Small Language Models for Education: Opportunities, Challenges, and a Shared Research Agenda

Yumou Wei[1][0009−0002−1364−8300], Steven Moore[2][0000−0002−5256−0339], Paulo F. Carvalho[1][0000−0002−0449−3733], John Stamper[1][0000−0002−2291−1468], Christopher Brooks[3][0000−0003−0875−0204], and Michael Liut[4][0000−0003−2965−5302]

[1] Carnegie Mellon University, Pittsburgh PA 15213, USA
{yumouw,pcarvalh,jstamper}@andrew.cmu.edu
[2] George Mason University, Fairfax, VA 22030, USA
smoore59@gmu.edu
[3] University of Michigan, Ann Arbor, MI 48109, USA
brooksch@umich.edu
[4] University of Toronto Mississauga, Mississauga, ON L5L 1C6, Canada
michael.liut@utoronto.ca

Abstract. Small language models (SLMs) are emerging as a promising alternative to large language models for educational applications. This workshop aims to explore the potential of SLMs in education, focusing on their unique advantages and the challenges they present. By bringing together researchers from the AIED, EDM, and L@S communities, we hope to foster interdisciplinary collaboration and identify new research directions that leverage the strengths of each community to improve learning outcomes. The expected results of the workshop include a comprehensive understanding of the current state of small language models, the identification of key research challenges and opportunities, and the development of a shared research agenda that guides future work.

Keywords: Small Language Models · Accessible EdTech · Shared Agenda

1 Introduction

We plan to organize a **full-day workshop**[1] that brings researchers from the Artificial Intelligence in Education (AIED), Educational Data Mining (EDM), and Learning at Scale (L@S) communities to **explore the potential of small language models (SLMs) in education**. We define SLMs as open-weight language models with fewer than 10 billion parameters [12], in contrast to large language models (LLMs) with hundreds of billions of parameters that have dominated recent discussions in the field. Capable of being deployed locally on

[1] Workshop website: https://slm4ed-workshop.github.io/

consumer-grade hardware, SLMs offer a unique opportunity for research communities to develop personalized, accessible, and privacy-preserving educational technologies. The purpose of this workshop is to foster interdisciplinary collaboration and develop a shared research agenda that takes advantage of the strengths of each community to address the unique challenges and opportunities presented by SLMs in education. By focusing on the intersection of these three communities, we hope to identify innovative applications of SLMs, develop new methodologies for their evaluation and deployment, and ultimately contribute to the development of effective educational technologies accessible to a wide range of learners.

We welcome researchers and practitioners from all three communities: AIED, EDM, and L@S. We believe that the integration of SLMs in education presents a multifaceted challenge that requires the combined expertise of these communities to address effectively. AIED researchers will bring insights into instructional design and the development of intelligent tutoring systems, while EDM researchers will contribute their expertise in data mining and analysis to understand how SLMs can be fine-tuned and evaluated using educational data. L@S researchers will provide valuable perspectives on the scalability and deployment of SLMs in large-scale educational settings. We expect to attract around 30 participants, including researchers, educators, and industry professionals interested in the emerging applications of SLMs in education. This workshop will provide a unique opportunity for attendees to engage in interdisciplinary discussions, share their research, and collaborate to develop a shared research agenda for SLMs in education.

2 Theme and Goals

It is an exciting time for the AIED, EDM, and L@S communities, as rapid advances in natural language processing (NLP) have opened up new possibilities for personalized learning, intelligent tutoring systems, and educational content generation. However, the integration of NLP technologies into education has largely been dominated by large language models (LLMs) with hundreds of billion parameters. The GPT model family [3], for example, has become nearly synonymous with "LLM" in public discourses, and the AIED community has been no exception [12]. A quick, conservative search of "large language model" or "LLM" in the AIED 2025 proceedings [1] reveals that 157 of 259 full and short papers (60.6%) contained these two keywords and that 93 of those 157 papers (59.2%) also mentioned "GPT". Although these LLMs have demonstrated impressive NLP capabilities, they are often proprietary and require significant computational resources to deploy. In addition, their adoption in education is constrained by privacy concerns, high inference costs, and latency issues, which can limit their accessibility and effectiveness in educational settings, especially in environments with limited resources [12].

