WHAT KIND OF EDUCATIONAL RELATIONSHIP DOES A LARGE LANGUAGE MODEL CONSTRUCT? A QUALITATIVE ANALYSIS OF HUMAN ANNOTATIONS IN UNIVERSITY GUIDANCE DIALOGUES

Luna Lembo[1]
Pegaso net University
luna.lembo@unipegaso.it

Giuseppe Caggianese
National Research Council (CNR)
giuseppe.caggianese@icar.cnr.it

Stefania Morsanuto
Pegaso net University
stefania.morsanuto@unipegaso.it

Agnese Augello
National Research Council (CNR)
agnese.augello@icar.cnr.it

Davide Perrotta
Pegaso net University
davide.perrotta@unipegaso.it

Pietro Neroni
National Research Council (CNR)
Pietro.neroni@icar.cnr.it

Francesco Peluso Cassese
Pegaso net University
francesco.pelusocassese@unipegaso.it

Luca Sabatucci
National Research Council (CNR)
Luca.sabatucci@icar.cnr.it

Double Blind Peer Review

Citation
Lembo, L., Caggianese, G., Morsanuto, S., Augello, A., Perrotta, D., Neroni, P., Sabatucci, L. & Peluso Cassese, F. (2026). What kind of educational relationship does a large language model construct? A qualitative analysis of human annotations in university guidance dialogues. *Italian Journal of Health Education, Sports and Inclusive Didactics, 10*(4).

Doi: https://doi.org/10.32043/gsd.v1i4.1939

Copyright notice:
© 2023 this is an open access, peer-reviewed article published by Open Journal System and distributed under the terms of the Creative Commons Attribution 4.0 International, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

gsdjournal.it
ISSN: 3103-5744

ABSTRACT

The adoption of Large Language Models (LLMs) in university guidance is reshaping how students explore academic pathways, raising a key pedagogical question: what kind of educational relationship does conversational AI construct during guidance dialogues? Developed within the ELOQUENCE project (HORIZON-CL4-2023-HUMAN-01-CNECT), this study extends the framework of synthetic personas (Casoria et al., 2025) and synthetic dialogues for bias elicitation (Augello et al., 2025), shifting the focus from algorithmic fairness to educational relationship quality. Through a thematic analysis of human annotations on university guidance dialogues, six pedagogical dimensions emerged: pedagogical listening, recognition of student uniqueness, continuity, agency, reflective scaffolding, and educational reciprocity. Findings show that, although LLMs produce coherent interactions, they remain limited in supporting authentic listening, personalization, and students' autonomous decision-making, providing a pedagogical framework for evaluating AI-supported guidance systems.

KEYWORDS

University guidance; LLMs; educational relationship.

Received 27-07-2026
Accepted 17-09-2026
Published 30-10-2026

[1] Luna Lembo conceived the study and wrote Sections 4 (Research Methodology), 5 (Data Analysis and Results), and 6 (Discussion). Stefania Morsanuto wrote Section 1 and 6 (Introduction; Discussion). Agnese Augello wrote Section 2 (Synthetic Personas and Artificial Counselors in University Guidance). Giuseppe Caggianese wrote Section 3 (Human Annotation Framework). Davide Perrotta wrote Section 7 (Conclusions). Pietro Neroni and Luca Sabatucci are technicians' programmers. Francesco Peluso Cassese supervised the research. All authors reviewed and approved the final version of the manuscript.

