://ailla.utexas.org/>

languages across Latin America. Working in close collaboration with indigenous communities, the archive safeguards linguistic records while making them accessible for both revitalization initiatives and academic research. Its extensive holdings include audio, video, and textual materials such as oral traditions, myths, songs, and conversations. Metadata plays a central role in structuring this collection, with emphasis on language, speaker characteristics (age, gender, role), recording context, transcription and translation status, and genre classification. By bridging scholarly inquiry with community engagement, AILLA provides a vital resource not only for linguistics, anthropology, and related disciplines but also for supporting indigenous speakers in reclaiming and celebrating their cultural and linguistic heritage, fostering broader cross-cultural understanding and collaboration.

The Endangered Languages Archive (ELAR) at SOAS (London)⁷ similarly preserves folklore materials such as songs, narratives, and rituals. Metadata here expands to include precise geographic location (often with GPS coordinates), cultural group affiliation, performance context, type of folklore, and technical details about the recording process.

The Finnish Folklore Archives (SKS)⁸ offer one of the most systematized approaches, with extensive digital collections that include Kalevala-meter poems, folktales, and charms. Metadata fields cover informant information (name, age, village/region), collector details, year of collection, folklore type, and motif classification according to the ATU index.

Analysis of several folklore corpora has allowed us to conclude that their metadata can be grouped into four categories: textual metadata, performer metadata, collection-related metadata, and technical metadata. The textual level metadata often includes information such as the title or identifier of the piece, its genre (for example, tale, song, proverb, legend, riddle, lament, prayer, incantation, or joke), as well as motif indexes like the ATU or Thompson classifications, and thematic keywords such as love, harvest, animals, weddings, initiation, or death. The performer metadata may record the name or anonymized identifier of the performer, their age, gender, social role (storyteller, singer, healer, elder), and dialect or ethnic affiliation. Collection-related metadata provides contextual details such as the name of the collector or researcher, the date of collection, the location (village, region, and ideally GPS coordinates), and the occasion of performance, whether it be a wedding, funeral, ritual, or casual storytelling event. Finally, technical metadata captures the form in which the material was preserved—whether audio, video, or transcription—along with the conventions used in transcription, the status of translation and annotation, and the license or access rights that govern its use.

⁷ <https://www.elar-archive.org/>

⁸ <https://www.finlit.fi/en/archives/>

As we can see, different folklore corpora develop their metadata structures according to their specific research and preservation needs, though many share common requirements. In our case, we took as a basis the so-called "passport of the folklore unit", which contains identification data, a brief description, data on the performance, a description of the performers, and other necessary information approved and used by folklore researchers in the Republic of Moldova. Thus, the metadata in our corpus should be aligned with those in passports. We observe that researcher-folklorists generally need metadata that serves both linguistic analysis and socio-cultural or ethnographic interpretation. For certain entries, these metadata may be partially or fully completed, depending on the information available in the literary source. Thus, textual metadata can usually be taken from the title of the book, its contents or annotation, or from the archive source. The other categories of metadata may or may not be present. We do not exclude the possibility that this classification may be extended or reduced in the future, depending on the specific needs.

6 AI-powered tools for generative representation of folk text

Folklore reflects the values, beliefs, and social structures of a community. Illustrating animated folk characters enables us to observe how these elements influence perceptions of men, women, and supernatural or anthropomorphized figures. In folkloric creations, such animated folk characters often include talking animals, spirits, mythological beings, or enchanted objects. We intend to include these illustrations in the FolkAI corpus, primarily using AI-generated images, supplemented, when possible, by volunteer illustrators. By examining their depictions in stories, songs, proverbs, or rituals, we gain insight into historical social norms and societal expectations. At the same time, visual representations not only make these patterns more memorable and engaging, but also reveal traditional decorative elements - such as embroidered curtains, clay jugs, village houses, or wooden fences - that enrich our understanding of everyday life, material culture, and the aesthetic dimension of folklore. Through such details, users can better grasp both cultural identities and the social roles embedded within them. In order to highlight the cultural authenticity of Bessarabian folklore, we will, where appropriate, integrate national decorative elements into the AI-generated illustrations - for instance, traditional embroidered textiles, painted pottery, etc. Thus, in such a way, we will enrich the visual narrative by connecting characters and scenes to the material and aesthetic heritage of the region.

To create these illustrations, it is essential to determine which tools are most suitable from the wide range of image generation platforms currently available. For this purpose, we examined the following tools: ReCraft, Sora, Canva, FreePik, Stable Diffusion, Fusion Brain, Pollo AI, Shedevrum, Supermaker, Google Gemini, ChatGPT Plus, NightCafe, CGDream, OpenArt.ai, and Nano Banana. In their study, Colesnicova *et al.* (2024) present a comparative analysis of current

AI-powered image generation platforms, including Google Gemini, ChatGPT Plus, NightCafe, CGDream, and OpenArt.ai, focusing on image quality, generation speed, economic models, and prompt interpretation accuracy. Results reveal significant differences: Google Gemini excels in speed and free access with consistent visual quality; ChatGPT Plus delivers highly aesthetic images and strong understanding of complex prompts but at slower speed; NightCafe offers fast generation and a hybrid credit model; CGDream struggles with prompt comprehension, often producing unrelated images; and OpenArt.ai provides stylistic versatility and customization, though occasionally with consistency and file management issues. These findings highlight the importance of selecting platforms based on specific needs, balancing quality, speed, cost, and interpretative reliability to ensure effective AI image generation.

