TechTrends (2025) 69:700–706
https://doi.org/10.1007/s11528-025-01116-7

COLUMN: RETHINKING CREATIVITY AND TECHNOLOGY IN EDUCATION

Between Imperative and Panic: Navigating AI's Educational Moment with Leon Furze

Lindsey McCaleb[1] · **Danah Henriksen**[1] · **Punya Mishra**[1] · **Nicole Oster**[1]

Accepted: 11 July 2025 / Published online: 29 July 2025

Keywords Artificial Intelligence · Technology · Education · Research · Learning · LLM · Chatbot

> Chatbots don't make sense—they make words — *Leon Furze*

> Every human is a unique intelligence. Developing a unique intelligence is a work of teaching and learning. And honoring that is the highest calling of a teacher — *John Warner*

Introduction

As we continue the exploration of the dynamic interplay between technology, education, and creativity, this article adds to our collection of conversations with influential voices around generative AI (GenAI) and education. Previous articles in this column series have explored the intertwined history of AI and education (Mishra et al., 2025a, 2025b), theoretical frameworks (Mishra et al., 2024), practical applications (Oster et al., 2024), and interviews with prominent voices in the field, such as Dr. Mairead Pratschke (Henriksen et al., 2025), Ethan Mollick (Henriksen et al., 2023), and Dr. Ronald Beghetto (Mishra & Henriksen, 2024). This article features a conversation with Leon Furze—a scholar, educator, and systems-oriented advocate for demystifying AI in education. As GenAI tools continue to reshape teaching and learning, Furze's insights move the conversation beyond the hype, toward thoughtful reflection on practice, infrastructure, educational values, and AI's broader social and cultural role. In speaking with Furze, we continue exploring what these shifts mean for the future of education.

Among the many scholarly practitioners within the AI-in-education conversation, Leon Furze stands out for his grounded, educator-centered perspective. A thoughtful and visible voice in the GenAI landscape, he connects practitioner insight, critical reflection, and pragmatic experimentation. With over fifteen years as a secondary English Language Arts teacher, his interdisciplinary background spans media studies, digital technology, drama, and STEM. Early in his career, Furze began blogging about writing and literacy, building a popular platform that became one of the first educator-facing sites to engage with GenAI. Currently a PhD candidate at Deakin University studying GenAI's implications for writing instruction, he is widely recognized as a leading voice of scholarly practice in the field.

Furze has also published books, articles, and online courses, initially centered on literacy and writing, and now focused on explicitly addressing the implications of GenAI on writing instruction. He is also known for co-developing the *AI Assessment Scale* (Perkins et al., 2024), which is now used by schools and educational institutions worldwide. Whether speaking, writing, researching, or designing theoretical frameworks, Furze approaches AI with a sensitivity to pedagogical context and an openness to ongoing revision, and these qualities have made him a timely and resonant voice in this evolving field. Before diving into specific frameworks and tools, it is worth grounding the conversation in how Furze thinks about AI itself: what it is, what it is not, and how these understandings shape educational practice.

✉ Danah Henriksen
danah.henriksen@asu.edu

Lindsey McCaleb
lmccaleb@asu.edu

Punya Mishra
punya.mishra@asu.edu

Nicole Oster
nicole.oster@asu.edu

[1] Mary Lou Fulton College for Teaching and Learning Innovation, Arizona State University, Phoenix, AZ 85069, USA

Demystifying GenAI

When discussing GenAI in education, Furze balances practical opportunities with deeper critical reflection, while also examining broader societal context. He pays close attention to the difference between perceptions and reality of what AI is and is not, and what AI can and cannot do. In his work, he has intentionally addressed the hype, discussing the metaphors of AI language and the myths surrounding the ever-changing technology. According to Furze:

> There's so much mythologizing of what it (AI) can do, which actually really helps the tech companies. It's great if you can put out a technology and then abstract and obfuscate it, to make everybody uncertain about what it actually does.