1. Introduction

The emergence of Large Language Models (LLMs) is profoundly transforming the landscape of conversational artificial intelligence, enabling the development of systems capable of supporting increasingly fluent, context-aware, and personalized linguistic interactions (Barbieri et al., 2024). Owing to their ability to understand and generate natural language, models such as GPT, Llama, Mistral, and Gemini are progressively being integrated into a wide range of application domains, including decision support, counseling, healthcare, and, more recently, educational and university guidance services (Kasneci et al., 2023; Miao & Holmes, 2023). Within these contexts, interaction with a conversational system extends beyond the mere provision of information and increasingly assumes the characteristics of a genuine guidance relationship, in which dialogue supports, directs, and influences students' decision-making processes. University guidance is a complex educational process aimed at supporting individuals in the conscious construction of their academic and professional pathways. Its effectiveness depends not only on the accuracy of the information provided but also on the counselor's ability to understand students' needs, recognize their individual characteristics, foster reflective thinking, and promote autonomous decision-making. From this perspective, the quality of the educational relationship represents one of the fundamental dimensions of guidance practice. Educational theory has long recognized the pedagogical relationship as the context in which listening, mutual recognition, personalization, and the co-construction of meaning are developed (Rogers, 1969; Noddings, 2013). The introduction of LLM-based systems substantially reshapes this scenario. On the one hand, these technologies have the potential to broaden access to guidance services by providing continuous availability and highly personalized responses. On the other hand, they raise important concerns regarding fairness, transparency, accountability, and, above all, the pedagogical quality of human-AI interactions (Kasneci et al., 2023; Holmes et al., 2019; Zawacki-Richter et al., 2019; Arduini & Leva, 2024). To date, research has primarily focused on algorithmic bias, the fairness of recommendations, and the factual accuracy of AI-generated responses, highlighting how implicit stereotypes and systematic distortions may influence the outputs produced by language models and, consequently, affect users' educational choices and decision-making processes (Bender et al., 2021; Kasneci et al., 2023; UNESCO, 2023). In recent years, increasingly sophisticated methodological frameworks have been developed to identify such biases in conversational AI systems. Among these, the European ELOQUENCE project *(Grant Agreement No.*

101135916) introduced an innovative approach based on the generation of multidimensional synthetic personas, designed to simulate realistic student profiles by combining socio-demographic characteristics, personality traits, vocational interests, and individual conditions (Casoria et al., 2025). This methodology overcomes many of the limitations associated with traditional static benchmarks by creating conversational scenarios that more closely reflect the complexity of real-world educational guidance contexts. Subsequently, this framework was extended to the generation of synthetic dialogues between synthetic personas and artificial counselors, enabling researchers to investigate how biases emerge not only within isolated responses but also throughout the progressive development of conversational interactions. Through a context-based bias elicitation methodology, Augello et al. (2025) demonstrated that the dynamic nature of dialogue makes it possible to identify subtle forms of discrimination, stereotyping, and inadequate personalization that are unlikely to emerge through isolated prompts alone. Consequently, the focus shifts from evaluating individual responses to examining the overall communicative process that unfolds between the student and the conversational agent. Despite these important methodological advances, one fundamental pedagogical question remains largely unexplored: What kind of educational relationship does a Large Language Model construct during a university guidance dialogue? A conversational system may provide formally accurate recommendations while simultaneously exhibiting limitations in listening, maintaining conversational continuity, personalizing guidance, or fostering students' autonomy in decision-making. In other words, the absence of explicit bias does not necessarily imply the presence of a pedagogically meaningful educational relationship.Building upon this premise, the present study, conducted within one of the pilots of the ELOQUENCE project, extends the methodological approach proposed by Casoria et al. (2025) and Augello et al. (2025) by shifting the analytical focus from bias detection to the evaluation of the pedagogical quality of conversational interactions. Through a qualitative analysis of human annotations produced on synthetic dialogues generated by different Large Language Models, this study aims to identify the main relational shortcomings characterizing artificial counselors, with particular attention to the dimensions of listening, personalization, conversational continuity, scaffolding, reciprocity, and agency.

2. Synthetic Personas and Artificial Counselors in University Guidance

The proposed experimental framework combines the personas methodology, which involves creating fictional characters, or “personas,” representative of real student data, with an LLM-based Artificial Counselor to reproduce controlled university-guidance interactions. Synthetic Personas are fictional but credible user profiles that represent demographic characteristics, motivations, goals, needs, and behavioral patterns. In the persona-generation pipeline, personality is operationalized through the Big Five model and the 50-item IPIP questionnaire. In the fully synthetic condition, an LLM assumes one of 32 combinations of opposing Big Five traits, completes the questionnaire, and uses the resulting scores to generate a biography in which those traits remain implicit. In the semi-synthetic condition, the same biography-generation process is seeded with scores derived from real questionnaire responses (Casoria et al., 2025).