We also experimented using ChatGPT, NightCafe, NanoBanana (which runs on Google's Gemini 2.5 Flash Image models), Canva, Freepik, and a human illustrator. Each was given the same prompt in English and Romanian to generate a visual representation of *doina* Doi Iubiți (the text of which is presented in Figure 2) and Figure 3 displays the six resulting illustrations. The given prompt in English was: "A mystical, romantic scene at a lake. On the shore of the lake, a young fir tree and a grapevine are planted side by side. The fir tree and the vine grow together, intertwining as if embracing each other. In the reflection of the lake's surface, the face of a young man can be seen within the fir tree, and the face of a young woman appears within the grapevine. Their faces are gentle, sorrowful, yet full of love, as if they are reaching out to embrace. The atmosphere is poetic and folkloric, with a touch of melancholy, like an old legend. Soft evening light, dreamy colours, mystical aura".



a) Canva, b) NanoBanana, c) FreePik, d) NightCafe, e) ChatGpt, f) Illustrator Doina Seiciuc

Figure 3: Visual representation of the poetry “Two Lovers” by AI tools and a human illustrator according to the English prompt

Correspondingly, the text in Romanian was as follows: "O scenă mistică, romantică, la marginea unui lac. Pe țărmul lacului, un brăduț tânăr și o viță-de-vie sunt plantate unul lângă celălalt. Bradul și vița cresc împreună, împletindu-se ca și cum s-ar îmbrățișa. În oglinda apei se reflectă chipul unui tânăr în ramurile bradului, iar în vița-de-vie se conturează chipul unei tinere. Fețele lor sunt blânde, triste, dar pline de iubire, ca și cum ar încerca să se atingă într-o îmbrățișare.

Atmosfera este poetică și folclorică, cu o notă de melancolie, asemenea unei legende vechi. Lumina blândă a serii, culori visătoare, aură mistică."




Figure 4: Visual representation of the poetry "Two Lovers" by AI tools and a human illustrator according to the Romanian prompt

Figure 4 does not include the result from NightCafe, which failed to execute the proposed task.

For the comparative analysis, we evaluated the platforms based on the following criteria: *prompt alignment, visual quality, stylistic and aesthetic quality, and technical performance*. Regarding the alignment of prompts with the resulting illustrations we observed that it varies by platform. Nano Banana delivers the highest fidelity, producing images that most closely resemble a surreal fusion of humans and trees. Canva, NightCafe, and ChatGPT offer a moderate interpretation, where both faces and trees are present but with noticeable variations. FreePik tends to prioritize realism, resulting in less imaginative or fantastical outputs. In terms of human interpretation, the illustration created by the artist stood out for its conceptual accuracy and profound expressiveness. When it comes to visual quality, most tools produce high-quality images, with Nano Banana and Canva standing out for their sharpness and attention to detail. The ChatGPT tool, by contrast, tends to create images with a softer, painterly effect that is slightly less crisp. For the human illustrator, the quality ultimately depends on the tools and techniques used; in this case, the resulting illustration is of very high quality.

In evaluating style and aesthetic, each platform brought a distinct artistic approach to the illustration. Nano Banana delivered a dreamy, fantasy-inspired visual, while Canva leaned toward surreal realism, evoking a misty and moody atmosphere. The ChatGPT tool produced images reminiscent of classic oil paintings or traditional illustrations. FreePik and NightCafe both favoured a more romantic, realistic, and cinematic style. The human illustrator chose a striking black-and-white aesthetic to mirror the sombre tone of the poem. By intertwining a fir tree with a grapevine, the artist used the dry branches to symbolize the lovers' death, while the living branches conveyed the eternal nature of their love. From a technical perspective, the integration of faces and trees was most successfully

achieved by Nano Banana. NightCafe exhibited minor artifacts and unnatural blending, while Canva showed only minimal issues in this regard. The human illustrator, however, demonstrated technical mastery by flawlessly rendering the faces of the man and woman, seamlessly weaving them into the branches of the fir tree and grapevine.

While each platform excels in different aspects, the choice between human and AI illustration ultimately depends on practical considerations such as time and resource availability. Even if the human illustration is more subtle and, in our opinion, more harmoniously reflects the spirit of the *doina*, we must nevertheless take into account the fact that human illustrators typically require a considerable amount of time to produce such refined work, whereas AI tools can generate images within a few minutes. Our comparison of the five AI platforms reveals that, with precise prompting, AI tools can produce results that closely approach the quality and expressiveness of human illustrators' work. In our analysis, we observed that different AI tools are particularly well-suited to specific genres and types of folkloric creations. For example, ChatGPT excels at creating caricatures, making it an excellent choice for humorous folk tales, or satirical stories. Nano Banana shines for legends, myths, and magical folk tales, especially those featuring transformations, enchanted forests, or supernatural elements. FreePik and NightCafe are ideal for realistic folk ballads, moral tales, or stories set in rural landscapes, offering a romantic, cinematic touch. Canva excels with lyrical folk poetry, tragic love stories, or emotionally charged tales, thanks to its moody, atmospheric surreal realism. Thus, for our research, we have decided to utilize all the AI tools we tested, and we may even expand our study to include additional AI platforms. We also remain open to volunteer contributions from human artists whenever possible.

