

Stop Blaming AI for What the Education System Abandoned

Basil C. Puglisi, MPA

A Human-AI Collaboration

A study of 372,793 college application essays finds that after ChatGPT becomes widely available, student writing becomes more lexically diverse while the underlying ideas become more uniform. The researchers call this “disjunctive homogenization.” The New York Times published a guest essay presenting the research as evidence that AI constricts creative thinking.

The data is sound, but the blame is wrong. AI does not homogenize student thinking; unstructured deployment without method governance does. The institution that deploys the tool without governing its use does not get to blame the tool for the outcome.

The Data Is Not in Dispute

The Moon, Kushlev, and colleagues’ study is one of the largest analyses of AI-era writing in a high-stakes admissions context. Three findings matter most.

Essays written after ChatGPT show increased vocabulary diversity but decreased idea diversity, both within individual essays and across the applicant pool. The surface language improves while the concepts converge. In a companion study of 2,200 essays, each additional human essay contributes more new ideas to the collective pool than each additional GPT essay, and the gap widens as more essays are added. Prompt modifications and parameter adjustments do not close it.

The third finding is the one that should end the conversation about blaming AI and start the movement to drive change within the institutions. Evaluators rate the AI-era essays as more creative, because polished vocabulary overrides the signals of conceptual sameness underneath. The evaluators are fooled by surface quality and cannot tell the difference between better words and better ideas. That is an evaluation failure, and the institution owns the evaluation.

The Discipline Failed First

Academic pedagogy has required “show your work” for generations, applying that standard across every domain it governs. Source citation, peer review, methodological transparency, evidence-based practice, and named human accountability for claims are the foundational standards the discipline produces, defends, and requires of credentialed members.

When AI enters classrooms, the discipline does not apply those same standards to deployment. No structured prompting protocol requires students to form ideas before consulting AI, and no metacognitive checkpoint requires them to evaluate what AI produces. Whether the deployment produces learning or formatted sameness goes unmeasured. The tools reach millions of students with less methodological rigor than many peer-reviewed journals require.

The institutions that spend decades teaching evidence standards deploy technology without meeting their own.

In January 2026, a cognitive neuroscientist testified before the Senate Commerce Committee that the problem is biological: humans evolved to learn from other humans, not from screens. He identified method governance as an alternative, then rejected it categorically, dismissing deployment quality and instructor training in sequence. His own data told a different story: Intelligent Tutoring Systems showed gains above his instructional benchmark while general one-to-one laptop deployment showed losses below it. The variance tracked deployment structure, with tightly governed systems above the benchmark and loosely distributed access below it. The discipline blamed the tool while the data blamed the method.

Some Users Are Lazy. That Is Still the Institution’s Problem.

The counterargument arrives on schedule: students do put in less effort when they can copy and paste AI output. Some users treat AI as a copy machine. Some don’t know how to prompt, sequence, or challenge AI output because no one taught them. These are fair observations, and pretending otherwise would weaken the argument.

But the observation doesn’t help the critics’ case, because it strengthens the institutional indictment.

Who is responsible for teaching students how to use the tool? Who deploys it without training or checkpoints, then measures the results and concludes the problem is the machine?

A 2026 validation study at the University of Hong Kong makes the institutional blind spot visible in the measurement itself. Researchers developed a “Metacognitive Laziness Scale” to measure student offloading of cognitive work to AI. The scale correlates with disaffection but shows a nonsignificant correlation with engagement. The measured pattern is withdrawal, and engagement shows no significant change. The researchers’ own limitations section acknowledges they cannot determine whether AI dependence diminishes metacognitive skills or whether students with preexisting deficits are simply more prone to offloading. The instrument measures students’ behavior without measuring whether the institution ever provided a governed deployment method. The construct is named “laziness,” and the deployment is never named at all.

When a hospital deploys a surgical tool without training residents to use it, and outcomes decline, the discipline does not blame the scalpel. When an aviation authority certifies a cockpit system without training pilots on its failure modes, and incidents increase, the regulator does not blame the avionics. The institution owns the method, and the method governs the outcome. Lazy use is a deployment failure before it is a user failure, and pretending otherwise lets the deployer walk away from the damage.

The Equity Finding Tells the Real Story

The research reports that homogenization hits hardest among minoritized applicants and linguistic minorities. The people whose perspectives are most distinct are the ones AI smooths toward the institutional mean.

