EPiC Series in Technology

Volume 3, 2026, Pages 93–97

Proceedings of AI for Accelerated Research Symposium

AI as an Accelerant for the Learning Sciences: Opportunities, Risks, and a Vision for the Future

Stephen Hutt

University of Minnesota Twin Cities, Minneapolis, MN, USA
shutt@umn.edu

Abstract

Artificial Intelligence (AI) offers the learning sciences new possibilities to advance not only educational practice, but also the study and theory of learning itself. AI enables richer insights from multimodal data, operationalizes long-standing educational theories in real-world contexts, and supports rapid experimentation, creating opportunities to address persistent inequities in research participation and representation. At the same time, these opportunities carry risks, including reinforcement of systemic inequalities, erosion of theoretical grounding, and privileging of narrow cultural perspectives. This vision paper argues for an interdisciplinary and human-centered approach that integrates the values of the learning sciences with the technical capabilities of AI. We present a vision grounded in participatory methods, responsible data science, and inclusive design, supported by privacy-preserving infrastructures.

1 Introduction

Technological innovation, particularly in computer science, has long influenced how we teach and learn, for example, intelligent tutoring systems (1) and addressing scale through Massive Open Online Courses (MOOCs)(6). As technology evolved, for example with development of areas such as affective computing (8), educational tools adapted in turn, with intelligent tutors incorporating emotional and metacognitive scaffolding (2, 4). However, these innovations have primarily shaped how we deliver education, not how we understand the learning process.

The Learning Sciences is an interdisciplinary field dedicated to investigating how people learn in both formal and informal settings. It draws from psychology, education, computer science, neuroscience, and design to study the cognitive, social, and contextual processes that underlie learning. By grounding research in theory and empirical evidence, the Learning Sciences aims to inform and support the practices of learning and teaching. AI has the potential not only to improve learning technologies, but also to fundamentally reshape how we *study and understand learning itself*. This paper argues that AI has the potential to catalyze discovery in the learning sciences by operationalizing theory, accelerating inquiry, and addressing persistent equity gaps. Realizing this promise requires developing and deploying AI through a critical and reflective lens.

2 AI as an Accelerant for Learning Sciences

AI presents a significant opportunity to accelerate discovery in the learning sciences. While technology has long enabled new forms of data collection and intervention, AI allows us to more efficiently analyze, interpret, and act upon that data. In this section, we focus on how AI can enhance both the depth of insights we derive from educational data and the speed at which we can translate those insights into meaningful, theory-informed advancements.

Modern learning environments generate vast amounts of data—from log files and other multimodal sources. Traditional statistical and rule-based approaches often fall short when it comes to capturing complex, dynamic relationships between behavior, cognition, and emotion. AI methods, particularly in machine learning and deep learning, offer new tools for discovering rich latent structures in these complex datasets such as subtle, temporally evolving patterns in student engagement, learning strategies, or emotional states—patterns. Just as importantly, AI allows these insights to be fed back into real-time educational systems to enable adaptive systems that respond to learners (4), while also providing researchers with new lenses through which to study learning processes at scale.

There are also potential theoretical benefits. The learning sciences have produced decades of rich theoretical models—constructs like metacognition, self-regulated learning, productive struggle, cognitive load, and motivational orientation. Although these frameworks are powerful, their translation into scalable educational tools has been slow and costly; AI can bridge this gap by facilitating operationalizing theory in real-world data. For example, recent work has begun to model constructs such as persistence or regulation using data captured from students' interactions with educational technologies (5, 9). These models, in turn, allow us to build adaptive systems that reflect and test our theoretical assumptions, while also iteratively refining those theories in response to real-world data. What was once prohibitively expensive or logistically unfeasible is increasingly achievable with AI-enhanced learning analytics. This opens new paths for theory-driven experimentation and validation at an unprecedented scale.

Taken together, the ability to extract meaningful insights from complex data and to embed theory into real-time systems changes the pace of educational research, bringing the "fail fast" ethos of the startup world into the (traditionally) slower cycles of educational research. This does not mean abandoning rigor. Instead, it means combining the rich academic history of the learning sciences with technological advances that allow us to iterate, test, and refine more efficiently. AI, when used responsibly and reflectively, can move us from multi-year research cycles to more agile, theory-informed cycles of discovery.

