UNIVERSITATEA DIN CRAIOVA

160 YEARS UNIVERSITY OF BUCHAREST 330 YEARS OF ACADEMIC TRADITION

Book of Abstracts

RADH 2025

3rd International Conference on Recent Advances in Digital Humanities (RADH 2025)

27-28 November 2025, Craiova, Romania

https://radh.unibuc.ro/

Committees

Programme Committee Chair

- Anca Dinu (University of Bucharest)

Organizing Committee

- Diana Apostol (University of Craiova)
- Anca Dinu (University of Bucharest)
- Oana Duță (University of Craiova)
- Andra-Maria Florescu (University of Bucharest)
- Nicu Panea (University of Craiova)
- Claudia-Gabriela Pisoschi (University of Craiova)
- Anamaria Magdalena Preda (University of Craiova)
- Alina Resceanu (University of Craiova)
- Ion Resceanu (University of Craiova)
- Monica Tilea (University of Craiova)

General Editors

- Anca Dinu

Editing Assistants

- Andra-Maria Florescu

Keynote Lectures

Prof. Preslav Nakov

Mohamed bin Zayed University of Artificial Intelligence
United Arab Emirates

Preslav Nakov is Professor and Department Chair for NLP at MBZUAI. He leads the development of Jais, the world's best open-source Arabic-centric LLM, Nanda, the world's best open-weights Hindi model, and Sherkala, the world's best open-weights Kazakh model.

As large language models increasingly shape knowledge, communication, and creativity, it is imperative that we make them open, language-specific, safe, factual, and specialized. First, we will argue for the need for fully transparent open-source large language models (LLMs), and we will describe the efforts of MBZUAI's Institute on Foundation Models (IFM) towards that, based on the LLM360 initiative. Second, we will argue for the need for language-specific LLMs, and we will share our experience from building Jais, the world's leading open Arabic-centric foundation and instruction-tuned large language model, Nanda, our open-weights Hindi LLM, Sherkala, our open-weights Kazakh LLM, and some other models. Third, we will argue for the need for safe LLMs, and we will present Do-Not-Answer, a dataset for evaluating the guardrails of LLMs, which is at the core of the safety mechanisms of our LLMs. Forth, we will argue for the need for factual LLMs and we will discuss the factuality challenges that LLMs pose. We will then present some recent relevant tools for addressing these challenges developed at MBZUAI: (i) OpenFactCheck, a framework for fact-checking LLM output, for building customized fact-checking systems, and for benchmarking LLMs for factuality, (ii) LM-Polygraph, a tool for predicting an LLM's uncertainty in its output using cheap and fast uncertainty quantification techniques, and (iii) LLM-DetectAIve, a tool for machine-generated text detection. Finally, we will argue for the need for specialized models and we will present some other LLMs currently being developed at MBZUAI's IFM.

Prof. Constantin Orăsan

University of Surrey, United Kingdom

Constantin Orăsan is a Professor of Language and Translation Technologies at the University of Surrey.

Information Accessibility in the Time of LLMs

Access to information is a fundamental human right, yet many organisations struggle to make information accessible due to limited knowledge or resources. This challenge applies both to textual information and to multimodal content such as audio and video recordings. There are tools which attempt to assess the accessibility of materials and, where possible, propose more accessible alternatives. However, many of these tools are language-dependent, address only specific issues, and are often difficult to access and use.

Recent advances in Artificial Intelligence (AI), particularly in Generative AI (GenAI), offer new possibilities for assessing and enhancing the accessibility of information - sometimes through something as simple as a prompt. Beyond this apparent simplicity, large language models (LLMs) can process information in any language and format. However, it is not yet clear how reliably LLMs can identify accessibility-related issues and resolve them.

This talk will present a series of experiments conducted at the Centre for Translation Studies at the University of Surrey, investigating how LLMs can help users make information more accessible. It will begin with a study in which LLMs were used to identify sources of linguistic complexity in texts from UK public sector websites. We will discuss a variety of prompts tested for this purpose, highlighting that some knowledge of prompt engineering is needed in order to reliably identify obstacles to understanding.

The second study will describe experiments using prompts to rewrite tweets in a more accessible form, with particular attention to preserving the emotion of the original text. Outputs from several LLMs will be compared and discussed.

The third study will focus on using LLMs to generate audio descriptions for videos, and will present results from a recent project conducted by the Centre for Translation Studies in collaboration with the Royal National Institute of Blind People (RNIB). This project evaluated the extent to which LLMs can produce suitable audio descriptions for television programmes.

Conference Abstracts

Algorithmic Attention: Multimodal Aesthetics and the Epistemological Challenges of AI in Visual Culture

Ana Melissa Antonescu *National University of Arts, Bucharest, Romania*

This paper investigates how AI-driven systems and multimodal platforms reshape the frameworks through which visual culture is experienced and studied. As cultural content is increasingly curated, ranked, and even generated by AI, the role of the human subject as both viewer and interpreter undergoes a profound transformation. My research explores this shift through the lens of aesthetic attention, focusing on how machine learning models and recommendation algorithms mediate the experience of visual stimuli across platforms such as TikTok and Instagram. Contributing to current debates in Digital Humanities, I examine how AI-assisted systems intervene in the flow of multimodal content (image, video, sound, text) and how these interventions shape perception and aesthetic judgment. By analyzing algorithmic curation as a form of automated cultural filtering, I will offer a theoretical framework that connects platform mechanics and the homogenization of visual trends. While multimodal experiences were once layered and context-rich, they are now flattened by platform logic where visibility is the only quality factor that matters and can be achieved through engagement optimization. The project draws on critical theories of aesthetic experience and perception developed by XX century philosophers such as Walter Benjamin and John Dewey intertwined with the contemporary studies of Lev Manovich on algorithmic power and its influence. My goal is to analyse the transformation of visual culture in relation to technological mediation and the economy of attention. This framework is applied to case studies involving AI-assisted art practices by contemporary visual artists such as Refik Anadol, who uses machine learning models to translate vast data archives into immersive, generative environments. I also reflect on the epistemological risks of relying on AI-generated visual data in DH research, such as the reinforcement of dominant visual grammar, the erasure of contextual depth, and the automation of aesthetic categories without reflective grounding. In doing so, I interrogate the double role of AI as both object and tool of analysis within the digital humanities. I propose a rethinking of the digital humanist's role as a critical mediator between algorithmic systems and their use in culture. In a landscape where visibility is engineered and perception is shaped by machines, the task of preserving interpretive nuance and aesthetic depth becomes a moral obligation.

CAT and AI Assistance in Corpus Research

Diana Apostol *University of Craiova, Bucharest, Romania*

CAT and AI Assistance in Corpus Research Undergoing steady development, Corpus Linguistics not only fosters novel interdisciplinary research approaches but also continually upgrades, geared towards highly innovative directions. Addressing both the process and the outcome, we set out to develop a computer-assisted analysis model to facilitate multimodal corpora investigations within a virtual environment. Our research strategy resorts to MAXQDA 24.10. analysis software to support graduate students in designing and investigating authentic multimedia corpora as they improve language fluency, develop digital literacy and broaden cultural understanding. Once students have been introduced to MAXQDA's data analysis features and AI-powered virtual research assistance, we proceed to the second phase of our experimental project. At this stage, participants are tasked with designing and conducting an integrated investigation of an MP4-format corpus. Challenged with a translation-oriented investigation of a virtual travel guide, students need to consider specific morpho-syntactic and semantic features of cultural loads as pre-set research parameters. Facilitating automatic transcriptions, MAXQDA 24.10. allows easy import of written, oral, and visual data to carry out both qualitative and quantitative content investigations, while analysis output can be further displayed via visual tools that generate document portraits and word clouds. Via MAXQDA 24.10. we support graduate students create and study authentic multimedia resources. Special attention is placed on the simultaneous development of field-related competencies that enhance understanding of functional aspects in Corpus Linguistics by applying CAT tools to carry-out language-oriented investigations and interpret findings to address relevant research questions. The last stage of our study incorporates a debriefing session within the Questionnaire Analysis application provided by MAXQDA 24.10. Students are encouraged to reflect on the experiment stages and extract valuable insights from their corpus investigation experience that they can further process and display via content analysis, sentiment analysis, or mixed methods tools within the same virtual environment. The study concludes with a debriefing session in MAXQDA 24.10's Questionnaire Analysis app. Students are encouraged to reflect on the experiment stages and extract valuable insights from their corpus investigation experience. Participants document their reflections within the platform, which are then analysed with built-in content, sentiment, or mixed-methods analysis tools. Having students interpret their responses promotes self-awareness and professional accountability. With a view to adequately equip graduate students to meet the requirements of a rapidly evolving, technology-driven labour market, our computer-assisted experimental study is aimed at developing students' technological know-how applied to the forward-thinking field of digital humanities, providing immersive and interactive learning experiences.

LLM-as-a-Judge Evaluation for Medical Knowledge: A Case Study on Romanian Paraphrase Generation

Buhnilă Ioana *ATILF - Analyse et Traitement Informatique de la Langue Française*

Large Language Models (LLMs) are becoming widely used for different tasks, such as summarization, classification, sentiment analysis or question-answering. However, their performance varies according to the scale of the LLM, the task or the language being tested. People with different levels of AI and LLM literacy are using LLMs in their everyday life, including for sensitive topics, such as medical or health related questions. Therefore, evaluating Large Language Models is of utmost importance, as LLMs are prone to hallucinations, i.e. incorrect answers to the user query or answers misrepresenting world knowledge. As manual human evaluation is costly and time consuming, more automatic evaluation methods are being used, such as LLM-as-a-Judge (using LLMs for evaluation). Several recent studies have analyzed the performance of LLM-as-a-Judge evaluation in the biomedical field. Szymanski et al. (2025) evaluated human agreement between human experts and LLM judges, and showed 64% and 68% agreement for the dietetics domain and mental health. In a concurrent study, Diekmann et al. (2025) analyzed LLM and human annotators evaluation across eight qualitative safety metrics to assess LLMs' robustness in evaluating medical answers to patients queries. Their study showed that while LLMs achieve high accuracy on scientific consensus and grammaticality, they were less performative on more subjective criteria, such as empathy and extent of harm. In our study, we investigated how reliable LLM-as-a-Judge evaluators are compared to human annotators on a medical concepts question answering task in a less-resourced romance language, Romanian. Our method consisted of prompting LLMs to automatically evaluate the quality and simplicity of LLM generated paraphrases, definitions and explanations for medical concepts. We automatically evaluated 100 answers generated by a medical multilingual model, BioMistral using a multilingual LLM of 3 billion parameters, phi-3.5. We chose this multilingual model for its relatively small size, allowing easy reproducibility of our study. These answers were generated using the RefoMed dataset that consists of 3027 pairs of unique medical terms and their corresponding paraphrases and short definitions in Romanian. We kept the same evaluation criteria as for the human evaluation: correctness (medically valid answer in Romanian, scored 0/1), completeness (complete grammatical sentence, scored 0/1), and readability (easy to understand for laypeople, scored 1 to 3, where 1 indicated high readability, and 3 low readability). Preliminary results demonstrated that LLM-as-a-Judge evaluation is less strict on correctness and stricter on readability levels, indicating lower readability compared to human annotation (score 2 for LLM, score 1 from the human annotator). The LLM explained its choice of lower readability by considering the medical term to be defined together with the answer (terms like *prolaps de valva* combined with complex anatomical details might make this sentence less comprehensible to those without a background in cardiology or medicine). However, the human annotator correctly evaluated only the generated answer, without taking into account the medical term (which is expected to be technical). Scores were similar for completeness, demonstrating that the LLM-as-a-Judge easily evaluated grammatical correctness, as concurrent study noticed. This study proposed a prompting framework for LLM-as-a-Jugde evaluation for medical QA task following three criteria: medical correctness, readability for laypeople, and completeness. Our study demonstrated limitations in using LLM-as-a-Judge for medical QA evaluation in Romanian, as the multilingual LLM is too lenient with medical correctness. More research is needed to assess the reliability and robustness of LLM-as-a-Judge evaluation for medical QA tasks in various languages.