Furze's approach to teacher training begins with an iceberg analogy, echoing Hemingway's idea that much of a story lies beneath the surface (Furze, 2023a, 2023b). In Furze's analogy: below the waterline is the vast, hidden foundation: massive training datasets, complex neural network architectures with billions of parameters, advanced pattern recognition algorithms, and intensive computational processes. And at the waterline sits the trained Large Language Model (LLM), a mathematical structure that captures statistical patterns from the training data. Above the surface is what users experience, which is the chat interface and generated text, images, or media that result from deep statistical pattern matching. According to Furze, helping educators see that what looks like intelligence is actually advanced pattern recognition helps demystify the technology, revealing both its impressive capabilities and its limits:

> A little bit of technical understanding goes a long way for teachers. They intuitively get it when you start talking about it in those terms. "Oh, okay, so this is why it's biased…And this is why it doesn't work like Google Search. And this is why I can't get it to do x, y, and z…" Whereas if you rely on a kind of mythical thinking, you fall into the trap of interacting with AI like a human.

In other words, Furze argues, "LLM don't make sense—they make words" (Furze, 2023a, 2023b). Understanding this distinction is critical since it helps us recognize that the most hyped practical applications of AI in education may not actually be the best uses of the technology. Moreover, understanding the basic construction of large language models reveals why they fabricate information, why their output lacks grounding in verified truth, and why their responses are fundamentally statistical predictions and not reasoned conclusions.

Keeping this in mind, Furze has found a range of uses for AI tools across both personal and professional contexts. He noted that he uses them so frequently that he does not always realize it, reflecting how deeply AI has become integrated into his workflow. At the same time, he draws clear boundaries around where the technology is most useful. In areas like business and consulting, he finds AI tools incredibly helpful for tasks like generating website code, transcribing audio for blog posts, and organizing early drafts. His process is iterative, shaped by carefully crafted prompts that help preserve his voice and intent. However, when it comes to his PhD research, Furze finds AI far less effective. He describes it as a "very frustrating, ambiguous, unpredictable technology" and "kind of hopeless" for tasks requiring critical synthesis and original insight. This kind of discernment is essential for knowing when AI is helpful and when it is not, especially in considering how GenAI might be used in educational contexts.

Practically speaking, Furze sees value in using AI for time-saving tasks, such as administrative work or drafting basic communications, but cautions about its flaws and frustrations if used uncritically. He notes that teachers, and their expertise, are still central to the meaningful work, and cautions against relying on GenAI for complex, pedagogically-rich tasks such as building rich curriculum resources. Despite the hype around GenAI as a standalone tutor, teachers remain essential to its effective use in education. Understanding how the technology works is key, but it is just as critical to build policies and tools that support wise engagement. Furze's practical frameworks reflect this shift from simply explaining AI to helping educators integrate it with intention.

Creating Frameworks for Nuanced Use

Concerns about student use of GenAI and plagiarism continue to be a top priority for educators. As a literacy and writing teacher, this issue is central to Furze's work and ongoing questions about AI integration in education. In response, he began developing the AI Assessment Scale in early 2023, aiming to move beyond binary, yes/no policies. The result is a more nuanced, five-point tool designed to support both students and educators in navigating AI use. It became evident, he says, that there was a need for "more than a binary…use it or don't use it" approach to AI and assessment. At the time, there were traffic light models, or straight bans on usage, and Furze and colleagues believed that a five-point scale could bring a more realistic, and nuanced, tool to the hands of educators. The initial design has evolved over time and the language has been revised for applicability to multiple disciplines and contexts, while keeping it specific enough to provide clarity to students

and teachers alike. In this context, Furze emphasizes that there is no "one-size-fits-all" approach in education, especially when navigating something as complex as GenAI. According to him:

> We've framed it carefully as a conversation starter with students and faculty around what appropriate use might look like, with the understanding that no-one really knows. Where some policies have gone too hard towards banning and blocking, or maybe even too permissive, we've said that we think there are still appropriate and inappropriate ways to use the technology. But we need to figure that out together.

