

Examining the Influence of Varied Levels of Domain Knowledge Base Inclusion in GPT-based Intelligent Tutors

Blake Castleman, Mehmet Kerem Turkcan¹

¹Department of Electrical Engineering, Columbia University, New York, NY
{bc3029,mkt2126}@columbia.edu

ABSTRACT

Recent advancements in large language models (LLMs) have facilitated the development of chatbots with sophisticated conversational capabilities. However, LLMs exhibit frequent inaccurate responses to queries, hindering applications in educational settings. In this paper, we investigate the effectiveness of integrating a knowledge base (KB) with LLM intelligent tutoring systems (ITSs) to increase response reliability. To achieve this, we design a scaleable KB that affords educational supervisors seamless integration of lesson curricula, which is automatically processed by the ITS. We then detail a preliminary study, where student participants were presented with questions about the artificial intelligence curriculum to respond to. GPT-4 ITSs with varying hierarchies of KB access and human domain experts then assessed these responses. Lastly, students cross-examined the ITSs' responses to the domain experts' and ranked their various pedagogical abilities. Results suggest that, although these ITSs still demonstrate a lower accuracy compared to domain experts, the accuracy of the ITSs increases when access to a KB is granted. We also observe that the ITSs with KB access exhibit better pedagogical abilities to speak like a teacher and understand students than those of domain experts, while their ability to help students remains lagging behind domain experts.

Keywords

intelligent tutoring system, GPT tutor, knowledge base inclusion, pedagogical abilities

1. INTRODUCTION

Asynchronous education has become increasingly more prevalent over the twenty-first century and is now crucial as a means to assist general education. The surge in online learning, however, has resulted in less student-to-student and student-to-teacher interaction, which has deterred academic engagement [1, 11].

B. Castleman and M. K. Turkcan. Examining the influence of varied levels of domain knowledge base inclusion in gpt-based intelligent tutors. In B. Paaßen and C. D. Epp, editors, *Proceedings of the 17th International Conference on Educational Data Mining*, pages 649–657, Atlanta, Georgia, USA, July 2024. International Educational Data Mining Society.

© 2024 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.
<https://doi.org/10.5281/zenodo.12729908>

The rise of artificial intelligence (AI) and, in particular, large language models (LLMs), have enabled the construction of intelligent tutoring system (ITS) chatbots, which hold promise in mitigating the latter issue by supplementing pedagogical interactions [6, 16]. GPT-3 [3] and GPT-4 [13] are some of the largest and most versatile LLMs trained to date, with impressive domain knowledge recall and natural language processing (NLP) abilities. This grants students the ability to pose queries for numerous subjects and receive tailored responses. The potential pedagogical impact on distance learners, therefore, is formidable, with some educational institutions and organizations already announcing the integration of GPT models in various programs [19, 22].

LLMs have several shortcomings, however, that discourage their use in educational settings. Primarily, with GPT models, a lack of domain knowledge or analytical reasoning frequently causes fictitious responses to be generated in their overzealous pursuit of a reply [16, 25]. As such, this potentially yields large misunderstandings in educational settings. As a result, LLM ITSs are controversial for present-day classroom application as they have not yet exhibited pedagogical abilities equal to or exceeding that of human teachers [7, 23]. To develop optimal LLM ITSs, it is crucial to keep LLM responses in alignment with educators' lesson goals and to prevent them from presenting artificial information to students.

This study is the first to examine the efficacy of integrating a knowledge base (KB) with LLM ITSs to achieve these goals. The direct installment of a KB with educators' subject matter may bring GPT to reconsider previously disregarded knowledge, mitigating these issues. Consequently, we aim to explore how integrating KBs may *enhance the response accuracy* of LLM ITSs, as well as how their *pedagogical abilities may improve* through such integration.

In this paper, we aim to address these open questions by creating GPT-4 ITSs with various KB access levels for accuracy and pedagogical ability comparison. A practice mode was developed for these ITSs in which they assess students' answers to questions about related lesson subject material. The tutors were granted the ability to use any information provided by the supervisor through their KB in the practice mode, given the various information constraints each ITS was assigned. Information constraints were abstracted into a hierarchy of three levels: the lesson's subject (no KB access), a topic description (partial KB access), and the overall

lecture material (full KB access).

