

Agentic PAL: Designing Human-Empowered AI Partnerships for Early Childhood Mathematics Learning

Anastasia Betts, Benjamin Ryon, Elida V. Laski, & Sunil Gunderia

Agentic AI

Early Childhood Mathematics

Family Engagement

Human-centered Design

learning engineering

Smart Learning Systems

The rapid expansion of AI-powered edtech has outpaced the field's capacity to ensure developmental appropriateness in early childhood mathematics. General-purpose large language models (LLMs) often generate activities that conflate concepts, skip developmental steps, or misalign with learning progressions, reflecting structural gaps in early childhood math expertise rather than limitations of AI itself. This paper introduces Agentic PAL, an AI-based system under development that uses a learning engineering approach to integrate learning science with a comprehensive, research-grounded early math knowledge infrastructure to support adult-child real-world math interactions. The theoretical foundations and emerging architecture behind Agentic PAL illustrate how thoughtfully designed AI-embedded edtech can strengthen adult expertise and learner outcomes in early childhood learning ecosystems.

Anastasia Betts¹, Benjamin Ryon², Elida V. Laski³, and Sunil Gunderia⁴

¹ Learnology Labs; anastasia@learnologylabs.org

² Learnology; ben@learnology.ai

³ Boston College; elida.laski@bc.edu

⁴ Mindset Copilot, sunil@mindsetcopilot.com

Abstract.

Keywords: Learning Engineering; Early Childhood Math; Agentic AI; Human-Centered Design; Smart Learning Systems; Family Engagement

Introduction

The rapid expansion of artificial intelligence (AI) in early childhood education has highlighted a growing tension between technological capability and the developmental nature of young children's learning (Hirsch-Pasek et al., 2015). Generic large language models (LLMs), now often embedded in "AI-powered" tools for learners, teachers, and families, frequently generate activities and explanations that misalign with early developmental learning progressions and fail to provide adults with the pedagogical supports needed to engage children in effective, developmentally grounded mathematics interactions (Betts,

2025). These challenges may not necessarily reflect the limitations of AI itself, but rather the absence of domain-specific knowledge infrastructures and pedagogical guardrails.

Agentic PAL (Personal Assistant for Learning) is an AI system that addresses this gap by integrating learning-science theory, learning engineering methods, and evidence from real-world use to empower adult-guided early math learning. Unlike traditional “AI tutors” that aim to instruct children directly on devices (Létourneau et al., 2025), Agentic PAL is designed to support (rather than replace) adults in helping young children learn through everyday real-world interactions. By functioning as a personalized learning concierge, Agentic PAL draws on a robust, continuously updated, research-grounded knowledge infrastructure to: (1) interpret informal adult observations and feedback captured during natural language conversations through a chat interface with an AI, (2) guide AI reasoning and system adaptivity, (3) support context-sensitive scaffolding, and (4) increase communication between home and school by sharing developmentally relevant, personalized data-driven insights for each learner.

Literature Review and Theoretical Grounding

Early mathematics learning emerges not from formal instruction alone, but through rich, frequent interactions with significant adults in everyday environments. Consider that children’s early exposure to mathematical talk, guided exploration, and hands-on math activity is strongly associated with later achievement (e.g., Levine et al., 2010; Gunderson & Levine, 2011). Moreover, high-quality adult–child interactions sharpen children’s attention to mathematical ideas (Daucourt et al., 2021). Effective systems aimed at supporting early mathematics should strive to strengthen the capacity and confidence of the adults who make these interactions possible, positioning adult expertise as a key leverage point for learning. Despite their critical role, many caregivers and early childhood educators report limited confidence in supporting young children’s mathematics learning (Ban et al., 2024). Math anxiety among parents is well-documented (Herts et al., 2019), and early childhood teachers often receive minimal preparation in early mathematics pedagogy or developmental learning trajectories (Clements & Sarama, 2020). These factors contribute to inconsistent or missed opportunities for mathematical learning in both home and school settings (Lu, Vasilyeva, & Laski, 2025).

Given this, the emergence of general-purpose large language models (LLMs) has fueled a rapid expansion of “AI-powered” tools for families and teachers. However, LLMs trained on broad internet data consume inconsistent and often inaccurate representations of early mathematics, including common confusions around foundational concepts (Betts, 2025). The resulting outputs include activities or explanations that are developmentally inappropriate or misaligned with learning progressions, and which provide limited guidance on how adults should scaffold, adapt, or interpret children’s thinking. These limitations highlight a critical learning engineering gap: early childhood mathematics requires AI systems that integrate learning-science principles, structured knowledge models, and explicit support for adult mediation in the real-world.