Thus, the scale is designed not as a rigid framework but rather to support collaborative conversations about appropriate AI use, which are essential as the technology evolves and our understanding of its capabilities develops. Recent updates to the scale have addressed visual design, shifting from traffic light colors to neutral pastels to avoid signaling "good" or "bad" use. A new circular layout also minimizes hierarchical connotations and underscores the value of each level in assessment design (Furze, 2024). The team released the scale as an open resource, encouraging institutions to adapt it locally and since its launch, the scale has been translated into over twenty languages, with versions tailored for primary and elementary students. More than a rubric, the AI Assessment Scale promotes critical thinking and supports shared reflection between students and educators on meaningful AI integration. Beyond providing the framework, Furze also shares examples of how he has applied the tool in his own practice (Furze, 2025a, 2025b, June).

When considering GenAI's impact on writing instruction, Furze encourages a shift toward viewing writing as a process of reflection, discovery, and meaning-making, rather than a fixed product or isolated assessment. While writing is often framed as a way to "make thinking visible," Furze argues that this "creative act has been turned into an academic checkpoint." He calls for a return to a process-oriented view of writing that centers the writer's voice and thinking. Similar to research that emphasizes the expressive and exploratory functions of writing (Warner, 2025), Furze advocates for helping students understand writing as a recursive process of thinking, where the act of putting words on paper (or a screen) helps to clarify and develop ideas. This view of writing as process-oriented, not strictly product-oriented, aligns with Murray's work on how writing is "the process of discovery through language" (Murray, 1972, p. 4). Furze argues that the question is not whether students should use GenAI in writing, but how they can do so while maintaining agency over their own thinking. He notes the importance of revisiting the writing process and exploring what meaningful AI use might look like at each stage:

> What does good AI use look like if you are generating ideas? How do you do it without rolling over your own human intuition and ideas? Or what does AI look like in the editing stage of writing? And how do you use it while sustaining your own voice? …So how do you use the technology at different points and moments of the writing process, in a way which still has your voice as the primary focal point?

The danger comes, he argues, when students outsource their thinking entirely to AI and short-circuit learning, as opposed to using AI to explore fresh perspectives. Educators, he suggests, should emphasize student agency and ownership of their ideas and voice throughout the writing process. While GenAI may be useful for steps such as idea-generation and revision, Furze argues that the responsibility for decision-making must remain with the student writer. This aligns with recent calls for human-centered AI use and design in education, particularly the importance of student agency and cognitive engagement during co-creative processes (Holmes et al., 2022; Luckin et al., 2022).

The dynamic interplay between human judgment and AI capabilities reflects what some researchers have described as creative friction: the productive resistance that arises when human values like nuance, depth, or intentionality meet the rapid output of generative tools (DeSchryver et al., in press). Across creative domains, this friction can invite reflection, slow thinking, and iterative engagement, to preserve human voice and meaning-making in Human-AI collaboration (Fu et al., 2025). Rather than rushing toward polished outputs, students benefit from technology collaborations where they must pause, critique, and rethink AI material, using it as a springboard rather than a shortcut.

To help students foster their own voice in writing, while using AI as a thought partner, Furze also recommends facilitating activities where students experiment with AI tools and then critically reflect on the results. This approach encourages learners to reflect on critical questions of voice, authority, bias, and more, not just in relation to how AI works, but also in terms of their own learning. Integrating AI into writing requires ongoing self-reflection as a writer. Beyond instructional frameworks, Furze applies this critical lens to broader claims about AI in education, particularly the promise of personalized learning, which he views with cautious skepticism.

Questioning the Promises of Personalization

Proponents of AI in education frequently argue that these tools can revolutionize learning through personalization, adapting content and pacing to individual student needs in ways that traditional classroom instruction cannot match. The promise

is compelling: AI tutors that adjust to each learner's style, chatbots that provide instant feedback, and systems that create customized learning pathways at scale for large numbers of students simultaneously.