To evaluate and quantify the effectiveness of each information level in the KB, we populated the LLM ITS with the AI curriculum and designed a three-stage evaluation pipeline to observe the effect of KB integration. First, we designated student participants with AI proficiency to respond to various AI questions with diverse accuracy. Next, human domain expert participants and ITSs assessed the students' responses and addressed each. Lastly, the students were allowed to observe the assessments and rated the pedagogical quality of each response. This allows us to observe the practicality of each knowledge hierarchy level in GPT-4 ITS implementations in contrast to human domain experts. Our code and data are open-sourced below for further research.¹.

2. RELATED WORK

2.1 Chatbot Tutor Pedagogical Abilities

Chatbot tutoring has been applied for the past two decades [12, 24] and recent research has investigated potential LLM frameworks. Demszky et al. (2021) [4] observed that an ITS based on BERT [5], trained on 2246 student-teacher exchanges, had remarkable conversational uptake to student queries. Tack and Piech (2022) [23] built upon Demszky et al. (2021) [4] by discovering that although LLMs such as Blender [18] and GPT-3 have considerable conversational uptake, core pedagogical abilities of these LLMs severely lag behind that of human teachers. They discovered these LLMs demonstrate insufficient capability in their pedagogical abilities to speak like a teacher, understand students, and help students. We aim to investigate if the substandard pedagogy measurements Tack and Piech (2022) [23] have found can be improved upon by expanding the knowledge these LLMs have to work within.

2.2 Chatbot Tutor Knowledge Base Incorporation

Using KBs for chatbot tutors to retrieve information is not a novel topic. An instance of such is AsasaraBot [10], an educational AI chatbot designed to teach high school students cultural content in foreign languages, which searches its own KB (or externally on the Web) for the most appropriate response and processes it within its dialogue manager.

BookBuddy [21], a virtual reading companion designed to teach English as a second language, uses a scaleable knowledge base of 20 books to choose from (depending on the child's interests). Trained on the Stanford Question Answering Dataset (SQuAD) [17], BookBuddy's generative model handles character and plot questions, while its rule-based algorithms produce vocabulary and arithmetic questions.

QuizBot [20], a dialogue-based agent that helps students learn factual knowledge, uses a system question pool populated by a supervisor's question-answer pairs that pertain to the lesson subject. It was designed to binary grade student answers via cosine similarity thresholding and was observed to rarely misclassify answers. In their studies, they evaluated the ability of the agent to generalize over three different

¹<https://github.com/b-castleman/llm-tutor-knowledge-base>

subjects and data suggested QuizBot's effectiveness to surpass that of flashcard techniques.

In this paper, we seek to quantify the efficacy of a GPT ITS implemented in conjunction with a knowledge base. Furthermore, we aim to observe the variations in our tutor's response quality from differing quantities of inputted supervisor information presented by the knowledge base through our access hierarchy framework. To our knowledge, there are currently no published chatbot tutors that have researched the effectiveness of teaching artificial intelligence course curricula or that have quantified the effect of disparate KB hierarchy levels on response quality.

3. FRAMEWORK

3.1 Prompting Structure

We use GPT-4 to create three separate LLM ITS instances. It is prompted to respond "as a teacher", being both accurate and enthusiastic. The tutor operates in a practice mode, where students are quizzed by the tutor about the material to reinforce and consolidate their understanding of what they have learned. After receiving a student's answer to a particular question, the tutor must respond appropriately depending on the student's answer choice(s) and reasoning. After the student responds to the question, GPT-4 ranks the responses on an integer scale from 1-5 and stores this value internally. Then, it is asked to qualitatively assess and address the student's answer given the previous rating, describing what is understood well by their answer and what may use improvement. This internal rating technique is inspired by Park et al. (2023) [14], who used situational GPT-4 ratings to simulate autonomous agent behavior.

3.2 Knowledge Base

The KB is designed for an educator to efficiently input their lesson material into JSON files for information processing at disparate knowledge hierarchies. We will now describe how each hierarchy level of information provided interacts separately with the system. Figure 1 illustrates the following described architecture to help understand the information processing pathways implemented.