3 AI and Equity in Learning Sciences

AI holds significant promise as a force for equity in the learning sciences and in education more broadly, though, as will be discussed in Section 4, its potential for harm must not be underestimated. At present, many underrepresented populations remain absent from in educational research datasets. This absence is not due to lack of importance, but rather to limited access to data collection opportunities, geographic and infrastructure barriers, and systemic inequities that shape who participates in research (3). AI offers new tools to address these gaps, including the ability to generate synthetic data that can augment primary datasets. When carefully constructed, such data can help create more balanced and fair models that better reflect the needs of minority groups.

Beyond synthetic augmentation, AI can also help address the long-standing "WEIRD" problem in the learning sciences, in which research has disproportionately focused on Western, Educated, Industrialized, Rich, and Democratic populations. AI-powered platforms can support data collection in more geographically and culturally diverse settings, enabling cross-cultural, multilingual, and contextually relevant research. Complex constructs such as self-regulated learning, collaboration, or affective states can now be measured in classrooms, homes, community centers, or entirely online learning environments. Furthermore, AI can support the creation of adaptive educational tools designed for a variety of platforms, ensuring that learning research extends into a variety of environments. However, it is important to note that AI iteself is inhernetly "WEIRD" (7), it is developed in the same nations that are typically studied, so while there is potential, we must also be conscious of a risk of perpetuating the problem.

AI may also serve as a mirror, revealing the assumptions and biases embedded within the learning sciences. By analyzing large, diverse datasets, it can surface inequities and patterns of exclusion that might otherwise go unnoticed. When coupled with critical interpretation, this reflexive capacity helps researchers recognize how their own perspectives, data choices, and theoretical frameworks shape findings—fostering a more equitable and representative foundation for future discovery.

4 Risks and Reflexivity

While AI holds great promise for accelerating discovery in the learning sciences, its integration into educational research requires a critical, reflective approach. Systems that overlook lived experience or exclude certain groups risk not only falling short, but deepening the very inequities they aim to address. A key concern is that AI may unintentionally homogenize or silence diverse learning experiences. Education is already shaped by historical and systemic inequities, and systems trained on biased or incomplete data risk reinforcing dominant narratives while neglecting learners whose needs fall outside the norm. As AI becomes more embedded in classrooms and research, the potential for pervasive surveillance also grows—threatening to constrain the exploration and risk-taking essential to deep learning. At the same time, the accessibility of large-scale data and powerful AI tools can encourage a mindset that prioritizes speed over depth, displacing the theoretical grounding that gives educational research its lasting value. Although earlier (Section 2) we highlighted the benefits of more agile experimentation, this must be balanced by reflective, theory-informed practice that safeguards diversity, privacy, and rigor in equal measure.

We must also confront the issue of epistemic injustice: Whose knowledge is being embedded in the systems we design? The training data for many large AI models is disproportionately shaped by Eurocentric, Western, and Anglophone perspectives, privileging certain ways of knowing while sidelining others (7). In educational contexts, this can define what "counts" as learning through narrow cultural lenses, excluding or diminishing the knowledge systems of historically marginalized communities. To build a truly global and inclusive science of learning, researchers must ask which voices are represented, and which are absent. Only through sustained reflection and inclusive practice can we ensure that AI in the learning sciences leads to more just, equitable, and expansive forms of understanding.

5 Toward a Human-Centered, Reflective AI in Learning Sciences

Realizing the potential of AI to advance the learning sciences requires more than technical innovation, it demands a shared vision that integrates the strengths, values, and methods of multiple disciplines. Bringing together AI and the learning sciences means uniting not only bodies of knowledge, but also the value systems that underpin each field. The human-centered ethos of the learning sciences offers a counterbalance to the efficiency and scale-oriented approaches that often dominate AI development. By embedding principles of responsible data science and design justice, we can ensure that AI-driven educational research serves the diverse needs of learners while also contributing back to AI itself, infusing its development with perspectives and priorities beyond computer science.

