

Preface

Yumou Wei^{1,*}, Steven Moore², Paulo F. Carvalho¹, John Stamper¹, Christopher Brooks³ and Michael Liut⁴

¹Carnegie Mellon University, Pittsburgh, PA 15213, United States

²George Mason University, Fairfax, VA 22030, United States

³University of Michigan, Ann Arbor, MI 48109, United States

⁴University of Toronto Mississauga, Mississauga, ON L5L 1C6, Canada

1. About the Workshop

This volume contains the accepted papers of the 1st Workshop of Small Language Models for Education (SLM4ED 2026), held on June 28, 2026 in Seoul, Republic of Korea, as a full-day workshop co-located with the Festival of Learning 2026—the joint meeting of the International Conference on Artificial Intelligence in Education (AIED), the International Conference on Educational Data Mining (EDM), and the ACM Conference on Learning @ Scale (L@S).

We define small language models (SLMs) as open-weight language models with fewer than roughly ten billion parameters that can be deployed locally or hosted cheaply. As more model providers release smaller and more efficient models, and as the research community develops better techniques for fine-tuning and deploying them, SLMs are emerging as a credible alternative to large language models (LLMs) for educational applications. While LLMs have demonstrated impressive natural language processing capabilities, their adoption in education is constrained by privacy concerns, high inference costs, and proprietary limitations. SLMs offer a more scalable and accessible path to integrating AI into education, given their capacity to be efficiently fine-tuned on high-quality, domain-specific data.

The workshop was designed to explore this potential and, by bringing together researchers from the AIED, EDM, and L@S communities, to foster interdisciplinary collaboration and identify research directions that leverage the distinctive strengths of SLMs to improve learning outcomes.

2. Topics

The call for papers invited original, unpublished work on small language models in educational settings, addressing but not limited to the following themes:

- Pedagogical frameworks and instructional design for SLM-based interventions;
- Methodologies for evaluating and fine-tuning SLMs using educational data;
- Scalability and infrastructure design for deploying SLMs in large-scale settings;
- Privacy, equity, and accessibility considerations of SLMs vs. LLMs;
- Innovative applications (e.g., knowledge tracing, question generation, or adaptive feedback);
- Comparative studies of SLMs and LLMs in educational contexts.

The call was open: submission was not restricted to members of any particular project, institution, or consortium, and was advertised to the AIED, EDM, and L@S communities as well as to the sister communities of Learning Analytics, Educational Data Science, and Human-Computer Interaction.

SLM4ED'26: The 1st Workshop of Small Language Models for Education (SLM4ED), June 28, 2026, Seoul, Republic of Korea

*Corresponding author.

✉ yumouw@andrew.cmu.edu (Y. Wei); smoores59@gmu.edu (S. Moore); pcarvalh@andrew.cmu.edu (P. F. Carvalho); jstamper@andrew.cmu.edu (J. Stamper); brooks@umich.edu (C. Brooks); michael.liut@utoronto.ca (M. Liut)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

3. Submission and Review Process

Submissions took the form of research papers of 5–7 pages, excluding appendices and references, formatted with the CEURART template. Papers were submitted through EasyChair.

To allow workshop participants to take advantage of the early-bird registration deadline (May 10, 2026) for the Festival of Learning, we operated a two-round submission and review process. The first round had a shorter turnaround time so that acceptance notifications could be sent before the early-bird registration deadline, while authors who needed more time could still submit to the second round. A paper that was not accepted in the first round could not be resubmitted to the second round. The deadlines were as follows:

- Submission portal opened: April 2, 2026
- First-round submission deadline: April 23, 2026
- First-round notification of acceptance: May 1, 2026
- Second-round submission deadline: May 15, 2026
- Second-round notification of acceptance: May 29, 2026
- Camera-ready deadline (both rounds): June 16, 2026

Reviewing was double-blind: submissions were anonymized, and reviewers did not see author identities. The program committee, listed in Section 7, was drawn from the organizing committee together with one further reviewer from outside it. The workshop chair administered the process—assigning each submission to committee members matched to their expertise, handling conflicts of interest, and making the final decisions—but did not serve as a reviewer. Every submission that entered review received at least 2 independent reviews. Reviewers assessed each paper on originality, technical soundness, relevance to the workshop themes, and clarity, and acceptance decisions were based solely on the quality and relevance of the work. Conflicts of interest were declared, and a conflicted reviewer took no part in either the review of or the decision on the paper concerned. No member of the organizing committee is an author of any paper in this volume. One program committee member is an author of a paper in this volume and was excluded from the review of, and from the decision on, that paper.