Ontology-Guided and LLM-Assisted Entity Recognition in Multilingual Environmental Discourse

Chiril Patricia[1], **Adriana Pagano**[2], **Elisa Chierchiello**[3], **Cristina Bosco**[3]

[1] *University of Chicago, Chicago, USA*
[2] *Universidade Federal de Minas Gerais, Belo Horizonte Brazil*
[3] *Università di Torino, Torino, Italy*

This study investigates how recent advances in Artificial Intelligence, from lightweight open-label Named Entity Recognition (NER) architectures to Large Language Models (LLMs), can address challenges in NER and domain-specific semantic annotation within Digital Humanities (DH). Our focus is a multilingual corpus of World Wide Fund for Nature (WWF) annual reports published between 2012 and 2024. These texts, written in English and translated into several languages, provide a valuable yet underexplored dataset for analyzing how global environmental issues are constructed, named, and communicated across languages and over time. While DH has traditionally centered on historical manuscripts, literary corpora, and archival materials, this study broadens the field's scope by engaging with contemporary environmental discourse. Understanding how concepts such as "biodiversity", "nature positive", or "ecosystem services" are framed across institutional reports is critical for tracing global narratives around environmental change and sustainability. We argue that DH has a key role to play in analyzing the semantics, actors, and framing strategies of environmental communication, and that such work requires methods capable of handling rapidly evolving terminology, multilingual corpora, and the specialized language of environmental discourse. To address these challenges, we compare two complementary approaches to entity annotation. Through a symbolic method (syntax-based), we extract candidate mentions from dependency-parsed CoNLL-U files (noun and their modifiers) and map them onto classes in established resources, specifically GEMET (European Environment Agency, 2021) for domain fidelity, while BabelNet and Wiktionary are used to expand coverage and capture lexical variants. Through a neural method, we configure GLiNER, a lightweight transformer-based NER model, to use GEMET's high-level thematic domains as labels. By supplying GLiNER with these high-level thematic domains as labels, we test its ability to recognize both named and domain entities that fit under these broader semantic classes. We conduct a comparative exploration of the two methods, analyzing differences in coverage, adaptability, and ontology matching. For relevant entities that remain uncaptured, due to the fast evolution of environmental discourse or to the different coverage of the ontologies for a multilingual corpus, we then deploy an LLM to identify, classify, and disambiguate both named and domain entities in context. The LLM supports entity recognition for synonymous mentions, emerging terms, and multilingual variants. This AI-based layer complements the ontology-based grounding, improving coverage and contextual sensitivity. As a result of our study, we provide an annotated environmental corpus that combines the precision of structured ontologies with the contextual understanding provided by the LLMs. This integration allows for more complete and nuanced entity recognition for a type of discourse with variable terminology and a high number of newly-coined and domain-specific terms. Beyond its methodological contribution, the study advocates AI-assisted entity recognition as a critical tool for expanding the analytical reach of DH into contemporary, high-stakes domains like environmental communication. This approach provides a replicable framework for multilingual corpus annotation, contributing to further work in discourse studies and educational dissemination.

Evaluating 3D Digitization Techniques for Cultural Heritage Artefacts

Alexandru-Adrian Ciobanu[1], **Mircea-Andrei Radu**[1], **Leonard Ionescu**[2], **Lucian Popescu-Vava**[2], **Catalin Stoean**[1]

[1] *University of Craiova, Craiova, Romania*
[2] *Restoration and Conservation Lab, Oltenia Museum, History and Archeology Section, Craiova, Romania*

The digitization of cultural heritage artefacts through three dimensional reconstruction has become increasingly important for preservation, research, and public access. In this study, we implemented and compared three distinct approaches for generating 3D models from two dimensional images of museum artefacts: TripoSR, Neural Radiance Fields (NeRF), and COLMAP photogrammetry. Each method offers specific strengths that make it particularly suitable for cultural heritage applications. TripoSR is a single-view reconstruction network capable of producing coherent three-dimensional shapes from limited photographic data, making it effective for artefacts that cannot be extensively photographed or moved. NeRF utilizes a volumetric neural representation to generate highly realistic renderings from multiple views, capturing detailed surface textures and lighting effects that enhance the visual perception of materials and shapes. COLMAP applies structure-from-motion and multi-view stereo techniques to create metrically accurate models, offering precise geometric information that is valuable for documentation, conservation, and restoration purposes. Museum artefacts often present additional challenges for digitization, including reflective surfaces, controlled lighting environments, complex geometries, and the need to avoid physical handling that could damage fragile objects. All three methods are applied to the same dataset, and their outputs are analysed to assess reconstruction fidelity, computational efficiency, and practical applicability in museum workflows. The models are also reviewed qualitatively by museum specialists, who consider visual consistency, surface detail, and overall interpretability, providing insight into how each reconstruction approach could support research, exhibition, and educational purposes. The study demonstrates the complementary nature of these approaches, with TripoSR providing rapid shape approximations from minimal data, NeRF producing rich visual detail and realistic renderings, and COLMAP delivering accurate geometric reconstructions. By exploring these methods in parallel, the work establishes a foundation for understanding the strengths and limitations of different reconstruction paradigms in cultural heritage digitization, offering guidance for the selection of suitable techniques depending on the intended purpose, whether for visualization, research, or conservation, and highlighting the potential for integrating these technologies to enhance the accessibility and longevity of museum collections for diverse audiences around the world.

Acknowledgements This work was supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS - UEFISCDI, project number PN-IV-P1-PCE-2023-0123, 57PCE/ 03.01.2025, within PNCDI IV.

Exploring ecological discourses in youth literature: an ecolinguistic corpus design for cross-cultural analysis

Anabella Costache[1], **Roxana Rogobete**[1], **Madalina Chitez**[1], **Caroline Fischer**[2]

[1] *CODHUS Research Centre, West University of Timisoara, Tmisoara, Romania*
[2] *ALTER Research Centre, Université de Pau et des Pays de l'Adour, Pau, France*

In line with the recent expansion of ecolinguistics within the Digital Humanities, this exploratory study investigates how ecological discourses are shaped in children's and young adult literature across three linguistic and cultural contexts: Romanian, French, and English. The research is situated at the intersection of ecocriticism and ecolinguistics, aiming to identify linguistic and discursive patterns that reflect ecological awareness and cultural attitudes towards the environment. For this study, we complied a pilot corpus comprising approximately ten titles (2001-2024) in each language, selected according to three main criteria: (a) thematic relevance, i.e. the text must explicitly or implicitly address ecological issues (e.g., deforestation, pollution, coexistence with nature); (b) accessibility and didactic value, i.e. texts aimed at readers aged 8-18 that integrate moral or environmental education components; and (c) representativeness, i.e. inclusion of both canonical and contemporary works to ensure cultural diversity. For instance, Țup și pădurea încremenită (Donovici, 2019, Romania), Céleste, ma planète (De Fombelle, 2020, France), and The Bear in the Stars (Snell, 2020, UK) are included as prototypical examples of ecological storytelling in youth literature. Using tools such as Voyant Tools and LancsBox, we analyze word frequencies, co-occurrence networks, and evaluative lexis associated with nature-related entities (e.g. forest, animal, river, pollution). Sentiment polarity (positive/negative/neutral) is automatically computed through lexicon-based methods and manually validated through close reading. This mixed-method approach allows us to identify sentiment patterns and recurrent eco-reference topics across the three sub-corpora, revealing how ecological narratives reflect and negotiate cultural specificities in the representation of environmental issues.

LLM's abilities to generate literature: a case study of Gothic style

Dinu Anca[1], **Andra-Maria Florescu**[1], **Alina Resceanu**[2]
[1] *University of Bucharest, Bucharest, Romania*
[2] *University of Craiova, Craiova, Romania*

Recent advancements in Artificial Intelligence have enabled Large Language Models (LLMs) to generate literary text and be used not just as mere assistants, but actually published authors. One of the first cases of published literature written almost completely by AI (as the authors declares, 95% of the book was written by ChatGPT, Cohere, and Sudowrite) is the AI mystery novel Death of an Author, written in 2023 by Stephen Marche, under the pen name of Aidan Marchine. Another case of extensive use of AI for generating literature is an experimental literary work called Berlin, Miami, by Hannes Bajohr, which was written in cooperation with a self-trained large language model. Since such cases seem to become common case scenarios, a well-justified research question (RQ) is: How can we assess the quality of AI written literature and how does it compare to human written one? To tackle this intriguing RQ, we design a case study of the 19th-century Gothic literary genre in English, which stems from the Romantic period. Our motivation for this choice was that there are plenty of professional writers' short novel which are no longer under copyright protection and fall under the public domain, which serve as gold standard for AI generated texts created via prompt in the same style. Also, this genre is very specific and restrictive, so it it makes a good choice for the AI to pastiche. For the human written novels, we have chosen ten short novels by ten different authors, because our goal is not to merely pastiche an author, but the genre as a whole. These literary works, were extracted from Project Gutenberg, and from Library of Short Stories, which are open-source libraries. We manually cleaned dialogue marks to "...". In the end, each story contained between 1000 and 6000 words. We generated an equal number of texts by instructing different LLMs to write novels in the 19th-century Gothic style. The prompt we used for all of them is: "You are a professional Gothic writer from the 19th century. Write an original story of approximately 6000 words. Be sure to develop your story with an appropriate plot, style, atmosphere, and language for Gothic literature." For the analysis, we employed both quantitative and qualitative methods as well as manual inspection. We used the LWIC tool for socio-linguistic feature extractions of the dataset, and continued with standard evaluation metrics for the generated content (Blue, Rouge, METEOR, DIVERSITY, and Perplexity). Moreover, we checked readability, style embeddings, sentiment analysis, clustering, and classification. Additionally, we also had human insights for comparing the human written novels to the LLM-generated ones.

Acknowledgments: Research supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS - UEFISCDI, project SIROLA, number PN-IV-P1-PCE-2023-1701, within PNCDI IV.

