The "Shortcut Pedagogy" Illusion: When Teachers Mistake the Speed of AI for the Quality of Instruction

Hafiz Hanif

Faculty of Human Development, Sultan Idris Education University · ORCID 0000-0001-9943-6557

CITE AS Hanif, H. (2026, September 28). The “Shortcut Pedagogy” Illusion: When Teachers Mistake the Speed of AI for the Quality of Instruction. Hack(edu).tech Chronicle. https://hackedu.tech/articles/the-shortcut-pedagogy-illusion-when-teachers-mistake-the-speed-of-ai-for-the-qua

PUBLISHED 28 September 2026

DOI 10.5281/zenodo.23154222

LICENCE CC-BY-4.0

AVAILABLE AT https://hackedu.tech/articles/the-shortcut-pedagogy-illusion-when-teachers-mistake-the-speed-of-ai-for-the-qua

ABSTRACT

Educators who master basic prompting quickly reach a peak of confidence, believing AI can automate planning, grading, and differentiation. Drawing on Dunning-Kruger research, the illusion of explanatory depth, and 2025 evidence that confidence in AI predicts less critical thinking while self-confidence predicts more, this article names the result shortcut pedagogy: generic, biased, pedagogically shallow content scaled at speed, with diagnostic feedback outsourced before the tool’s limits are understood. Training must build the descent: calibrated trust through deliberate failure-finding.

KEYWORDS

Pedagogy, Artificial Intelligence, Dunning-Kruger, Psychology

Two hours. That is roughly how long it takes, in my experience running AI workshops for educators, before somebody in the room has their moment. The prompts start working. Lesson plans appear in seconds. Quiz items, differentiated worksheets, feedback comments, semua siap. By the coffee break the person is glowing, and by the following week I will receive the message I have learned to expect: "Dr., saya dah siapkan bahan satu semester. AI settle semua."

I understand the joy. I remember my own version of it. But this column is about why that message worries me more than it impresses me, because I have come to believe we are watching the birth of something I call shortcut pedagogy: teaching decisions made at the speed of generation, by educators who have climbed to a peak of confidence with the tool long before they have any map of its cliffs.

The psychology behind this is old, well-documented, and painfully relevant. In 1999, Kruger and Dunning published their famous studies showing that the people who perform worst at a task tend to overestimate themselves the most, and the reason is cruel in its simplicity: the skills you need to do a thing well are the same skills you need to judge whether you are doing it well (Kruger & Dunning 1999). A beginner is not just unskilled; they are missing the very

instrument that would tell them so. Now apply that to prompting. Prompting has the lowest floor of any professional skill I know. Type a sentence, receive a miracle. Nothing in that loop tells you what a deeply-briefed, pedagogically-anchored prompt would have produced instead, so nothing tells you that what you received was the generic version. The staircase to confidence is two steps tall. The staircase to competence is much, much longer, and the tragedy is that from the second step, both staircases look the same.

There is a second illusion stacked on top, and cognitive scientists have a lovely name for it: the illusion of explanatory depth. People consistently believe they understand how complex systems work, zippers, helicopters, toilets, until they are asked to actually explain the mechanism step by step, at which point their self-rating collapses (Rozenblit & Keil 2002). AI is the perfect host for this illusion. The tool responds so fluently, so obligingly, that using it feels like understanding it. Ask that same confident teacher to explain why the AI's "differentiated" worksheet actually matches the three ability levels in their own classroom, or what the model did with the word "differentiate" in the first place, and the explanation runs out fast. They understand the tool the way I understand a helicopter: enough to point at it, not enough to fly.

If this were only about feelings, no harm done. But in 2025, researchers surveying 319 knowledge workers across 936 real AI use cases found the mechanism that turns misplaced confidence into degraded work: the higher a person's confidence in the AI, the less critical thinking they applied to its outputs, while, notably, the higher their confidence in their own expertise, the more critical thinking they applied (Lee et al. 2025). Read that pair of findings again, slowly, because it is the whole story of shortcut pedagogy in one sentence. Trust in the tool switches the checking off. Trust in yourself switches it on. And whom does the two-hour workshop produce? A person with maximum new trust in the tool and no new trust in themselves. We are manufacturing, at scale, the exact psychological profile that the research says checks the least.