3.2.1 No KB Access: Lesson Subject

The base-level hierarchy, hereon referred to as having no KB access, is told only what the lesson subject is (i.e. introduction to artificial intelligence, basic derivatives, prokaryotic cells). As a result, any ITS that uses only this KB level is only using GPT-4 training with minimal supervisor influence to generate responses. The lesson subject is given to the tutor when defining the initial system prompt and is repeated when internally rating any answer.

3.2.2 Partial KB Access: Topic Description

The topic description is a short definition of what the given topic is about, consisting of no more than a few sentences. Any tutor using KB levels up to this extent is referred to as having partial KB access as it can observe high-level information on the lesson subject without being presented with any intricacies. The topic description is used when defining the initial system prompt and is given to the response ranker.

each subtopic being randomly and uniquely picked for each student. Every participant was then asked to respond to each question presented to them with varying levels of accuracy. This was realized by asking participants to answer a third of their questions with full accuracy, a third of their questions with partial accuracy or partial inaccuracy, and a third of their questions very inaccurately. We prompted participants for a variety of accuracy levels in order to ensure a similar corpus of response qualities was achieved, which was necessary to properly assess whether or not the KB could react proficiently to students’ answer slips and guesses. No other restrictions were imposed on answers. No time limit was imposed for participants’ answers.

4.3.2 Stage #2: Domain Expert and Intelligent Tutor Addressal

In stage two, we recruited two professional AI curriculum domain expert volunteers who agreed to the anonymized release of their collected data. Both domain experts were professionals in AI education with multi-year experience in curating and teaching AI course curricula for secondary and post-secondary students. Additionally, both domain experts agreed they had the necessary expertise to properly assess and address potential student responses to the lesson material and provided questions. Before beginning the evaluation, we secured explicit confirmation from the domain experts on their agreement to the correctness of the video lectures that would be provided to students. This validation process ensured a robust foundation for the subsequent phases of our research, reinforcing the reliability of data collected. We utilized these two domain experts as well as the three ITSs (all five henceforth collectively denoted as “assessors”) to assess and respond to students’ answers. Specifics on assessors’ instructions can be found in Appendix B.

4.3.3 Stage #3: Ranking of Assessor’s Responses

For the third stage, the student participants from Stage #1 evaluated which assessor’s response (after processing as per Appendix E) was of the best quality for each question-answer pair provided. The assessments were anonymized to prevent biases, and the corresponding question-answer pairs were presented with each assessment as contextual background. Due to an overwhelming amount of data to present to students, we uniformly and randomly selected only one domain expert’s answer to present to the students for comparison against the ITSs’ responses (Appendix D).

The triples (question-answer-assessments) given to each participant consisted of those containing the answers they provided to questions in Stage #1 along with the triples to an additional, random participant’s question-answer pairs. This allowed for each assessor’s response to have two independent evaluator’s rankings and mitigated personal biases.

As every question-answer pair had four separate replies (three ITSs and one randomized domain expert) to rank, we formatted this evaluation as a pairwise comparison between each ITS versus the domain expert. This allowed us to observe how students’ assessor preferences changed with KB differences. We utilized the same criteria implemented in past literature and asked participants which response is more likely to be said like a teacher, understanding of a student,

and helpful to a student [23], which builds on key pedagogical requirements set forth for AI models being applied as AI teachers [2]. For each pedagogical ability criterion, we asked participants to indicate whether either assessment was of better quality or if they were too similar to differentiate.

Furthermore, a considerable amount of assessments suggested inaccurate answers as truthful information. To ensure participants’ familiarity with which choices were valid and rational, a “Factual Information” section was included for each question. This section gave a brief delineation, based on rigorous research, for the answer choices that most accurately addressed the given question. The assessors’ replies were never directly mentioned in this section to reduce biases.

5. RESULTS

5.1 Stage #2 Results: Agreement and Accuracy

We cross-compare the domain experts’ assessments against one another to find an agreement level of 63% for what each domain expert believed the correct answer choices were whereas the ITSs only agreed 43% of the time on what they believed the correct answer choices were. Some of the answer choices were ambiguous depending on the specific context that each assessor interpreted (e.g. the potential impact of a technology to have indirectly influenced events), meaning assessor groups would not always be in agreement in their responses. In calculating these agreement level percentages, we choose for agreements to be strict and for any assessor response deviance to not count as agreement.