7 Conclusions

In this paper, we introduced the FolkAI corpus, the first diachronic-parallel digital collection of Bessarabian folklore, developed to serve both academic and cultural purposes. Its construction builds upon the tools and resources previously utilized in creation of the HeDy platform and the MolDoRo Treebank, enriched with components specifically designed for the processing of folkloric texts. To enhance accessibility and engagement, at least 1000 illustrations will accompany the corpus entries, making the content more visually informative. In addition, we analysed a variety of existing folklore corpora and examined their metadata structures. Our study revealed that metadata standards vary considerably across different research teams; however, several common elements can be grouped into four main categories described in section 5. In future work, we plan to perform a more detailed analysis of these metadata schemes and adapt them to the specific needs of the FolkAI corpus, ensuring better alignment with the "passports" of literary folklore used by researchers in the field.

AI Use Statement

Generating images using AI is one of the subjects of study in this article; the tools used to illustrate an example are listed in section 6. Additionally, AI technologies were used to enhance readability and language only. AI tools were not used to generate text, analyse, or interpret data.

Acknowledgements

The research for this paper was supported by the project 25.80012.0807.50SE, “Digitization and creation of the FolkAI corpus with texts of Bessarabian folklore from the 19th–20th centuries”.

References

- Brian Ó Raghallaigh, Andrea Palandri, and Críostóir Mac Cárthaigh. 2022. Handwritten Text Recognition (HTR) for Irish-Language Folklore. In Proceedings of the 4th Celtic Language Technology Workshop within LREC2022, Marseille, France. European Language Resources Association, pp 121–126.
- Cătălina Mărânduc. 2022. *Instruments for Processing the Old Romanian Language*. PhD thesis, Iași, 2022.
- Dan Cristea, Cristian Păduraru, Petru Rebeja, Andrei Scutelnicu, and Mihaela Onofrei. 2021. *Data Structure and Acquisition in DeLORo – a Technology for Deciphering Old Cyrillic-Romanian Documents* In: Proceedings of the 16th International Conference on Linguistic Resources and Tools for NLP, online, 13-14 december 2021, pp. 59-74.
- Dan Tufiş, Verginica Barbu Mititelu, Elena Irimia, Vasile Păiş, Radu Ion, Nils Diewald, Maria Mitrofan, and Mihaela Onofrei. 2019. *Little Strokes Fell Great Oaks. Creating CoRoLa, The Reference Corpus of Contemporary Romanian*. Revue roumaine de linguistique, No./Issue 3. pp. 227-240.
- Grigore Botezatu. 1973. Folklore from the Codru region (in Romanian with Cyrillic alphabet), Chisinau, 362 p.
- Grigore Botezatu. 1981. A folk creation: proverbs and sayings. (in Romanian with Cyrillic alphabet), Chisinau, 309p.
- Joshua Hagedorn and Saandor Darányi. 2022. *Bearing a Bag-of-Tales. An Open Corpus of Annotated Folktales for Reproducible Research*. In: Journal of Open Humanities Data 8,16 (2022), DOI: <http://doi.org/10.5334/johd.78>.
- Mari Sarv and Risto Järv. 2023. *Layers of Folkloric Variation: Computational Explorations of Poetic and Narrative Text Corpora*. Folklore: Electronic Journal of Folklore, 90, 2023, pp. 233–266. DOI: 10.7592/FEJF2023.90.sarv_jarv.
- Maud Ehrmann, Matteo Romanello, Simon Clematide, Phillip Benjamin Ströbel, and Raphaël Barman. 2020. *Language Resources for Historical Newspapers: the Impresso Collection*. In Proceedings of the Twelfth Language Resources and Evaluation Conference, Marseille, France. European Language Resources Association, pp. 958–968.
- Minghan Wang, Viet-Thanh Pham, Farhad Moghimifar, and Thuy-Trang Vu. 2025. *Proverbs Run in Pairs: Evaluating Proverb Translation Capability of Large Language*

- Model*, In: Findings of the Association for Computational Linguistics: ACL 2025, July 27- August 1, 2025, pp. 1646–1662.
- Mircea Petic, Olesea Caftanатов, Svetlana Cojocar, Ludmila Malahov, and Atul Kr. Ojha. 2025. *Digitizing of Bessarabian Folklore: From Textual Heritage to Digital Corpus*. In: Proceedings of the Workshop on Intelligent Information Systems WIIS2025, October 16-18, 2025, Chisinau, Republic of Moldova (in print).
- Stefanie Dipper and Ronja Laarmann-Quante. 2024. *UD for German Poetry*. In Proceedings of the 4th International Conference on Natural Language Processing for Digital Humanities, Miami, USA. Association for Computational Linguistics, pp. 177–188.
- Tudor Bumbu, Lyudmila Burtseva, Svetlana Cojocar, Alexandru Colesnicov, and Ludmila Malahov. 2023. *Digitization of Romanian Historical Printings*, Chisinau: Vladimir Andrunachievici Institute of Mathematics and Computer Science, Moldova State University, 176 p. ISBN: 978-9975-68-497-2.
- Vanja Kljajević and Ljiljana Šarić. 2025. ‘*Corpus-Based Investigation of Proverbs: Challenges and New Directions*’ (2025) 9(2) *Corpus pragmatics: international journal of corpus linguistics and pragmatics* 135.
- Vlada Colesnicova, Olesea Caftanатов, Svetlana Cojocar, Alexandru Colesnicov, and Ludmila Malahov. 2024. *Digitization of Moldovan Folklore Texts using the HeDy platform*. In: Proceedings IMCS-60, October 10-13, 2024 Chisinau, pp. 272-275. ISBN 978-9975-68-515-3.
- William Lamb, Beatrice Alex, and Mark Sinclair. 2022. *Handwriting recognition for Scottish Gaelic*. In Proceedings of the 4th Celtic Language Technology Workshop within LREC2022, Marseille, France. European Language Resources Association, pp. 60–70