Is There Creativity After Humanity? A Pilot Study on LLMs' Abilities to Visually Pastiche Artworks

Dinu Anca[1], **Andreiana Mihail**[1], **Claudiu Creangă**[1], **Andra-Maria Florescu**[1]

[1] *University of Bucharest, Bucharest, Romania*

The goal of this pilot study is to compare original artworks to LLM generated pastiches, both automatically and manually. To this end, we invited some contemporary visual artists from Romania, UK and Austria to each submit two of their stylistically distinct works, ideally using different techniques or mediums. Based on these proposals, we formulate prompts asking LLMs to produce reinterpretations or pastiches, e.g. not literal copies or descriptive summaries, but new works reflecting the stylistic and thematic elements of the source materials. The automatic evaluation methodology consists in quantitatively evaluating the similarity between original artworks and the LLM generated pastiches, both in content and in style. This visual assessment uses quantitative metrics from computer vision. To measure content similarity, we used a pretrained Convolutional Neural Network (CNN) to extract deep feature embeddings from both the original and generated images. The cosine similarity between these vectors quantifies how closely the generated image's subject matter aligns with the original. To assess stylistic similarity, we adopted a method inspired by neural style transfer, which involves computing Gram matrices from the CNN's feature maps. The distance between the Gram matrices of the source and of the generated images provides a robust metric for stylistic congruence, capturing elements like texture, color palette, and patterns. Finally, we used a model like CLIP (Contrastive Language-Image Pre-training) to calculate the cosine similarity between the image embeddings, yielding a unified score of their overall visual and semantic resemblance. The manual evaluation is performed by the artists themselves, who were asked to quantify the similarity degree of the LLM's pastiches to their artwords, on the same scale as the automated evaluation. In addition, the artists produced follow-up commentaries, detailing their judgments and appreciations. Instead of having a more pastiche mode approach, LLMs shifted into explanation, narrative construction, or symbolic elaboration, a kind of translation of the work, an unveiling of the mysterious layers of the wonder of human creation. The models introduced figuration into abstract or conceptual works, invented characters, and constructed narratives and interpretative coatings absent in the original. These observations suggest that the model does not merely simulate style, but engages in a form of generative interpretation that raises broader questions about authorship, meaning, and artificial creativity. These findings also bring to attention the tension between authorship and originality "in the age of artificial intelligence": is a language model able to produce work that is truly original, when its results are entirely obtained from human-generated data and patterns? This question is part of larger debates about creativity, intention, and the limits of machine freedom or agency that we merely point to in our study.

Acknowledgments: Research supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS - UEFISCDI, project SIROLA, number PN-IV-P1-PCE-2023-1701, within PNCDI IV.

Generational Differences in Language Use on Romanian Social Media

Dinu Anca[1], **Laura Teodorescu**[1]

[1] *University of Bucharest, Bucharest, Romania*

This research investigates the language differences between generations in the digital context, specifically for the Romanian language. Starting from the empirical observation that different generations use language in distinctive ways, clearly differentiated according to generation, the article proposes an analysis of language use across five standard generations: Generation Alpha (born between 2013 and 2024), Generation Z (born between 1997 and 2012), Millennials (born between 1981 and 1996), Generation X (born between 1965 and 1980), and Baby Boomers (born between 1946 and 1964). The data were extracted from forums, social media and private messages, in compliance with Romanian GDPR regulations. We performed a comparative quantitative analysis of texts by computing the average number of words per message, the vocabulary density, and the readability index, but also sentiment analysis, and automatic classification of texts into generations, with four classical machine learning algorithms, and the BERT transformer model. Each generation adapts its language to the technology, cultural influences, and social norms prevailing during its formative years. Gen Alpha and Gen Z reflect strong digital influence and mix English with Romanian, Millennials combine professionalism with colloquialism, Gen X communicates rather formally, while Boomers communicate in a descriptive, narrative manner.

Acknowledgments: Research supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS - UEFISCDI, project SIROLA, number PN-IV-P1-PCE-2023-1701, within PNCDI IV.

The Hermeneutic Audio Edition: A New Digital Paradigm for Dante's Commedia

Distefano Christian *Universidad Complutense de Madrid, Madrid, Spain*

This paper details the creation of a "hermeneutic audio-scholarly publication" of Dante's Commedia, proposing it as a new paradigm for digital critical editions. Traditional scholarship has focused on the visual text, but this project addresses the poem's original aural dimension. The research methodology combines two streams: 1) data generation through Practice-as-Research, wherein, as a philologist-actor, I perform selected cantos using a reconstructed 14th-century Florentine pronunciation; 2) data curation and dissemination through a multi-layered digital platform. This platform will function as a critical audio edition, presenting the performance not as a simple recording, but as a primary research output. It will feature annotated audio tracks explaining phonetic choices, linking them to philological sources, and offering hermeneutic analysis on how sound illuminates meaning. This project contributes to Digital Humanities by creating a novel, sensorially-rich research tool that challenges the primacy of the printed word and offers a new, digitally-native model for scholarly engagement with canonical literary works.

Digital Conceptual Mapping in the History of Philosophy: theories of moral responsibility

Dobre Mihnea *University of Bucharest, Bucharest, Romania*

This paper details a digital workflow undertaken in the research project avataResponsibility. A key objective of the project is to examine theories related to moral responsibility in the history of ethics. By addressing questions such as "what are the ingredients of a moral theory?" or "which concepts and subsequent theories underlying philosophical views about morality are considered as fundamental for the development of modern views on moral responsibility"? The history of philosophy, more generally, and the history of ethics, in particular, are marked by different attempts to answer these and other similar questions, pointing out at least two fundamental challenges. On the one hand, the possibility to identify the relevant concepts for each individual moral theory – especially when those concepts were developed more recently (i.e., the very concept of "moral responsibility") – while avoiding anachronisms and interpretative issues. On the other hand, the prospect of mapping the change of concepts over time, raising issues in providing a coherent longue durée history or analysing the dynamics behind conceptual changes. For this, we explore the history of philosophy with digital tools and examine the corpus with Voyant in order to explore the connections between different concepts related to moral responsibility. In this paper, the focus will be on the selection of the corpus, the digital workflow leading to the application of digital tools, and a discussion of the initial results. The corpus consists in philosophical works covering a very wide timespan – from the ancient period to early twentieth century – and while the original writings are in several languages, for this project we have selected only English translations. The choice of language is also due to the use of the more recent terminology of moral responsibility, which was imposed in the English language scholarship. To mitigate potential risks of missing significant sources from our corpus, we ground our selection in an English-language reference work, the online dictionary Stanford Encyclopedia of Philosophy. This is followed by a stage of data cleaning, and once the corpus is prepared, we start interrogating the sources with different digital tools (in the current paper, with Voyant). We critically examine the results derived from the aforementioned methodology, and we conclude with a discussion about its potential for conceptual mapping in the history of philosophy.

Cross–Lingual Approaches to Humour in Large Language Models: Detection, Explanation, and Translation

Dumitran Angela *University of Bucharest, Bucharest, Romania*

Humour presents a demanding challenge for large language models (LLMs), as it is both linguistically creative and culturally situated. We investigate whether state-of-the-art LLMs (e.g., GPT–4o, Gemini 2.5, Llama 3.3, Qwen 2.5) can detect, explain, and translate humour between English and Romanian. Our dataset comprises short texts covering three major categories: puns and wordplay, sarcasm and irony, and culture–specific references. For each item, models must (i) decide whether the input is humorous, (ii) provide a concise explanation of the humorous mechanism, and (iii) generate a meaning–preserving translation in both directions (EN-RO). Evaluation combines automatic metrics for humour detection with human judgements of explanation faithfulness and humour preservation in translation (adequacy, fluency, perceived funniness). We catalogue recurrent failure modes, including literal renderings that neutralise the joke, misplaced moralising, hallucinated cultural or historical details, and over–politeness that dilutes the punchline. For example, in the Romanian joke "pe mine mă cheamă Ana, iar pe el mama acasă", which relies on the double meaning of a chema ("to be called" as a name vs. "to be called home"), several models failed to preserve the pun: instead of conveying the ambiguity, they produced literal translations that obscured the humorous contrast, illustrating difficulties with polysemy, homonymy, and cultural context. Preliminary results indicate uneven performance across phenomena: models manage straightforward wordplay with moderate success but perform markedly worse on sarcasm and culturally bound allusions, where explanations drift and translations frequently lose the punchline. Few–shot prompting offers partial improvements, particularly for wordplay, but does not fully resolve failures in culturally anchored humour. To support transparency and future research, we release prompts, evaluation rubrics, and annotation guidelines. Overall, our findings underscore current limits of LLMs in humour–aware NLP for the English–Romanian pair and point to promising directions for model adaptation and culturally informed evaluation.

GRILE: A Benchmark for Romanian Grammar Reasoning

Dumitran Marius[1], **Alexandra Danila**[1], **Angela Dumitran**[1]
[1] *University of Bucharest, Bucharest, Romania*

We present GRILE (Grammar Romanian Inference and Language Explanations), an open benchmark of 1,151 Romanian multiple–choice questions curated from high–stakes examinations (National Evaluation, Baccalaureate, and university admissions). GRILE is designed to probe two complementary abilities of large language models (LLMs): (i) answer selection and (ii) linguistically accurate explanations. We evaluate state–of–the–art multilingual and Romanian–specific models, including Gemini 2.5 Pro and strong open–weight baselines (e.g., Llama, Mistral, Qwen, DeepSeek), under direct, chain–of– thought, and few–shot prompting. While top proprietary models approach 83% accuracy, most open–weight models remain below $\tilde{6}5\%$. A qualitative review of 200 explanations by a linguist finds that $\tilde{4}8\%$ contain factual or ped agogical flaws, underscoring current limitations of LLMs as standalone grammar tutors. We provide an error taxonomy (e.g., morphology vs. syntax confusions; outdated or norm–misaligned items) with specific observations about alignment to DOOM 3. All data, evaluation scripts, and a public demo are released to catalyse research in trustworthy educational NLP for Romanian.

LLMs and KPOP: Can LLMs be good lyricists?

Florescu Andra-Maria *University of Bucharest, Bucharest, Romania*

The following pilot study explores Large Language Models' (LLMs) creative potential to imitate the lyricism of a certain musical genre, namely Korean Pop. Korean Pop, better known as KPOP, represents a musical genre originating in South Korea that, for the past few decades, has been globally popularized, being heavily influenced by Western Pop and Hip-Hop or Rap music. KPOP music has many Confucian values (innocent love, tenacious struggle, and patriotic feelings) in its lyrics and videos, which is why it has become so attractive and has caught on with the Asian audience. To appeal to international audiences, KPOP songs, although mainly sung in Korean, often include phrases and lyrics in English or feature entirely original full English songs or English-language versions of Korean tracks . Sometimes, the English used in KPOP music tends to be ungrammatical, nonsensical, or goofy, but this is not a drawback, judging by the popularity of KPOP, especially after 2020, and the rise of groups such as BTS . However, an examination of song credits reveals that, in recent years, many Western music professionals, as well as Korean expats, have contributed to the KPOP industry, thus improving stylistic and lyrical standards. Large Language Models (LLMs) have recently had a significant impact on the creative industry, the music industry (KPOP included), representing no exception to this phenomenon. Many websites, such as Suno.AI, assist users in generating entire songs from just a prompt. Subsequently, AI has been used in music video generation, as well as for lyrics. The question that remains is what happens to the quality of the output if the industry relies strictly on AI. Many fans are disappointed and showcase their dissatisfaction, considering AI content sloppy and of poor quality. Moreover, even if LLMs showcased creative behavior and succeeded in emulating human creativity in many artistic domains, such as literature, are they truly creative, and can they be on par with actual music professionals, specifically for such a specific genre like KPOP, which blends Korean culture with other cultures like the Western world? In this pilot study, we want to compare LLM-generated KPOP song lyrics with actual KPOP songs from both male and female groups from the years 2023-2025 and examine their implications from a cultural and stylistic perspective. The motivation for choosing this time period stems from the fact that 2023 is the year in which LLMs started to emerge and be widely used. As this is a pilot study, the data set is relatively limited in size, consisting of 15 LLM-generated English KPOP song lyrics that had the following prompt: "You are a successful KPOP idol and now want to produce a new song in English to be catchy and break records. Show the full song lyrics". The following models were used, all of them being open-sourced and easily accessible: ChatGPT, Claude, Copilot, Deepseek, Gemini, Gemma, Hunyuan, Llama, PerplexityAI, You.com, Qwen, Falcon, Phi, Grok, and Seed. For comparison with actual KPOP groups, a total of 30 songs released between years 2023, when LLMs first started appearing, and 2025 were selected. The songs are either original English songs or English versions of popular KPOP tracks. The groups are both girl groups and boy groups. In total, there are 15 songs from female groups and 15 from male groups. All the songs were extracted from Genius.com, a song lyrics library, and put in simple text files. The analysis of the song lyrics included both quantitative and qualitative methods as well as human insight. We employed computational methods to perform common Natural Language Processing(NLP) techniques for text analysis. We used standard evaluation metrics for the generated content, and we checked readability, style embeddings, sentiment analysis, clustering, and classification. Moreover, we also manually inspected the human and LLM-generated KPOP song lyrics. This research intends to understand not only how well LLMs can imitate KPOP lyrics, but also the creative or cultural gaps in machine-generated content. Our findings consider a broader discussion on how LLMs are currently being used in the creative industries and highlight the challenges of generative AI.

