Can LLMs be Effective in-class Tutors?

Sebastian Kobler and **Jonathan Kummerfeld**
School of Computer Science
University of Sydney
skob7546@uni.sydney.edu.au jonathan.kummerfeld@sydney.edu.au

Abstract

Traditional tutoring approaches are known to be effective yet expensive. This paper introduces a personalised tutoring system using large language models (LLMs) to guide students through first-year university tutorials. The system offers real-time, adaptive feedback based on individual student needs, enhancing engagement and promoting deeper learning. Unlike existing LLM tutors, this approach delivers tailored support, adjusting explanations and challenges dynamically. We plan to measure its effectiveness through student performance data, engagement metrics, and feedback surveys. We anticipate improved learning outcomes, increased student satisfaction, and reduced instructor workload, demonstrating the system's potential to create a more personalised and efficient learning environment.

1 The Problem

Studies on tutoring effectiveness (Cohen et al., 1982; Chi et al., 2001) have demonstrated the significant benefits of personalised instruction. While recent research has focused on improving Q&A tutors (Chevalier et al., 2024; Park et al., 2024; Schmucker et al., 2024; Chen et al., 2024), and LLM-based systems such as (Khanmigo, 2024) and Google's LearnLM (Jurenka et al., 2024) offer useful support, these approaches lack the classroom-specific personalisation needed for effective individualised learning. Existing systems struggle to adjust dynamically to the unique learning paces and needs of individual students, limiting their effectiveness in large, varied student groups. Extending LLM tutor capabilities to provide more responsive, personalised guidance in real-time could significantly enhance both student outcomes and the efficiency of instruction in university environments.

2 Proposed Solution

To solve the discussed issues, we plan to develop an in-class tutoring system with the following capabilities:

2.1 Personalisation

To maximise each student's learning potential, it is essential that our models are personalised. We achieve this through two key methods: initial profiling and dynamic adjustment. Initial profiling gathers information about students' preferences and prior knowledge through surveys or quizzes before the course begins. This helps to assess their familiarity with key concepts. Dynamic adjustment occurs in real time, as the LLM monitors student responses and adapts accordingly. Advanced students may receive challenge problems, while those struggling are provided with foundational examples to reinforce understanding.

Additionally, the LLM offers unprompted encouragement throughout the lesson, recognising when a student is on the right track but hesitant, and providing positive reinforcement to boost confidence. As the lesson progresses, the system tracks areas where students face difficulties and uses this data to customise homework and suggested readings. This ensures that students focus on their own specific challenges, rather than following a generic, one-size-fits-all approach to homework.

2.2 Teacher Controls

While the LLM guides the lesson, it works in collaboration with the teacher, who provides the lesson plan, prompts the LLM to move on to the next topic when necessary, and addresses student questions that require human intervention.

The lesson plan is generated using GPT-4o by incorporating the tutorial sheet as context and emphasizing the lesson's key learning objectives, which are later used to assess student understanding. The plan is structured with timed sections based on Gagné's nine events of instruction (Gagné et al., 1992). These events create a systematic teaching process, starting with capturing students' attention and presenting content, followed by guided practice, feedback, and strategies to reinforce retention, ensuring an effective learning experience.

The teacher also has access to a dashboard featuring visual analytics that track the class's progress. It allows the teacher to view individual student scores or class averages, organized by topic. This provides valuable insights into the overall lesson progress and helps guide instructional decisions.

3 Ideal Scenario

In an ideal 60-minute class using the system, the instructor begins with a brief lecture while each student interacts with their LLM-based tutor. The system monitors comprehension, providing extra explanations or more advanced material based on individual needs. Students then work through problems with real-time support from the tutor, which adapts to their progress, offering hints or advanced challenges as necessary. The instructor accesses a dashboard with real-time analytics, allowing for targeted interventions where needed. In the final 10 minutes, the system generates personalised summaries and customised homework, ensuring each student continues learning at their own pace after class. This creates a highly adaptive, student-focused learning environment.

4 Progress

Work so far has primarily involved designing lesson plans for the personalised tutors to use, as well as building the front end teacher and student system. By consulting past research (Fan et al., 2024; Hu et al., 2024) as well as tutorial sheets given for a first year-chemistry class, we designed simple lesson plans the LLM uses as context. Further experimenting will involve implementing the ideas mentioned throughout Section 2.

5 Conclusion

Our work covers a gap in current research, extending the benefits of conversation-based tutoring systems using LLMs into the classroom. Upon completion of this system, educators would be less burdened by having to address individual student difficulties in larger classes, allowing them to focus on higher-level instruction and management. In addition, students would receive real-time, adapted feedback suitable for their specific needs fostering deeper engagement and thus improved academic outcomes.

References

Yulin Chen, Ning Ding, Hai-Tao Zheng, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2024. Empowering private tutoring by chaining large language models. *Preprint*, arXiv:2309.08112.

Alexis Chevalier, Jiayi Geng, Alexander Wettig, Howard Chen, and Sebastian Mizera et al. 2024. Language models as science tutors. *Preprint*, arXiv:2402.11111.

Michelene Chi, Stephanie Siler, Heisawn Jeong, Takashi Yamauchi, and Robert Hausmann. 2001. Learning from human tutoring. *Cognitive Science*, 25.

Peter A Cohen, James A. Kulik, and Chen-Lin C. Kulik. 1982. Educational Outcomes of Tutoring: A Meta-analysis of Findings. *American Educational Research Journal*, 19(2).

Haoxiang Fan, Guanzheng Chen, Xingbo Wang, and Zhenhui Peng. 2024. LessonPlanner: Assisting Novice Teachers to Prepare Pedagogy-Driven Lesson Plans with Large Language Models. ArXiv:2408.01102 [cs].

R.M. Gagné, L.J. Briggs, and W.W. Wager. 1992. *Principles of Instructional Design*. Harcourt Brace Jovanovich College Publishers.

Bihao Hu, Longwei Zheng, Jiayi Zhu, Lishan Ding, Yilei Wang, and Xiaoqing Gu. 2024. Teaching Plan Generation and Evaluation With GPT-4: Unleashing the Potential of LLM in Instructional Design. *IEEE Transactions on Learning Technologies*, 17.

Irina Jurenka, Markus Kunesch, Kevin R. McKee, Daniel Gillick, and Shaojian Zhu et al. 2024. Towards responsible development of generative ai for education: An evaluation-driven approach. *Preprint*, arXiv:2407.12687.

Khanmigo. 2024. Khanmigo.

Minju Park, Sojung Kim, Seunghyun Lee, Soonwoo Kwon, and Kyuseok Kim. 2024. Empowering personalized learning through a conversation-based tutoring system with student modeling. CHI '24.

Robin Schmucker, Meng Xia, Amos Azaria, and Tom Mitchell. 2024. Ruffleriley: Insights from designing and evaluating a large language model-based conversational tutoring system. *Preprint*, arXiv:2404.17460.