THE MAPPING GREETING FORMULAE IN THE HISTORICAL MOLDOVA PROJECT: AN OVERVIEW

ANCA-DIANA BIBIRI¹, MIHAELA MOCANU¹,
LIVIU-ANDREI SCUTELNICU², LIVIU PETCU¹, VIORICA POPA³,
ELENA LACUSTA³, ANGELA BEJAN³

¹*Institute of Interdisciplinary Research, "Alexandru Ioan Cuza" University of Iași*
{anca.bibiri, mocanu.mihaela, liviu.petcu}@uaic.ro,

²*Faculty of Computer Science, "Alexandru Ioan Cuza" University of Iași*
andrei.scutelnicu@info.uaic.ro

³*"Alecu Russo" Bălți State University*
{viorica.popa, elena.lacusta, angela.bejan}@usarb.md

Abstract

The paradigm of greetings and farewells is dynamic, as social relations significantly influence it, with changes affecting these relations and prevailing norms of politeness. Studying greetings requires a sociolinguistic approach, since extralinguistic factors guide their selection. A speaker cannot communicate without referencing the addressee. Both the greeter and the respondent must adapt to the context's parameters to behave appropriately. Moreover, greetings and farewells reflect realities such as tradition, religion, lifestyle, and outlook on life, which persist as long as they are imbued with community significance. Through this project, we aim to expand the greetings inventory by including a wider range of contexts and expressions. Using a comprehensive sociolinguistic questionnaire, the project will document everyday greeting and farewell formulae as well as those used during religious holidays and important life events (such as birth, baptism, engagement, marriage, and periods of mourning, including death, vigils, and funeral ceremonies). The Greetmap_Ro-Md project seeks to develop a linguistic resource (an annotated corpus) of greetings from the territory of historical Moldova (on both sides of the Prut River). This database will include greetings annotated with relevant sociolinguistic variables from various perspectives: diatopic (both sides of the Prut River), diastratic (social, cultural, or educational differences), diaphasic (stylistic variations), and diamesic (communication contexts). Such a resource will map the formal and informal aspects of Romanian greeting formulae across different contexts and periods, offering a nuanced view of how cultural and situational factors influence their usage. The result will be a valuable contribution to understanding the sociolinguistic landscape of greetings and address forms in historical Moldova,

with potential applications in cross-cultural studies and language technology. This linguistic resource will serve as an essential tool for specialists and provide substantial material for research and studies in the field.

Keywords: Greeting formulae, Romanian language, linguistic resource, cultural heritage

1 Introduction

Language continuously develops as societies change. While the phonological and grammatical systems of a language develop very slowly, changes in politeness systems and behaviours usually happen more quickly, reflecting shifts in the socio-cultural environment. The wave of global change has already wiped out many customs and cultural practices, due especially to Western European and American influences. The extended family is weakening, rituals and customs are observed less seriously, and traditional practices are becoming increasingly rare, passing from one generation to the next. In line with these modern societal changes, politeness behaviour and the related set of expressions are also evolving through reduction, simplification, and secularisation. Researchers have observed that in our contemporary world, older individuals (particularly in villages) tend to use politeness formulae more often than the younger generation in daily interactions or special occasions, such as life events or religious holidays. Caused by urbanisation, secularisation, migration, and exposure to Western cultures, younger speakers are more likely to adopt a simpler approach to politeness behaviour, characterised by fewer formulas, less complex patterns, fewer honorific terms, fewer religious-bound expressions, and a preference for more egalitarian or simply ‘cool’ forms (Zafiu, 2011).

As human beings, we encounter various situations in the social world, as demonstrated by our recorded speech acts. As Spolsky (1998) considers: “The existence of language variation, therefore, is not accidental or meaningless. It adds a vital set of social dimensions, making it possible for language to reflect and record an individual’s demographic, geographic, sociological, educational and religious background. It helps constitute identity; it claims solidarity; it expresses attitudes towards power and prestige.” Social variables such as relative status, age, and gender influence language choices and affect one’s politeness behaviour. These variables interact in complex ways, along with other factors such as the formality of the situation and the personal relationship between the speaker and the addressee.

Although regarded as relatively fixed linguistic structures, standardised formulae, polite conventional phrases, and mandatory elements in specific communication contexts, greetings and farewells remain a vital aspect of verbal behaviour, carrying deep social, cultural, and historical significance. As expressions of politeness, greetings and farewells serve to open and close nearly any communication; their selection and use are guided by generally accepted

norms within a particular linguistic community. Violating these norms affects not only the communication itself but also the social relationships between speakers. Moreover, greetings and farewells are part of a textual tradition that extends beyond individual languages.

2 State of the Art

Research on greetings and farewells requires a sociolinguistic approach as extra-linguistic factors influence their choice. A speaker cannot communicate without considering their addressee. Both participants in a greeting exchange must adapt to the context's parameters to understand the nuances of the situation, such as social distance, formality, and context, so that their behaviour aligns with the cultural 'rules' of interaction. Conversely, greeting and farewell formulae reflect realities as tradition, religion, lifestyle, and perspectives on life, which remain meaningful within the community. The paradigm of greetings and farewells is dynamic, as it is influenced by social relationships, changes impacting these relationships, and current norms of politeness (Duranti, 1997; Coulmas, 2005).

