

AI in the Classroom, Inequality in the Background: The Dual Effects of Generative AI on Educational Equity Across Linguistic and Cultural Lines in Higher Education—A Systematic Review and Critical Synthesis

Yuezhou Liu

Moray House School of Education and Sport, Edinburgh, UK
yuezhou_liu@163.com

Abstract. Amid the rapid proliferation of GenAI in higher education, claims about its democratizing potential have outpaced empirical scrutiny of its equity consequences. This study critically reviews generative artificial intelligence (GenAI) and its implications for equity in higher education. Drawing on literature published between 2023 and 2026, it develops a three-dimensional analytical framework encompassing opportunity, process, and outcome equity to examine the dual role of GenAI as both an inclusion-enabling and inequality-reproducing force. The findings suggest that GenAI enhances educational inclusion through multilingual support, personalized learning, and expanded participation opportunities. However, these benefits are accompanied by structural risks, including algorithmic bias, unequal digital literacy, and systematic misclassification in AI-driven assessment systems. These dynamics collectively reinforce existing educational inequalities rather than eliminate them. The study further demonstrates that current governance frameworks remain conceptually fragmented, lacking operational specificity, contextual adaptability, and empirical validation. As a result, the equity outcomes of GenAI remain highly uneven across institutional and regional contexts. Localized governance, informed by sustained empirical tracking, is recommended over generic institutional frameworks for mitigating the uneven impacts of GenAI.

Keywords: generative artificial intelligence, higher education equity, algorithmic bias, digital divide, educational governance framework

1. Introduction

Since the release of ChatGPT in late 2022, generative artificial intelligence (GenAI) has rapidly entered the teaching, learning, and assessment practices of higher education. This transformation, however, is unfolding within a global higher education system characterized by persistent and multidimensional inequalities, including the dominance of English as an academic lingua franca, unequal digital infrastructure, and culturally embedded assumptions within evaluation systems.

Against this backdrop, existing research has unfolded along three main lines. The first positions GenAI as an educational equalizer, emphasizing its capacity to lower writing anxiety for non-native English speakers and bridge language gaps. The second reveals the same technology functioning as an invisible exclusionary device: AI detection tools systematically misclassify non-native writing, large language models carry cultural and linguistic biases, and uneven digital literacy leads to new forms of usage stratification, all of which quietly reproduce existing hierarchies. From this tension a third strand has emerged, exploring structural interventions at institutional, policy, and instructional design levels that seek to move GenAI from possible equity to actual equity. However, these three lines of work expose a dual role that has yet to be organized in an integrated way.

To address this, the study approaches educational equity along three interconnected dimensions: opportunity, process, and outcome. The analysis is organized around four sets of issues: the multilingual and accessibility-related affordances of GenAI; algorithmic bias and structural inequities in AI-mediated evaluation; the reconfiguration of the digital divide in GenAI contexts; and emerging governance frameworks and their limitations. Drawing on a systematic review of research between 2023 and 2026, the study adopts a critical synthesis to identify recurring mechanisms through which GenAI shapes educational equity across these four domains. The study aims to clarify how GenAI simultaneously enables and constrains educational equity, and to identify directions for more structurally responsive policy and practice.

2. GenAI as a potential facilitator of equity in higher education

2.1. Multilingual assistance and cross-cultural learning support

The strong performance of large language models (LLMs) in English has been extensively documented, while their applicability in non-English linguistic contexts has attracted increasing scholarly attention. Empirical studies suggest that GenAI can provide timely and scalable academic support for learners from diverse linguistic backgrounds, including translation, grammar correction, and content clarification. These functions partially mitigate language-related barriers that are deeply embedded in traditional higher education settings.

Albedah et al., in a review of 161 studies, highlight that models such as Aya and LLaMA demonstrate emerging multilingual capabilities, offering baseline academic support for users of low-resource languages [1]. Similarly, Gamboa et al. show that although LLMs may exhibit systematic disparities when processing sociolinguistically diverse inputs, language-specific fine-tuned models can still provide meaningful educational support, including translation, paraphrasing, and basic question answering [2]. These findings suggest that, compared to learning environments without AI assistance, GenAI represents a substantive shift from absence to partial support, even if its performance remains uneven across contexts.