In contrast, small language models, which we define as **open-weight language models with fewer than 10 billion parameters**, have only recently

begun to gain attention in the community. SLMs offer several advantages over LLMs, including lower computational requirements, reduced inference costs, and the ability to be deployed locally on consumer-grade hardware [12]. This makes SLMs particularly attractive for educational applications [10], where privacy concerns and resource constraints are often significant barriers to adoption [9]. In addition, SLMs can be efficiently fine-tuned on high-quality domain-specific educational data to provide accurate curriculum-aligned responses [5] and have great potential to build flexible AI agents [2].

This workshop aims to explore the potential of SLMs in education, focusing on their unique advantages and the challenges they present. By bringing together experts from the AIED, EDM, and L@S communities, we hope to foster interdisciplinary collaboration and identify new research directions that leverage the strengths of SLMs to improve learning outcomes. The expected results of the workshop include a comprehensive understanding of the current state of SLMs in education, the identification of key research challenges and opportunities, and the development of a shared research agenda that can guide future work.

3 Expected Outcomes and Contributions

The expected outcomes of this workshop include:

- A comprehensive understanding of the current state of SLM research in education, including recent advancements, ongoing challenges, and potential applications.
- Identification of key research challenges and opportunities related to SLMs in education, informed by interdisciplinary discussions among AIED, EDM, and L@S researchers.
- Development of a shared research agenda that can guide future work in this area, which will be documented in a white paper co-authored by all participants.
- Strengthened interdisciplinary collaborations between researchers from the AIED, EDM, and L@S communities, as well as with participants from related fields such as human-computer interaction and machine learning.
- Dissemination of the workshop's findings and outputs through publications in the CEUR Workshop Proceedings[2] and a relevant journal, as well as through the workshop's website.

We believe that these outcomes will make significant contributions to the research communities by providing a clear agenda for the integration of SLMs into education. The insights gained from this workshop will help inform the design, development, and deployment of SLM-based educational applications accessible to a wide range of learners. In addition, the interdisciplinary collaborations that are fostered by this workshop will help drive innovation and advance research in this burgeoning area.

[2] https://ceur-ws.org/

4 Program Committee

The program committee of this workshop includes a PhD student, an early-career researcher, and four senior researchers from the AIED, EDM, and L@S communities. We have a track record of organizing successful workshops at major conferences in the field [7,4,8,6], and we are committed to fostering an inclusive and collaborative environment that encourages participation from researchers at all stages of their career. Our combined expertise in educational technology, data mining, and learning analytics will enable us to facilitate meaningful discussions and guide the development of a comprehensive research agenda for SLMs in education. The biographies of the organizers are as follows:

- **Yumou Wei** (Workshop Chair) is a PhD student in the Human-Computer Interaction Institute at Carnegie Mellon University. His research focuses on building educational technologies that support mastery learning at scale. He has published papers in the EDM and Learning Analytics and Knowledge (LAK) conferences describing innovative uses of SLMs for KC modeling [11] and question generation [13].
- **Steven Moore** is an Assistant Professor in the Department of Information Sciences and Technology at George Mason University. He studies how to design educational technologies that improve student learning and how people use AI to learn.
- **Paulo F. Carvalho** is an Assistant Professor in the Human-Computer Interaction Institute at Carnegie Mellon University. His research sits at the intersection of learning science and educational technology, investigating how the cognitive, metacognitive, and motivational processes of learners can inform the design of more effective, practice-first learning technologies.
- **John Stamper** is an Associate Professor in the Human-Computer Interaction Institute at Carnegie Mellon University. His research focuses on using big data collected from educational systems to improve student learning.
- **Christopher Brooks** is an Associate Professor in the School of Information at the University of Michigan. He builds and studies the effects of educational technologies in higher education and informal learning environments, with a particular domain focus on data science education and methodological interests in predictive modeling, learning analytics, and collaborative learning.
- **Michael Liut** is an Assistant Professor, Teaching Stream (Computer Science) in the Department of Mathematical and Computational Sciences, University of Toronto Mississauga. His research focuses on the design and development of educational technologies that support learning in computer science, with a particular interest in the use of AI and machine learning to enhance student learning outcomes.

This is the first edition of the workshop, and we hope to establish it as an annual event that continues to foster interdisciplinary collaboration and advance research on small language models in education.