These changes stem from internal sources (modifications of traditional greetings and farewells, their recombination, or functional shifts in words and phrases) and external sources (borrowings from other languages). Greeting and farewell formulae evolve in structure and meaning, acquiring new semantic and functional values driven by societal changes. Consequently, some greetings and farewells transcend their original relational context and continue to be used in different circumstances from those originally intended (Pietreanu, 1984).

Located in southeastern Europe, at the crossroads of pathways to the West and East and threatened from all sides, Romania has remained a keeper of cultural traditions and values throughout its tumultuous history, which have been passed down and inherited over the centuries. Apart from the region it belongs to (the Balkans), what mainly characterises it is its ethnic and cultural diversity, incorporating elements from Central and Eastern Europe with Latin, Slavic, Hungarian, Germanic, and Turkish influences. Although Romanian offers a rich and diverse array of greeting formulae, research on this subject remains limited. Some aspects are sporadically discussed in dialectology, and more recently in stylistic research. Linguistic atlases include isolated entries related to greeting and farewell formulae, while Romanian grammar and linguistic treatises address this issue only peripherally. The first comprehensive collection of politeness formulae, along with their structural and functional analysis, was presented in 1983 by Tatiana Slama-Cazacu. Methodologically, approaches to Romanian greetings and farewells include: psycholinguistic, which explores the cognitive and psychological processes behind how people understand, produce, and interpret these expressions, often highlighting how social context influences language use (Slama-Cazacu 1983; 1985); stylistic, focusing on the stylistic and rhetorical features of greetings and farewells, analysing the varied forms and expressions in

different settings, such as formal versus informal contexts (Zafiu, 2003); pragmatic, which investigates language use in relation to social norms and intentions, studying how greetings and farewells function as tools for social interaction and relationship management (Mladin, 2003; Constantinescu, 2018); semantic, examining the meanings and nuances, including degrees of politeness, familiarity, or affection, of different greeting and farewell formulae (Polloni, 1993); contrastive or comparative analyses with other Romance languages, positioning Romanian greetings and farewells alongside those of other Romance languages to identify similarities, differences, and cross-linguistic influences that highlight the unique aspects of Romanian social expressions within the Romance language family (Mladin, 2003; Aioane, 2005); and sociolinguistic, as noted by Pietreanu (Pietreanu, 1983).

The volume *Politeness in Europe* (Hickey and Stewart, 2005), which provides a broad picture of the concept across different countries, does not include Romania. The chapter ‘Eastern Europe’ includes only Estonia, Hungary, Poland and the Czech Republic, but they are not treated as greeting formulae. Following Hickey and Stewart’s example, Rebecca Leugner, in her study *Politeness in Romania. Address Forms* (2007) examines the use of address forms, especially the pronouns of address in the Romanian language.

Any analysis of greetings is usually linked to concepts such as politeness and decency, which sociolinguistics associates with a human need for social interaction. The concept of politeness aligns with the ‘interaction norms’ in sociolinguistics, rooted in American ethnology and sociology (Hymes and Gumperz, 1972; Brown and Levinson, 1978). Guided by generally accepted norms within a community, social interactions are shaped by the image speakers wish to promote about themselves (Goffman, 1967). In this particular context, the pragmatic values characterising greetings become complex and subtle: on one hand, greeting formulae enhance the speaker’s image, while on the other, they aim to establish positive communicative interaction with the addressee, achieving a form of *captatio benevolentiae* intended to lay the groundwork for effective communication.

3 Scientific Relevance of the Project

Greeting and address formulae play an important role in shaping community relationships, the social status or rank of interlocutors, and their identities. Researchers show that these formulae reflect realities such as tradition, religion, lifestyle, perspective on life, and so on. The proposed research aims to systematically map the greeting formulae and address forms used in historical Moldova, creating a valuable linguistic resource that enhances understanding and use of the Romanian language within the context of European culture, where other Eastern European countries are better represented. The research will make significant contributions to linguistic studies and fill the gap in research dedicated

to this topic, representing Romanian within the broader context of politeness studies from an interdisciplinary perspective. We aim to build, through this project, an electronic tool that is easily accessible for users interested in these formulae. In our previous projects (AMPROM and SoRoEs), we recorded formal and informal greetings in a limited number of contexts, considering three sociolinguistic factors: age, gender, and level of education (Mocanu and Bibiri, 2019). Additionally, we analysed a rich inventory of questions (particularly partial interrogative questions/WH-Q) used as greetings, either alone or, more commonly, accompanying proper greetings (Bibiri *et al.*, 2016). A statistical approach to this mini-corpus of recorded greetings was undertaken in a previous study that aimed to map variations in commonly used greetings across Romanian regions (Bibiri *et al.*, 2021). All these form the foundational basis for the proposed research theme to explore further specific sociolinguistic factors, especially with data from broader contexts.

Through this project, we aim to expand the greetings inventory by including a wider range of contexts and expressions. Using a comprehensive sociolinguistic questionnaire, the project will document everyday greetings and farewells, as well as those used during religious holidays and significant life events (such as birth, baptism, engagement, marriage, and periods of mourning, including death, vigil, and funeral). With an experienced interdisciplinary team (comprising specialists such as linguistics, pragmatics, semiotics, and computational linguistics), we will succeed in creating a valuable linguistic resource. Collected data will map the formal and informal aspects of Romanian greetings across different contexts and periods and provide a nuanced view of how cultural and situational factors influence the use of these formulae. The outcome will be a meaningful contribution to understanding the sociolinguistic landscape of greetings and address forms in historical Moldova, with future applications for cross-cultural studies and language technology.