From an equity perspective, these capabilities are particularly significant in higher education systems where English functions as the dominant academic language. Non-native English-speaking students often face structural barriers in accessing knowledge and participating in scholarly discourse. GenAI-mediated multilingual assistance can partially lower these barriers by facilitating access to content and enabling broader academic participation.

Stiegelbauer et al. compare Google Translate, DeepL, and ChatGPT-4 in higher education and healthcare contexts, finding that ChatGPT-4 demonstrates superior fluency and semantic accuracy in cross-linguistic communication, thereby supporting its applicability in multilingual academic environments [3]. However, important limitations remain. Low-resource languages, such as Malay and Indonesian, are often underrepresented in training corpora and may therefore receive lower-

quality outputs due to data scarcity [1]. This results in uneven learning support quality across linguistic groups, which itself constitutes a form of process-level inequity.

Furthermore, evidence suggests that GenAI systems may struggle to preserve pragmatic nuance, emotional tone, and culturally embedded meaning in translation tasks, particularly in high-stakes contexts such as clinical or academic documentation [3]. Over-reliance on machine translation may therefore risk flattening the cultural depth of language, underscoring the necessity of human oversight and cross-cultural mediation in educational applications.

2.2. Personalized learning and universal access to educational resources

One of the key affordances of GenAI is its capacity to deliver adaptive, individualized learning support tailored to learners' diverse needs. For students in resource-constrained contexts, learners with difficulties, and non-traditional student populations, such personalization partially compensates for the limitations of conventional classroom instruction, which often struggles to accommodate heterogeneity at scale.

Okonkwo and Ade-Ibijola demonstrate that GenAI systems can generate learning materials calibrated to different knowledge levels and learning trajectories, provide interactive one-on-one tutoring, and dynamically adjust instructional content based on learner feedback [4]. Similarly, Kohnke et al. argue that tools such as ChatGPT function as "always-available virtual tutors," offering differentiated scaffolding aligned with learners' linguistic proficiency and cognitive profiles. This capability is particularly consequential in contexts characterized by teacher shortages and unequal access to qualified instructional support [5].

Beyond individualized instruction, GenAI also helps redistribute the redistribution of educational resources across geographic boundaries. The DIKSHA platform in India represents a notable case of large-scale digital public infrastructure, reaching 1.48 million schools across 35 states and union territories and providing educational content in 36 languages. Recent integration of AI-enabled learning assistants into this ecosystem has further extended access to instructional resources for students in remote regions, enabling them to engage with learning materials previously concentrated in urban, high-resource institutions. This development illustrates how GenAI can function as a technological mediator in reducing spatial inequalities in educational provision.

2.3. Expanding opportunities for participation in higher education

GenAI also expands opportunities for participation in higher education among marginalized and underrepresented groups. For students with disabilities, AI-enabled tools such as speech recognition, text-to-speech systems, and real-time captioning significantly reduce barriers to learning, thereby enabling more equitable access to classroom instruction and academic content. Empirical studies suggest that such technologies can improve both engagement and academic outcomes for disabled learners [4].

For example, Microsoft's Seeing AI application has been widely implemented in higher education settings to assist visually impaired students by converting printed materials, lecture notes, and textbooks into synthesized speech while also supporting image-based interpretation of diagrams and visual content. Similarly, Google's Live Transcribe provides real-time captioning that enables deaf and hard-of-hearing students to access spoken lectures synchronously, thereby enhancing classroom inclusivity.

In addition, GenAI supports non-traditional learners, including working adults and lifelong learners, by offering flexible, on-demand learning pathways. This flexibility reduces both temporal

and opportunity costs associated with higher education participation, thereby enabling broader engagement across different life stages and socioeconomic conditions. Collectively, these developments indicate that GenAI may function as an enabling infrastructure for expanding participation, although its effectiveness remains contingent on equitable access to underlying technological resources.

3. Multiple dimensions of inequality exacerbated by GenAI in higher education

3.1. Algorithmic bias and process inequity stemming from linguistic disparities

The training datasets underlying large language models (LLMs) are predominantly composed of Western-centric linguistic and cultural corpora. As a result, these systems embed implicit normative assumptions that tend to reproduce dominant epistemologies while marginalizing non-Western knowledge systems. Nordt distinguishes between two analytically separable but interrelated forms of bias in LLMs: linguistic bias, embedded in language structures and model architectures, and cultural bias, which emerges through culturally grounded scripts and normative expectations activated during prompting [6].