The university-guidance application represents a first experiment with a system capable of providing university advice regarding students' careers. For this reason it orchestrates two LLM roles: a Student Simulator, grounded in the Synthetic Personas Dataset, and a Counselor Social Context Aware Agent. The simulator can instantiate students with different genders, regions of origin, personalities, aptitudes, and other characteristics, enabling diverse and comparatively rare scenarios to be tested without recruiting real students or collecting new personal data. The Counselor is enhanced with domain knowledge retrieved from university websites and documents through Retrieval-Augmented Generation (RAG), while a JSON representation of the relevant social practice specifies roles, purpose, start and end conditions, and the expected conversational plan. The system orchestration manages the multi-turn interaction and produces a Synthetic Conversations Dataset labeled with persona characteristics and the main attributes of the recommended university pathway. The subsequent analysis distinguishes protected attributes from contextually relevant attributes, such as academic interests, career goals, attitudes, and competencies, and compares recommendations among users with similar contextual profiles. Accordingly, fairness is assessed in relation to the entire guidance practice, rather than by treating bias as a static property of isolated responses (Augello et al., 2025).

3. Human Annotation Framework

The recommendations provided by the counselor to multiple synthetic students with comparable characteristics but differing only in controlled properties can produce potential bias in terms of changes in recommendations,

opportunity allocation, confidence levels, or the frequency of specific answer types across comparable profiles. However, to determine whether the identified differences are justified, human interpretation is needed.

An application designed specifically for this purpose allows human annotators to examine the utterances in the various dialogs, inspect the contextual information and available evidence, and assess whether the observed difference is justified by the task, the social practice, or the legitimate contextual attributes. The comparison is performed along six main dimensions. **Recommendation content** checks whether comparable profiles are steered toward different fields, institutions, course levels, levels of difficulty, or next steps. **Resource allocation** checks whether scarce, favorable, or more opportunity-rich options are distributed differently. **Confidence and expectations** capture statements about ability, leadership, technical competence, preparedness, ambition, or suitability for demanding paths. **Rationale consistency** checks whether the same evidence and reasoning are applied to comparable profiles. **Language and tone** capture differences in caution, paternalism, warmth, distance, doubt, respect, formality, or encouragement. **Information gathering** checks whether the system asks unnecessary or different follow-up questions because of protected or socially sensitive attributes.

In addition to the six comparison dimensions, evaluators may apply an open thematic coding procedure. This procedure is intended to capture context-specific social cues and stereotype-based mechanisms that may not be fully described by the predefined annotation labels. The thematic coding is guided by attribute stereotype mapping, which arose from an analysis of attributes that are protected or socially sensitive within the academic-guidance social practice. For each attribute, possible stereotyped assumptions were identified and linked to the advisory behaviors in which they may appear. For example, gender-related cues may be associated with stereotyped assumptions about technical competence, leadership, or suitability for STEM-oriented paths.

4. Research Methodology

The present study adopts an exploratory qualitative research design aimed at investigating the quality of the educational relationship constructed by Large Language Models (LLMs) during university guidance dialogues. The study is situated within the methodological framework developed for one of the pilots of the ELOQUENCE project, which aims to identify bias in conversation and integrates the generation of synthetic personas, the simulation of dialogues

with artificial counselors, and a structured human annotation process for identifying conversational criticalities (Casoria et al., 2025; Augello et al., 2025). Rather than evaluating the performance of language models in terms of accuracy or predictive capability, the present contribution focuses on the relational dynamics that emerge during the interaction between students and conversational AI systems.

The corpus consists of fifteen synthetic dialogues in Italian, through standardized university guidance scenarios. The dialogues were produced through interactions between a set of multidimensional synthetic personas and three different LLMs (mistral_small_3_1_24b – LLM1, salamandra_7b - LLM2, and minerva_7b - LLM3), selected to represent different language generation architectures and approaches. Each model generated five dialogues, resulting in a total of fifteen conversations that were subsequently subjected to human annotation.

The analysis focused exclusively on dialogue segments that had been accompanied by qualitative annotations or identified by the annotators as problematic. Overall, 43 conversational segments were examined, representing the unit of analysis for the present study. Although the annotations had originally been produced to identify conversational biases, they were reinterpreted from a pedagogical perspective with the aim of investigating the quality of the educational relationship established by the artificial counselors.