The learning sciences emphasize that young children progress through conceptual pathways in which later ideas depend on effective mastery of earlier ones (Bloom, 1984; Doignon & Falmagne, 1985), and that efficient progress along these pathways is supported through guided participation with a more knowledgeable other (Vygotsky, 1986; Zaretskii, 2009). Put simply, children learn most efficiently when adults offer timely scaffolds that extend their thinking just beyond what they can do independently, keeping learners in their Zone of Proximal Development (ZPD). This requires adults to recognize not only what the child currently understands, but what they are most ready to learn next (ZPD), and how to adjust support in the moment. In practice, reliably supporting children in this way places substantial cognitive and pedagogical demands on adults—demands that are rarely systematically or effectively supported in early childhood contexts.

These issues are further amplified by a “triple expertise” gap in early childhood mathematics (Betts, 2025): (1) a practitioner gap, in which many early childhood educators receive limited training in mathematical development (Clements & Sarama, 2020; Laski et al., 2013); (2) a perception gap, in which adults underestimate the complexity and importance of early math (Bittner & Bull, 2022); and (3) an infrastructure gap, in which the field lacks widely available, fine-grained knowledge models that map the developmental dependencies and common misconceptions (Clements & Sarama, 2024). Addressing these gaps

requires extensive research-based knowledge infrastructures (Doignon & Falmagne, 1985) that include systematic representations of concepts, learning trajectories, prerequisite relationships, and pedagogical strategies that encode what is known from decades of early mathematics research (Clements & Sarama, 2020). When operationalized through a learning engineering approach, such infrastructures prevent generative AIs from perpetuating the limitations of their training data, and enable them to effectively and accurately reason about learner readiness, interpret informal evidence from adult-child interaction feedback, and generate developmentally aligned personalized guidance for each learner.

A Learning Engineering Approach to Agentic PAL

Agentic PAL's design approach reflects the core commitments of learning engineering: integrating the learning sciences, using human-centered design methodologies, and using data-informed decision making to iteratively refine PAL's learning system (Goodell, 2023). The evolution of Agentic PAL is grounded in a nested learning engineering cycle (Craig et al., 2025) in which caregiver–AI interactions generate learning evidence that enables PAL to update individual children's knowledge models, while aggregate patterns of use inform ongoing system design and refinement.

PAL's first iteration was intentionally low-fi: a lightweight, text-message–based system that delivered research-aligned early mathematics activities to families, paired with simple feedback prompts for parents to report how the activity went. The system was built on a comprehensive, granular knowledge model containing hundreds of early math learning objectives organized by domain and developmental progressions that mapped precursor and successor relationships (Clements & Sarama, 2020). Parents received targeted activities matched to their child's "readiness" to learn key mathematical concepts and skills, along with brief explanations of why each activity mattered and how to best support their child's mathematical thinking (i.e., "Learning Science for Parents").

A small exploratory pilot study ($n = 19$) conducted during the Spring of 2025 showed that children whose families used PAL had greater gains on the TEMA-3 (Ginsburg & Baroody, 2003), with an average increase of 5.6 points compared to 1.9 points in a business-as-usual control group ($d > 1.0$). All PAL children showed improvement (100% vs. 57% in the control group), and 82% progressed from below to above grade level, with nearly all reaching kindergarten readiness and beyond by the posttest. Moreover, parent engagement was strongly associated with child learning gains ($r = .624, p < .01$; Betts, Ryon, & Laski, 2026). Building on this directional evidence, PAL is presently in use with ~500 students (from preschool through first grade), their families, and teachers across multiple sites in the United States. This larger-scale deployment is generating a rich stream of data and continues to affirm that adults benefit from regular, targeted support and that children's math learning increases when families engage. However, this broader real-world deployment has highlighted the limitations of a one-directional, prescriptive interface.