This dual-layered bias structure suggests that LLM outputs are not merely technical artifacts but culturally situated productions of knowledge. Consequently, when applied in educational contexts, these systems may systematically privilege Western rhetorical norms while rendering non-Western linguistic and epistemic expressions less visible or less legitimate. This has direct implications for students from diverse linguistic backgrounds, whose interaction with GenAI systems may be subtly shaped toward dominant linguistic standards.

Recent scholarship further challenges the "tool-neutrality" assumption that underpins much of the GenAI literature. Ngo et al. argue that GenAI in education is embedded within broader sociotechnical infrastructures, and may exacerbate inequality through four interconnected pathways: infrastructural asymmetries, disparities in digital literacy, algorithmic bias, and English-dominant linguistic norms [7]. In particular, the privileging of English-native discourse conventions in LLM outputs creates a normative pressure on non-native English users to conform to standardized academic English, thereby potentially suppressing linguistic diversity and reinforcing asymmetrical power relations in academic communication.

3.2. Linguistic discrimination and identity inference in academic publishing

Linguistic inequality in academic publishing represents a more deeply embedded and less visible form of stratification than classroom-level language disparities. Using a large-scale dataset of approximately 80,000 peer review reports from the International Conference on Learning Representations (ICLR), Lepp and Smith demonstrate that authors affiliated with institutions in low English proficiency contexts face systematic disadvantages in peer evaluation processes [8].

Importantly, the diffusion of GenAI has not eliminated such biases; rather, it has reconfigured their expression. Reviewers increasingly rely on indirect linguistic cues, such as perceived "ChatGPT-like writing style" or other markers, to infer authors' linguistic background and, in some cases, to adjust evaluative judgments accordingly. This process reflects a shift from explicit language-based discrimination to more implicit forms of identity inference. As a result, authors report heightened self-monitoring and repeated revision of linguistic features that might signal non-native status, indicating the emergence of anticipatory adaptation behaviors within academic writing practices [8]. Rather than mitigating inequality, GenAI thus appears to have displaced the locus of

discrimination from surface-level linguistic forms to deeper inferential judgments about author identity.

From a theoretical perspective, this finding challenges the assumption that LLMs function as neutralizing agents of linguistic inequality. Instead, it suggests that the structural roots of academic stratification lie not in language per se, but in evaluative regimes that equate "standardized English" with epistemic quality and scientific rigor. Such ideological coupling continues to shape gatekeeping practices in scholarly communication, particularly affecting early-career researchers and scholars from non-dominant linguistic backgrounds.

3.3. The GenAI divide: from digital access to the compound stratification of literacy and resources

The concept of the "digital divide" has evolved significantly in the GenAI era. Whereas earlier formulations primarily emphasized disparities in physical access to digital infrastructure, recent scholarship highlights a more complex, multidimensional form of stratification. Matjie et al. identify three interrelated dimensions of the GenAI-driven divide in education: linguistic adaptability of AI systems, cultural misalignment between developers and users, and algorithmic bias embedded in system design [9]. This suggests that inequality is increasingly produced not only through differential access, but also through the structural properties of the technologies themselves.

Empirical evidence further demonstrates how these structural inequalities manifest at the user level. Beckman et al., using a mixed-methods design with 194 university students, identify three distinct user profiles: novice, cautious, and enthusiastic users, revealing systematic variation in engagement with GenAI [10]. Students with lower AI literacy tend to resist GenAI and are more likely to avoid its use, whereas those with higher literacy levels integrate it more effectively into learning processes. These findings indicate that GenAI may function as a stratifying mechanism that amplifies pre-existing disparities in digital competence, rather than mitigating them.

At the institutional level, Wong, Aristidou, and Scheuermann demonstrate that significant variation exists in GenAI readiness across universities, particularly within Asian higher education systems. These disparities are shaped by national policy environments and institutional resource endowments, leading to uneven implementation of GenAI-related curricula and competencies [11]. The authors identify four core competencies required in the GenAI era: AI ethics, AI literacy, human-machine collaboration, and uniquely human skills, yet access to these competencies remains unevenly distributed across institutional contexts, thereby reinforcing systemic inequality in educational outcomes.

