Is Human–LLM Interaction Culture-Dependent? A Cross-Linguistic NLP Analysis of Student Interviews on AI-Assisted Thesis Writing

Madalina Chitez[1], Karla Csuros[1], Dejana Jelena Milićević[2], Petya Osenova[3], Stefan Marinov[3], Teodor Valchev[3], Nikolay Paev[3], Otto Kruse[2], Christian Rapp[2], Andreea Dinca[1], Roxana Rogobete[1], Claudia Doroholschi[1], Loredana Punga[1], Anabella Costache[1], Dumitru Tucan[1], Cristina Baniceru[1]

[1]West University of Timisoara
[2]Zurich University of Applied Sciences
[3]Institute of Information and Communication Technologies, Bulgarian Academy of Sciences
madalina.chitez@e-uvt.ro

Abstract

This study investigates whether human–LLM interaction in academic writing exhibits cross-cultural variation. Using NLP-informed corpus methods, we analyze nine semi-structured student interviews from three national contexts (Romania, Bulgaria, Switzerland) to examine how AI use is linguistically constructed across three dimensions of epistemic positioning: agency strength, authority dynamics, and discourse-level stance. Results show a strong predominance of distancing and hedging strategies, with AI consistently framed as a functional writing support tool rather than an epistemic authority. At the same time, modest but systematic cross-country differences indicate culturally embedded variation in how students discursively negotiate epistemic responsibility and evaluation in AI-assisted writing practices.

Keywords: Human–LLM interaction, Epistemic positioning, Cross-linguistic discourse analysis, AI-assisted writing, Student interviews, Academic thesis writing, Corpus-informed NLP

1. Introduction

Academic thesis writing is a central higher-education practice through which students construct and justify knowledge rather than reproduce information. Writing functions not only as communication but also as a cognitive process that supports claim formulation, evidence evaluation, and epistemic responsibility (Bean, 2011; Hofer and Pintrich, 1997). At the same time, academic writing conventions are shaped by cultural and educational traditions described in contrastive rhetoric (Kaplan, 1966) and subsequent discourse research (Connor, 2002). The rapid integration of large language models (LLMs) introduces a new dimension to these culturally embedded practices. While AI tools increasingly support idea generation and text formulation, they lack epistemic awareness and cannot assume responsibility for knowledge claims. Human–AI collaboration may therefore reflect culturally situated norms of authorship, intellectual autonomy, and knowledge evaluation. Despite growing research on AI-assisted writing, little is known about whether human–LLM interaction itself exhibits cross-cultural variation in academic discourse, as most existing studies rely primarily on surveys or experimental tasks rather than natural language data. To address this gap, the study proposes semi-structured student interviews as a cross-linguistic discourse resource suitable for corpus-informed NLP analysis. We examine how human–LLM interaction is linguistically constructed across three dimensions of epistemic positioning: epistemic agency, authority dynamics, and discourse-level stance toward AI outputs. By focusing on observable linguistic patterns, the study provides preliminary corpus-based insights into how students discursively position AI within academic knowledge construction across national contexts.

2. Related Work

2.1. Cross-Cultural Academic Writing and Digital Writing Contexts

Research in contrastive rhetoric shows that academic writing conventions vary across linguistic and educational traditions (Kaplan, 1966). Subsequent discourse studies document differences in textual organization, authorial positioning, and argumentative structure across academic cultures (Clyne, 1987; Connor, 2002). Corpus-informed work further highlights distinctive writing practices in Central and Eastern Europe shaped by historical educational traditions (Chitez et al., 2018). In the Romanian context, analyses of BA theses reveal systematic differences in rhetorical structuring and epistemic positioning (Băniceru et al., 2012), which continue to influence students' English academic writing, particularly in stance expression (Bercuci and Chitez, 2023). At the same time, research on digital writing environments shows that writing technologies reshape knowledge construction processes by supporting iterative drafting, feedback in-

Proceedings of Shaping Multilingual, Multimodal AI for the Social Sciences and Humanities (LLMs4SSH) @ LREC 2026, pages 39–44
11 May 2026.