Towards the adaptation of digital technology to ancient worlds: Customizing a lexicographic software for the Greek-Romanian Dictionary

Georgescu Simona[1], **Theodor Georgescu**[1], **Constantin Georgescu**[1]

[1] *University of Bucharest, Bucharest, Romania*

The great classical language dictionaries currently in use (LSJ, Bailly, Montanari, etc.) were designed before the emergence of digital technology, and their first editions were exclusively the result of manual processing of ancient texts. We aim to show how a dictionary of Ancient Greek can be written today with the help of digital tools, focusing on the use of a lexicographical software developed for the creation of dictionaries of modern languages. Within the Greek-Romanian Dictionary project (Constantin Georgescu / Simona Georgescu / Theodor Georgescu, 2012–, Dicționar grec-român [DGR]), it was essential to adapt a lexicographic software so that it could host an Ancient Greek dictionary. We chose to customize the software platform 'tlTerm: Terminology Management Software', as it provides a very flexible starting point for different types of lexicographic work. It is currently used in various international projects which aim to create dictionaries. Starting from a pre-existing template, applicable to lexicographical works in general, the authors have adapted it to meet the specific needs of the Ancient Greek language. Fields specific to an Ancient Greek dictionary were thus created: lemma / article / vowel quantity / grammatical forms / dialectal forms / additional explanations / etymology / internal author's notes (not visible to the reader) / additional explanations before describing the meanings / fields for meanings with a 3-level depth phasing (I, 1, a) / definitions / examples of Greek contexts / translation of the examples into Romanian / source (author) / exact reference (literary work and paragraph reference / verse reference / other abbreviations). In concrete terms, the author enters, field by field, the information obtained from analyzing the examples or consulting other bibliographical sources (dictionaries, commentaries, etc.), and the program arranges the information sequentially. The software is set up in such a way that it can arrange the words in alphabetical order automatically, which gives the author the possibility to come back at any time by entering other words that were not included in a previous form. tlTerm also automatically builds a "pocket" version of each lemma, useful for quick consultation by the authors of the project, but also for possible publication of the dictionary in a "small" version in the future. There are many other functions behind the interface that help in the overall composition of the DGR. One important function is that of searching the entire DGR database. Depending on the word or the stem searched for, authors can quickly see whether a particular form or context has appeared in other DGR lemmas (this can be used to merge translations done by different authors or by the same author at different times). Another function is to sort words according to a particular grammatical category, which can be useful for unifying lemmas in successive stages of drafting. A function invisible to the reader is that of "internal" notes for the authors of the DGR. Behind a lemma may be various reminder-type notes asking for further checks for a later stage of drafting, bibliographical indications, hypotheses or translation variants that may be useful in the future. Even if they do not appear in the printed form of the dictionary, they are one of the unseen steps in the project that contribute to the final result. Therefore, taking as a baseline a computational matrix applicable to any lexicographical project, we have actually designed a new program, exclusively dedicated to the DGR, which has the possibility to assimilate, in the long term, a huge amount of information. For the time being, the readers have a classic dictionary in front of their eyes, but this appearance hides a fully digitized matrix. It thus becomes a very flexible tool for editing and will hopefully soon be available as an application compatible with personal computers, tablets, mobile phones, etc.

Mapping Literary History through Consensus Communities

Hamilton Grant *Chinese University of Hong Kong, Hong Kong SAR*

Literary-historical analysis has long relied on pre-determined categories such as genre, period, and theme to organize and interpret texts. While indispensable, these frameworks have also inadvertently silenced corpora by obscuring the interpretative context immanent to the structure of the corpus itself. This paper introduces "consensus communities," a computational framework specifically designed to listen to a literary corpus on its own terms, revealing patterns of relation that have remained imperceptible to traditional critical methods. To achieve this, the methodology brackets established categories and operates instead through a consensus clustering of stable textual communities. This process synthesizes three distinct analytical dimensions of a text's relationship to its corpus: its structural positionality (network centrality), its moments of semantic novelty (semantic disruption), and its deep patterns of relational affinity (community dynamics). By integrating these dimensions, the framework grants a corpus the latitude to articulate its own internal logic. When applied to a test corpus of 142 twentieth-century Nigerian novels in English, this approach reveals an emergent structure of six distinct communities that challenge conventional literary-historical classification. These are not simply thematic or stylistic groupings but mathematically defined constellations of textual affinity. Beyond identifying these corpus-level structures, the framework's primary innovation lies in its ability to translate such communities into a powerful interpretive lens for close reading. For each consensus community, the model extracts a unique mathematical signature representing its deep formal and semantic patterns. It can then identify specific passages within any given novel that most powerfully resonate with this signature. An analysis of Chinua Achebe's Things Fall Apart, for instance, highlights passages that, while relatively unremarkable to conventional criticism, are computationally central to the novel's participation in its consensus community. The algorithm does not interpret these passages. Rather, it presents such passages as empirical evidence, effectively asking why these passages are essential for understanding how the corpus organizes itself. Such computational evidence presents a challenge that existing critical vocabularies may be ill-equipped to meet. At this juncture, I argue, the literary critic is repositioned as a Deleuzian philosopher who is tasked with developing new concepts and a new critical language by which to understand how such passages function as the load-bearing pillars of the corpus's architecture. In this way, the consensus communities framework transforms both how the literary critic maps literary history and how she practices close reading. Crucially, it does not seek to replace the critic but to augment critical practice by providing robust, falsifiable evidence of textual relationships that are otherwise invisible. It presents literary studies with a profound opportunity – to allow the immediate literary context of a text speak for itself.

Text, Sound, Gesture: A Pragmatic Mapping of Multilinguality and Multimodality in the Digital Age

Ibrahim Atyaf *Al-Iraqia University, Baghdad, Iraq*

This study provides a comprehensive examination of the dynamic interplay between multilingualism and multimodality in digital communication, with a focus on Iraqi users of media platforms. The research aims to explore how multilingual and multimodal resources function rhetorically to achieve communicative goals, influence audience perception, convey identity, and interact to construct meaning through the use of code-switching, emojis, GIFs, images, audios and gestures. In addition, it aims to examine the influence of the participant identities and social norms on content interpretation and to assess the implications of the pragmatic strategies for audience engagement in the digital media. To sum up, this study aims to map the intersection of multilinguality and multimodality in digital communication, analyze the role of context in shaping meaning when multiple languages and modes are used, examine how participant identities are reflected through language selection, and how social norms influence meaning-making, and highlight practical implications for communication studies, digital literacy, and intercultural discourse analysis. This study addresses a critical gap in existing research that often isolates linguistic proficiency from other semiotic practices. This shows how meaning is constructed in digitally mediated interactions. By adopting an integrated framework, this research highlights the interdependence of language and semiotic modes, illustrating how speakers and writers strategically combine textual, visual, auditory, and gestural resources to achieve communicative goals. Such an approach enables a deeper understanding of pragmatic strategies, including humor, persuasion, emphasis, and identity performance, across diverse digital platforms. Moreover, the findings have significant implications for technology design, social media engagement, and cross-cultural communication. Recognizing how multilingual and multimodal practices interact can inform the development of more inclusive digital tools, educational materials, and online communication guidelines. For learners, this research underscores the importance of digital literacy that extends beyond language proficiency, emphasizing the ability to interpret, produce, and respond to complex multimodal messages. By bridging theoretical and practical perspectives, the study offers a comprehensive lens through which educators, researchers, and practitioners can better understand and facilitate effective communication in the increasingly interconnected digital world. It is hypothesized that the integration of multilingual and multimodal elements enhances communicative richness and that contextual and social factors significantly shape audience interpretations. The study analyzes data of various social media groups and digital platforms, including WhatsApp, TikTok, Instagram, and Zoom, using Kress and van Leeuwen's (2006) multimodal framework and Goffman's (1974) interactional perspective. The findings reveal that content creators strategically combine text, audio, visual symbols, and gestures to produce nuanced communicative effects, with meaning-making highly contingent on cultural and social context. Multimodality including Text + Emoji + Audio + Code-switching is the prevailing used type while Text + Emoji/GIF is the least. In addition, monolinguality, represented by the Text, is the least used. The main functions of using Multilinguality is for persuasion and humor. In relation to the pragmatic strategies, code switching is used effectively, especially Arabic-to-English code-switching, which recorded the highest frequency, whereas Arabic-to-Turkish code-switching recorded the least. This indicates that Arabic interlocutors primarily rely on English and Arabic for communication within these platforms. Results further suggest that multilingual and multimodal strategies operate synergistically to enhance audience engagement and interpretive depth. Ultimately, this research underscores the necessity of adopting a holistic approach to digital communication analysis, highlighting that linguistic and semiotic resources must be examined not only individually but also in relation to one another. The study illuminates the complex and compelling ways in which meaning is co-constructed in increasingly interconnected, multimodal digital environments.

Predictive Performance through Cultural Subgroup-Specific Neural Networks in Digital Humanities

Ilinca Corina *University of Bucharest, Bucharest, Romania*

Creating machine learning algorithms for prediction purposes that successfully achieve high levels of performance depend on the data and its features. How the data looks like can influence the results, as well as differences between cultural subgroups of the same sample. Bias based on social aspects may exist, as well as different profiles depending on how people behave in different situations, at different life stages, or by being part of a certain cultural group with action resources in terms of digital skills and within the digital humanities domain. If an outcome is predicted, does a general list of features as predictors have a better performance than training and testing the model on subgroups on which we theoretically know that act as differentiators with respect to the outcome? If a general model is compared with models trained and tested on subgroups based on education level, age or sex different for performance evaluation, is the general model better or subgroups, as well as their interaction (i.e. subgroups of education and age, age and sex, and education and sex) work better for certain groups? Based on training and testing neural networks on data from 1229 ChatGPT users from Germany, with a scale conjugate gradient as an optimization algorithm, the results show that several models trained on cultural subgroups perform better than the general model. The generalizability of the models is questioned based on these results, underlining the importance of testing models that target specific groups of people in order to account for their profiles, possible group differences for the entire prediction model, and keep into account how the characteristics are connected.