As a result of this inflexible and conservative definition, we also calculate the overall accuracy level for all assessors, taking into account these ambiguous situations. In this definition for accuracy, we declare that ambiguous situations are acceptable so long as no evidence renders the response incorrect or unreliable beyond a reasonable doubt. With this definition, we found that domain experts demonstrated a combined accuracy rate of approximately 87% while ITSs exhibited a combined accuracy rate of only 70%. In particular, the no KB, partial KB access, and full KB access ITSs had accuracy scores of 67%, 77%, and 74%, respectively. Information on ITS token usage is available in Appendix F.

5.2 Stage #3 Results: Pedagogical Abilities

We examine the preliminary results by calculating the percentage of replies classified as a particular pedagogical ability for each ITS (Appendix G). Table 1 exhibits the results of this computation and Appendix Figure 3 showcases a bar graph of the rankings to illustrate the visual change in reply selection quantity with KB additions.

We observe, for the ITS with no KB access, that the domain experts attained a relative majority for all pedagogical abilities, particularly in the dimension of being helpful to students. This agrees with Tack and Piech (2022) [23], who found similar trends in pedagogy for GPT-3 AI teachers responding to student dialogues. Though their LLM (GPT-3) and our own (GPT-4) are disparate, as well as the domain of education, the exhibition of their similar pedagogical qualities underscores the consistency of these didactic models across different iterations and applications.

Intelligent Tutor	Pedagogical Ability	GPT-4 Preference	Domain Expert Preference	No Clear Preference
No Knowledge Base	Talking like a teacher	27.14%	50.00%	22.86%
	Understanding the student	31.43%	38.57%	30.00%
	Helpful to the student	25.71%	57.14%	17.14%
Partial Knowledge Base	Talking like a teacher	44.29%	38.57%	17.14%
	Understanding the student	41.43%	35.71%	22.86%
	Helpful to the student	38.57%	51.43%	10.00%
Full Knowledge Base	Talking like a teacher	38.57%	31.43%	30.00%
	Understanding the student	41.43%	31.43%	27.14%
	Helpful to the student	34.29%	45.71%	20.00%

Table 1: Preference rates for different tutors based on evaluations of four participant annotators for each pedagogical ability.

Upon integration with a KB, the ITSs’ pedagogical abilities exhibited increases in preference for GPT-4 responses. The partial KB ITS achieved 17.15%, 10.00%, and 12.86% preference increases for talking like a teacher, understanding the student, and being helpful to the student respectively. Meanwhile, the full KB ITS achieved 11.43%, 10.00%, and 8.58% preference increases, respectively. This suggests that having both partial KB access and full KB access increases the ITSs’ pedagogical dimensions.

Notably, the reply percentage distributions for talking like a teacher and understanding the student realigned in favor of both ITSs with KB access. No ITS’s ability to be helpful to the student overcame the domain experts’, despite the other two abilities accomplishing this task.

The partial KB ITS was particularly preferred for talking like a teacher and tied with the full KB ITS for understanding the student. However, the domain expert reply preferences were consistently lower for the full KB ITS as opposed to the partial KB ITS for all three pedagogical abilities.

6. DISCUSSION AND FUTURE WORK

Our study faced a core limitation in the evaluation due to a shortage of AI curriculum experts with adequate knowledge and time allocations. Consequently, we had to reduce the number of student participants and prompt students for a variety of response accuracy to ensure we received assessor responses on incorrect answers. This approach purposefully ensured a generation of sufficient inaccurate student answers for assessors to evaluate. However, trying to mimic naturally inaccurate student answers (i.e. slips and guesses) with artificial inaccuracies could have potentially introduced experimental bias. Ideally, in the scenario of more personnel resources becoming available, future studies of KB integration should have more AI curriculum domain experts and student participants to resolve the aforementioned issues.

We also chose to include participants who have completed courses in AI instead of participants engaged in an AI course for the first time. This afforded participants a preferable capability to contrast the pedagogical abilities presented in Stage #3, where seemingly correct arguments from assessor responses could mislead students who lack prior, formal education in AI. Moreover, this decision raised the likelihood that the inaccurate student answers provided in Stage #1 were plausible and well-articulated.