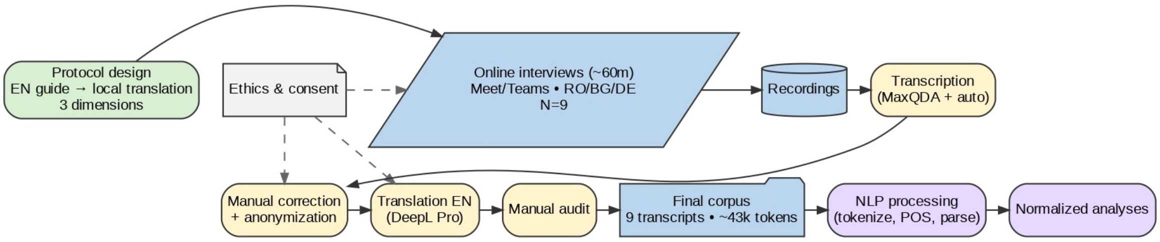

Figure 1: Pipeline

tegration, and interaction with information sources (Kruse and Rapp, 2019). Broader syntheses indicate that digital tools influence epistemic engagement by mediating how students evaluate and assume responsibility for knowledge claims (Kruse et al., 2023).

2.2. Epistemic Positioning and AI-Assisted Writing

Corpus-based research demonstrates that epistemic positioning is systematically encoded through recurrent lexico-grammatical patterns such as stance verbs, modality, and complement structures (Biber et al., 1999). These patterns have been operationalized in corpus and NLP studies to model certainty, evaluation, and authority in academic discourse (Biber, 2006; Biber et al., 2004). Such approaches provide a methodological basis for computational analyses of epistemic agency and stance. Recent work on AI-assisted writing shows that generative systems increasingly support planning, drafting, and revision processes in higher education (Strobl et al., 2019; Imran and Almusharraf, 2023). Empirical syntheses further indicate that generative AI reshapes authorship dynamics and knowledge-construction practices (Sanz-Tejeda et al., 2026). However, NLP research highlights persistent epistemic challenges, as language models may express certainty without reliable evidential grounding (Ghafouri et al., 2024).

3. Data and Methods

3.1. Data collection pipeline

We conducted nine semi-structured interviews across three national contexts (Romania, Bulgaria, Switzerland; three participants per site). Interviews were carried out in participants' native languages (Romanian, Bulgarian, German) to support accurate reflection on writing practices. To ensure cross-context comparability, the interview guide was collaboratively designed in English and then translated locally; it covered (i) intellectual development, (ii) thesis writing practices, and (iii) the role of generative AI. Interviews were held online (Google Meet / Microsoft Teams), lasting approximately one hour each. Recordings were transcribed using MaxQDA (GmbH, 2026), for the Romanian and Swiss datasets, and an in-house model based on the Open AI Whisper speech-to-text model[1] for the Bulgarian dataset, with subsequent intensive manual correction (approx. 4 hours per interview), during which all identifying information was removed to ensure full anonymization. The verified transcripts were translated into English using DeepL Pro (DeepL, 2026), and translations were manually audited to preserve meaning and reduce interpretive drift. Ethical safeguards (see Figure 1) included written informed consent, secure storage and restricted access to recordings/transcripts, and the use of privacy-compliant tools for transcription and translation.

3.2. Dataset

The dataset consists of nine anonymized semi-structured interview transcripts, totaling approximately 43,000 tokens in English translation. The corpus is evenly distributed across three national contexts, Bulgaria, Switzerland, and Romania, with three interviews included in each sub-corpus. However, substantial variation exists in document length. Bulgarian interviews range from approximately 4,100 to 4,600 tokens, Swiss interviews from about 3,200 to 5,200 tokens, and Romanian interviews from roughly 3,400 to 9,600 tokens, the latter showing the greatest internal variability and containing the longest individual transcript in the dataset. Given this heterogeneity in transcript size, all quantitative analyses were conducted using normalized frequency measures to ensure comparability across interviews and national sub-corpora.

3.3. Computational Pipeline and Operationalization

All computational linguistic processing was conducted using the Stanza NLP library (version 1.11.0) configured for English. We utilized the `combined` models for tokenization and

[1] https://github.com/openai/whisper

multi-word token (MWT) expansion, the `combined_nocharlm` model for lemmatization, and the `combined_charlm` models for part-of-speech (POS) tagging and dependency parsing.