References

1. Artificial Intelligence in Education: 26th International Conference, AIED 2025, Palermo, Italy, July 22–26, 2025, Proceedings, Part VI. Springer-Verlag, Berlin, Heidelberg (2025)
2. Belcak, P., Heinrich, G., Diao, S., Fu, Y., Dong, X., Muralidharan, S., Lin, Y.C., Molchanov, P.: Small language models are the future of agentic ai (2025). https://doi.org/10.48550/arXiv.2506.02153, https://arxiv.org/abs/2506.02153
3. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language models are few-shot learners. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) Advances in Neural Information Processing Systems. vol. 33, pp. 1877–1901. Curran Associates, Inc. (2020)
4. Kizilcec, R.F., Baker, R.S., Brooks, C., Geathers, J., Hicke, Y., Moore, S., Wu, B.: Applications of generative ai to support teaching and learning in higher education: A half-day workshop. In: Cristea, A.I., Walker, E., Lu, Y., Santos, O.C., Isotani, S. (eds.) Artificial Intelligence in Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium, Blue Sky, and WideAIED. pp. 267–274. Springer Nature Switzerland, Cham (2025)
5. Koutcheme, C., Woodrow, J., Piech, C.: Aligning small language models for programming feedback: Towards scalable coding support in a massive global course. In: Proceedings of the 57th ACM Technical Symposium on Computer Science Education V.1. p. 610–616. SIGCSE TS 2026, Association for Computing Machinery, New York, NY, USA (2026). https://doi.org/10.1145/3770762.3772539, https://doi.org/10.1145/3770762.3772539
6. Moore, S., Singh, A., Khosravi, H., Denny, P., Brooks, C., Stamper, J.: Partnerships for cocreating educational content. In: LAK23 Workshop on Partnerships for Cocreating Educational Content (2023), https://stevenjamesmoore.com/assets/papers/lak23_workshop_moore.pdf
7. Moore, S., Singh, A., Lu, X., Jin, H., Khosravi, H., Denny, P., Brooks, C., Wang, X., Kim, J., Stamper, J.: Learnersourcing: Student-generated content @ scale: 3rd annual workshop. In: Proceedings of the Twelfth ACM Conference on Learning @ Scale. p. 410–413. L@S '25, Association for Computing Machinery, New York, NY, USA (2025). https://doi.org/10.1145/3698205.3733963, https://doi.org/10.1145/3698205.3733963
8. Moore, S., Tong, R., Singh, A., Liu, Z., Hu, X., Lu, Y., Liang, J., Cao, C., Khosravi, H., Denny, P., Brooks, C., Stamper, J.: Empowering education with llms - the next-gen interface and content generation. In: Wang, N., Rebolledo-Mendez, G., Dimitrova, V., Matsuda, N., Santos, O.C. (eds.) Artificial Intelligence in Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium and Blue Sky. pp. 32–37. Springer Nature Switzerland, Cham (2023)
9. Reich, J., Ito, M.: From good intentions to real outcomes: Equity by design in learning technologies. In: Fitzgerald, W., Burns, J., Sonwalkar, N., Urry, J. (eds.) The Digital Learning Challenge: Obstacles to Educational Uses of Copyrighted Material in the Digital Age, pp. 1–42. The Digital Media and

Learning Research Hub, Irvine, CA (2017), https://clalliance.org/publications/good-intentions-real-outcomes-equity-design-learning-technologies/
10. UNESCO: Small language models (slms): A cheaper, greener route into ai (March 2024), https://www.unesco.org/en/articles/small-language-models-slms-cheaper-greener-route-ai, last updated 11 November 2025
11. Wei, Y., Carvalho, P., Stamper, J.: Kcluster: An llm-based clustering approach to knowledge component discovery. In: Proceedings of the 18th International Conference on Educational Data Mining. pp. 228–240. International Educational Data Mining Society (Jul 2025). https://doi.org/10.5281/zenodo.15870197, https://doi.org/10.5281/zenodo.15870197
12. Wei, Y., Carvalho, P., Stamper, J.: Small but significant: On the promise of small language models for accessible aied (2025), https://arxiv.org/abs/2505.08588
13. Wei, Y., Stamper, J., Carvalho, P.F.: Generate-then-validate: A novel question generation approach using small language models (2026), https://arxiv.org/abs/2512.10110