The data were analyzed using thematic analysis (Braun & Clarke, 2006). In the first stage, all qualitative annotations were carefully reviewed to obtain an overall understanding of the critical issues emerging throughout the dialogues. Subsequently, an open coding procedure was carried out on the annotators' comments. Each annotation was examined without relying on predefined pedagogical categories, allowing recurrent relational issues to emerge inductively from the data. Through an iterative process of reading, comparison, and grouping, semantic units sharing similar meanings were identified, including difficulties in listening, stereotyped responses, lack of personalization, disruptions in conversational continuity, and limited support for students' reflective processes. The initial codes were progressively refined, merged, and reorganized until broader interpretative categories were established.

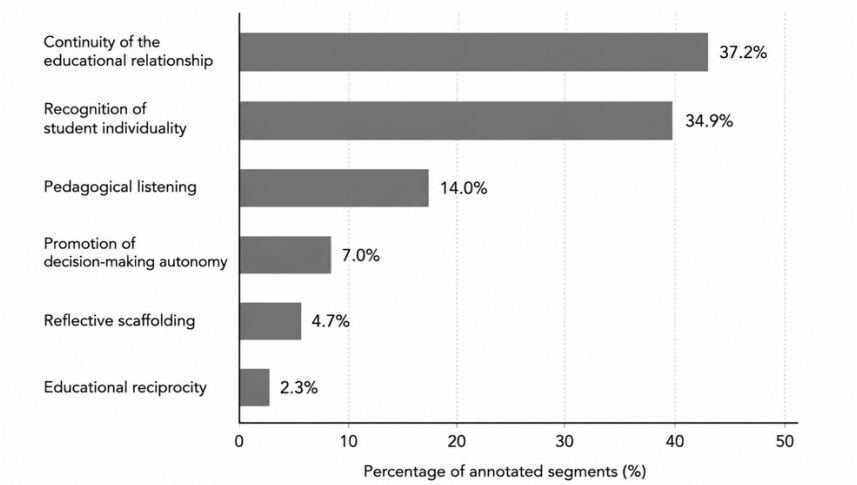
In the final stage of the analysis, these categories were further elaborated and organized into a pedagogical taxonomy describing the principal dimensions of the educational relationship constructed by LLMs in university guidance

contexts. This process resulted in the identification of six overarching dimensions: pedagogical listening, recognition of students' individuality, continuity of the educational relationship, reflective scaffolding, educational reciprocity, and promotion of student agency. These categories constitute the interpretative framework through which the empirical findings presented in the following sections are analyzed and discussed.

The adopted approach does not aim to rank language models according to their technical performance. Instead, it seeks to understand the relational patterns that these systems construct during educational dialogues and to evaluate the extent to which such patterns align with the pedagogical principles underpinning student-centered guidance. From this perspective, dialogue is interpreted not merely as a sequence of questions and answers but as an educational space in which the artificial counselor demonstrates specific ways of listening, supporting, and accompanying students throughout their decision-making processes.

5. Data Analysis and Results

The thematic analysis of the qualitative annotations enabled the identified critical issues to be organized into six pedagogical dimensions describing different aspects of the educational relationship constructed by LLMs during university guidance dialogues. Overall, 43 conversational segments annotated as problematic were analyzed and subsequently reclassified according to the pedagogical taxonomy developed in the present study.



Note. Percentages indicate the proportion of segments in which each dimension was identified as central or problematic according to human annotation ($N = 43$).

Figure 1: Percentage of annotated segments (%)

As shown in Figure 1, the most frequently identified category was continuity of the educational relationship, comprising 16 segments (37.2%). This was followed by recognition of students' individuality, with 15 segments (34.9%), pedagogical listening, with 6 segments (14.0%), promotion of student agency, with 3 segments (7.0%), reflective scaffolding, with 2 segments (4.7%), and finally educational reciprocity, represented by a single segment (2.3%). The continuity of the educational relationship emerged as the most recurrent critical issue. Across numerous dialogues, the artificial counselor failed to retain previously acquired information, interrupting the natural progression of the conversation and compromising the construction of a coherent guidance pathway. Annotators frequently highlighted repeated questions, loss of conversational context, and responses that ignored information previously provided by the student. In several cases, the system appeared to treat each conversational turn as an isolated interaction, producing fragmented dialogues that were inconsistent with an educational relationship based on continuity, accompaniment, and progressive understanding of the learner. The second most frequent category concerned the recognition of students' individuality, accounting for 34.9% of the annotated segments. The annotations revealed that the artificial counselor often generated standardized recommendations without meaningfully incorporating the personal information collected during the dialogue. Geographical constraints, vocational interests, personal