Furze, however, has a somewhat critical perspective on AI and personalized learning, stating that "personalization in the tech world is often about isolation." When we consider how people all learn differently and have varied needs, Furze raises an important but often overlooked point: the perspectives and assumptions of those developing AI tools inevitably shape how the tools are designed and envisioned for use. While there may be a small percentage of learners who can learn successfully through a chatbot in a self-paced manner, that does not necessarily apply to the greater student population. Furze shares:

> The better approach from my experience working in schools is understanding that everybody learns in totally different ways, and that sometimes a student might respond well to having a little bit of additional support, whether it comes from AI or a human tutor or whatever, but those interactions with the teacher in the classroom are where actual personalization happens.

Furze points out that when students are not learning, the issue is often deeper learning disengagement rather than a lack of technological tools or resources. Simply throwing AI at the problem will not address these underlying causes. We must consider pre-existing issues related to curriculum structure and the overall educational system (e.g., siloed content blocks and tightly scheduled learning) that often limit learner agency. This also connects with long-standing critiques of traditional educational structures, including rigid hierarchies, standardized curricula, and teacher-centered instruction (Freire, 1970; Illich, 1971), as opposed to more student-centered, emancipatory models of learning. Furze also emphasizes the need to think critically about how to support students in making their learning and thinking visible (Hattie, 2008), and how that might shift with the use of emerging technologies. In balancing student autonomy with instructional guidance, educators remain essential to teaching and learning. As experts in pedagogy and content, they understand that while GenAI can support the process, it will not solve every problem. This tension between what AI purports to offer and what learners need becomes even clearer when we consider who benefits most from these tools, raising a paradox that complicates assumptions about AI-enabled learning.

The Subject Expertise Paradox

Perhaps counterintuitively, AI learning tools may work best for those who already have a solid grasp of the subject matter. Users' ability to benefit from AI tools often depends on both their domain knowledge and their understanding of how AI functions. These factors influence how accurately they interpret outputs, how useful they find the tool, and how effectively they can apply it to support meaningful learning. This pattern has for some time been observed in work on AI literacy and learner cognition. For example, Holmes et al. (2023) argue that AI is problematic for learners who do not possess enough background knowledge to critically evaluate outputs.

Mishra (2025) recently proposed a 4-quadrant framework to illustrate the interplay of domain and AI knowledge (Fig. 1). If we focus on two key contrasting quadrants in the diagram: users with low knowledge in both areas, AI and the domain to be learned, face "The Novice's Dilemma." They are most vulnerable to AI errors, lacking both the domain expertise to evaluate outputs and the AI literacy to recognize potential failure modes. In contrast, users in the upper right corner, with expertise in both areas experience "The Dual Expertise Challenge," achieving the most reliable outcomes while requiring the highest cognitive effort as they simultaneously evaluate both content accuracy and AI process reliability. Given that most learners sit in the bottom left quadrant (The Novice's Dilemma), we can start to understand why chatbot-based learning simply will not work for many.

Consistent with this framework, Furze argues that students need foundational skills and content knowledge first before receiving additional support through AI platforms:

> It's a paradox…you get this learning technology, which is 'adaptive' and 'personalized', and all of those fancy words. But in order to drive it properly, you already need to know what you're doing in the subject that you're trying to learn. And you know, [I am] speaking from personal experience…I don't always know the right question to ask, or I don't have some fundamental piece of knowledge.