That finding doesn’t say what the critics want it to say.

If AI smooths diverse voices toward a mean, the mean was always there. The institution always had a convergence target, always rewarded movement toward it, and always penalized deviation from it. The admissions counselors who rate AI-polished essays higher are not being tricked by the technology. They apply the standard they have always applied, and that standard has always been institutional voice.

For decades, polish was a proxy for privilege. The student with private tutors, editorial support, and college counselors produced institutional-grade writing. The student

without those resources produced writing that sounded different, that carried the texture of a life lived outside the credentialing pipeline. AI compresses that gap, and what the institution mourns is the gatekeeping function that polish once served.

Method Governs Outcome

Cognitive science has supported the answer for decades: method-governed deployment where cognitive demand remains on the learner.

The sequence is human idea formation first and AI engagement second, which means brainstorming before prompting and drafting before polishing, with each delegation decision made only after a human position exists. Structured interaction requires the student to articulate a position before the AI responds to it. Metacognitive checkpoints require the student to evaluate, accept, reject, or modify AI output rather than absorbing it wholesale.

This is not a new invention. It is the standard the discipline already requires of every other educational practice it governs. The discipline that requires evidence-based and method-governed practice with named human accountability for every other domain of its work does not get a pass on this one because the tool is new and the failure is embarrassing.

AI did not erase the human mind, it exposed whether the human mind showed up. The tool is a mirror. Blaming the mirror for what it shows is how institutions avoid the question they do not want to answer: what was the process producing before the mirror arrived?

References

- Alvero, A. J., Lee, J., Regla-Vargas, A., Kizilcec, R. F., Joachims, T., & Antonio, A. L. (2024). Large language models, social demography, and hegemony: Comparing authorship in human and synthetic text. *Journal of Big Data*, 11, Article 138. <https://doi.org/10.1186/s40537-024-00986-7>
- Dizon, J. I. W. T., Mendoza, N. B., Gašević, D., & Ganotice, F. A., Jr. (2026). Assessing AI-driven metacognitive offloading: Initial development and validation of the Metacognitive Laziness Scale. *ECNU Review of Education*, 9(2). <https://doi.org/10.1177/20965311261450994>
- Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28), Article eadn5290. <https://doi.org/10.1126/sciadv.adn5290>
- Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, 56(2), 489–530. <https://doi.org/10.1111/bjet.13544>
- Hattie, J. (2009). *Visible learning: A synthesis of over 800 meta-analyses relating to achievement*. Routledge. <https://doi.org/10.4324/9780203887332>
- Horvath, J. C. (2026, January 15). Written testimony of Dr. Jared Cooney Horvath, PhD, MEd, neuroscientist and educator, before the U.S. Senate Committee on Commerce, Science, and Transportation. U.S. Senate Committee on Commerce, Science, and Transportation. https://www.commerce.senate.gov/wp-content/uploads/media/doc/Horvath_Written%20Testimony.pdf
- Moon, K., Green, A. E., & Kushlev, K. (2025). Homogenizing effect of large language models (LLMs) on creative diversity: An empirical comparison of human and ChatGPT writing. *Computers in Human Behavior: Artificial Humans*, 6, Article 100207. <https://doi.org/10.1016/j.chbah.2025.100207>
- Moon, K., Kushlev, K., Bank, A., Lira Luttges, B., Viskontas, I., Kaufman, J. C., Johnson, D. R., Duckworth, A. L., & Green, A. E. (2026). The creative link between words and ideas is weakening in the AI era. *PsyArXiv*. https://doi.org/10.31234/osf.io/jsz58_v6
- U.S. Senate Committee on Commerce, Science, and Transportation. (2026, January 15). Experts tell committee AI presents greater risk to children than social media [Press release]. <https://www.commerce.senate.gov/press/dem/release/experts-tell-committee-ai-presents-greater-risk-to-children-than-social-media/>
- Winthrop, R. (2026, May 27). What 370,000 college essays tell us about A.I.'s effects on creativity. *The New York Times*. <https://www.nytimes.com/2026/05/27/opinion/writing-creativity-ai.html>

Original articles on cognitive research, education and AI Method Governance at BasilPuglisi.com

The Other AI: Audio Briefings on Augmented Intelligence and AI Governance
[Spotify](#) · [Apple Podcasts](#) · [Amazon Music](#) · [YouTube](#)

#AIassisted using HAIA Ecosystem