3.4. Further widening of urban-rural and regional educational resource gaps

Beyond individual and institutional stratification, GenAI also contributes to the spatial reconfiguration of educational inequality. Research on rural and regional higher education in Australia highlights a dual effect: while GenAI enhances learning flexibility and access to educational support for geographically dispersed students, it simultaneously risks intensifying existing disparities in digital infrastructure, cultural representation, and institutional capacity [12]. Infrastructural constraints in less-developed regions further compound this spatial inequality. Matjie et al. emphasize that in contexts such as parts of Africa and Eastern Europe, limited network connectivity and insufficient computational resources constitute structural barriers to meaningful GenAI adoption [9]. As a result, nominally equal access to GenAI technologies does not translate

into equivalent learning conditions, producing a gap between formal availability and effective usability.

At the systemic level, these disparities are reinforced by uneven resource distribution across higher education institutions. Universities in urban and economically developed regions typically possess greater financial capacity, more advanced technological infrastructure, and higher concentrations of digitally skilled faculty, enabling more rapid and effective integration of GenAI into teaching and learning. By contrast, institutions in rural and under-resourced areas face persistent constraints in funding, infrastructure, and human capital, limiting their ability to fully leverage GenAI. Over time, these cumulative disparities risk entrenching and reproducing spatially structured inequalities in higher education systems.

4. Bias in AI detection tools and the institutional predicament of non-native English speakers

4.1. Systematic misclassification: quantitative evidence and causal analysis

The most immediate equity concern associated with GenAI in higher education does not stem from the use of AI writing tools themselves, but rather from the institutionalization of AI detection systems as arbiters of academic integrity. In many higher education contexts, these systems are increasingly used to determine whether students have engaged in AI-assisted writing, thereby introducing significant risks of systematic misclassification, particularly for non-native English speakers.

A systematic review by Yvan et al. synthesizing 27 studies provides the most comprehensive empirical evidence to date [13]. Using Kane's validity framework, the study evaluates whether AI-based assessment tools exhibit construct-irrelevant variance, whereby linguistic proficiency—an unrelated construct to the intended assessment domain—systematically distorts evaluation outcomes. The findings are striking: experimental studies consistently show that AI detection systems misclassify non-native English writing as AI-generated at rates ranging from 50.2% to 61.3%, whereas false positive rates for native English writing remain below 5%. In addition, automated scoring systems exhibit a systematic penalty of 0.5 to 1.2 standard deviations against non-native English writers.

This pattern is further reinforced by Zarrouk's analysis of multiple AI detection systems applied to authentic academic manuscripts [14]. Even texts that were entirely human-written and only minimally edited were frequently flagged as AI-generated. Notably, highly formalized and structurally rigid academic genres, such as quantitative research articles, were more likely to be misclassified. These findings suggest that linguistic features commonly associated with non-native academic writing, including formulaic expression, reduced lexical variation, and strict adherence to formal structures, are systematically interpreted as indicators of machine-generated text.

Taken together, these results indicate that AI detection systems do not merely reflect technical limitations, but instead operationalize a form of algorithmic bias in which linguistic conformity is equated with authenticity. This produces a structural paradox: the more closely non-native speakers adhere to academic conventions, the more likely they are to be penalized by automated detection systems.

4.2. Institutional consequences

The institutional consequences of such misclassification extend far beyond technical error rates. For students, false accusations of AI-assisted misconduct may result in formal academic sanctions,

psychological distress, and a broader erosion of trust in assessment systems. Survey evidence involving 2,847 students indicates that 79% of non-native English speakers express substantial concern regarding the fairness of AI detection systems, while 63% identify this issue as a primary academic anxiety [13].

Importantly, GenAI introduces a dual role in this context. While it can be regarded as an assistive tool for overcoming linguistic barriers, it can also trigger suspicion and disciplinary action when its use is inferred rather than explicitly detected. This duality places non-native English speakers in a structurally vulnerable position, where their reliance on linguistic support tools may paradoxically increase their exposure to institutional scrutiny.

