Internship offer: Designing NLP tools for Student Digital Twins

Michele Linardi

July 2026

1 Research Context

At CY Cergy Paris University and the University of Toulouse, our research aims to develop trustworthy, robust, and human-centered Generative AI systems that adapt to students' learning habits and individual needs.

Generative AI tutors are rapidly becoming pervasive across both secondary and higher education. Students increasingly rely on AI-powered assistants for explanations, problem-solving, revision, and personalized feedback, making these systems an integral part of the modern learning experience. As these tools become embedded in education, it is essential that they move beyond generic conversational capabilities and provide support that is pedagogically sound, personalized, and trustworthy.

We argue that, despite the remarkable capabilities of large language models (LLMs), they are not inherently designed to account for the diverse requirements of individual learners. Current systems often overlook important factors such as a student's prior knowledge, learning preferences, misconceptions, and emotional state. Negative emotions such as frustration, anxiety, fear of failure, cognitive overload, and cognitive biases can significantly influence students' perception of the generated feedback. Addressing these limitations is therefore essential for the next generation of intelligent tutoring systems.

Our research investigates the integration of small language models (SLMs) language models [6] with fewer than one billion parameters that can be deployed efficiently on edge devices within an agentic AI educational tutor [3].

By leveraging compact, privacy-preserving models, we aim to provide adaptive tutoring while reducing computational costs and enabling on-device personalization. Within the tutoring process, the AI agent follows an agentic reasoning framework. It begins by receiving the student's input and establishing a clear instructional objective, which guides all subsequent decisions throughout the interaction.

The agent then analyzes the student's responses and contextual information to identify relevant patterns, infer the learner's current level of understanding, and assess the educational context. This reasoning process enables the tutor

to select appropriate pedagogical strategies and generate personalized interventions.

Our broader research objective is to analyze the current state of the art in AI-based tutoring systems [5, 1, 7, 4, 2] and develop a novel adaptive tutoring framework that more accurately models individual learners. To achieve this, we propose the integration of a student digital twin: a generative model that acts as a computational proxy for a specific learner. Such a model captures the learner's knowledge, behavioral patterns, preferences, and interaction history, allowing it to predict how the student is likely to respond to new questions or learning situations. By incorporating this digital twin into the tutoring loop, the AI tutor can anticipate student needs, personalize instructional strategies, and continuously adapt its guidance to maximize learning effectiveness while maintaining trustworthiness.

2 Internship Offer

Such an internship aims to investigate how linguistic cues and prompt specifications influence the quality and pedagogical effectiveness of student-demand content generation of agentic AI tutors. We hypothesize that analyzing both the student's language and the prompts provided to AI agents will enable us to assess the usefulness, appropriateness, and reliability of AI-generated educational responses.

A key component of this research is the development of a student digital twin capable of realistically simulating a learner's behavior and language in different pedagogical scenarios, such as exam preparation, revision sessions, or concept discovery. Beyond reproducing the types of questions that a student is likely to ask, the digital twin should also evaluate whether the AI tutor's responses align with the learner's expectations, knowledge level, and instructional objectives. Furthermore, it should ideally model factors that significantly influence learning, including emotional states such as frustration, confusion, anxiety, or loss of motivation. By incorporating these dimensions, it can assess not only the pedagogical relevance of generated content but also whether the responses are likely to support or hinder the student's learning process.

The methodology we propose to investigate the integration of Natural Language Processing (NLP) techniques to enhance this evaluation framework. In particular, the selected methods should enable the automatic detection of linguistic signals related to student engagement and emotional state, as well as the identification of hallucinations, biases, and other indicators of unreliable or pedagogically inappropriate AI-generated content. These analyses will contribute to the development of trustworthy, adaptive, and human-centered agentic tutoring systems.

The selected student will benefit from an existing research platform, including an implemented agentic AI tutoring prototype and an established evaluation benchmark. These resources will allow the intern to focus on designing, integrating, and evaluating novel NLP components, integrating the student's digital

twin.

3 Contact Persons

Michele linardi (Assistant Professor - CY Cergy Paris University) : michele.linardi@cyu.fr.

Farah Benamara Zitoune (Professor in Computer Science, Université de Toulouse IRIT and IPAL-CNRS Singapore) : farah.benamara@irit.fr.

References

[1] Sihan Chen, John P. Lalor, Yi Yang, and Ahmed Abbasi. PersonaTwin: A multi-tier prompt conditioning framework for generating and evaluating personalized digital twins. In Ofir Arviv, Miruna Clinciu, Kaustubh Dhole, Rotem Dror, Sebastian Gehrmann, Eliya Habba, Itay Itzhak, Simon Mille, Yotam Perlitz, Enrico Santus, João Sedoc, Michal Shmueli Scheuer, Gabriel Stanovsky, and Oyvind Tafjord, editors, *Proceedings of the Fourth Workshop on Generation, Evaluation and Metrics (GEM²)*, pages 774–788, Vienna, Austria and virtual meeting, July 2025. Association for Computational Linguistics.

[2] COSchool. A comprehensive review of ai-based intelligent tutoring systems. `https://www.coschool.ai/feeds/blog/intelligent-tutoring-systems`, May 2026. Accessed: 2026-08-19.

[3] Firuz Kamalov, David Santandreu Calonge, Linda Smail, Dilshod Azizov, Dimple R. Thadani, Theresa Kwong, and Amara Atif. Evolution of ai in education: Agentic workflows. *arXiv preprint arXiv:2504.20082*, 2025.

[4] Angélique Létourneau, Marion Deslandes Martineau, Patrick Charland, John Alexander Karran, Jared Boasen, and Pierre Majorique Léger. A systematic review of ai-driven intelligent tutoring systems (its) in k-12 education. *npj Science of Learning*, 10, 2025.

[5] Rui Li, Heming Xia, Xinfeng Yuan, Qingxiu Dong, Lei Sha, Wenjie Li, and Zhifang Sui. How far are LLMs from being our digital twins? a benchmark for persona-based behavior chain simulation. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 15738–15763, Vienna, Austria, July 2025. Association for Computational Linguistics.

[6] Amelie Smith and Michele Linardi. Poster. `https://perso.etis-lab.fr/linardi/POSTER_AI_LEARNING_2026/Poster_Smith_Linardi.pdf`, 2026. Poster presented at Learning in the Age of AI: Epistemic Practices, Evaluation and Governance, June 11, 2026, Campus de l'Arche, La Défense, Paris, France.

[7] Haojie Xie, Yirong Chen, Xiaofen Xing, Jingkai Lin, and Xiangmin Xu. PsyDT: Using LLMs to construct the digital twin of psychological counselor

with personalized counseling style for psychological counseling. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1081–1115, Vienna, Austria, July 2025. Association for Computational Linguistics.