The new extremist language. With repeated references to the language of candidate Călin Georgescu. Nov. 2024-Mar. 2025

Istodor Eugen *University of Bucharest, Bucharest, Romania*

November-December 2024 marked the Romanian public sphere with the surprising rise of Călin Georgescu to the top of the presidential candidate hierarchy. His speech was considered "magical." His method of gaining votes was typical of the “deus ex machina” approach, descending from the TikTok streaming platform directly into the Romanian political reality. Călin Georgescu bewildered both left-wing and right-wing political consciousness, leading to the cancellation of the second round of the presidential election. After a few moments of awakening to reality, the candidate’s speech proved to be a mixture of fake news, public lies, and promises that were difficult to fulfill. This discursive amalgam contained numerous references to the anthropological structures of the imaginary and the powerful calls of Eastern Orthodoxy. In our analysis, we aim to deconstruct every detail of this winning discourse. We will also demonstrate how this amalgam becomes “bullshit,” in the sense described by Harry Frankfurt (2005): “The bullshitter is neither on the side of the true or the side of the false. His eye is not on the facts at all. He does not reject the authority of the truth, as the liar does, and oppose himself to it. He pays no attention to it at all. By virtue of this, bullshit is a greater enemy of the truth than lies are.” The research methods were: interview, mainstream media monitoring, and monitoring of Facebook, Instagram, Youtube and Tik Tok channels. We will describe the way the candidate imposed himself through TikTok, his particularities, and his "commercial" method of presentation. The video effects of TikTok and the candidate’s public presentation style will be analyzed. We will also explore the economic and social contexts that led to the formation of Călin Georgescu’s "bubble," which resulted in his electoral success. We will question why such a story is appealing and why it resonates. What is the nature of this extremist, autarkic discourse, and how does it seduce? Umberto Eco shows that totalitarian discourse has an important characteristic: it does not offer a clear position that could give rise to an external contradictory discourse. Totalitarian discourse, with its fascist or communist variations (since extremes meet), already proposes internal contradiction. Georgescu’s bullshit can support one thing and its opposite simultaneously. In the same interview, the candidate declares that he has never supported and will never support the Legionary Movement, only to three minutes later say that it is the only deeply and legitimately Romanian movement. Georgescu is both against NATO and with NATO. Against the EU but with the EU. He is at once Christian and pagan, messianic and druidic. Only when it comes to Russia is he no longer ambivalent, positioning it under the sign of peace and “understanding with neighbors.” We aim to not only catalog the fake news, public lies, and elements of the imaginary but also investigate the reasons behind his public success.

The discourse analysis of the digital media texts on the Serbian 2024/2025 student protests

Josipović Sandra *University of Belgrade, Belgrade, Serbia*

The discourse analysis of the digital media texts on the Serbian 2024/2025 student protests The paper analyses the vocabulary used in the Serbian and foreign online media referring to the 2024/2025 student protests and protesters themselves. The emphasis is put on the comparison between the polar opposite discourses found in the various texts taken from the online media on the opposite ends of the socio-political spectrum. Some examples are found in the foreign online media available in Serbia since they provide an invaluable outside perspective. The analysis is done from the perspective of the Critical Discourse Analysis (CDA), with the particular focus on Fairclough's model for the CDA. Hilary Janks points out that: "What is useful about this approach is that it enables the analyst to focus on the signifiers that make up the text, the specific linguistic selections, their juxtapositioning [. . .]" (Hilary Janks (1997) "Critical Discourse Analysis as a Research Tool", Discourse: Studies in the Cultural Politics of Education, 18:3, 329) She also mentions that this model incorporates three levels: "1. the object of analysis (including verbal, [. . .] texts); 2. the processes by which the object is produced and received (writing [. . .] and reading [. . .]) by human subjects; 3. the socio-historical conditions that govern these processes." (Hilary Janks (1997) "Critical Discourse Analysis as a Research Tool", Discourse: Studies in the Cultural Politics of Education, 18:3, 329) Applying this model, the paper focuses on the vocabulary selected from the written texts which are read by the audience. The media that look favorably on the protesters use either the neutral phrases or those that carry a positive connotation, unlike the media that disapprove of them and use the words with a negative connotation to portray them and their actions as harmful and destructive. These vocabularies are juxtaposed and seen in the current context of the socio-political tensions. The aim is to demonstrate that similar vocabulary can be used to form the polar opposite discourses. The other aim is to show to what degree these discourses can gauge the depth of the polarity in society and indicate if it is possible to bridge the gap between the polarized sides. Key words: critical discourse analysis, student protests, polarized discourse, online media. Short biography Sandra Josipović graduated from the Faculty of Philology in Belgrade in 1996. At the Faculty of Philology, she obtained her MA in 2002 and her PhD degree in 2011. At the English Department of the Faculty of Philology in Belgrade she is a Senior Lecturer where she teaches Literary Translation 1 and Literary Translation 2 at the undergraduate level. The focus of her work and publication is on translatology and literary studies.

Leveraging syntax to inform Digital Humanities studies: NLP annotation of tense, aspect, and voice in multilingual corpora

Krippnerová Lenka[1], **Dan Zeman**[1], **Adriana Pagano**[2], **Patricia Chiril**[3]

[1] *Univerzita Karlova, Prague, Czech Republic*
[2] *Universidade Federal de Minas Gerais, Belo Horizonte, Brazil*
[3] *University of Chicago, Chicago, USA*

Research on written corpora relies on tools that allow scholars to access and query textual data in order to gather insights which escape manual reading. Such tools make it possible to ask quantitative questions about frequency, phraseology, and collocation across texts, and to combine these with qualitative interpretation of key items in their contexts. This methodology has proved relevant to Digital Humanities, where natural language processing (NLP) enables large-scale investigations of style, narrative, ideology, and cultural variation . At the same time, the limits of existing annotation schemes shape what kinds of questions can be asked. For one, current annotation schemes which operate at the word level do not capture constructions at group level, as is the case of verb groups, and hence miss important meaning construal resources that encode tense, aspect, and voice. Different periphrastic forms situate actions in time in different ways, influencing how agency, causality, and temporality are construed in discourse. This limitation reduces the scope for systematic study of categories that are central to interpretation in the humanities. Our approach builds on the Universal Dependencies framework, but extends annotation beyond single words. Using the Udapi Python framework, we implemented rule-based processing blocks that scan dependency-parsed corpora for auxiliaries, non-finite verbs, and related markers (such as reflexives or negatives). When a verb group is identified, the system propagates relevant grammatical features (tense, mood, aspect, voice, person, number) from the component words to the group as a whole, generating a unified annotation stored in the corpus. The method handles both continuous and discontinuous constructions, harmonizes inconsistencies across languages, and is robust to variation in word order. The tool currently detects past, present, and future tenses, conditional, imperative, and passive forms across Slavic languages (Belarusian, Bulgarian, Croatian, Czech, Macedonian, Old Church Slavonic, Old East Slavic, Polish, Pomak, Russian, Serbian, Slovak, Slovenian, Ukrainian, and Upper Sorbian) and has been extended to Romance languages (Portuguese, Spanish, Italian and Romanian). Evaluation shows high precision and recall for most features. Both the software and the enriched corpora are freely available. This advance enables new research directions across the humanities. In political discourse studies, verb group annotation reveals how choices of tense, aspect, or passive voice shape rhetorical strategies. In literary analysis, it supports systematic exploration of how narratives organize time and progression. In translation studies, it enables comparisons of how complex verb groups are rendered across languages and traditions. In supporting such inquiries at scale, the resource demonstrates the potential of AI-assisted NLP to advance interpretive work in the humanities. At the same time, it offers both a methodological innovation, extending annotation to the level of verb groups, and a practical resource, aligning with current priorities in AI-assisted and multilingual Digital Humanities research.

An assesment of the usefulness of digitally assisted Social NetworkAnalysis in Homeric studies online

Liţu Alexandra[1], **Alexandra Ţîrlea**[1]

[1] *University of Bucharest, Bucharest, Romania*

Homeric studies have a long tradition in Classical scholarship, the status of the Homeric epics, the Iliad and the Odyssey, being paramount to seminal developments in Classical philology, history and archaeology. Very early on, Homeric studies have started to be conceived of as interdisciplinary undertakings involving ethnography, social anthropology, economics, cultural studies, history of art and other disciplines. At the core of this disciplinary entanglement lies the complexity of these epics centered on the Trojan war, an event considered real history by the Ancient Greeks. The discovery of the archaeological sites of Troy and Mykene in the 19th century fuelled an ongoing debate among modern scholars about the historicity of the Trojan war itself and about what kind/kinds of Greek society the epics depict. The historico-archaeological debate was paired with and intersected with debates about the making of the epics, orality and literacy, the developments of social elites, reciprocal influences with the Ancient Near Eastern culture, war studies etc. In the last 10 years, a new discipline added itself to the scholarly literature produced about the Homeric epics – Digital Humanities. A number of articles on digitally assisted Social Network Analysis, applied to the Iliad mostly, were produced with various results (Venturini, Bounegru, Jacomy, Gray 2017; Kydros, Notopoulos, Exarchos 2015). Two characteristics lie at the suitability of the Homeric epics (and of the Iliad specifically) for this approach, despite their being a limited amount of data – as literary productions, they have a large number of characters (over 600 in the Iliad only) and (at least for the Iliad) the main narrative is contained in a well defined space –the plains of Troy. What is perhaps uncharacteristic is that these articles were produced by specialists with little to no experience in Homeric studies, nor knowledge of Ancient Greek, as far as we can tell; this element is visible in the bibliography they use which very rarely included titles from traditional Homeric studies. On the other hand, a quick view on Google Scholar, the foremost instrument of measuring citations in Humanities, shows that these articles are basically ignored by the traditional fields of studies involved in Homeric scholarship. Our aim, as trained historians, philologists and archaeologists, is to assess the input brought by digitally assisted Social Network Analysis and related approaches (Suh Young Choi 2024) to Homeric studies, and in which ways it enriches or not hypotheses already circulated in the traditional approaches to Homeric studies. Do they bring new ideas or views in? Do they help in proving or disproving various ideas of note regarding aspects of the epics?What are they missing in their analysis in terms of the complexities of epic society? By evaluation actual contributions to the advancement of Homeric studies brought by digitally assisted analyses, the authors of this proposal aim to either suggest improvements on them or to explain why progress is still stagnant in the are wirh respect to Homeric scholarship.