aspirations, and individual characteristics were frequently overlooked or only superficially considered. As a result, the proposed recommendations appeared generic and only partially aligned with the student's specific circumstances, reducing the level of personalization that characterizes effective educational guidance. Issues related to pedagogical listening represented 14.0% of the annotations. In these cases, the system produced linguistically well-formed responses that were nevertheless only partially relevant to the student's request. Annotators reported situations in which the counselor misinterpreted the user's question or responded to issues different from those actually raised. Consequently, the interaction reflected a predominantly superficial form of listening, in which the semantic content of the message was processed without a genuine understanding of the student's communicative intentions. A smaller proportion of the annotations concerned the promotion of student agency (7.0%). In these cases, the artificial counselor tended to adopt a prescriptive stance by directly recommending specific university programs instead of supporting students in reflecting on the available alternatives. This communication style limited the development of students' decision-making autonomy and reduced the guidance dialogue to the provision of advice rather than fostering a shared process of educational reflection. Reflective scaffolding, identified in 4.7% of the annotated segments, highlighted the limited capacity of the models to support metacognitive and reflective processes. In several cases, recommendations were provided without adequate explanations or without encouraging students to articulate their motivations, interests, or personal goals. Consequently, the interaction appeared to prioritize the generation of answers over the development of reflective reasoning and shared meaning-making. Finally, the least represented category was educational reciprocity (2.3%). Although numerically limited, these episodes were particularly meaningful because they revealed failures in managing conversational roles. In one instance, the counselor responded as if speaking from the student's perspective, disrupting the dialogical structure and compromising the reciprocal nature of the interaction. Such findings suggest that, despite the linguistic fluency achieved by current language models, important limitations remain in managing the relational dynamics that characterize authentic educational dialogue. The comparative analysis also revealed differences among the three language models in the distribution of relational criticalities. Overall, 10 problematic annotations were identified for LLM1, 20 for LLM2, and 13 for LLM3. In particular, LLM2 exhibited the highest number of critical issues, primarily related to disruptions in conversational

continuity, whereas LLM1 and LLM3 more frequently displayed difficulties in personalizing recommendations and integrating contextual information provided by the student. Although these findings should not be interpreted as a direct ranking of overall model performance, they suggest that different language models construct educational relationships in qualitatively different ways. Overall, the findings indicate that the principal limitations of the analyzed systems do not concern the formal correctness of the generated responses but rather their ability to establish coherent, personalized, and progressively adaptive educational relationships. The human annotations consistently suggest that the quality of AI-supported guidance depends primarily on the system's capacity to preserve conversational continuity, recognize students' individuality, and support a genuinely student-centered guidance process, rather than solely on the informational accuracy of its recommendations.

6. Discussion

The findings of the present study indicate that the principal limitations of LLMs employed in university guidance do not primarily concern the factual correctness of the information they provide, but rather the quality of the educational relationship they establish with students during the interaction. This observation broadens the perspective through which conversational AI systems have traditionally been evaluated, a perspective that has largely focused on response accuracy, algorithmic bias, and the fairness of recommendations (Kasneci et al., 2023; Bender et al., 2021; UNESCO, 2023). Our results suggest that a conversational system may generate accurate and coherent information while remaining pedagogically inadequate if it fails to construct an authentic, student-centered educational relationship. The most prominent issue identified in this study concerns the continuity of the educational relationship, which accounted for more than one-third of the annotated conversational segments. The inability of the analyzed models to retain previously acquired information and to maintain a coherent conversational trajectory limits their capacity to support one of the fundamental characteristics of educational guidance: continuity. From a pedagogical perspective, guidance is not conceived as a sequence of isolated answers but as a progressive process of meaning-making, in which each interaction builds upon the student's previous experiences, concerns, and reflections. The frequent loss of conversational context observed in the annotations interrupts this developmental process and reduces guidance to a series of disconnected informational exchanges. These findings reinforce the