As the system has scaled, several new insights have emerged. Many families receive nudges at moments when they cannot engage, and although they may intend to return later, daily demands often intervene. Others complete activities but do not return to report how they went, leaving gaps in the evidence PAL needs to effectively adapt to a child's learning needs. Families also report frequent contextual mismatches; for example, an activity designed for the kitchen may arrive while they are at the grocery store, at the park, or waiting in the car, making it difficult to implement. In addition, many parents and early childhood teachers experience math anxiety or limited confidence in their own mathematical understanding (Herts et al., 2019; Lu, Vasilyeva, & Laski, 2025), which can make even accessible, informal activities feel intimidating.

Taken together, these realities highlight a fundamental limitation of traditional, structured input channels for capturing early learning in everyday contexts, and point to the kinds of challenges that thoughtfully designed AI systems are uniquely positioned to address. Systems intended to strengthen adult–child interactions must adapt not only to children's learning trajectories, but also to the affordances, constraints, and emotional landscapes of the adults who support them. Given this, a set of design requirements were identified based on close observation of how families and teachers interacted with PAL in real contexts (Table 1).

Table 1: Data-driven design requirements for Agentic PAL

Design Requirement	Description
R1: Flexible Timing	Adults must be able to share insights whenever they have capacity, not only when a prompt arrives.
R2: Natural Language Input	Parents and teachers need to describe interactions in their own words (in their own language), without structured reporting.
R3: Context Adaptation	Activities should adapt to real contexts (e.g., grocery store, car, park), not assume fixed settings.
R4: Interpretable Informal Evidence	The system must extract reliable learning signals from messy, everyday descriptions.
R5: Reduce Math Anxiety	Interactions should feel supportive and non-evaluative, building adult confidence.

Future Directions and Planned Research on Agentic PAL

Agentic PAL is currently moving through multiple phases of iterative design focused on developing a more dynamic and responsive solution, including: refining the conversational interface, strengthening natural-language evidence extraction, and operationalizing probabilistic inference mechanisms that map adult input onto the child's individualized knowledge graph in PAL's knowledge infrastructure. The next phase of work will compare the efficacy of Agentic PAL with the previous PAL (1.0) system. Across conditions, activities will remain aligned to the same underlying knowledge model, enabling the effects of the conversational, agentic layer on learning and family engagement to be examined. Data will include parent and teacher interactions with the system, process-level indicators of adult-child mathematical engagement, children's learning trajectories within the knowledge model, and standardized assessment outcomes. This research will address several key questions:

1. Does Agentic PAL support more frequent or higher-quality adult-child mathematical interactions in everyday contexts?
2. Does natural-language evidence improve the accuracy and stability of the child's evolving knowledge graph?
3. Do children whose parents use Agentic PAL demonstrate accelerated progress along early mathematics learning pathways compared to those supported by PAL 1.0?

Findings from this work will inform the iterative refinement of Agentic PAL going forward and contribute insights for others designing AI systems that empower humans, are grounded in the learning sciences, and who use learning engineering approaches to support learning across dynamic real-world contexts.

Conclusion

Rapidly evolving AI-powered tools in early childhood edtech underscore the need for a stronger, more explicit commitment to learning engineering approaches. Designing systems for young learners requires more than technical sophistication or generative flexibility; it demands grounding in the learning sciences and intentional design for how learning actually unfolds in everyday adult-child interactions. Effective systems must integrate theory, evidence, and iterative design in ways that strengthen the adults who mediate and guide young children's early learning. Agentic PAL illustrates how these learning engineering commitments can be operationalized in practice. By combining a research-based knowledge infrastructure with a conversational, knowledge-guided interface, Agentic PAL is designed to help adults recognize what a child is ready to learn, how to scaffold effectively, and when and how to introduce appropriate levels of support and challenge. Importantly, Agentic PAL builds directly on evidence and insights generated through earlier implementations, demonstrating how learning engineering cycles connect field data, theory refinement, and system redesign. In this way, Agentic PAL serves not as a

singular solution, but as an example of how learning engineering approaches can help us move from insight to implementation, in direct alignment with the LERN Convening theme.