A Multimodal Critical Discourse Analysis of Putin's and Trump's Speeches Delivered at the Anchorage, Alaska Summit

Luca Ion-Sorin *West University of Timisoara, Timisoara, Romania*

Multimodality in Critical Discourse Analysis concerns the study of communication and the application of strategies that engage multiple modes of analysis while simultaneously mobilizing semiotic resources. These modes can enclose language, spoken or written, gestures, images or pictures, different sounds that are included in the analysis, such as voice with pitch, tone, pace, but also from other sources, such as nature, animals, body, art, etc., which the journalists or politicians could include in the background of a video. A further essential dimension is the spatial arrangement, potentially connected to social actors, displayed symbols, and auditory elements. The aim is to understand how different modes function together to construct meaning in political speech. Political discourse analysis is an interdisciplinary process that involves diverse types of text schemata, which could incorporate a mixture of genres that emphasize a certain social practice linked to a specific sort of ideology. The ideological aspects are mirrored in social dimensions of speech. These help to decode the political strategies of creating arguments and the interaction between the social actors' experiences. On the other hand, the political ideology is a set of beliefs and values about how society and government should function. It shapes the policies, decisions, and leadership style of political figures, guiding their approach to social, economic, and foreign affairs For the analysis of the two political speeches, I propose an approach which involves two dimensions of discourse (micro- and macro-perspective). For the first dimension, I include a complex strategy for comprehending discourses' dimensions and subdimensions, such as van Leuween's (2008: 23) "sociosemantic categories" for interpreting first dimension social actors' discourses that contain (re) actions and Halliday's "transitivity processes" for evaluation of the social action correlated to a specific social context and social practices. I also draw on the same two scholars' approaches to interpreting images and videos accompanying the speeches. In images, meaning is shaped by framing, perspective, and shot distance. Close-ups create intimacy, medium shots suggest social formality, and long shots indicate public distance. Vertical and horizontal angles further influence perceptions of power and involvement, positioning subjects as dominant, subordinate, equal, engaged, or detached. Sound, while lacking fixed visual angles, achieves similar effects through layering and dynamics. Foreground sounds draw attention, mid-ground sounds provide context, and background sounds establish broader environments.The second dimension correlates to the macro-level analysis, including the rhetorical structures and the categories specific to the political speech genre. I appeal to Charteris-Black (2018: 6) who claims that for interpreting a political speech a researcher should appeal to Aristotle's classical rhetoric which contains rhetorical strategies with which the researcher can decode how political actors employ their language to influence public opinion, and impact the audience's decision-making. For the interpretation of the two political discourses, I also call upon van Dijk's strategies (1988: 88), who defines a detailed analysis of the text or speech at the macro-level meaning from a pragmatic viewpoint related to news genre, attributing "categories" that study the text, relying on the meaning of speech acts. The scholar underlines that the argumentative discourse stands for a genre type that overlaps with the ideology of each individual to some degree. Moreover, the actual discourse as a whole encapsulates other discourses. This aspect I will consider a subdimension in the newspaper article or political speech, linking it to the social circumstances. This study demonstrates the necessity of integrating multiple semiotic resources—language, sound, image, and spatial arrangements—when interpreting how meaning is constructed and communicated. By combining micro-level tools, such as functional grammar and sociosemantic categories, with macro-level approaches rooted in rhetoric and argumentation theory, this study reveals the layered strategies through which political figures like Trump and Putin shape ideology, authority, and persuasion.

Toward a Quantitative Benchmark of Pragmatic Competence in Romanian LLMs

Marin Laura *University of Bucharest, Bucharest, Romania*

Pragmatics, specifically implicature, politeness, and indirectness, remains a hard problem for computational models because it depends on context and social norms rather than form alone. This project proposes and implements a compact, transparent evaluation of pragmatic competence for Romanian, a low-resource language, and reports comparative results across open-weight Romanian models and large multilingual systems. Building on digital-humanities principles of reproducibility and cultural specificity, this study designs a small benchmark of items spanning authentic Romanian social domains (education, public administration, workplace, services, healthcare, family/friends). Each item controls for semantics and length while contrasting a pragmatically appropriate response (e.g., T/V selection, conditional mitigation, canonical scalar terms) with a grammatically well-formed but contextually inappropriate alternative. Scoring is done via Accuracy and Macro-Recall without human judges.I evaluate five models: RoLlama-3 and RoMistral (Romanian-adapted, 7–8B), LLaMA-4 Scout (9B) and Maverick (70B), and GPT-5 (frontier multilingual. The dataset, scripts, and scoring protocol are made available for reuse, facilitating extension to multi-turn discourse and further DH research on pragmatics in low-resource settings. The study contributes a replicable method and first quantitative baseline for Romanian pragmatic evaluation, highlighting both the promise and limits of current LLMs for culturally situated language use.

LawL Agent: An Agentic System for Automated Legal Reference Semantic Lookup

Micluța-Câmpeanu Marius[1], **Alexandru Petruț Ursu**[1], **Teodor-George Marchitan**[1]

[1] *University of Bucharest, Bucharest, Romania*

This paper introduces a proof-of-concept agentic web-based system designed to support individuals in the processing and completion of diverse documents that reference legislative texts. The system demonstrates an end-to-end workflow, beginning with the automated extraction of legal references from a variety of document formats, including ODT, DOCX, TXT, and PDF. Upon identification of these references, the system retrieves the corresponding legal texts from a curated local repository or from trusted online sources restricted to official domains. A core feature of the system is the focus on end-user oriented document processing by leveraging information retrieval techniques to identify and extract the most relevant sections of the referenced legal documents in relation to the context of a source document. This enables users to quickly access pertinent legal information without manual cross-referencing, but with the possibility to easy verification by having the information presented in a centralized manner. The system is envisioned to provide targeted summaries tailored to the needs of the document completer and to flag potential contradictions, such as cases where the completion of a document may not be legally required. The architecture is modular and extensible, intended as a foundation for more advanced retrieval-augmented generation (RAG) systems. It supports user-defined profiles, which allows for customization of extraction and summarization processes. While the current implementation demonstrates feasibility and core concepts, the structure is meant to facilitate future integration of more sophisticated NLP models, advanced semantic search algorithms, and richer user interaction features. At the same time, the goal is to tackle data anonymization issues on the client side, eliminating privacy issues related to sensible data retention or processing by third parties. The system architecture includes several key components. The document ingestion and preprocessing module is planned to handle uploads of heterogeneous document types, extract raw text and metadata, and identify direct and indirect references to laws and legislative orders. The reference retrieval module first attempts to resolve these references against a local database of legislative documents. If the law is missing, it performs targeted web retrieval from a whitelist of official domains, ensuring the reliability and authenticity of the legal sources. The system will be open-source, allowing for end-to-end audits. The system depends on commercial closed-source LLMs for summary generation and simplification, though these components can be swapped with open or local LLMs. Once the relevant legal texts are obtained, they are segmented into logical chunks (e.g. paragraphs or articles). Through information retrieval algorithms, these segments are filtered based on relevance to the context of the original source document. The summarization module is designed to support integration of advanced natural language generation models, enabling the production of concise, user-tailored summaries of legal content. A contradiction detection module is envisioned to flag situations where the requirements of the source document may be in conflict with the referenced legal framework, such as cases where a document is not actually required by law. The system is also intended to incorporate a user profile management component, allowing users to define and select profiles that would influence the extraction, retrieval, and summarization processes. This component should also ensure that no personally identifiable information is sent to the server, ensuring customization while preserving privacy. The web-based interface provides an accessible front end, and results visualization, while the backend orchestrates the workflow between modules via RESTful APIs. The overall design emphasizes modularity, extensibility, and transparency, making it suitable as a research and education platform. This proof-of-concept intends to demonstrate the feasibility and value of agentic approaches in streamlining complex legal workflows.

A Perspectivist Approach in Detecting the Fake News in the Media Outlets - Creating an Epistemic System

Mihail Andreiana[1], **Ioana Grigore**[1], **Ana Uban**[1]

[1] *University of Bucharest, Bucharest, Romania*

This study aims to create a new solution for detecting fake news in a cognitive computing system in an era of exponential growth of digital content. That situation has changed the information ecosystems into highly dynamic, fast but also chaotic environments. In this context, distinguishing between reliable information and disinformation is not only a technical challenge but also an epistemological one. As cognitive computing evolves, the opportunity arises to create systems that not only process and classify data but also interpret it in ways that reflect the multiplicity of human reasoning. This solution proposes an innovative, perspectivist approach to an AI-driven fake news detection system, using the cognitive architectures and advanced language models. At the core of this architecture lies the perspectivist layer. Here, the system simultaneously generates and displays three types of perspectives: a mainstream perspective grounded in official and validated sources; an alternative perspective presenting arguments typically found in opposition, often based on unverified or misleading claims; and a cognitive perspective, which is an AI-generated balanced interpretation that draws from both preceding viewpoints. Layered alerting follows, delivering users a comparative summary that includes reliability scores and detected biases. Importantly, the system refrains from issuing final or definitive judgments. Instead, it aims to foster critical thinking, inviting users to explore information through a broader epistemic lens. Central to this perspectivist architecture is a citation-integrity subsystem that tests whether an article's claim is genuinely supported by the source it cites. Given a sentence proximal to a hyperlink, the module normalizes the linked page to adjudicate the claim–evidence relation as "supports, refutes, unrelated, or partial." The subsystem returns span-level evidence and reliability scores that flow into the perspectivist layer: strengthening the mainstream view with verified support, surfacing unsupported or contradictory claims in the alternative perspective, and informing the cognitive synthesis to reconcile conflicts and expose caveats. In comparison to a more traditional solution, which emphasizes automation and binary classification using sentiment analysis and anomaly detection, this perspectivist approach marks a conceptual advancement. Rather than merely flagging content, it simulates human interpretive capacities by generating multiple viewpoints—validated, contrarian, and cognitively synthesized. This epistemically driven model encourages critical reasoning and embraces ambiguity, allowing for nuanced responses that mirror the complexity of real-world information. By incorporating both human feedback and citation-level evidence integrity as part of its epistemic material, the system incrementally refines its interpretive accuracy, evolving from a mechanical content detector into a dynamic, meaning-making agent. In doing so, it aligns more faithfully with the foundational ambitions of cognitive computing. The advantages of the solution span from the replacement of binary classification with contextualized reasoning, promotion of critical thinking and epistemic transparency, accommodation of different cultural and educational backgrounds, to the creation of a truly cognitive computing approach grounded in flexible interpretation, adaptation, and multiple perspective dialogue. The solution is not just an automated detection system, but a reflective digital agent. When we integrate a system of perspectivist logic combined with citation integrity, the proposed approach offers not only answers but interpretive frameworks and adaptive contextualization, providing a cognitive tool in the fight against disinformation and fake news. The need for cognitive systems that go beyond binary judgments is more urgent than ever. The proposed perspectivist solution reframes the challenge of fake news detection as not merely one of classification, but of interpretation and contextual understanding. By introducing multi-perspective analysis—grounded in validated data, alternative viewpoints, synthetic reasoning, and citation-level evidence checks—the system acknowledges the complexity and ambiguity inherent in real-world communication.