To mitigate parsing and tagging errors inherent in translated conversational discourse, we implemented a multi-stage validation protocol. First, spoken-language artifacts (e.g., timestamps, speaker labels, platform artifacts) were programmatically removed using regular expressions to prevent dependency parsing disruption. The transcripts, which were translated into English using DeepL Pro, were manually audited prior to NLP processing to preserve semantic meaning and reduce interpretive drift. Following the automated extraction of first-person agency constructions, a random sample ($N = 15$) of the dependency-parsed outputs was extracted. This subset underwent manual verification by the research team to confirm the accuracy of the first-person subject tagging, verb lemmatization, and the subsequent semantic class mapping, ensuring algorithmic reliability before conducting the full corpus analysis.

Within this pipeline, AI-related contexts were operationalized using a dictionary-based sentence-window approach. We defined an explicit lexicon of target terms: *chatgpt, gpt, gpt-4, openai, copilot, claude, gemini, bard, llm,* and *generative ai*. The transcript texts were lowercased, and sentences were flagged if they contained any of these strings. To accurately link these mentions to student agency, we employed a co-occurrence window of ± 1 sentence. Specifically, a first-person agency verb was classified as "AI-related" if a target term appeared in the exact same sentence, the immediately preceding sentence, or the immediately following sentence. While this approach prioritizes precision over recall, it introduces specific edge cases. Primarily, it misses broader co-references; if a participant names an AI tool and subsequently refers to it as "the tool" or "it" outside the ± 1 sentence window, the associated agency verbs are not captured. Furthermore, because the algorithm relies on direct string matching, there is a minor risk of substring collisions, though the specificity of the chosen lexicon largely mitigates this risk.

4. Results

4.1. Agency Strength in Human–LLM Interaction

Weak or distancing agency markers overwhelmingly dominate the corpus, accounting for between 86.9% and 90.3% of all constructions across national contexts. These include hedging verbs such as *think* and procedural verbs such as *use*, which primarily describe AI interaction in terms of assistance rather than knowledge evaluation. Students frequently frame their interaction through statements such as “I use AI mainly for translation,” emphasizing functional support rather than epistemic authority. Moderate agency markers, representing between 4.8% and 7.6% of instances, include inquiry and evaluative verbs such as *ask* and *consider*, which signal active engagement without strong epistemic commitment. These constructions typically describe exploratory interaction, as illustrated by expressions such as “I ask whether the text makes sense.” Strong agency markers remain relatively rare, accounting for only 4.9% to 6.8% of constructions. These include verification and control verbs such as *check*, *decide*, and *adjust*, which explicitly signal responsibility for evaluating AI outputs. When such verbs occur, they consistently emphasize human oversight, for example in statements such as “I check whether the information is accurate.”

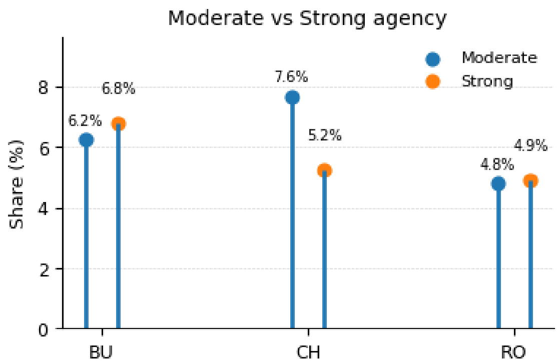

Figure 2: Epistemic positioning

Cross-country differences are present but modest. Swiss interviews show the highest proportion of moderate agency markers (7.6%), suggesting a slightly stronger orientation toward inquiry and evaluative engagement. Bulgarian interviews exhibit the highest proportion of strong agency markers (6.8%), indicating a somewhat greater tendency to linguistically express verification or decision-making processes. Romanian interviews display the lowest levels of both moderate (4.8%) and strong (4.9%) agency markers, reflecting a stronger reliance on distancing or procedural descriptions.