4 Project Objectives

The project aims to develop a linguistic resource, an annotated corpus, that records and analyses greeting and address formulae in the historical region of Moldova, covering areas on both sides of the Prut River. This corpus will be an essential tool for experts and a valuable database for advancing research on the Romanian language. Additionally, it will serve as useful didactic material for teaching, studying, and using the Romanian language.

To achieve the aim of the Greet Map project, the following objectives are considered:

O.1: Building a corpus that includes greeting and address formulae used in different contexts and periods of the year. This includes everyday greetings and those related to religious holidays and significant life events, which reflect the linguistic richness of historical Moldova.

O.2: Annotation of recorded formulae through the questionnaire applied across the historical region of Moldova (in the two countries). The annotation phase will classify and tag the greeting and address formulae according to context, sociolinguistic factors, and linguistic features.

O.3: Interdisciplinary approach of greeting and address formulae used in different contexts. A multi-disciplinary analysis will be conducted, leveraging different approaches from sociolinguistics, pragmatics, and semiotics to explore greetings variations across different contexts in historical Moldova.

O.4: Interactive map with digital corpus in historical Moldova. The project will provide users with access to the annotated corpus, enabling a visual exploration of greeting practices across historical Moldova. This map will make the data accessible and engaging for researchers, educators, and students.

Through this project, we develop an inventory and a database of greeting and address formulae, annotated with sociolinguistic variables considered from diatopic (geographical differences), diastratic (social, cultural, or educational differences), diaphasic (stylistic or register variations), and diamesic (context of communication) perspectives. This approach considers that speakers reveal their life and identity in their responses, alongside their sociocultural, economic, and geographical backgrounds across time and space. The GreetMap corpus will be analysed at various levels: lexical (formal versus informal, urban versus rural), morpho-syntactic (POS tagging annotation), and semantic-pragmatic (using NVivo software).

5 Methodology

The proposed project aims to develop a lexical resource focused on greeting formulae across different socio-cultural contexts. It also seeks to identify the social and cultural patterns behind Romanian speakers' greeting choices within the framework of sociolinguistics (considering various variables in the communication context), pragmatics (taking into account social influences on language use during interaction), and semiotics (relating to linguistic expression, gestural structures, ceremonies, etc.). Data collection methods involve an online sociolinguistic questionnaire, followed by analysis methods that employ statistical approaches to identify speakers' strategies in using recurring greeting formulae linked to sociolinguistic variables. The questionnaire comprises open-answer items depicting communication situations; multiple-choice questions for quick demographic analysis, including personal background data (age, gender, degree, geographic origin, religious affiliation, etc.). Its structure is informed by the review of bibliographic literature and automatic lexical extraction of greeting formulae from The Reference Corpus of Contemporary Romanian – CoRoLa (<https://corola.racai.ro>).

The CoRoLa corpus, developed under the auspices of the Romanian Academy, includes texts published after 1945 and serves as the reference corpus for

contemporary Romanian. It is the largest one, containing 1,257,752,712 tokens for the written component and nearly 152 hours of recordings for the oral component (Barbu Mititelu *et al.*, 2018; Tufiş *et al.*, 2019). The texts cover all language styles across four major domains with 71 subdomains, ensuring broad and representative vocabulary coverage. CoRoLa is open for research and is accessed using the KorAP platform (<https://korap.racai.ro/>) for corpus interrogation.

The research will be conducted through several consecutive stages. In the initial phase, we will review relevant literature in the field to establish the theoretical and methodological framework of the study. Following this, an inventory of greeting and address formulae will be compiled through automatic extraction from the CoRoLa corpus. In the subsequent images, we present two examples of interrogation results using different queries: a wedding wish as *Casă de piatră!*




Figure 1: Interrogation result for the wedding wish *Casă de piatră!*

Another example is a religious greeting used during the Easter period: *Hristos a Învia!*




Figure 2: Interrogation result for the religious greeting *Hristos a înviat!*

Using this inventory, we will design the questionnaire for online administration. After drafting, the questionnaire will be pre-tested in the two geo-linguistic regions. The pre-test will help us identify potential ambiguities, inconsistencies, or shortcomings. Afterwards, the questionnaire will be revised and improved accordingly. These steps mark the first phase of the project's implementation.

Next year, we will develop the corpus and annotate the collected data. The improved version of the questionnaire will be distributed online to participants from the two geo-linguistic regions. Special attention will be given to ensuring that the sample respects relevant sociolinguistic variables, thus enabling a representative and comparative analysis of greeting and address formulae across the two contexts.

Further, we will focus on the sociolinguistic analysis of the greeting and address formulae used in everyday life, during major life events and religious holidays. A digital corpus is depicted on an interactive map to identify specific patterns in the diatopic distribution.

6 Preliminary Results

Greetings represent one of the most universal and socially significant forms of linguistic interaction. Far from being merely formulaic or routine, greeting exchanges encode complex social meanings while reflecting speakers' identity, relationships toward politeness and formality. In sociolinguistic research, greeting practices have been shown to index broader social structures and value systems, functioning as sites where power relations, gender roles, and generational identities are negotiated and reproduced (Mills, 2005). Despite their apparent

simplicity, greeting exchanges encapsulate intricate patterns of social meaning and provide valuable insight into the norms governing interpersonal communication (For an overview, see also Bibiri *et al.*, 2016; Mocanu, 2018; Mocanu & Bibiri, 2019; Bibiri *et al.*, 2021).