Further evidence from Eaton highlights that AI detection systems disproportionately affect neurodivergent students, particularly those with dyslexia, whose writing patterns may differ systematically from standardized academic norms [15]. Similarly, Liang et al. demonstrate that over 60% of TOEFL essays written by non-native English speakers are incorrectly flagged as AI-generated. In contrast, essays written by native English-speaking students show negligible false positive rates [16].

Collectively, these findings suggest that bias in AI detection systems is not an isolated technological flaw but a structural mechanism that translates linguistic difference into perceived academic misconduct. This constitutes a substantive threat to procedural justice in higher education assessment systems.

5. Existing governance frameworks, theoretical models, and critical assessment

5.1. Predominant ethical frameworks and theoretical models

In response to the growing challenges posed by GenAI in higher education, a range of governance frameworks and ethical models have been proposed to guide equitable integration. These frameworks generally attempt to reconcile technological adoption with principles of fairness, inclusivity, and educational justice.

The Equitable AI in Language Education Model (EAIEM) proposed by Delmadorou conceptualizes GenAI as a culturally embedded but non-neutral educational tool. Grounded in sociocultural theory, second language acquisition, and constructivist learning theory, the model emphasizes four core principles: inclusivity, transparency, human–AI collaboration, and participatory design. Its central argument is that GenAI integration in language education must actively support linguistic diversity while fostering intercultural competence and advancing educational equity [17].

Similarly, the EPI-AI framework proposed by Ngo et al. adopts a systemic perspective, conceptualizing equitable AI implementation across five interdependent dimensions: access and infrastructure, AI literacy, pedagogy, governance, and sociocultural responsiveness. The framework highlights that equitable GenAI adoption is not a technological deployment issue, but rather a relational and context-dependent process embedded within broader educational systems [7].

From an implementation perspective, Matjie et al. emphasize the importance of participatory design approaches that incorporate local languages, cultural contexts, and community stakeholders into AI development processes. They further highlight the need for capacity building among educators to bridge both technological and cultural gaps [9].

5.2. A critical reflection on existing governance approaches

Despite their conceptual contributions, existing governance frameworks exhibit three fundamental limitations when evaluated from an implementation-oriented perspective.

First, most frameworks remain highly abstract and lack operational specificity. While principles such as "participatory design" and "transparency" are widely endorsed, they are rarely translated into concrete institutional mechanisms. For instance, existing models do not clearly define how different stakeholders, students, educators, developers, and policymakers, should be organized, nor do they specify decision-making structures or accountability mechanisms.

Second, these frameworks are predominantly developed within Western or high-resource educational contexts, limiting their transferability to low-resource settings. In particular, recommendations involving multilingual integration and localized AI development require substantial financial, technical, and human capital, which may not be readily available in developing regions.

Third, there is a lack of robust longitudinal evidence evaluating the long-term effectiveness of these governance interventions. Most existing studies rely on conceptual analysis or small-scale pilots, with limited empirical validation of whether such approaches meaningfully reduce educational inequality over time. As a result, the actual equity impacts of "inclusive AI" and participatory governance models remain largely speculative.

5.3. Future research directions

Building on the analysis above, future research on GenAI and educational equity should advance in four key directions.

First, comparative cross-cultural and cross-regional studies are needed to identify context-sensitive governance models that reflect the diversity of institutional and sociotechnical environments. Second, there is a need for rigorous empirical evaluation of existing governance interventions, particularly through longitudinal designs that capture both intended and unintended consequences over time.

Third, interdisciplinary collaboration should be strengthened by integrating insights from education, computer science, sociology, linguistics, and ethics in order to develop more comprehensive theoretical and analytical frameworks. Finally, future research should prioritize the lived experiences of marginalized groups, ensuring that GenAI systems are evaluated not only in terms of technical performance but also in relation to their distributive and procedural justice implications.

6. Conclusion

This study has provided a critical synthesis of the dual role of generative artificial intelligence (GenAI) in shaping educational equity in higher education. The analysis demonstrates that GenAI holds significant potential to enhance inclusion through multilingual support, adaptive learning systems, and expanded access to educational participation. However, these affordances coexist with structural risks, including algorithmic bias, unequal digital literacy, and systemic misclassification in AI-based assessment systems.