This insight reframes how we think about scaffolding within Vygotsky's Zone of Proximal Development (ZPD), which describes the space between what a learner can do independently and what they can accomplish with support from a more knowledgeable other (Vygotsky, 1978); GenAI complicates this model. While it might function as a more knowledgeable other, it is also unreliable and lacks contextual awareness and pedagogical judgment. For expert users, it can act as a productive scaffold that extends thinking. For novices, however, GenAI blurs the boundaries of the ZPD by offering information that appears helpful but possibly inaccurate or misaligned with the learner's current understanding. Without sufficient prior knowledge, students may not clearly evaluate, interpret, or apply what the AI provides. This can lead to a false illusion of learning or even create or enhance misconceptions.

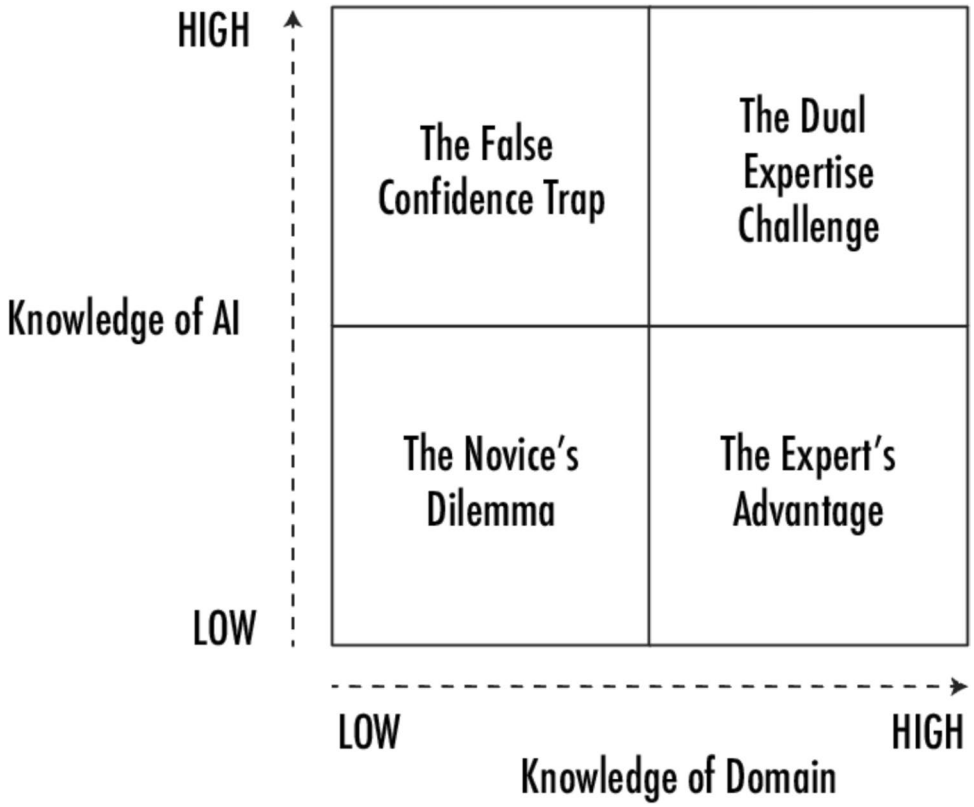

Fig. 1 The AI-Domain Knowledge Framework (Mishra, 2025)

Furze emphasizes the need to think critically about the broader structures of education, including how we design learning environments and balance technology use. His view aligns with frameworks such as Technological Pedagogical Content Knowledge (TPCK) (Mishra & Koehler, 2006), which call for thoughtful integration of content, pedagogy, and technology. Furze describes effective AI use as a balance between subject matter expertise, an engaged learner, and appropriate tools. If AI's value in education depends so heavily on context and critical use, then it must also be examined through ecological and ethical lenses. To this end, Furze offers a striking metaphor that captures the long-term implications of unchecked digital production.