The preliminary analysis reveals that the Romanian language exhibits a rich and diverse inventory of greeting formulae, which reflect complex sociolinguistic patterns. Variation in greeting practices appears to be systematically influenced by social factors such as age, gender, and educational background. Variations across age groups are particularly pronounced: younger speakers tend to use shorter, more innovative, and informal greeting expressions, frequently drawing on elements from popular culture, digital communication, or slang. Their greetings display a tendency toward linguistic economy and social creativity, reflecting shifting norms of informality and solidarity. By contrast, older speakers have a preference for more conventional and formulaic greetings characterised by formality and politeness, and adhere closely to traditional norms of social interaction. This contrast suggests that age represents a key sociolinguistic variable, marking different attitudes toward linguistic innovation and politeness conventions.

Gender also emerges as a relevant sociolinguistic variable, reflecting gender-specific patterns of communication and social expectation within the speech community. This pattern shows broader gendered norms in language use, where men's greetings tend to emphasise solidarity and brevity. In contrast, women's greetings often display greater attentiveness to politeness and interpersonal warmth.

Furthermore, educational background is an important factor influencing greeting behaviour. Speakers with higher education display greater awareness of politeness norms, often opting for more elaborate and contextually appropriate formulae.

Thus, the sociolinguistic variables underline the dynamic relationships between language, identity and social structure, as greetings reflect and reinforce the sociocultural organisation of everyday interaction.

7 Impact of the Project

The project GreetMap_Ro-Md intends to develop a linguistic resource — an annotated corpus — of greetings from the territory of historical Moldova, encompassing areas on both sides of the Prut River. The database includes greetings, annotated with relevant sociolinguistic variables, from a diatopic (covering both sides of the Prut), diastratic (reflecting social, cultural, or educational variations), diaphasic (different levels of style), and diamesic (varied communication contexts) perspectives. This linguistic resource will serve as an essential tool for specialists and offer extensive material for future studies and research in the field.

The project GreetMap_Ro-Md also aims to connect Moldovan researchers with international research networks where Romanian scholars actively participate, further strengthening Moldova's presence in the global research community. Moreover, this collaborative effort marks the initial phase for joint project applications in European funding competitions, utilising the strengths of both teams. This will broaden the impact of Romanian and Moldovan research within the wider European academic scene, paving the way for future accomplishments in linguistic and cultural heritage.

AI Use Statement

This study did not use AI-generated input. The originality of the ideas, the critical analysis of the theoretical background, the process of data collection, the methodology used and described, the research implementation, the data validation and interpretation, and the conclusions were carried out by the authors.

Acknowledgements

„This work was supported by a grant of the Ministry of Research, Innovation and Digitisation, CNCS – UEFISCDI, project number 4PCBROMD/2025, PN-IV-PCB-RO-MD-2024-0383, within PNCDI IV”.

References

- Alessandro Duranti. 1997. Universal and Culture-Specific Properties of Greetings. *Journal of Linguistic Anthropology*, 7(1): 63-97, American Anthropological Association.
- Anca-Diana Bibiri, Mihaela Colhon, Mihaela Mocanu. 2021. Statistics of Some Regional Varieties of Romanian Greetings. In *Proceedings of the 16th International Conference on Linguistic Resources and Tools for Natural Language Processing*, edited by Petru Rebeja, Mihaela Onofrei, Dan Cristea, Dan Tufiş, Iaşi, 13-14 December 2021, „Alexandru Ioan Cuza“ University of Iaşi Publishing House, pages 181-198.
- Anca-Diana Bibiri, Mihaela Mocanu, and Adrian Turculeţ. 2016. The intonation of Romanian greetings: a sociolinguistics approach. *International Journal of Social, Behavioral, Educational, Economic, Business and Industrial Engineering*, 10(4): 653–661.
- Bernard Spolsky. 1998. *Sociolinguistics*, Oxford, Oxford University Press.
- Dan Tufiş, Verginica Barbu Mititelu, Elena Irimia, Vasile Păiş, Radu Ion, Nils Diewald, Maria Mitrofan, and Mihaela Onofrei. 2019. Little Strokes Fell Great Oaks. Creating Corola, the Reference Corpus of Contemporary Romanian. *Revue Roumaine de Linguistique*, 64(3): 227-240.
- Dániel Z. Kádár. 2016. Historical intercultural socio-pragmatics: A study on ritualisation. *Journal of Historical Pragmatics*.
- Florian Coulmas. 2005. *Sociolinguistics – The Study of Speakers' Choices*. Cambridge University Press.