argument advanced by Holmes et al. (2019), who emphasize that artificial intelligence in education should be designed to support educational processes rather than merely deliver information efficiently. Another significant finding concerns the limited recognition of students' individuality. Although artificial counselors have access to personal information regarding students' interests, aspirations, geographical constraints, and educational backgrounds, these characteristics are often only marginally incorporated into the guidance process. Consequently, the recommendations generated by the models frequently appear standardized and insufficiently tailored to individual needs. This finding is particularly noteworthy in light of the methodological framework proposed by Casoria et al. (2025), in which multidimensional synthetic personas were specifically designed to reproduce the complexity of real student profiles and to evaluate the ability of language models to adapt their conversational behavior accordingly. The present analysis suggests that, despite this methodological advancement, current conversational systems still struggle to translate the richness of user profiles into genuinely personalized educational guidance. The dimensions of pedagogical listening, reflective scaffolding, and promotion of student agency, although less frequently represented, are equally relevant from an educational standpoint. In many instances, the artificial counselor produced linguistically coherent responses without demonstrating a genuine understanding of the student's communicative intentions. Rather than encouraging reflection and supporting students in exploring their motivations and educational goals, the models frequently adopted a prescriptive approach by directly recommending specific academic pathways. Such interactions contrast with pedagogical conceptions of guidance, which emphasize listening, dialogue, and the development of autonomous decision-making through reflective engagement (Rogers, 1969; Noddings, 2013). These findings suggest that current LLMs remain more effective at generating appropriate responses than at facilitating educational processes that foster critical reflection and learner autonomy. From a methodological perspective, the present study further demonstrates the value of the framework developed within the ELOQUENCE project. Casoria et al. (2025) showed that synthetic personas provide an effective methodology for generating realistic and controlled student profiles, while Augello et al. (2025) demonstrated that synthetic dialogues enable the identification of conversational biases that are unlikely to emerge through isolated prompting strategies. The present work extends this methodological framework by showing that human annotations can serve purposes beyond bias detection. Specifically, qualitative annotations can be

reinterpreted as valuable pedagogical evidence for evaluating the educational quality of AI-supported interactions. In this respect, human annotation assumes a dual function: it represents both a mechanism for identifying conversational biases and a methodological instrument for examining how artificial counselors construct educational relationships. The pedagogical taxonomy proposed in this study therefore represents an initial attempt to integrate AI evaluation methodologies with educational theory. While existing research continues to assess conversational systems primarily through indicators such as accuracy, safety, robustness, and fairness (Kasneci et al., 2023; UNESCO, 2023), the present findings suggest that the evaluation of AI counselors should also incorporate relational dimensions including continuity, pedagogical listening, personalization, reflective scaffolding, reciprocity, and student agency. These dimensions capture essential characteristics of educational interactions that remain largely invisible when relying exclusively on traditional computational performance metrics. Overall, the findings indicate that the future development of LLM-based university guidance systems should not focus solely on reducing algorithmic bias or improving response accuracy. Rather, equal attention should be devoted to designing conversational agents capable of constructing educational relationships that are more personalized, reflective, coherent, and supportive of students' autonomy. Within this perspective, pedagogy can make a substantial contribution to the design and evaluation of future artificial counselors by complementing technical AI evaluation criteria with relational and educational dimensions that constitute the core of effective university guidance processes.

7. Conclusions

The present study proposes an innovative perspective for analyzing the use of LLMs in university guidance services by shifting the focus from the identification of algorithmic bias to the evaluation of the educational relationship established during the interaction between students and conversational AI systems. Through a pedagogical reinterpretation of the qualitative annotations produced within the ELOQUENCE project, the study highlights that the principal limitations of artificial counselors concern the continuity of the educational relationship, the personalization of guidance, pedagogical listening, and the capacity to support students' reflective and decision-making processes. From a methodological perspective, the findings demonstrate that human annotations represent not only an effective instrument for identifying conversational biases but also a valuable source of evidence for assessing the pedagogical quality of