GenAI Output as "Digital Plastic"

As generative AI tools flood the internet with synthetic content, sometimes referred to as "AI slop" (Mahdawi, 2025), concerns are mounting about the long-term effects on our digital ecosystems. From auto-generated clickbait and low-quality e-books to synthetic images and videos, these outputs are increasingly difficult to distinguish from authentic, human-created media. While some of this content may seem benign, it contributes to a broader shift in how we create, assess, and trust information online. This erosion of trust is especially troubling in a landscape already destabilized by years of social media amplification, misinformation, and declining confidence in traditional sources of news and expertise (Lewandowsky et al., 2017). For educators, these shifts raise urgent questions about how to prepare students for a world where the line between credible and artificial content is increasingly blurred. In response, scholars, journalists, and educators have begun turning to new metaphors and frameworks to make sense of this evolving digital terrain.

While metaphors for AI are increasingly common, Furze offers one that is especially compelling: that of seeing AI outputs as "digital plastic." The term refers to "a mass-produced, synthetic form of data that, like its physical counterpart, doesn't degrade much over time" (Furze, 2023a, 2023b). Like plastic, AI-generated content can be useful, but it also risks creating clutter and long-term damage. As Furze observes, "by the late 2020s, the amount of digital plastic in the online strata outweighs the earlier shades of human-generated content" (Furze, 2023a, 2023b). Just as we once underestimated the environmental consequences of physical plastic, we must now consider the risks of digital debris and algorithmic saturation (Crawford, 2021). A critical stance on AI use in education calls for intentionality not only in how we use these tools, but in how we teach students to be informed and reflective users. Media literacy, digital balance, and ethical technology design all play a role in helping learners navigate these risks. Metaphors like "digital

plastic" can be powerful tools in making abstract issues tangible. More recently, Furze and colleagues have argued that the digital plastic metaphor may also provide a theoretical foundation for understanding both the affordances and challenges of GenAI to help learners navigate this synthetic media landscape (Roe et al., 2025). Furze notes:

> "Digital plastic" is a useful metaphor because it speaks to the physical, real-world environmental impact of AI which we know is a problem….It's flooding the digital ecosystem, and like plastic, it's largely unregulated and pushed by a handful of small, already powerful companies. But, like real-world plastics, it can be useful as well as harmful…AI may be a really useful accessibility technology; it might be useful for creative people in many different ways. So it's not all negative, but it does need more cautious oversight.

Establishing effective oversight requires creating educational environments where conversations about AI remain productive, open, and collaborative. Furze emphasizes the importance of maintaining community cohesion during times of technological transition, especially given the rapid growth and limited public understanding of GenAI. This goes beyond traditional communities of practice (Wenger, 1998), calling for sustained collaboration even when views on technology diverge. Furze also notes that much of the current discourse is shaped by multinational corporations with a vested interest in accelerating adoption, often without considering educational needs or ethical concerns. He emphasized that educational community is essential, because:

> What I hate to see is where educators get very divisive…what we need to be doing is pointing that anger where it belongs, which is the companies responsible for the technologies. Be aware of the positives and negatives, and the strengths and limitations, and where you find limitations and things that don't ethically sit right, direct that anger upwards and continue to work with colleagues and students in a way which is more open, more accepting.

If educational leaders can build a culture of collaboration and team learning, teachers will be better prepared to guide students in appropriate uses of AI. When it comes to school policy, Furze notes the importance of asking vendors hard questions about their products and practices. At the same time, he cautions against overly specific policies, given how quickly GenAI is evolving. While there is a temptation to define terms like ethics, bias, and transparency, it may be more useful to establish flexible guidelines that invite ongoing dialogue rather than rigid compliance with rules that might soon be outdated. Just as the "digital plastic" metaphor encourages us to look beyond the surface appeal of innovation, Furze's broader stance calls for intentional pause. In this moment of rapid change, he urges us not to react, but to respond—with discernment, curiosity, and care. In doing so, he invites a deeper reckoning with what it means to lead, teach, and learn through uncertainty, offering a vision of education grounded not in fear or futurism, but in thoughtful engagement and collective sensemaking.