4.2. Authority Dynamics in AI-Related Discourse

To examine how epistemic authority is linguistically constructed when students explicitly refer to AI, we analyzed agency markers occurring in AI-related contexts and grouped them into delegation, hedging, and supervision functions. Across all national sub-corpora (Figure 3), delegation markers overwhelmingly dominate AI-related discourse. They account for approximately 65% of epistemic expressions in Romania, 70% in Bulgaria, and 74% in Switzerland. These patterns indicate that students primarily describe AI interaction through task-oriented verbs such as *use*, *write*, and *find*, which frame AI as a functional writing assistant rather than a knowledge source requiring evaluation.

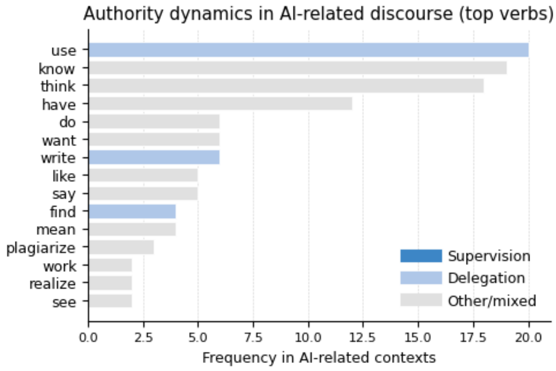

Figure 3: Authority verbs (all data)

Hedging markers constitute a secondary layer of epistemic positioning, representing roughly 27–30% of AI-related expressions across countries. This suggests that when discussing AI, students frequently maintain linguistic distance from the reliability of its outputs, signaling uncertainty or caution rather than authority. In contrast, supervision markers remain consistently low across contexts. They account for only 7.7% of AI-related agency constructions in Romania, 9.8% in Bulgaria, and 7.5% in Switzerland. These markers include verbs associated with verification, correction, and decision-making, and their limited presence indicates that explicit linguistic expressions of epistemic control over AI outputs are relatively rare. Cross-country variation (Figure 4) is modest but systematic. Bulgarian interviews show the highest relative proportion of supervision markers, suggesting a slightly stronger orientation toward linguistic expressions of oversight. Swiss interviews display the strongest dominance of delegation markers, indicating a more pronounced framing of AI as a task-execution tool. Romanian interviews fall between these patterns but show the highest relative presence of hedging, reflecting a somewhat stronger tendency toward epistemic distancing.

4.3. Discourse-Level Epistemic Stance Toward AI

The epistemic stance categories in Figure 5 capture three complementary dimensions of how students linguistically position knowledge in AI-related discourse. *Hedging markers* include modal verbs and stance expressions (e.g., *may*, *might*, *perhaps*, *I think*) that signal epistemic uncertainty and reduced

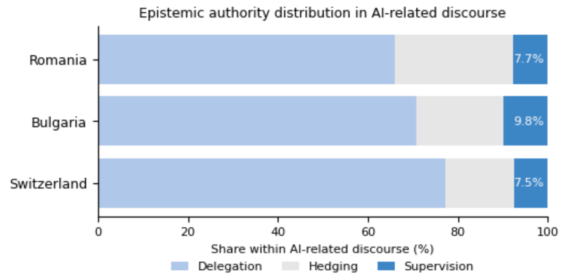

Figure 4: Authority verbs (per country)

commitment, consistent with corpus-linguistic accounts of modality (Biber et al., 1999; Biber, 2006). *Attribution markers* assign information to external agents (e.g., “AI says,” “it suggests”), thereby shifting epistemic authority away from the speaker. *Evaluation markers* consist of adjectives expressing judgments of informational quality (e.g., *accurate*, *helpful*, *problematic*). Importantly, hedging is operationalized here at the discourse level rather than as first-person epistemic verbs used in earlier agency analyses. Thus, the two measures capture complementary aspects of epistemic positioning: agency-related hedging reflects responsibility for knowledge claims, whereas modality-based hedging reflects degrees of certainty toward information.

Discourse-level stance markers show a clear predominance of hedging across all national subcorpora (normalized per 1,000 words). Hedging occurs at approximately 6.9 in Bulgarian interviews and 9.3–9.4 in Swiss and Romanian datasets, making it the most frequent epistemic feature. Attribution markers remain rare (0.23–1.29), while evaluation markers occur at intermediate levels (4.33–6.31). These patterns indicate that AI-related discourse is characterized primarily by epistemic uncertainty rather than explicit attribution or evaluation. Cross-country differences are modest but consistent: Swiss interviews show the highest levels of hedging and evaluation, whereas Bulgarian interviews display the lowest attribution rates, suggesting that students generally frame AI-supported writing through caution rather than trust or authority.