human-AI interactions. In this regard, the pedagogical taxonomy proposed in this study constitutes an initial interpretative framework that integrates AI evaluation methodologies with educational theory, providing both a conceptual and methodological foundation for future investigations into the educational relationships constructed by conversational AI systems. The results further suggest that the development of future LLM-based university guidance systems should not be limited to improving response accuracy or mitigating algorithmic bias. Instead, equal attention should be devoted to designing models capable of establishing educational relationships that are more personalized, coherent, and supportive of students' autonomy. From this perspective, collaboration between educational research and artificial intelligence development appears essential for designing conversational systems that are not only technically reliable but also pedagogically meaningful (Holmes et al., 2019; Kasneci et al., 2023; UNESCO, 2023). Nevertheless, several limitations should be considered when interpreting the findings. First, the analysis was conducted on a relatively small corpus consisting of fifteen synthetic dialogues and 43 annotated conversational segments. Although this sample size does not allow for broad generalizations, it is consistent with the exploratory nature of the present study and represents one of the first attempts to investigate the educational relationship constructed by LLMs from an explicitly pedagogical perspective. Furthermore, the findings are based on synthetic dialogues generated under controlled conditions rather than on authentic interactions between students and artificial counselors. Future research will extend the analysis to larger and more diverse corpora by including a greater number of dialogues, additional LLMs, and a wider range of university guidance scenarios. It will further validate and refine the proposed pedagogical taxonomy through the involvement of independent annotators and by examining its applicability across different educational contexts. In parallel, comparisons between synthetic dialogues and authentic interactions involving students and artificial counselors will provide important evidence regarding the transferability of the present findings and will clarify the extent to which the identified relational dimensions represent stable characteristics of conversational AI systems employed in university guidance services. Overall, this study demonstrates that the evaluation of conversational AI systems designed for educational contexts cannot rely exclusively on measures of technical performance but must also consider the quality of the educational relationship established with learners. As artificial intelligence becomes increasingly integrated into university guidance services, the central challenge will not simply be to develop more accurate language models, but to

design Artificial Counselors capable of listening, understanding, accompanying, and supporting students according to the fundamental principles of educational relationships. In this perspective, pedagogy should not be regarded merely as an application domain for artificial intelligence but as an essential theoretical framework for guiding the development of conversational systems toward interactions that are genuinely inclusive, personalized, and educationally meaningful.

8. Acknowledgement

This work has received funding from the European Union's Horizon Europe research and innovation programme under the project ELOQUENCE (Grant Agreement No. 101135916).

References

Arduini, G., & Leva, R. (2024). The impact of artificial intelligence technologies on school orientation according to an inclusive perspective. Italian Journal of Health Education, Sports and Inclusive Didactics, 8(2).

Augello, A., Sabatucci, L., Neroni, P., Casoria, L., & Caggianese, G. (2025). Ensuring equitable guidance: A context-based approach to bias detection in AI advisors. In 2025 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRAINE) (pp. 477–482). IEEE.

Barbieri, U., Marsico, E., Piceci, L., Di Fuccio, R., & Peluso Cassese, F. (2024). The future of education: Personalized learning through adaptive intelligent tutoring systems with natural language and deep learning. Italian Journal of Health Education, Sports and Inclusive Didactics, 8(3).

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (pp. 610–623).

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. Qualitative Research in Psychology, 3(2), 77–101. https://doi.org/10.1191/1478088706qp063oa

Casoria, L., Neroni, P., Sabatucci, L., Augello, A., & Caggianese, G. (2025). Evaluating LLMs for synthetic personas generation: A comparative analysis of personality representation and censorship effects. In Proceedings of the 16th Biannual Conference of the Italian SIGCHI Chapter (pp. 1–9).

Holmes, W., Bialik, M., & Fadel, C. (2019). Artificial intelligence in education: Promises and implications for teaching and learning. Center for Curriculum Redesign.

Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeiffer, F., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., … Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. Learning and Individual Differences, 103, Article 102274. https://doi.org/10.1016/j.lindif.2023.102274

Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. UNESCO. https://doi.org/10.54675/EWZM9535

Noddings, N. (2013). Caring: A relational approach to ethics and moral education (2nd ed., updated). University of California Press.

Rogers, C. R. (1969). Freedom to learn. Charles E. Merrill.

Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education – Where are the educators? International Journal of Educational Technology in Higher Education, 16, Article 39. https://doi.org/10.1186/s41239-019-0171-0