- George Caragață. 1940. *Formulele de salutare în limba română*, Iași, Tipografia Alexandru A. Țerek, 17 p.
- Jennifer Coates. 2016. *Women, Men, and Language. A Sociolinguistic Account of Gender Differences in Language*, third edition, Routledge Linguistics Classics.
- Leo Hickey and Miranda Stewart (eds.). 2005. *Politeness in Europe*, volume 127. Multilingual Matters Ltd, Clevedon, Buffalo, Toronto.
- Li Wei. 2010. The Function and Use of Greetings, *Canadian Social Science*, 6(4): 56-62.
- Marica Pietreanu. 1984. *Salutul în limba română. Studiu sociolingvistic*, București, Editura Științifică și Enciclopedică.
- Mihaela Mocanu, and Anca-Diana Bibiri. 2019. Greetings and Farewells in Contemporary Romanian: a sociolinguistic approach. *Zeitschrift für romanische Philologie* (ZrP), 135(3): 866-886.
- Mihaela Mocanu. 2018. The Greeting in the Romanian Language. Gender Differences in the Use of Greetings. *Philologica Jassyensia*, 14(1), 83-95.
- Mihaela-Viorica Constantinescu. 2018. The Historical Variation of the Pragmeme 'Greeting' in Romanian. *Revue Roumaine de Linguistique*, LXIII (1-2): 73-97.
- Mirela Aioane. 2015. Formule de salut în limbile romanice, *Analele Universității "Alexandru Ioan Cuza" din Iași*. Secțiunea IIIe. Lingvistică 51, 15-19.
- Penelope Eckert and Sally McConnell-Ginet. 2003. *Language and Gender. The Social Construction of Meaning*. Cambridge, Cambridge University Press.
- Penelope Eckert. 1997. Age as a Sociolinguistic Variable. In Florian Coulmas (ed.) *The Handbook of Sociolinguistics*, Cambridge University Press.
- Rebecca Leugner. 2009. *Politeness in Romania, Address forms*, Seminar Paper, Grin Verlag, Germany.
- Rodica Zafiu. 2011, Formule de salut. *România literară*, Anul XLIII, p. 15.
- Sara Mills. 2005. Gender and Impoliteness. *Journal of Politeness Research*, I, Walter de Gruyter, p. 263-280.
- Tatiana Slama-Cazacu. 1983. Formule de politețe în limba română: structură și funcții (I). *Studii și cercetări lingvistice*, XXXIV, nr. 3.
- Tatiana Slama-Cazacu. 1985. Conceptul de «bunăcuviință» și formulele care îl exprimă în limba Română. *Anuar de Lingvistică și Istorie Literară*, XXX, 413-419.
- Verginica Barbu Mititelu, Dan Tufiş, and Elena Irimia. 2018. The Reference Corpus of the Contemporary Romanian Language (CoRoLa). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, pages 1178-1185, Miyazaki, Japan. European Language Resources Association (ELRA).

INDEX OF AUTHORS

- Adrian Marius Dumitran, 185, 222
 Alexandr Parahonco, 249, 265
 Alexandra Ilie, 199
 Alina Turculet, 162
 Altanbek Zulkhazhav, 55
 Amalia Todiraşcu, 113
 Ana Vulpe, 87
 Anca Vasilescu, 162
 Anca-Diana Bibiri, 346
 Andreea I. Niculescu, 13
 Andreea Nechifor, 162
 Andrei Păun, 235
 Angela Bejan, 346
 Anisia Popescu, 31
 Ayaulym Sairanbekova, 55
 Carmen Vasile, 113
 Cătălin Mihăilă, 113
 Cătălin Mititelu, 126
 Claudius Teodorescu, 87
 Cristina Caterev, 87
 Dan Bumbu, 265
 Dan Cristea, 98
 Dan Tufiş, 235
 Daniela Gîfu, 145
 Daniela Istrati, 249
 Eduard Barbu, 222
 Elena Irimia, 113, 235, 249
 Elena Isabelle Tamba, 87
 Elena Lacusta, 346
 Elena Pelican, 303
 Eric Curea, 235
 Florin-Teodor Olariu, 67
 Gulmira Bekmanova, 55
 Horia Cucu, 126
 Inga Titchiev, 319
 Ioana Buhnîlă, 113
 Ioana Vasilescu, 31
 Ioana-Maria Biolan, 113
 Ion Bostan, 265
 Isabella Şinca, 113
 Iurie Ciala, 265
 Jochen Ehnes, 13
 Laura Orynbay, 55
 Liliana Botnaru, 87
 Livia Căruntu-Caraman, 87
 Liviu Petcu, 346
 Liviu-Andrei Scutelnicu, 346
 Lori Lamel, 31
 Ludmila Malahov, 249, 332
 Maria Mitrofan, 235, 249
 Maria Ungureanu, 1
 Maria-Corina Dimitriu, 199
 Mariana Nastasia, 87
 Marius Cerescu, 265
 Marius-Radu Clim, 87
 Matei-Iulian Cocu, 185
 Mihaela Cristescu, 113
 Mihaela Mocanu, 346
 Minghui Dong, 13
 Mircea Petic, 332
 Nichta Degteariov, 265
 Olesca Caftanatov, 249, 265, 319, 332
 Paula Alexandra Crucianu, 145
 Petru Rebeja, 98
 Radu Gabriel Chivoreanu, 40
 Radu Ion, 235, 249
 Ramona Luca, 67
 Răzvan-Cosmin Cristia, 185
 Rodica Branişte, 249
 Sergiu Nisioi, 249
 Silviu-Ioan Bejinariu, 67, 145
 Speranţa Cecilia Bolea, 145
 Svetlana Cojocaru, 332
 Tiberiu Boroş, 40
 Tudor Bumbu, 87, 265
 Vasile Păiş, 235, 249
 Vasile Vancea, 281
 Verginica Barbu Mititelu, 113, 235, 249
 Victoria Alexei, 249
 Victoria Bobicev, 249
 Viorica Popa, 346
 Vlad Belciug, 303
 Yaru Wu, 13