



Who benefits? Equity and the digital divide in Generative AI-supported learning in higher education

Swaminathan Govind

Department of Education, Bharathiar University, Coimbatore, Tamil Nadu, India

Abstract

Generative artificial intelligence (GenAI) is frequently promoted as a democratizing force in higher education, capable of extending personalized instructional support to students who otherwise lack access to it. This review examines whether the empirical literature (2023–2026) supports that claim or whether it instead points toward a widening gap between students who benefit from GenAI and those who do not. Drawing on recent meta-analyses of learning outcomes alongside digital-equity scholarship, the review argues that GenAI's documented benefits for academic performance are conditioned on forms of access that are themselves unevenly distributed: reliable connectivity and devices, AI and digital literacy, and instructor capacity to design effective GenAI-integrated tasks. It organizes these conditions into a three-level framework of access, skill, and outcome divides, and reviews evidence on how each operates in higher education contexts. The review concludes that without deliberate institutional and policy intervention, GenAI risks reinforcing rather than narrowing existing educational inequities, and outlines equity-oriented implications for pedagogy, policy, and future research.

Keywords: Generative AI, digital divide, digital equity, higher education, AI literacy, educational inequality

Introduction

Generative AI (GenAI) tools such as ChatGPT are often described as having democratizing potential in higher education: capable of providing on-demand tutoring, feedback, and personalized explanation to students who might not otherwise have access to such support. This framing implies that GenAI could narrow gaps between students with differing levels of prior preparation or access to academic resources. Yet the same technology requires reliable devices and connectivity, a baseline level of digital and AI literacy, and instructional environments capable of integrating it productively, all of which are themselves unevenly distributed across and within institutions. This review asks whether the empirical literature on GenAI's effect on learning outcomes supports its democratizing promise, or whether the barriers surrounding its use are more likely to concentrate its benefits among already-advantaged students.

Review Approach

This narrative review draws on recent meta-analyses and systematic reviews of GenAI's effect on higher education learning outcomes, cross-referenced with digital-equity and digital-divide scholarship addressing GenAI specifically, and supplemented by targeted searches for studies published between 2023 and 2026. Search terms combined "generative AI" or "ChatGPT" with "digital divide," "digital equity," "AI literacy," and "educational inequality." Sources were prioritized where they addressed differential access, literacy, or outcomes associated with GenAI use, or where they provided outcome data that could be examined for equity-relevant moderators such as institutional resourcing or prior digital access.

A Three-Level Framework for the GenAI Divide

Digital-equity scholarship applied to GenAI describes unequal benefit as operating across three interrelated levels.

The first-level divide concerns basic access: devices, reliable internet connectivity, and, increasingly, access to higher-performing GenAI tools that may sit behind paywalls or require institutional licensing. The second-level divide concerns skills: the digital and AI literacy required to use GenAI tools effectively, including the ability to write effective prompts, critically evaluate AI outputs, and recognize the limitations and potential biases of generated content. The third-level divide concerns outcomes: even where access and basic skills are present, differences in how GenAI is integrated into instruction can produce unequal gains in learning and, ultimately, unequal digital agency. This framework suggests that closing the first-level access gap, while necessary, is not sufficient to ensure equitable benefit from GenAI in digital learning.

Evidence on Effectiveness and Its Equity-Relevant Boundaries

Meta-analytic evidence indicates that GenAI use is, on average, associated with moderate-to-large positive effects on academic performance. A systematic review and meta-analysis of 69 experimental studies conducted between 2022 and 2024 reported strong positive effects on academic achievement and higher-order thinking, while a meta-analysis of 22 studies reported a moderate overall effect on learning achievement. A further meta-analysis restricted to randomized controlled trials reported a large pooled effect on learning performance with low evidence of publication bias.

These average effects, however, are drawn overwhelmingly from studies conducted at well-resourced institutions with reliable technology access and, in many cases, structured instructional integration of GenAI as part of a designed intervention. Systematic reviews of GenAI attitudes and behaviors in higher education note that most existing research concentrates on students and instructors already engaged with the technology, leaving comparatively little

evidence on outcomes for students facing first- or second-level access barriers. This creates an important asymmetry: the literature is well positioned to describe GenAI's benefit under conditions of adequate access and structured use, but poorly positioned to describe its effect, if any, for students who lack either.

Access, Literacy, and Instructional Capacity as Equity Barriers

Device and connectivity access remains an uneven foundation for GenAI-supported learning. Digital-equity commentary emphasizes that access is not simply binary; students may have intermittent connectivity, older devices poorly suited to running AI-integrated applications, or informal access via shared or public devices, all of which shape how consistently they can incorporate GenAI into their learning routines compared to students with dedicated, reliable access.