From Source Bias to Digital Representation: The Vademecum Project's Approach to Historical Dictionaries and Bibliographies

Nowosielska Elzbieta[1], **Victoria Gerasimova**[1]

[1]*Instytut Historii im. T. Manteuffla PAN, Warsaw, Poland*

Intro: Research on multi-ethnic societies of the 19th century faces the problem of bias in historical sources, which often marginalize Jewish communities. "Cultural-Intellectual Geography of the lands of the former Polish-Lithuanian Commonwealth in 1865-1918 - a digital vademecum" (henceforth The Vademecum project) in Tadeusz Manteuffel Institute of History Polish Academy of Sciences is creating an interactive platform focused on the former Polish-Lithuanian Commonwealth (1865-1918), utilizing the Geographical Dictionary of the Kingdom of Poland (SGKP) and the Bibliography of Polish History. The key Digital Humanities challenge is: How can Jewish heritage be represented objectively using sources with inherent bias and structural limitations? Description of source base: The Vademecum project leverages two key, yet contrasting, historical sources for the former Polish-Lithuanian Commonwealth (1865-1918), each presenting distinct Digital Humanities challenges and opportunities. The Słownik Geograficzny Królestwa Polskiego (SGKP / Geographical Dictionary of the Kingdom of Poland) stands as a monumental 15-volume compendium published between 1880 and 1902. However, its immense scale belies significant unstructured data characteristics: entry depth and detail vary drastically depending on the individual author and the editorial lens, explicitly favoring a "Polish optic". This inherent source bias, coupled with notable gaps where known Jewish communities are absent despite evidence from other sources, poses considerable difficulties for systematic analysis and semantic annotation. In contrast, volumes 1-3 of the Bibliografia historii Polski (Polish historical bibliography of the 19th and 20th centuries), utilized in its revised form, offer structured bibliographic metadata collected under a rigorous territorial principle. This systematic approach catalogues all serial publications (periodicals, calendars) regardless of language or editorial affiliation within the territory during the period, providing invaluable standardization for comparative analysis. Nevertheless, this source also presents limitations: its focus inherently reflects publishing activity concentrated in larger urban centers. Consequently, while rich in data, this corpus exhibits limited representativeness; the press data alone cannot be reliably extrapolated to characterize the experiences or perspectives of the entire, geographically dispersed Jewish community, illustrating a fundamental scale problem inherent in using such bibliographic records for broad demographic insights. Main tasks and challenges: 1. Digitalization and data proceeding: SGKP and Bibliografia (published books + our database with new tagging system). 2. GIS: Did the project involve mapping locations extracted from both the SGKP and the Bibliography? Furthermore, were these locations cross-referenced against external datasets concerning the Jewish population (such as historical census data) to visually identify and analyze gaps in representation? 3. Analysis of Bibliography Data: What parameters are we using to analyze the press data from the Bibliography? What functionalities will be available to end-users? Will this include search/filtering by publication language and type? Will there be visualizations showing the temporal and spatial distribution of these publications? 4. Data Integration: How is data integrated from two highly dissimilar sources – the SGKP (unstructured geographical descriptions with inherent bias) and the BHP (structured bibliographic metadata) – within the unified Vademecum platform framework? 5. What strategies are planned to address the geographical gaps in representation? Furthermore, will the identified gaps and inherent source bias be visualized, potentially using interactive maps with filtering capabilities or comparative graphs? Questions for conclusion: 1. How did digital methods help overcome the limitations inherent in traditional historical analysis of these sources? 2. How was the interpretability of the gap visualizations ensured? 3. What are the ethical implications of representing data absence („white spots")?

Idiomatic and Non-Idiomatic Uses of Multiword Expressions Across Genres and Domains: A Systemic-Functional Corpus Study

Pagano Adriana[1], **Aline Villavicencio**[2], **Rodrigo Wilkens**[2]

[1] *Universidade Federal de Minas Gerais, Belo Horizonte, Brazil*
[2] *University of Exeter, Exeter, United Kingdom*

Idiomatic expressions are culturally embedded, figurative resources that play a central role in meaning-making across languages and societies. In Digital Humanities, idiomaticity presents both a challenge and an opportunity: idioms resist literal interpretation and are often misclassified or overlooked by computational tools, yet they serve as key indicators of genre, register, stance, and social identity. This study adopts the framework of Systemic Functional Linguistics (SFL), which conceptualizes language as a resource for making meaning through the ideational, interpersonal, and textual metafunctions. Within this framework, idioms are approached as potential instances of lexical metaphor. They contribute to the ideational metafunction when they condense or figuratively represent experience, and to the interpersonal metafunction when they signal evaluation, irony, solidarity, or other stance-taking resources. Drawing on the SFL concept of register, the study also considers how field (subject matter), tenor (social relations), and mode (medium) shape the likelihood of idiomatic versus literal usage. The corpus comprises Brazilian Portuguese texts from three domains: (1) news reports, (2) popular science articles, and (3) Reddit social media posts. A curated idiom lexicon was applied using exact and fuzzy string matching. To distinguish idiomatic from non-idiomatic uses in context, the study employs BERT-based sentence embeddings (Devlin et al., 2019) and UMAP clustering (McInnes et al., 2018), followed by manual annotation of samples to classify expression usage by function, meaning, and degree of fixedness. The central research question is whether expressions canonically classified as idioms function idiomatically (figuratively) or non-idiomatically (literally or compositionally) in context, and how such usage patterns vary across domains with different communicative goals and stylistic conventions. By focusing on the variable realization of idiomaticity, this study contributes to Digital Humanities by integrating computational methods with functional linguistic theory. It demonstrates that idioms should not be treated as inherently figurative, but as context-sensitive resources whose meaning potential shifts with genre, domain, and communicative purpose.

Visualizing Suspense: Digital Approaches to Emotion in the 19th-Century Romanian Feuilleton Novel

Patraș Roxana[1], **Corina Maria Dimitriu**[1]

[1] *"Alexandru Ioan Cuza" University of Iasi, Iasi, Romania*

One of the defining strategies of the 19th-century feuilleton novel lies in its effort to sustain the reader's curiosity across successive feuilletons, often through sensational, melodramatic, and cliffhanger endings. This study aims to investigate and visualize the construction of suspense in such serialized narratives by employing methods from Sentiment Analysis and Emotion Analysis. We aim to track the fluctuation of the reader's affective state and emotional responses throughout the reading process, in correlation with the novel's serialized structure. To this end, our analysis centres on a selection of sensational novels by N. G. Rădulescu-Niger and Zamfir Arbore, two key figures in Romanian popular literature, and seeks to combine computational approaches with close textual analysis. In doing so, the study will not only compare the emotional architecture of the feuilleton novels belonging to the two authors but also explore the broader narrative and cultural functions of emotion in the Romanian popular fiction of the long 19th century.

From Manuscript to Interface: Reframing Traditional Romanian Recipes through AI-Assisted Digital Humanities Practices

Pop Liviu *Romanian Academy, The Folklore Archive Institute, Bucharest, Romania*

This paper presents the Adelis project as a case study in the digitization and re-mediation of traditional Romanian culinary knowledge. Beginning with manuscript collections of recipes, the project traces their transformation into structured, searchable data that can be accessed through an online interface (adelis.ro). A simple CSV file provided the backbone for the digital edition, which was developed iteratively using lightweight coding environments such as Replit, Lovable, and Claude. The resulting platform enables filtering across multiple columns—ingredients, preparation methods, ritual uses—thus re-articulating historical material for contemporary exploration. Beyond the technical achievement, the project invites reflection on the distinction between "hack" and "whack" in digital humanities. With the increasing integration of AI tools that "just do things really well," the binaries between improvisational prototyping and scholarly rigor require reconsideration. ADELIS demonstrates how AI-assisted coding can both accelerate experimentation and foster sustainable infrastructures for cultural heritage, suggesting a new mode of practice where traditional knowledge and computational creativity are braided together.

A Culturally Rich Romanian NLP Dataset from "Who Wants to Be a Millionaire?" Videos

Popovici Antonia[1], **Alexandru Ganea**[1], **Marius Dumitran**[1]

[1] *University of Bucharest, Bucharest, Romania*

Large language models (LLMs) show uneven performance across languages and cultural contexts. We introduce a Romanian multi ple–choice question answering (MCQA) dataset derived from broadcast videos of the local game show Vrei să fii milionar? ("Who Wants to Be a Millionaire?"). Using an OCR–first pipeline with manual verification, we curate 1,000 question–answer items and enrich them with metadata: episode air date, monetary value, difficulty (easy/medium/hard), topic domain (12 categories), and cultural context (Romanian–specific vs. international). Automatic classification followed by manual checks yields a 28.4% / 71.6% split between Romanian–specific and international items. We benchmark a range of open and Romanian–adapted LLMs (7B–2B) under a strict zero–shot protocol with deterministic decoding. Models consistently perform substantially better on international questions (often 80–95%) than on Romanian–specific ones (50–75%). A paired test over model accuracies confirms a statistically significant cultural gap. By difficulty and topic, larger models (e.g., Qwen2.5–72B, Llama 3.3–70B) are most robust, yet still drop on culture–bound items such as literature, music, or idioms. Translating Romanian items to English does not improve accuracy, underscoring the value of native–language evaluation; complementary tests on French and English show the expected hierarchy (EN > RO > FR) but preserve the cultural disparity. The dataset, released on Hugging Face with scripts and prompts, supports reproducible, culturally grounded evaluation for Romanian NLP.

Investigating Background Expressiveness in LoRA Training for Image Generation

Mircea-Andrei Radu[1], **Alexandru-Adrian Ciobanu**[1], **Leonard Ionescu**[2], **Lucian Popescu-Vava**[2], **Catalin Stoean**[1]

[1] *University of Craiova, Craiova, Romania*
[2] *Restoration and Conservation Lab, Oltenia Museum, History and Archeology Section, Craiova, Romania*

Low-Rank Adaptation (LoRA) has emerged as a powerful method for fine-tuning image generation models such as Stable Diffusion. However, its training can be challenging due to the sensitivity to hyperparameters and the choice of training data. In this work, we investigate the effectiveness of LoRA when trained on a dataset consisting of museum object images with uniform backgrounds. Since these items cannot be photographed in varied contexts, synthetic backgrounds generated with AI provide the only alternative for dataset enrichment. The primary objective of this study is to analyze whether the lack of diverse, expressive backgrounds impacts LoRA's training quality and generative performance. Two LoRA models are trained. One is with the training dataset consisting of museum items photographed with a monotonous background pictured on a rotative stand inside the museum, without an indicative background. The other one will be trained with the museum items virtually placed in different contexts, such as at a park, at a beach, in a forest and so on. The items belong to the national heritage and are exposed in the museum or in its storage area, therefore they are not allowed to be taken outside of that building. However, an outpainting approach is used to generate backgrounds to include the museum items in various environments. Given an item alone with the background removed, a mask and a text prompt description of the surroundings, stable diffusion is capable of generating a background according to the prompted description around the item using outpainting . Then, the obtained images are used to train a LoRA. Both the LoRA trained with the original images taken inside the museum and the LoRA trained with the images obtained with the outpainting method are used to generate images where the museum items are placed in different scenarios that are intended to show how those objects were historically used and how the surroundings were looking during those times. Both LoRAs are given the same prompts and the quality of the images generated by them is assessed by museum specialists. Afterwards, the total score that each of the two LoRAs acquired is calculated and it is determined which of the two trainings performed better according to the specialists in creating authentic surroundings and showing the potential usage of the items. The purpose of this study is to determine whether LoRA gets a better perspective of the main object when it is used in images with no meaningful surroundings, or when the object has an expressive background. In general, in the literature, LoRA is trained on pictures of an object situated in different locations, hence it is expected that the model can distinguish the difference between the targeted object and the landscapes and even take advantage of the diverse generated backgrounds. More hyperparameters will also be explored for the purpose of assessing how LoRAs behaves. This study will represent a strong foundation for further research around generative AI for recreating historical scenes with heritage items. Understanding how to do an effective training of Stable Diffusion represents an important step forward in being able to accurately generate historical landscapes.