Using “check all that apply” questions were used to make assessors required to evaluate multiple components of every question in Stage #2. As a result, it should be noted that other question types may produce different results. Dissimilar curricula may also produce different results as question complexity and response variation depend on them.

Lastly, though KBs external to LLMs possess multifaceted benefits (e.g. ease of updates, direct prompting) fine-tuning these models on educational subject matter may prove incredibly resourceful for accuracy and pedagogical ability gains. We hope to investigate this possibility in future research.

7. CONCLUSION

Our study is the first to delve into the extent to which an integrated KB can enhance the pedagogical abilities of LLM chatbot tutors. This technique seeks to mitigate inaccuracies and irrelevant responses from LLMs, thereby potentially increasing pedagogical abilities. We gather preliminary data on such systems by exploring, at varying KB hierarchies, how effective and accurate the ITSs are compared to domain experts. An evaluation was conducted on AI curriculum utilizing both ITSs (with varying KB access constraints) and domain experts to assess student answers to various “check all that apply” question types for a given curriculum.

The assessor responses retrieved from our evaluation pipeline indicated that all ITSs had lower agreement and accuracy rates than domain experts. Both ITSs with KB access, however, had notably higher accuracy rates as compared to the ITS without KB access. After quantifying the ITSs’ pedagogical abilities, we observe, for the aforementioned curriculum and question types, that the KB integration increases the preference for ITS responses and decreases the preference for domain expert responses. Furthermore, our data suggested that the abilities of ITSs to speak like a teacher and understand students surpassed that of domain experts when KB capabilities were used. The ability to be helpful to students, though increasing with KB inclusion, still failed to exceed that of domain experts.

This comprehensive evaluation highlights the practicality of implementing KBs with chatbot ITS frameworks using LLMs. Though LLM ITSs still have necessary development for application, this methodology spotlights a scaleable technique for instructors to include their course information with chatbot ITSs that results in accuracy and pedagogical ability gains.

8. ACKNOWLEDGMENTS

We would like to thank the organization Aiphabet for providing their prerelease AI curriculum for our use. We also would like to thank the anonymous domain experts and student participants for their invaluable contributions to the study.