Conclusion

Underlying this conversation was perhaps the most challenging aspect of AI in education: its relentless pace of change. The constant and unbounded release of new models, features, and capabilities creates a feeling of technological whiplash. Furze offers a historical lens through which to view this moment of apparent chaos and uncertainty:

> The trajectory of all of these technologies—from phone lines to electricity, the Internet, or smartphones—is that they arrive on the scene with a big explosion. There's a lot of moral panic. There's a lot of imperative...and then, technology just becomes part of day-to-day life. It's important that teachers not get swept up in the imperative, but consider how we can make the most of this time, where we can critically engage before it disappears and just becomes a part of the way we do things.

This pattern reveals a great deal about our current moment with AI: we are living through the explosive phase, with urgent demands for immediate adoption or rejection. Rather than getting overwhelmed by the constant shapeshifting of AI, we can question the myth of technological inevitability (Furze, 2025a, 2025b April), and as educators, critically engage using thoughtful experimentation and evaluation.

When asked about the future trajectory of AI development, Furze takes a measured perspective that characterizes his approach to the field. Rather than trying to make sweeping predictions about GenAI and educational transformation, he focuses on what he calls the "near future," or the next six months. One key development he anticipates is the continued rise of what he terms "pseudo-reasoning" models: tools that convincingly simulate deep thinking while operating through complex pattern recognition behind the scenes. This recalls Bender et al.'s (2021) critique of large language models as "stochastic parrots," capable of mimicking human text without true understanding. The distinction between genuine reasoning and sophisticated pattern matching echoes Furze's earlier emphasis that educators should understand what GenAI actually does rather than what it appears to do; i.e., we need to demystify the magic.

Most importantly, Furze's advice to educators navigating this fast-changing landscape is about experimentation over anxiety, creating space between uncritical adoption

and reflexive rejection. Through his AI Assessment Scale, the "digital plastic" metaphor, his emphasis on community dialogue, and his engagement with other scholars and practitioners, Furze models what thoughtful engagement around GenAI might look like. His work suggests that the most valuable contributions to AI in education may not come from those who adopt every new feature or platform, but from those who pause to understand, question, and judiciously integrate tools while keeping the focus on educational goals and the needs of our students and educators. As AI continues its path toward becoming part of the educational "woodwork," Furze's measured approach offers a roadmap for educators to harness its potential while staying grounded in the core purposes of teaching and learning.

References

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021, March). On the dangers of stochastic parrots: Can language models be too big?. In Proceedings of the 2021 ACM conference on fairness, accountability, and transparency (pp. 610–623).

Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

DeSchryver, M., Henriksen, D., & Leahy, S. (in press). From friction to synergy: The complex interplay of human creativity and AI. *Possibility Studies and Society.*

Freire, P. (1970). *Pedagogy of the oppressed* (M. B. Ramos, Trans.). Herder and Herder.

Fu, Y., Newman, M., Going, L., Feng, Q., & Lee, J. H. (2025). Exploring the collaborative co-creation process with AI: A case study in novice music production. *arXiv preprint* arXiv:2501.15276.

Furze, L. (2023, May 18). The AI Iceberg: Understanding ChatGPT. *Leon Furze*. https://leonfurze.com/2023/05/18/the-ai-iceberg-understanding-chatgpt/

Furze, L. (2023, November 22). Chatbots don't make sense—they make words. *Leon Furze*. https://leonfurze.com/2023/11/22/chatbots-dont-make-sense-they-make-words/

Furze, L. (2024, August 28). Updating the AI Assessment Scale. *Leon Furze*. https://leonfurze.com/2024/08/28/updating-the-ai-assessment-scale/

Furze, L. (2025, April 28). The myth of inevitable AI. *Leon Furze*. https://leonfurze.com/2025/04/28/the-myth-of-inevitable-ai/

Furze, L. (2025, June 11) How I use the AI Assessment Scale: Part 1. *Leon Furze.* https://leonfurze.com/2025/06/11/how-i-use-the-ai-assessment-scale-part-1/

Hattie, J. (2008). *Visible learning*. Routledge.