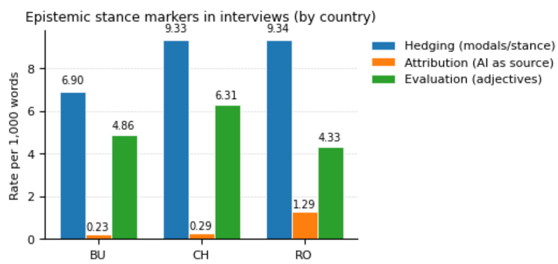

Figure 5: Epistemic stance markers

5. Discussion

This exploratory study provides corpus-based evidence that human–LLM interaction in academic writing is linguistically constructed through systematic patterns of epistemic positioning. Across contexts, students predominantly use distancing and hedging strategies, framing AI as a functional support tool rather than an epistemic authority. At the same time, modest but consistent cross-country differences suggest culturally shaped variation in how epistemic responsibility is discursively negotiated. Higher levels of verification language in the Bulgarian data and stronger evaluative engagement in the Swiss interviews may reflect differing academic socialization practices regarding knowledge control, autonomy, and critical assessment. In contrast, the Romanian interviews show a stronger reliance on procedural and distancing formulations, which may indicate a greater tendency to frame AI as a technical writing aid rather than an epistemic partner, consistent with previous findings on locally shaped academic discourse conventions (Kaplan, 1966; Connor, 2002).

Methodologically, the study demonstrates the value of combining qualitative interview data with corpus-informed, computationally tractable feature extraction to analyze epistemic positioning in natural discourse rather than self-reported attitudes. However, the findings remain exploratory due to the small sample size, reliance on translated transcripts, and the focus on reported rather than observed writing practices. Future research should extend this approach to larger multilingual datasets, integrate writing-process data, and develop automated stance-detection methods to further investigate culturally embedded patterns of human–AI collaboration.

5.1. Limitations and Robustness

The findings regarding agency and epistemic authority are highly sensitive to the strict syntactic operationalization employed in this study. The extraction algorithm is constrained to explicit first-person pronouns (e.g., *I, we, my*) that hold an exact nominal subject (`nsubj`) dependency relation to a head token tagged as a verb (`VERB`) or auxiliary (`AUX`). Consequently, this operationalization is conservative. It successfully isolates explicit claims of agency but does not parse complex evaluative multi-word expressions (e.g., "I am of the opinion that") or passive constructions ("It was checked by me"), which may result in an underestimation of total epistemic positioning. Furthermore, shifting to a modality-only operationalization, which captures adverbs or modal verbs (e.g., *may, might, perhaps*) regardless of the grammatical subject, yields distinct frequency distributions. This complementary dynamic is evident in our data, where discourse-level stance markers highlight uncertainty, while first-person agency markers emphasize functional delegation.

Additionally, because the current computational pipeline processes the English translations exclusively, systemic shifts in modality or hedging introduced by the DeepL translation algorithm constitute a limitation of the current findings. To rigorously estimate these translation effects, future iterations of this methodology should leverage multilingual NLP capabilities. By instantiating native language pipelines for Romanian, Bulgarian, and German, researchers could execute the identical dependency extraction on the original source-language transcripts. Mapping these native verbs to our epistemic classes and comparing their normalized frequencies against the English-translated results would allow for a quantitative assessment of whether machine translation processes artificially inflate or deflate specific markers of stance and agency across national sub-corpora.

6. Conclusions

This study shows that human–LLM interaction in academic writing is systematically shaped by patterns of epistemic positioning, with students across contexts predominantly framing AI as a functional support tool rather than an epistemic authority. At the same time, cross-country differences indicate that AI-supported writing practices remain embedded in culturally specific norms of knowledge construction, responsibility, and evaluation. These findings highlight the importance of analyzing human–AI interaction through multilingual discourse data and interpretable linguistic features, demonstrating how language technologies can support empirical investigations of socially situated communication practices. By operationalizing epistemic positioning in naturalistic interview discourse, the study illustrates how computational methods can provide theoretically grounded tools for examining emerging forms of AI-mediated knowledge work in academic and research contexts. Future research can extend this approach to larger multilingual datasets, incorporate writing-process and multimodal evidence, and develop automated stance-detection methods to further support interdisciplinary investigations of human–AI collaboration.