References

- [1] C. M. Amerstorfer and C. Freiin von Münster-Kistner. Student perceptions of academic engagement and student-teacher relationships in problem-based learning. *Frontiers in Psychology*, 12, 2021.
- [2] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, E. Brynjolfsson, S. Buch, D. Card, R. Castellon, N. S. Chatterji, A. S. Chen, K. A. Creel, J. Davis, D. Demszky, C. Donahue, M. Doumbouya, E. Durmus, S. Ermon, J. Etchemendy, K. Ethayarajh, L. Fei-Fei, C. Finn, T. Gale, L. E. Gillespie, K. Goel, N. D. Goodman, S. Grossman, N. Guha, T. Hashimoto, P. Henderson, J. Hewitt, D. E. Ho, J. Hong, K. Hsu, J. Huang, T. F. Icard, S. Jain, D. Jurafsky, P. Kalluri, S. Karamcheti, G. Keeling, F. Khani, O. Khattab, P. W. Koh, M. S. Krass, R. Krishna, R. Kuditipudi, A. Kumar, F. Ladhak, M. Lee, T. Lee, J. Leskovec, I. Levent, X. L. Li, X. Li, T. Ma, A. Malik, C. D. Manning, S. P. Mirchandani, E. Mitchell, Z. Muniyikwa, S. Nair, A. Narayan, D. Narayanan, B. Newman, A. Nie, J. C. Niebles, H. Nilforoshan, J. F. Nyarko, G. Ogut, L. Orr, I. Papadimitriou, J. S. Park, C. Piech, E. Portelance, C. Potts, A. Raghunathan, R. Reich, H. Ren, F. Rong, Y. H. Roohani, C. Ruiz, J. Ryan, C. R’e, D. Sadigh, S. Sagawa, K. Santhanam, A. Shih, K. P. Srinivasan, A. Tamkin, R. Taori, A. W. Thomas, F. Tramèr, R. E. Wang, W. Wang, B. Wu, J. Wu, Y. Wu, S. M. Xie, M. Yasunaga, J. You, M. A. Zaharia, M. Zhang, T. Zhang, X. Zhang, Y. Zhang, L. Zheng, K. Zhou, and P. Liang. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [3] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [4] D. Demszky, J. Liu, Z. Mancenido, J. Cohen, H. Hill, D. Jurafsky, and T. Hashimoto. Measuring conversational uptake: A case study on student-teacher interactions. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1638–1653, Online, 2021. Association for Computational Linguistics.
- [5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, 2019. Association for Computational Linguistics.
- [6] Y. K. Dwivedi, N. Kshetri, L. Hughes, E. L. Slade, A. Jeyaraj, A. K. Kar, A. M. Baabdullah, A. Koochang, V. Raghavan, M. Ahuja, H. Albanna, M. A. Al-bashrawi, A. S. Al-Busaidi, J. Balakrishnan, Y. Barlette, S. Basu, I. Bose, L. Brooks, D. Buhalis, L. Carter, S. Chowdhury, T. Crick, S. W. Cunningham, G. H. Davies, R. M. Davison, R. Dé, D. Dennehy, Y. Duan, R. Dubey, R. Dwivedi, J. S. Edwards, C. Flavián, R. Gauld, V. Grover, M.-C. Hu, M. Janssen, P. Jones, I. Junglas, S. Khorana, S. Kraus, K. R. Larsen, P. Latreille, S. Laumer, F. T. Malik, A. Mardani, M. Mariani, S. Mithas, E. Mogaji, J. H. Nord, S. O’Connor, F. Okumus, M. Pagani, N. Pandey, S. Papagiannidis, I. O. Pappas, N. Pathak, J. Pries-Heje, R. Raman, N. P. Rana, S.-V. Rehm, S. Ribeiro-Navarrete, A. Richter, F. Rowe, S. Sarker, B. C. Stahl, M. K. Tiwari, W. van der Aalst, V. Venkatesh, G. Viglia, M. Wade, P. Walton, J. Wirtz, and R. Wright. Opinion paper: “so what if ChatGPT wrote it?” multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71:102642, 2023.
- [7] L. Kohnke, B. L. Moorhouse, and D. Zou. ChatGPT for language teaching and learning. *RELC Journal*, 54(2):537–550, 2023.
- [8] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [9] U. Macar, B. Castleman, N. Mauchly, M. Jiang, A. Aouissi, S. Aouissi, X. Maayah, K. Erdem, R. Ravindranath, A. Clark-Sevilla, and A. Salleb-Aouissi. Teenagers and artificial intelligence: Bootcamp experience and lessons learned. *arXiv preprint arXiv:2312.10067*, 2023.
- [10] K. Mageira, D. Pittou, A. Papasalouros, K. Kotis, P. Zangogianni, and A. Daradoumis. Educational AI chatbots for content and language integrated learning. *Applied Sciences*, 12(7), 2022.
- [11] G. Maimaiti, C. Jia, and K. F. Hew. Student disengagement in web-based videoconferencing supported online learning: An activity theory perspective. *Interactive Learning Environments*, 0(0):1–20, 2021.
- [12] B. D. Nye, A. C. Graesser, and X. Hu. AutoTutor and family: A review of 17 years of natural language tutoring. *International Journal of Artificial Intelligence in Education*, 24(4):427–469, 2014.
- [13] OpenAI. GPT-4 technical report, 2023. Accessed: 2023-09-05.