Henriksen, D., Oster, N., Mishra, P., & McCaleb, L. (2025). Generative AI, Creativity, Culture, and the Future of Learning: A Conversation with Mairéad Pratschke. *TechTrends, 69*(1), 3–9.

Henriksen, D., Woo, L. J., & Mishra, P. (2023). Creative uses of ChatGPT for education: A conversation with Ethan Mollick. *TechTrends, 67*(4), 595–600.

Holmes, W. (2023). *The unintended consequences of artificial intelligence and education.* Education International.

Holmes, W., Bialik, M., & Fadel, C. (2022). *Artificial intelligence in education: Promises and implications for teaching and learning* (2nd ed.). Center for Curriculum Redesign.

Illich, I. (1971). *Deschooling society*. Harper & Row.

Lewandowsky, S., Ecker, U. K. H., & Cook, J. (2017). Beyond misinformation: Understanding and coping with the "post-truth" era. *Journal of Applied Research in Memory and Cognition, 6*(4), 353–369. https://doi.org/10.1016/j.jarmac.2017.07.008

Luckin, R., Cukurova, M., Kent, C., & Du Boulay, B. (2022). Empowering educators to be AI-ready. *Computers and Education: Artificial Intelligence, 3*, Article 100076.

Mahdawi, A. (2025, January 8). AI-generated 'slop' is slowly killing the internet, so why is nobody trying to stop it? *The Guardian*. https://www.theguardian.com/global/commentisfree/2025/jan/08/ai-generated-slop-slowly-killing-internet-nobody-trying-to-stop-it

Mishra, P. (2025, February 13). The GenAI and expertise paradox. *Punya Mishra*. https://punyamishra.com/2025/02/13/the-genai-and-expertise-paradox-why-it-makes-expert-work-more-important-but-harder/

Mishra, P., & Henriksen, D. (2024). Creative dialogue with generative AI: Exploring the possible with Ron Beghetto. *TechTrends, 68*(3), 395–401.

Mishra, P., & Koehler, M. J. (2006). Technological pedagogical content knowledge: A framework for teacher knowledge. *Teachers College Record, 108*(6), 1017–1054.

Mishra, P., Henriksen, D., & Dunnigan, J. (2025b). From symbols to statistics: The parallel histories of machine and human learning. *TechTrends, 69*, 459–466. https://doi.org/10.1007/s11528-025-01083-z

Mishra, P., Henriksen, D., Woo, L. J., et al. (2025a). Control vs. agency: Exploring the history of AI in education. *TechTrends, 69*, 247–253. https://doi.org/10.1007/s11528-025-01064-2

Mishra, P., Oster, N., & Henriksen, D. (2024). To Thine Own Mind Be True: Understanding Cultural Technologies, From Cave Walls to ChatGPT. *TechTrends, 68*, 975–982.

Murray, D. (1972). Teach writing as a process not product. *The Leaflet, 71*(3), 11–14.

Oster, N., Henriksen, D., & Mishra, P. (2024). Chatgpt for teachers: Insights from online discussions. *TechTrends, 68*(4), 640–646.

Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A Framework for ethical integration of Generative AI in educational assessment. *Journal of University Teaching and Learning Practice*, *21*(06). https://doi.org/10.53761/q3azde36

Roe, J., Furze, L., & Perkins, M. (2025). GenAI as Digital Plastic: Understanding Synthetic Media Through Critical AI Literacy. arXiv preprint arXiv:2502.08249.

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes* (M. Cole, V. John-Steiner, S. Scribner & E. Souberman, Eds.; A. R. Luria, M. Lopez-Morillas & M. Cole, Trans.). Harvard University Press.

Warner, J. (2025). *More than words: How to think about writing in the age of AI*. Basic Books.

Wenger, E. (1998). Communities of practice: Learning as a social system. *Systems Thinker, 9*(5), 2–3.