Importantly, these outcomes are not incidental or transitional effects of emerging technology. Rather, they reflect deeper sociotechnical structures in which GenAI systems are embedded, including linguistic hierarchies, institutional evaluation norms, and uneven distributions of digital

infrastructure. As such, GenAI should not be conceptualized as a neutral tool, but as a constitutive component of contemporary educational governance regimes.

While existing governance frameworks offer important normative and conceptual guidance, they remain limited in their operational specificity and empirical grounding. Addressing these limitations requires a shift toward context-sensitive governance models, interdisciplinary collaboration, and longitudinal research capable of capturing the evolving and distributed impacts of GenAI across diverse educational settings.

Ultimately, the central challenge is not simply how to integrate GenAI into higher education, but how to ensure that its integration does not reproduce or deepen existing inequities. This requires rethinking the relationship between technology, pedagogy, and justice in ways that move beyond technological optimism toward structurally informed educational governance.

References

- [1] Albedah, F. (2025). Artificial intelligence in language education: A systematic review of multilingual applications, large language models, and emerging challenges. *Language Teaching Research Quarterly*, 49, 247-268.
- [2] Gamboa, L. C. L., Feng, Y., & Lee, M. (2025, November). Social bias in multilingual language models: A survey. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing* (pp. 27845-27868).
- [3] Stiegelbauer, L. R. (2025). AI and multilingual communication: Accuracy, trust and linguistic equity in higher education and healthcare. *Incursions into the Imaginary*, 16(1), 339–358.
- [4] Okonkwo, C., & Ade-Ibijola, A. (2021). ChatGPT and the future of education. *Journal of Artificial Intelligence and Technology*, 2(3), 115-123.
- [5] Kohnke, L., Moorhouse, B. L., & Zou, D. (2023). ChatGPT for language teaching and learning. *RELC Journal*, 54(2), 537-550.
- [6] Nordt, P. (2025, June). Disentangling linguistic and cultural biases in large language models: A comparative study. In *Proceedings of the MEI: CogSci Conference (Vol. 19, No. 1)*.
- [7] Ngo, C. L. et al. (2025). Bridging or breaking? A systematic review of how generative AI shapes equity in foreign language education. *International Journal of Applied Linguistics*.
- [8] Lepp, H., & Smith, D. S. (2025, June). "You Cannot Sound Like GPT": Signs of language discrimination and resistance in computer science publishing. In *Proceedings of the 2025 ACM conference on fairness, accountability, and transparency* (pp. 3162-3181).
- [9] Matjie, M. A., Nethavhani, A., & Matlakala, M. (2026). AI and the digital divide in education. *Frontiers in Computer Science*, 8, 1759027.
- [10] Beckman, K., Apps, T., Howard, S. K., Rogerson, C., Rogerson, A., & Tondeur, J. (2025). The GenAI divide among university students: A call for action. *Internet High. Educ.*, 67, 101036. <https://doi.org/10.1016/j.iheduc.2025.101036>.
- [11] Wong, W., Aristidou, A., & Scheuermann, K. (2025). The future of learning or the future of dividing? Exploring the impact of generative artificial intelligence on higher education. *Data & Policy*, 7, e44.
- [12] Kornhaber, R., Bragg, S., Matheson, A., & Irwin, P. (2026). Qualitative inquiry into generative artificial intelligence impact on rural and regional higher education. *Australian and International Journal of Rural Education*, 1–23.
- [13] Yvan, N. J. C., Charles, K., & Lucy, A. C. (2026). The impact of generative artificial intelligence tools on assessment equity for non-native English-speaking medical students: A systematic review. *BMC Medical Education*.
- [14] Zarrouk, R. (2025). False alarms: How linguistic features of non-native academic writing trigger false positives in AI detection tools. In *Proceedings of the Middle East University International Conference on AI in Education*.
- [15] Eaton, S. E. (2026). Neurodiversity and academic integrity: Toward epistemic plurality in a postplagiarism era. *Teaching in Higher Education*, 31(4), 685-704.
- [16] Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. *Patterns*, 4(7).
- [17] Delmadorou, E. I. (2026). Generative artificial intelligence in language education: Ethical dimensions and the equitable AI in language education model. *Research Papers in Language Teaching & Learning*, 16(1), 96.