References

- Ban, J., Msall, C., Douglas, A. A., Rittle-Johnson, B., & Laski, E. V. (2024). Knowing what they know: Preschool teachers' knowledge of math skills and its relation to instruction. *Journal of Experimental Child Psychology*, 246, 105996.
- Betts, A. (2025, July 23). The AI mirror: How GenAI reflects and amplifies gaps in early math. The Learning Agency.
- Betts, A., Ryon, B., & Laski, E. (2026, April). "Smart" nudges, real gains: Family engagement for kindergarten math readiness. Paper to be presented at the Annual Meeting of the American Educational Research Association (AERA), Los Angeles, CA.
- Bittner, K., & Bull, R. (2022). Understanding educator attitudes towards and perceptions of mathematics in early childhood. *Journal of Early Childhood Research*, 20(3), 341-356.
- Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational researcher*, 13(6), 4-16.
- Clements, D. H., & Sarama, J. (2024). Systematic review of learning trajectories in early mathematics. *ZDM – Mathematics Education*. Advance online publication.
- Clements, D. H., & Sarama, J. (2020). *Learning and teaching early math: The learning trajectories approach*. Routledge.
- Craig, S. D., Avancha, K., Malhotra, P., C., J., Verma, V., Likamwa, R., Gary, K., Spain, R., & Goldberg, B. (2025). Using a Nested Learning Engineering Methodology to Develop a Team Dynamic Measurement Framework for a Virtual Training Environment. In International Consortium for Innovation and Collaboration in Learning Engineering (ICICLE) 2024 Conference Proceedings: Solving for Complexity at Scale (pp. 115-132).
- Daucourt, M. C., Schatschneider, C., & Hart, S. A. (2021). The home math environment and math achievement: A meta-analysis. *Journal of Experimental Child Psychology*, 205, 105073.
- Doignon, J. P., & Falmagne, J. C. (1985). Spaces for the assessment of knowledge. *International journal of man-machine studies*, 23(2), 175-196.
- Ginsburg, H. P., & Baroody, A. J. (2003). *Test of Early Mathematics Ability (3rd ed.)*. PRO-ED.
- Goodell, J. (2023). Chapter introduction: What is learning engineering? In J. Goodell & J. Kolodner (Eds.), *Learning engineering toolkit: Evidence-based practices from the learning sciences, instructional design, and beyond* (pp. 5–13). Taylor & Francis.
- Gunderson, E. A., & Levine, S. C. (2011). Some types of parent number talk count more than others: Relations between parents' input and children's cardinal-number knowledge. *Developmental Science*, 14(5), 1021–1032.
- Guss, S. S., Clements, D. H., & Sarama, J. (2022). High-quality early math: Learning and teaching with trajectories and technologies. In A. Betts & K.-P. Thai (Eds.), *Handbook of Research on Innovative Approaches to Early Childhood Development and School Readiness* (pp. 349-373). IGI Global.

- Herts, J. B., Beilock, S. L., & Levine, S. C. (2019). The role of parents' and teachers' math anxiety in children's math learning and attitudes. In *Mathematics anxiety* (pp. 190-210). Routledge.
- Hirsh-Pasek, K., Zosh, J. M., Golinkoff, R. M., Gray, J. H., Robb, M. B., & Kaufman, J. (2015). Putting education in "educational" apps: Lessons from the science of learning. *Psychological science in the public interest*, 16(1), 3-34.
- Laski, E. V., Reeves, T. D., Ganley, C. M., & Mitchell, R. (2013). Mathematics teacher educators' perceptions and use of cognitive research. *Mind, Brain, and Education*, 7(1), 63-74.
- Létourneau, A., Deslandes Martineau, M., Charland, P., Karran, J. A., Boasen, J., & Léger, P. M. (2025). A systematic review of AI-driven intelligent tutoring systems (ITS) in K-12 education. *npj Science of Learning*, 10(1), 29.
- Levine, S. C., Suriyakham, L. W., Rowe, M. L., Huttenlocher, J., & Gunderson, E. A. (2010). What counts in the development of young children's number knowledge? *Developmental Psychology*, 46(5), 1309–1319.
- Lu, L., Vasilyeva, M., & Laski, E. V. (2025). Home math environment as a mediator of socioeconomic differences in early math skills: A study of Chinese families from disparate backgrounds. *Developmental Psychology*.
- Thai, K. P., Craig, S. D., Goodell, J., Lis, J., Schoenherr, J. R., & Kolodner, J. L. (2023). Learning engineering is human-centered. In J. Goodell (Ed.), *The learning engineering toolkit: Evidence-based practices from the learning sciences, instructional design, and beyond* (pp. 83–124). Routledge.
- Vygotsky, L. S. (1986). *Thought and language*. MIT press.
- Zaretskii, V. K. (2009). The zone of proximal development: What Vygotsky did not have time to write. *Journal of Russian & East European Psychology*, 47(6), 70-93.