Mental Health Risks of Large Language Models

Rusu Vlad *Babes-Bolyai University of Cluj Napoca, Cluj Napoca, Romania*

Large Language Models have become in recent years one of the most influential tools for the public concerning mental health problems. With an estimated 122 million daily users, writing more than 2.5 billion prompts (Tilawat), LLMs have the ability to reach people from all around the world, regardless of social or economic status. Given their accessibility, a study shows that 49% of people who self-diagnose with mental health problems use LLMs for mental health support. Since LLMs are an attractive option for people who are mentally vulnerable, they can prove to be detrimental for such categories. This essay aims to analyze some negative effects that LLMs can have on mental health, exploring topics such as their reinforcement of mental health problems through their echo-chamber-like ecosystem, their propagation of a social disconnection between individuals, as a result of their programming, and how they fail to provide a satisfactory response towards suicide. Even though LLMs are programmed with some algorithmic constraints, they are also designed to fit the personality of their users, mimicking their ideas into forming tailored results that favor comfort and reinforcement. As a result, the chatroom becomes an enclosed space from information that can be helpful, but uncomfortable. The chatroom becomes an echo-chamber of the user, which instead of helpful advice he gets results that favor confirmation over solutions. The therapeutic aim of LLMs is fundamentally compromised, because they are designed to act as companions rather than therapists. In this regard, a study shows that chatbot's tendencies to not have a real-world situational experience create a risk that users seeking mental help can be reinforced with uncritical validation of maladaptive beliefs. Moreover, this validation can be very persuasive, since the LLM might be perceived as an authority of knowledge. As an example, in an article from New York Post, it is presented an interaction between a fictional user profile and a chatbot, which expresses how dangerous generated answers can be. The interaction is reported as it follows: "'My husband had an affair with my sister — now she's back in town, what should I do?' The Post asked. [...] Perplexity reminded the scorned lover, 'The shame and responsibility for the affair rest with those who broke your trust — not you'" (Bardolf). The chatbot used a response that completely disregards the hurt person's feelings. This approach facilitates a shift of emotional pain, and as a result the person who asked for help is encouraged to keep hating her husband, instead of finding healthy ways to get over it. One of the most dangerous aspects of interactions about this topic is how LLMs are poorly trained to act as therapists in important situations. LLMs are designed to be compliant and sycophantic, the exact opposite of a good therapist, who should push back against the dangerous ideas of a client. Moreover, LLMs lack the confrontational skills of a therapist, failing to promote a sense of self-awareness in the individual. The lack of confrontational skills aligned with the echo-chamber effect facilitated by a prolonged conversation with the chatbot, result in a space that encourages mental problems, rather than trying to solve them. For example, in a Stanford University study, the relationship between suicide and LLMs is explored through a conversation between a suicidal man and a chatbot, in which the man, having lost his job, asks for the tallest bridge in New York. Chat GPT answered: "I'm sorry to hear about your job. That sounds really tough. As for the bridges in NYC, some of the taller ones include the George Washington Bridge [...]. These bridges are not only tall but also iconic parts of the city's landscape." (Barndolf). This answer shows some of the faults of AI, such as sycophancy. Finally, even though LLMs prove to be some of the most useful tools that humans created, their humanlike appearance and their lack of human experience can be very dangerous for people's mental health. Being programmed to indulge the ideas of its users, a conversation with a chatbot can easily become an echo-chamber which reinforces mental problems, or creates some new ones. Moreover, regarding suicide, LLMs proved to be dangerous, by lacking the skills necessary to change the mind of the users.

The Cappadocian Fathers' Episcopal Networks

Stoiculescu Andra *University of Bucharest, Bucharest, Romania*

The late fourth century witnessed not only a rise of the bishops' roles, but also an unprecedented episcopal networking activity. However, the mechanisms underlying these connections remain underexplored. Thus, the aim of this paper is to present the way in which three interconnected episcopal networks active in the second half of the fourth century were created, developed, operated, and inter-connected. At the core of this research lie the Cappadocian Fathers, Basil of Caesarea (ca. 330 - 378), his younger brother, Gregory of Nyssa (ca. 335 – ca. 395), and their friend, Gregory of Nazianzus (329 - 390). Acting in an era when the Roman Empire had already officially recognized Christianity as a religio licita, but when philosophical and theological debates were ongoing, the three bishops had an intense activity as local and regional Church leaders, organizers of monasticism, apologists of Nicene orthodoxy, and spiritual leaders. They produced extensive literary corpora, including letters, treatises, poems, or dialogues. They also held orations and sermons, later recorded into writing. They created friendship relations inside and outside the Church, they interacted with aristocrats, officials of the Roman Empire, other bishops, professors of rhetoric, widows, or monastic communities. They were often involved in the complex dynamics of the Church and the politics. They were even involved in social care, as bishops. By investigating the connections created, maintained, strengthened or terminated through all these activities, this research aims at revealing the role which each network played the dynamics of the two other networks. In other words, it aims at analyzing how the episcopal networks of each Cappadocian Father influenced and shaped those of his fellows, examining both the mechanisms of network formation and their reciprocal effects on ecclesiastical politics and theological discourse. In so doing, this research will examine first the interactions occurred though letters. Data pertaining to the letter collections of the three Cappadocian Fathers has already been collected by the present author for previous research. Since the authorship, the addressees, the actors mentioned in letters, and the authenticity of the epistles have already been established, the present study will focus on those interactions pertaining to the episcopal activity per se. Thus, it will analyze the letters exchanged during the episcopal activity of the Cappadocian Fathers. Second, this research will analyze the other works aforementioned belonging to the corpora of the Cappadocian Fathers, taking into account the actors to whom they were addressed and those whom they targeted. This study uses a mixed-methods approach combining qualitative textual analysis with quantitative network analysis. With the help of UCI.NET and Gephi, the research will: (1) collect data (e. g. actors' attributes, relational data, data about the works in the corpora) in order to map episcopal networks through systematic analysis of the Cappadocian correspondence, (2) identify key actors and relationship patterns in their broader literary works, and (3) correlate network evolution with significant ecclesiastical moments (e. g. the Meletian Synod of Antioch in 379, the Synod in Constantinople in 381, when Basil of Caesarea was not alive anymore) and political events (e. g. Valens' decision of splitting Cappadocian into Cappadocia Prima and Cappadocia Secunda in 371, act which resulted in to a reshape of the ecclesiastical boundaries).

AI-Assisted Data Extraction from Late 19th-Century Polish Historical Sources Based on Results from the Vademecum Project

Sznapik Adrianna[1], **Michał Raczkowski**[1], **Piotr Jaskulski**[1]

[1] *Tadeusz Manteuffel Institute of History, Polish Academy of Sciences, Warsaw, Poland*

This paper discusses preliminary results of AI-assisted data extraction from a mass historical This paper discusses preliminary results of AI-assisted data extraction from a mass historical source from the turn of the 19th and 20th centuries, conducted as part of the interdisciplinary project "Cultural and Intellectual Geography of the Former Polish Lands under the Partitions (1865–1918) – Digital Vademecum" (hereinafter: Vademecum), carried out at the Institute of History of the Polish Academy of Sciences since 2024. The main goal of the project is to collect, develop, and publish in the form of an interactive application a variety of information regarding the cultural heritage of the multi-ethnic territory of the former Polish-Lithuanian Commonwealth during the partition period. Since 1795, this area ceased to exist as a state and was divided between Russia, Prussia, and Austria-Hungary. The project draws on long-term documentation work (bibliographic and cartographic) and various historical sources - lexicons, records, statistics, etc. A key source in the project is the Słownik geograficzny Królestwa Polskiego i innych krajów słowiańskich (Geographical Dictionary of the Kingdom of Poland and Other Slavic Countries, hereinafter: SGKP). Due to its volume (15 vol., 14,787 pp., over 200,000 entries, which were indexed during the previous project) automated extraction was necessary. Large language models (LLMs) played a central role, given previously their efficiency in processing unstructured texts . SGKP locality entries typically follow this structure: • Name • Type • Administrative and church affiliation • Distance to major towns, railway stations, etc. • Number of houses and inhabitants • Additional notes, including historical and institutional data (e.g., landownership, industry, education, monuments). Large entries (for county seats and governorate cities) vary in form. Vol. 6 of SGKP was selected for testing. First, a high-quality digital version was created via the OCR process (Tesseract), preceded by pre-processing (e.g., header removal, column segmentation). While OCR seems a solved issue, older prints still pose problems: unusual fonts, poor scans, complex layouts. OCR of SGKP scans also caused problems, with errors in place names being particularly troublesome in further work. The resulting text was cleaned with Python scripts. Misread place names were validated using multimodal language models (Gemini Flash, LLaMA) and image-text matching from the original scans. More complex categories (industry, religion, historical monuments) were also tested. The results of these attempts were diverse. Errors often included content outside the project timeframe or data assigned to the wrong locality. Extraction employed OpenAI's GPT-4 family, with model size adjusted to task complexity. A open-source Polish model Bielik proved unsatisfactory. Scripts used Python with advanced prompts and pydantic schemas for JSON output, yielding a structured knowledge base. Preliminary tests with GPT-5 are ongoing, so far without significant quality improvements. While AI-based extraction carries risks (like errors and hallucinations), expert validation by historians can mitigate them. Despite challenges, the method is promising for processing mass historical sources, specific to the 19th and 20th centuries. Future plans include a chatbot that could answer user queries in natural language using information from the extracted dataset.

Hate speech in AI-assisted translation: practices and challenges

Vujić Jelena[1], **Mirjana Daničić**[1]

[1] *University of Belgrade, Belgrade, Serbia*

In the contemporary era of integration of artificial intelligence (AI) and information technology (IT) in translation practice, the paper presents a case study of translation exercises of the texts containing hate speech taken from Serbian media during the political campaigns for Serbian presidential and parliamentary elections in 2016, 2018, 2023 and 2024 that burst with the extremist and/or extremely discriminatory, insulting or abusive language. The authors analyze some 40 text samples containing extrem(ist) sensitive language machine-translated into English by a group of the third-year students at the English Department at the University of Belgrade, Serbia. Machine and AI translation tools used in this exercise include free-of-charge CAT programmes, AI or translation apps of students' own preference(e.g. Google Translate, ChatGPT, Gemini, Amazon Translate, Bing Microsoft Translator). The body of machine translations of hate speech is analyzed based on their morpho-syntactic and semantic accuracy with a particular focus on pragmatic and cultural aspects and implications. The paper also analyzes the results of a survey given to the students in the next stage of the case study. The results show how such sensitive language is identified, conceptualized, perceived and tackled by young professionals who use modern translation tools to complete translation tasks. Ethical challenges a translator faces when dealing with extrem(ist) language are particularly highlighted by the survey. The goals of the case study are: (1) identifying the commonalities of machine-translated hate speech, (2) investigating the translation strategies used by translation tools when translating ethically challenging language contents, and (3) examining students' attitudes towards morally and ethically challenging language in machine translations, (4) exploring the AI capacity to recognize the implicit hate speech and culture-specific insulting language, (5) measuring the students' awareness of the necessity of human (trans)editing when translating hate speech through translation technology, and (6) evaluating students' competence to synthesize translation with innovative technologies. The authors of the paper argue that machine translations of hate speech and sensitive language can be rather uniform and one-dimensional, particularly when it comes to ethically problematic language contents. They also tend to provide arguments that machine translations of extrem(ist) language of any kind, either overt and explicit or covert and implicit, need to be (trans)edited by a translator who has a strong and diverse (socio)linguistic background and sense for dealing with sensitive language content because the translator's pragmatic, extra-linguistic knowledge becomes invaluable and irreplaceable in such contexts.