- [14] J. S. Park, J. C. O'Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein. Generative agents: Interactive simulacra of human behavior. *arXiv preprint arXiv:2304.03442*, 2023.
- [15] M. Pietsch, T. Möller, B. Kostic, J. Risch, M. Pippi, M. Jobanputra, S. Zanzottera, S. Cerza, V. Blagojevic, T. Stadelmann, T. Soni, and S. Lee. Haystack: The end-to-end NLP framework for pragmatic builders, 2019. <https://github.com/deepset-ai/haystack>.
- [16] J. Qadir. Engineering education in the era of ChatGPT: Promise and pitfalls of generative AI for education. In *2023 IEEE Global Engineering Education Conference (EDUCON)*, pages 1–9, 2023.
- [17] P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas, 2016. Association for Computational Linguistics.
- [18] S. Roller, E. Dinan, N. Goyal, D. Ju, M. Williamson, Y. Liu, J. Xu, M. Ott, E. M. Smith, Y.-L. Boureau, and J. Weston. Recipes for building an open-domain chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, Online, 2021. Association for Computational Linguistics.
- [19] K. Roose. GPT-4 is exciting and scary. *The New York Times*, 2023. <https://www.nytimes.com/2023/03/15/technology/gpt-4-artificial-intelligence-openai.html>. Accessed: 2023-09-05.
- [20] S. Ruan, L. Jiang, J. Xu, B. J.-K. Tham, Z. Qiu, Y. Zhu, E. L. Murnane, E. Brunskill, and J. A. Landay. QuizBot: A dialogue-based adaptive learning system for factual knowledge. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pages 1–13, New York, NY, USA, 2019. Association for Computing Machinery.
- [21] S. Ruan, A. Willis, Q. Xu, G. M. Davis, L. Jiang, E. Brunskill, and J. A. Landay. BookBuddy: Turning digital materials into interactive foreign language lessons through a voice chatbot. In *Proceedings of the Sixth (2019) ACM Conference on Learning @ Scale*, L@S '19, New York, NY, USA, 2019. Association for Computing Machinery.
- [22] J. Shaw. Embracing AI. *Harvard Magazine*, 2023. <https://www.harvardmagazine.com/2023/08/ai-in-education>. Accessed: 2023-09-05.
- [23] A. Tack and C. Piech. The AI teacher test: Measuring the pedagogical ability of blender and GPT-3 in educational dialogues. In A. Mitrovic and N. Bosch, editors, *Proceedings of the 15th International Conference on Educational Data Mining*, pages 522–529, Durham, United Kingdom, 2022. International Educational Data Mining Society.
- [24] S. Wollny, J. Schneider, D. Di Mitri, J. Weidlich, M. Rittberger, and H. Drachsler. Are we there yet? - a systematic literature review on chatbots in education. *Frontiers in Artificial Intelligence*, 4, 2021.
- [25] W. Zhong, R. Cui, Y. Guo, Y. Liang, S. Lu, Y. Wang, A. Saied, W. Chen, and N. Duan. AGIEval: A human-centric benchmark for evaluating foundation models. *arXiv preprint arXiv:2304.06364*, 2023.

APPENDIX

A. STAGE #1: QUESTION CREATION

To create impartial yet relevant questions, we used GPT-4 to generate a set of 36 “check all that apply” questions. This question type was chosen as it yields an exponential amount of answer possibilities, increasing the likelihood that any of the assessors may evaluate an answer inaccurately. With the expansive question bank, the influence of outliers is minimized and correlations in assessor agreements can be observed.

As this information alone may not have been sufficient for understanding a student’s weaknesses from incorrect responses, we prompted participants to explain their answers for every answer given. The question relevancy was equally divided amongst the four subtopics, and GPT-4 was given full KB access to the particular AI subtopic as well as the topic description for generating each question. This best mimicked question creation by various educators, who often have abundant resources (i.e. textbooks, internet) for creating intricate questions.

Before proceeding, we went through all generated questions to ensure they were valid for different levels of difficulty, pedagogically sound, and relevant to the video lecture’s material. Five questions were chosen to be regenerated for being either repeats of other questions or for lacking relevance to the given lecture material.

B. STAGE #2: ASSESSORS’ PROMPT AND INSTRUCTIONS

We utilized the two domain experts as well as the three ITSs to assess and respond to students’ answers. The five assessors were given the same prompt and were asked to do the following:

- Respond to the student’s answer given the current question
- Explain what about their answer is accurate or inaccurate
- Suggest what about the answer can use improvement
- Help clear any misconceptions

Assessors were limited to a 75-word constraint and had no time limit. The only difference in the prompting for the ITSs opposite to the domain experts was the ITS’s access to an internal ranking of the answer.