The workshop received 15 submissions in total. Of these, 3 were desk-rejected before peer review, and the remaining 12 were sent out for review (3 in the first round and 9 in the second). Eight papers were accepted for presentation and for inclusion in this volume (2 from the first round and 6 from the second). No invited or non-peer-reviewed papers are included in this volume. All accepted papers were presented as lightning talks at the workshop.

4. Accepted Papers

The eight accepted papers span the evaluation, adaptation, and deployment of small language models across a range of educational tasks.

Three papers examine how well small models perform in tutoring and assessment roles. *Can Small Language Models Teach Effectively? Evaluating Explanations and Analogies Strategies* evaluates three sub-1.2B models on explanation-and-analogy strategies in mathematics tutoring dialogues, and characterizes their distinct pedagogical failure patterns. *CS TutorBench: Benchmarking Small Language Models as Tutors for Block-Based Programming* introduces a scenario-based benchmark with a pedagogical rubric for CS tutoring in a block-based robotics environment, and finds that model family and instruction tuning predict tutoring quality better than parameter count. *Bounding Boxes to Improve Small Language Model Performance on Vision-Based Grading Tasks* shows that cropping student responses with bounding boxes substantially improves both accuracy and computational cost when grading handwritten exam responses.

Two papers address the generation and tracing of learning content. *DiffiQG: Difficulty-Controlled Question Generation Using Small Language Models* conditions question generation on Bloom’s Taxonomy using fine-tuned Flan-T5 models to achieve reliable difficulty control without large-scale models.

Compression-First Knowledge Tracing with a Small Language Model compresses each student interaction into two embedding tokens, allowing a 150M-parameter backbone to process thousands of interactions per student and match established knowledge-tracing baselines.

The remaining three papers concern fine-tuning, deployment architecture, and domain adaptation. *Fine-Tuning A 1.2B Large Language Model for Systematic Review Screening in Educational Technology* demonstrates large performance gains from fine-tuning a small open-weight model for title-and-abstract screening in a systematic review. *Browser-Sized Agents: Small Language Models and WebLLM for Tool-Augmented AI-Assisted Programming Education* argues for deploying SLMs as tool-augmented agents rather than monolithic oracles, and sketches an in-browser architecture for programming education. *Structured Feedback for Arabic Competitive Debates Using a Small Language Model* combines task-specific argument mining with an open-weight 9B Arabic model to produce rubric-aligned, evidence-grounded feedback for debate adjudicators and coaches.

5. Workshop Program

The workshop was organized around a “context and connection” morning phase and a “consensus and output” afternoon phase. Morning and afternoon paper sessions were interleaved with open discussion and small-group synthesis to stimulate the interchange of ideas across the AIED, EDM, and L@S communities. Workshop participants are planning to co-author a white paper outlining a shared research agenda for small language models in education based on the discussions and insights generated during the workshop.

6. Organizing Committee

- Yumou Wei (Workshop Chair), Carnegie Mellon University, United States
- Steven Moore, George Mason University, United States
- Paulo F. Carvalho, Carnegie Mellon University, United States
- John Stamper, Carnegie Mellon University, United States
- Christopher Brooks, University of Michigan, United States
- Michael Liut, University of Toronto Mississauga, Canada

7. Program Committee

The program committee for SLM4ED 2026 comprised four members of the organizing committee together with one further reviewer from outside it. Its members were widely recognized researchers in the AIED, EDM, L@S, learning analytics, and computing education research communities that the call for papers addressed:

- Steven Moore, George Mason University, United States
- Paulo F. Carvalho, Carnegie Mellon University, United States
- Christopher Brooks, University of Michigan, United States
- Michael Liut, University of Toronto Mississauga, Canada
- H. Chad Lane, University of Illinois Urbana-Champaign, United States

Papers were selected on the quality and relevance of the work alone. Review assignments, conflict-of-interest handling, and final acceptance decisions were coordinated by the workshop chair under the double-blind process described in Section 3. The call for papers, the review process, and the program committee were specific to SLM4ED 2026 and separate from those of the Festival of Learning 2026 host conferences.

8. Acknowledgments

We thank the authors for their contributions, and the organizers of the Festival of Learning 2026 for hosting the workshop.

Yumou Wei, Steven Moore, Paulo F. Carvalho, John Stamper, Christopher Brooks, and Michael Liut
Editors

September 2026

Declaration on Generative AI

The authors did not use generative-AI tools in the preparation of this preface.