AI literacy represents a second, less visible barrier. Reviews of AI literacy conceptualization describe it as encompassing technical competence, critical evaluation skills, and ethical understanding of how GenAI systems function and where they are prone to error or bias. Students without structured instruction in these competencies are more likely to use GenAI in ways that either underuse its potential or lead to problematic reliance, such as accepting outputs uncritically. Because AI literacy instruction is not uniformly embedded in curricula, its distribution tends to track existing patterns of institutional resourcing and prior educational advantage. Instructor capacity constitutes a third barrier operating at the level of instructional design rather than individual student characteristics. Survey-based research on university instructor attitudes toward GenAI finds substantial variation in confidence and beliefs about its educational value, which in turn shapes how consistently instructors model effective, structured GenAI use for students. Where instructors lack training or institutional support for GenAI-integrated pedagogy, students are left to develop their own, often unstructured, practices, a pattern more consequential for students without prior exposure to effective AI use from other sources such as family or informal networks. Together, these barriers imply a compounding effect: students facing access limitations are also more likely to be enrolled in under-resourced institutions with less capacity to provide structured AI literacy instruction and instructor support, meaning that access, skill, and instructional-capacity gaps are unlikely to be independent of one another in practice.

Discussion and Implications

The evidence reviewed here suggests that GenAI's positive average effect on learning outcomes should not be read as evidence of its democratizing potential without qualification. That average effect is disproportionately generated under conditions of adequate access, developed AI literacy, and structured instructional integration, conditions that are themselves unevenly distributed. Left unaddressed, this pattern risks a divergence in which well-resourced institutions use GenAI to further strengthen already strong instruction, while under-resourced institutions either lack meaningful access or adopt GenAI in unstructured ways with more limited or inconsistent benefit. For institutional policy, this implies that equity cannot be treated as a downstream consideration addressed after

GenAI adoption; it needs to be built into adoption planning from the outset, including subsidized or institutionally licensed access to capable tools, embedded AI literacy instruction across curricula rather than as an optional add-on, and dedicated instructor training and support for GenAI-integrated task design. For pedagogy, closing the second- and third-level divides likely matters more than closing the first-level access divide alone; providing a device or a tool without accompanying literacy instruction and structured task design is unlikely to produce the learning benefits documented in the existing meta-analytic literature. For research, there is a clear need for studies that explicitly sample across institutional resourcing levels and report outcomes disaggregated by access and literacy conditions, since the current evidence base is heavily weighted toward already-advantaged learning environments and may overstate the benefit typically available to less-resourced students.

Conclusion

Generative AI's promise as a democratizing force in higher education rests on an assumption that access, literacy, and instructional capacity are either already in place or will follow naturally from technology adoption. The literature reviewed here suggests this assumption is not well supported: the conditions under which GenAI reliably improves learning outcomes are themselves unevenly distributed, and the evidence base disproportionately reflects outcomes for students already positioned to benefit. Realizing GenAI's potential to narrow rather than widen educational gaps will require deliberate, equity-focused investment in access, AI literacy, and instructor capacity, rather than treating equitable benefit as an automatic byproduct of technology availability.

References

1. Ahmed F. The digital divide and AI in education: Addressing equity and accessibility. *AI EDIFY Journal*, 2024;1(2):12–23.
2. Almatrafi O, Johri A, Lee H. A systematic review of AI literacy conceptualization, constructs, and implementation and assessment efforts (2019–2023). *Computers and Education Open*, 2024;6:100173.
3. BERA. Digital equity in the age of generative AI: Bridging the divide in educational technology. British Educational Research Association blog, 2024.
4. Cabellos B, Sánchez-Reina R, Fernández-Sanjurjo M. University teachers' beliefs about the use of generative artificial intelligence in education. *Frontiers in Psychology*, 2024, 15.
5. Chiu TK, Ahmad Z, Ismailov M, Sanusi IT. What are artificial intelligence literacy and competency? *Computers and Education Open*, 2024;6:100171.
6. Deng R, et al. Does ChatGPT enhance student learning? A systematic review and meta-analysis of experimental studies. *Computers & Education*, 2025.
7. Liu Y, et al. The impact of ChatGPT on students' academic achievement: A meta-analysis. *Journal of Computer Assisted Learning*, 2025.

8. Regulating the digital divide: Policy interventions for equitable access to AI-enabled education. *International Journal of Education, Management and Technology*, 2025, 11(4).
9. Wu F, Dang Y, Li M. A systematic review of responses, attitudes, and utilization behaviors on generative AI for teaching and learning in higher education. *Behavioral Sciences*, 2025;15(4):467. <https://doi.org/10.3390/bs15040467>
10. Zilibotti F, et al. [As cited in Regulating the digital divide, 2025, on second- and third-level digital divides], 2022.