C. EVALUATION PIPELINE STAGES

<p><i>Question Generation</i></p> <p>Question: What does a neuron output if the sum of its inputs does not exceed the threshold value?</p> <p>a) 0 b) 1 c) -1 d) The threshold value</p>
<p><i>Stage #1</i></p> <p>Student Answer: a - the mcculloch pitts artificial neuron only fires if its input sum exceeds the threshold value, otherwise it remains a zero</p>
<p><i>Stage #2</i></p> <p>Reply A: That’s correct! Your understanding of the McCulloch-Pitts artificial neuron is accurate. The neuron outputs a 0 if the input sum does not exceed the threshold value.</p> <p>Reply B: Great job, that’s correct! In the context of a neuron with a threshold, assuming the McCulloch and Pitts model, both the inputs to the neuron and the response from the neuron are all binary (0 or 1). 1 means the neuron fires (on) and 0 means the neuron doesn’t fire (off).</p> <p>Factual Information: A is the only correct answer as artificial neurons (for McCulloch & Pitts) output a zero for insufficient aggregate values.</p>
<p><i>Stage #3</i></p> <p>Which of these replies is more likely said by a teacher?</p> <p style="text-align: center;">I cannot tell</p> <p>In which reply is the teacher understanding the student more?</p> <p style="text-align: center;">A</p> <p>In which reply is the teacher helping the student more?</p> <p style="text-align: center;">B</p>

Figure 2: A depiction of the evaluation pipeline stages alongside a sample of a student participant’s rating for assessor replies.

D. STAGE #3: OVERWHELMING DATA FOR COMPARISON

For the third stage, we uniformly and randomly selected only one domain expert’s answer to present to the students for comparison against the ITSS’ responses. We roughly estimated that each pairwise comparison would take one minute to analyze and interpret for meaningful results, meaning that our participant volunteers would likely spend closer to two hours for the third stage evaluations in addition to their earlier contributions in the first stage. Concerns grew that an estimated time requirement exceeding an hour may perhaps induce boredom in the participants and would yield more haphazard data. Therefore, our effort to minimize comparisons attempts to preserve our data’s integrity.

E. STAGE #3: ASSESSOR RESPONSE PROCESSING

Before relaying the assessors’ replies to the participants, all replies were processed as according to the following list:

- Fixed all spelling and grammar mistakes
- Removed all sentences of pure affirmation or motivation (i.e. “Keep up the good work!”)
- Removed all declarations affirming the student’s knowledge level in AI as a whole (i.e. “Your understanding aligns well with the biological basis of artificial neurons.”)
- Removed impractical assessment scenarios (i.e. “How would you change your explanation? [Pause and give time to answer]”)

Though student participants were never informed that any reply stemmed from GPT, we suspected they may intuitively be able to recognize common GPT-generated sentence structures against human domain expert sentence structures. Thus, these changes were targeted in order to further anonymize the study and conceal which replies were from domain experts or from LLMs.

F. STAGE #2 RESULTS: TOKEN USAGE

We find the computational price for the described architectures of the no KB, partial KB access, and full KB access ITSS to call GPT-4 to be approximately 400, 600, and 2000 input tokens per student answer assessment, respectively. Additionally, all ITSSs output approximately 60 tokens per student answer assessment. This corresponds to total monetary prices of about \$0.016, \$0.022, and \$0.064 per student answer assessment, respectively, as present-day GPT-4 token costs being \$0.03 per 1000 input tokens and \$0.06 per 1000 output tokens.

G. STAGE #3 DATA PROCESSING

We examine pedagogical ability results by calculating the percentage of replies classified as a particular pedagogical ability for each ITS. This allows us to isolate the probability an ITSS’ responses will be preferred or unfavorably by a student. The calculation is performed by finding the percentage of times each reply option (GPT preference, domain

expert preference, no clear preference) was selected within a pedagogical ability category (speaking like a teacher, understanding the student, helpful to the student) for a specific ITS (no KB access, partial KB access, full KB access). Mathematically, this is derived by dividing the total count of selections for a reply option for an ITS-pedagogical ability pair by the total count of all rated replies for that particular ITS. This number is then converted into a percentage for the particular ITS-pedagogical ability so that it can be.

For example, for the ITS with no KB access in the talking like a teacher pedagogical category, GPT-4 preference was chosen 19 times, and a total of 70 ratings were given for that ITS in that category. This yields a percentage of $\frac{19}{70} \times 100\% = 27.14\%$

H. RANKINGS GRAPH VISUALIZATION




Figure 3: Visualization of ITSs' pedagogical ability improvements as knowledge base access increases for a) speaking like a teacher, b) understanding the student, and c) being helpful to the student.