To make this vision a reality, the learning sciences community must embrace deep, sustained collaboration. AI's future should not be shaped solely by technologists; it must be co-created with educators, psychologists, ethicists, community stakeholders, and learners themselves. Such collaboration will allow AI to transcend narrow disciplinary boundaries, supporting research and practice that are both technologically sophisticated and socially grounded.

This vision also requires robust infrastructure, including secure, privacy-preserving platforms for experimentation; accessible and interoperable data infrastructures; and inclusive research designs that allow broad participation. Given that educational data is protected by privacy regulations worldwide, we must actively engage with legal, ethical, and technical strategies for safeguarding learner information while enabling meaningful research. Without intentional investment in these infrastructures, many of the most transformative possibilities for AI in the learning sciences will remain out of reach.

Looking ahead, the goal is not simply to create systems that adapt to learners, but to design systems that also adapt to the evolving theories of learning themselves. As our understanding of cognition, motivation, and social interaction deepens, AI systems should be capable of integrating these shifts in real time, effectively becoming partners in the scientific process. Achieving this will require reimagining AI not as a static tool, but as a dynamic collaborator, one that grows alongside our theories, our evidence base, and our commitment to equity. In doing so, we can chart a course where AI accelerates discovery, amplifies inclusion, and strengthens the foundations of the learning sciences.

6 Conclusion

AI offers the learning sciences a transformative opportunity, not merely to enhance educational tools, but to accelerate theory-driven discovery, deepen our understanding of learning, and broaden participation in that discovery. As argued throughout, AI can support the operationalization of constructs like self-regulation and metacognition at scale, enable faster and more iterative cycles of experimentation, and help address long-standing inequities in representation. Realizing this promise requires vigilance toward its risks, such as, the reinforcement of bias, erosion of theoretical grounding, and privileging of narrow perspectives, as well as a commitment to human-centered, interdisciplinary collaboration. By investing in privacy-preserving infrastructures, adopting participatory design practices that include educators and learners, and embedding reflexive evaluation frameworks that foreground equity and rigor, the field can shape AI as a genuine collaborator in advancing theory and inclusion. Ultimately, our task is not simply to apply AI, but to co-evolve with it, designing systems that adapt alongside evolving theories of learning and reflect the full diversity of human experience.

References

[1] J. R. Anderson, C. F. Boyle, A. T. Corbett, and M. W. Lewis. Cognitive modeling and intelligent tutoring. *Artificial intelligence*, 42(1):7–49, 1990.

[2] R. Azevedo and A. C. Strain. Integrating Cognitive, Metacognitive, and Affective Regulatory Processes with MetaTutor. In *New Perspectives on Affect and Learning Technologies*. 2011.

[3] C. Baek and T. Doleck. Learning analytics: A comparison of western, educated, industrialized, rich, and democratic (weird) and non-weird research. *Knowledge Management & E-Learning*, 16(2):217, 2024.

[4] S. Hutt, K. Krasich, J. R. Brockmole, and S. K. D'Mello. Breaking out of the Lab: Mitigating Mind Wandering with Gaze-Based Attention-Aware Technology in Classrooms. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 2021.

[5] S. Hutt, J. Ocumpaugh, J. M. A. L. Andres, N. Bosch, L. Paquette, G. Biswas, and R. S. Baker. Investigating SMART Models of Self-Regulation and their Impact on Learning. In *Proceedings of the International Conference on Educational Data Mining*, pages 580–587. International Educational Data Mining Society, 2021.

[6] T. R. Liyanagunawardena, A. A. Adams, and S. A. Williams. MOOCs: A systematic study of the published literature 2008-2012. *The International Review of Research in Open and Distributed Learning*, 14(3):202–227, 2013.

[7] R. Mihalcea, O. Ignat, L. Bai, A. Borah, L. Chiruzzo, Z. Jin, C. Kwizera, J. Nwatu, S. Poria, and T. Solorio. Why ai is weird and shouldn't be this way: Towards ai for everyone, with everyone, by everyone. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 28657–28670, 2025.

[8] R. W. Picard. Affective computing: Challenges. 59(1):55–64, 2003.

[9] M. Siadaty, D. Gasevic, and M. Hatala. Trace-Based Microanalytic Measurement of Self-Regulated Learning Processes. 3(1):183–214, 2016.