Acknowledgements

This work was supported by the Swiss National Science Foundation (SNSF) under the Multilateral Academic Projects (MAPS) grant (No. IZ11Z0_230816), awarded for the project “Exploring the Triad of Academic Writing, Critical Thinking and AI Literacy: Interdisciplinary Perspectives on

Undergraduate Thesis Writing" (TRAI), co-financed by the Romanian Ministry of Education and Research through UEFISCDI (Grant No. 15ROCH / 01/07/2025) and by the Bulgarian Ministry of Education and Science (Reg N D03-34/30.07.2025).

7. Bibliographical References

John C. Bean. 2011. *Engaging Ideas: The Professor's Guide to Integrating Writing, Critical Thinking, and Active Learning in the Classroom*, 2 edition. Jossey-Bass.

L. Bercuci and M. Chitez. 2023. The influence of the romanian academic style on student writing in english. In M. Vránová, editor, *New to the LSP classroom? A selection of monographs on successful practices*, pages 169–185. Vernon Press.

D. Biber. 2006. *University language: A corpus-based study of spoken and written registers*. John Benjamins.

D. Biber, S. Conrad, and V. Cortes. 2004. If you look at…: Lexical bundles in university teaching and textbooks. *Applied Linguistics*, 25(3):371–405.

D. Biber, S. Johansson, G. Leech, S. Conrad, and E. Finegan. 1999. *Longman grammar of spoken and written English*. Longman.

C. Băniceru, M. I. Borchin, C. Doroholschi, and D. Tucan. 2012. Academic writing in romania: A contrastive analysis of ba thesis introductions in romanian and english. *Quaestiones Romanicae*, 1:331–344.

M. Chitez, C. I. Doroholschi, O. Kruse, Ł. Salski, and D. Tucan, editors. 2018. *University writing in Central and Eastern Europe: Tradition, transition, and innovation*. Springer.

Michael Clyne. 1987. Cultural differences in the organization of academic texts: English and German. *Journal of Pragmatics*, 11(2):211–241.

Ulla Connor. 2002. New directions in contrastive rhetoric. *TESOL Quarterly*, 36(4):493–510.

DeepL. 2026. Deepl translate.

B. Ghafouri, S. Mohammadzadeh, J. Zhou, P. Nair, J. J. Tian, H. Tsujimura, and K. Pelrine. 2024. Epistemic integrity in large language models.

VERBI GmbH. 2026. Maxqda.

Barbara K. Hofer and Paul R. Pintrich. 1997. The development of epistemological theories: Beliefs about knowledge and knowing and their relation to learning. *Review of Educational Research*, 67(1):88–140.

Muhammad Imran and Norah Almusharraf. 2023. Analyzing the role of chatgpt as a writing assistant at higher education level: A systematic review. *Contemporary Educational Technology*, 15(4):ep464.

Robert B. Kaplan. 1966. Cultural thought patterns in inter-cultural education. *Language Learning*, 16(1–2):1–20.

O. Kruse and C. Rapp. 2019. Seamless writing: How the digitisation of writing transforms thinking, communication, and student learning. In *Seamless learning: Perspectives, challenges and opportunities*, pages 191–208. Springer.

O. Kruse, C. Rapp, C. M. Anson, K. Benetos, E. Cotos, A. Devitt, and A. Shibani. 2023. *Digital writing technologies in higher education: Theory, research, and practice*. Springer.

Aránzazu Sanz-Tejeda, Juana Celia Domínguez-Oller, Josep Maria Baldaquí-Escandell, Raquel Gómez-Díaz, and Araceli García-Rodríguez. 2026. The impact of generative AI on academic reading and writing: a synthesis of recent evidence (2023–2025). *Frontiers in Education*, 10:1511718.

Carolin Strobl, Emily Ailhaud, Katerina Benetos, Amy Devitt, Otto Kruse, Antje Proske, and Christian Rapp. 2019. Digital support for academic writing: A review of technologies and pedagogies. *Computers & Education*, 131:33–48.