So what actually gets scaled when a confident beginner starts generating a semester in a weekend? Three things, mostly invisible to them.

Generic content, first. As I argued in an earlier column, a prompt is compressed pedagogy, and research on teachers' actual prompting shows the quality of what comes out tracks the specificity of what goes in (Moorhouse et al. 2025). The shallow prompt produces the lesson for a classroom that exists nowhere, and speed simply multiplies it. Speed, we must remember, is a multiplier, not an upgrader: it scales whatever quality you already have, including the absence of it.

Bias, second, quietly inherited from the machine's training diet and unexamined because the confident user is precisely the one not examining.

And third, the one that troubles me most as a researcher of learning: pedagogical shallowness wearing the costume of completeness. AI outputs look finished. The worksheet has instructions, sections, even a cheerful rubric. But "looks like teaching material" and "produces learning" are different properties, exactly the way plausible and true are different properties, and a teacher who cannot yet articulate the difference will ship the costume.

Which brings me to the outsourcing that should genuinely frighten us: diagnostic feedback. Of all the tasks teachers are rushing to hand over, marking with comments is the most popular, and I understand why; it is the most exhausting. But here is what the shortcut hides. Marking was never merely clerical. **Reading a student's wrong answer is how a teacher learns the class**: this one confuses area with perimeter, that one can compute but cannot explain, the back row all made the same error so my Tuesday explanation failed. Feedback is the teacher's diagnostic instrument pointed at their own teaching. Outsource it before you understand the tool's limits, and two things break at once: the student receives comments that are praise-shaped but diagnostically empty, and the teacher stops receiving the information stream that made them better every year. The emerging evidence on cognitive offloading, that heavy unguided

reliance on AI is associated with weakened critical thinking (Gerlich, 2025), lands hardest here, because the teacher is offloading the very cognition their profession is built on. A doctor who outsources the examination and keeps only the prescription is not more efficient. He is less of a doctor, faster.

I want to be clear, as always, that the villain here is not the tool and not the teacher. It is the shape of our training. We run workshops that teach prompting mechanics and stop exactly at the peak of confidence, then wonder at what follows. The fix is to build the descent into the training itself, deliberately, and the research points to how. Make teachers push the tool until it fails in front of them: generate the worksheet, then hunt its errors as a group activity, the way the illusion of explanatory depth is punctured, by demanding the explanation. Require the justification step: before any AI material enters a classroom, the teacher explains, in one minute, why this fits their actual students; if the explanation cannot be produced, the material is not ready, and neither is the adoption. Keep a calibration habit: mark one set of books by hand each term, not for nostalgia, but the way pilots hand-fly approaches, to keep the diagnostic muscle alive and to have a baseline against which the machine's feedback can be judged. And, following the Lee finding, invest training where it actually protects quality: in the teacher's confidence in their own pedagogical judgment, because that is the confidence that switches the checking on.

The teachers who will thrive with AI are the ones who pass through the valley, who learn the tool's limits the honest way and come up the far slope with calibrated trust. My worry is for the many we are leaving happily on the first peak, generating faster every week.

Because a shortcut you do not understand is not a shortcut. It is a detour that has not introduced itself yet, and in education, the ones who eventually pay the toll are not the ones who took the wrong turn.

References

Gerlich, M. (2025). AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking. Societies, 15(1), 6. https://doi.org/10.3390/soc15010006

Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. Journal of Personality and Social Psychology, 77(6), 1121–1134. https://doi.org/10.1037/0022-3514.77.6.1121

Lee, H.-P. (Hank), Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (pp. 1–22). ACM. https://doi.org/10.1145/3706598.3713778

Moorhouse, B. L., Ho, T. Y., Wu, C., & Wan, Y. (2025). Pre-service Language Teachers' Task-specific Large Language Model Prompting Practices. RELC Journal. https://doi.org/10.1177/00336882251313701

Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: an illusion of explanatory depth. Cognitive Science, 26(5), 521–562. https://doi.org/10.1207/s15516709cog2605_1