

Volume 18 (2026), pp. 75-85
*American Journal of STEM Education:
Issues and Perspectives*
Star Scholars Press
<https://doi.org/10.32674/rnayrn64>

Cultural Diversity and Artificial Intelligence: Toward a Pluralistic Framework

Muddassir Siddiqi
College of DuPage, USA

ABSTRACT

AI tools are reshaping classrooms, yet many default to Western norms, which can lead to bias, language exclusion, and cultural erasure. This paper presents a workable framework for STEM educators and leaders, which involves blending ethical pluralism with inclusive, participatory design, partnering with language communities, embedding culturally grounded governance, investing in capacity building, and conducting ongoing cultural audits. Practically, this means co-creating curricula and datasets with local stakeholders, requiring multilingual and accessible AI resources, aligning policy with community values, and auditing tools before and after adoption. The result is AI use that enhances belonging, accuracy, and fairness for diverse learners while strengthening institutional accountability.

Keywords: Cultural diversity; AI ethics; Ethical pluralism; Participatory design; Language preservation; Algorithmic bias; STEM education

INTRODUCTION

The global proliferation of artificial intelligence (AI) technologies, especially large language models (LLMs), has far-reaching implications for cultural diversity, defined here as the presence and interaction of various traditions, belief systems, languages, social identities, and ways of life across human societies (Lin, 2020). While Western cultural frameworks predominantly develop these technologies, their worldwide adoption has complex social, political, and cultural implications. As AI

increasingly influences many important aspects of human life, such as education, healthcare, governance, and communication, it involves the cultural assumptions of its builders. It sometimes conflicts with, or misrepresents, the values and lived experiences of non-Western societies (Youvan, 2024). The ethical frameworks surrounding AI, which also largely mirror Global North priorities, raise concerns about their universal applicability and the disadvantages they may impose on other religious, cultural, and philosophical traditions (Roche et al., 2023). Since AI permeates the lives of millions of people, addressing these concerns is crucial to ensuring equitable outcomes of AI technologies for humanity.

The builders of AI and the datasets they use to train AI models influence outcomes. If developers of AI and their datasets fail to consider cultural pluralism deliberately, they risk reproducing ethical imperialism and digital colonization. These risks are particularly severe for historically marginalized groups, which already face unequal access to digital infrastructures, such as people with disabilities, members of racial and ethnic minorities, LGBTQ+ people, and low-income or rural communities (Roche et al., 2022). Large language models (LLMs) often train from data that reflect historical prejudices, systemic biases, and the disproportionate influence of dominant cultural narratives. When these patterns are automated and scaled through algorithmic processes, they risk reinforcing existing inequities rather than disrupting them (O'Neil, 2016). As a result, individuals and communities who deserve greater resources and recognition may continue to be marginalized, and the historical processes of suppression and erasure may be further entrenched through algorithmic bias. Therefore, ensuring that the development of AI is equitable and inclusive of diverse cultures constitutes both an ethical imperative and a necessary condition for advancing social justice in technological innovation.

At its core, the scholarly debate on cultural diversity and artificial intelligence centers around the dual nature of AI. On the one hand, AI poses risks of reinforcing global inequality and cultural homogenization; however, it also offers transformative opportunities for preserving, revitalizing, and promoting cultural diversity when an inclusive framework and practices are prioritized in its design and deployment (Koc, 2025). While embracing a plurality of ethical viewpoints can help counteract these risks, the continued predominance of Western cultural presumptions in AI development threatens to solidify existing inequalities and cultural bias. As a result, it is imperative to evaluate and ensure that AI systems value and promote cultural diversity. To this end, this paper proposes six strategies for culturally inclusive AI: integrating ethical pluralism, inclusive design and development, participatory design for linguistic and cultural equity, culturally grounded governance, education and capacity building, and cultural auditing and accountability.

Note on Terms

Cultural diversity: the coexistence and interaction of multiple artistic traditions, languages, and identities.

Cultural inclusivity: the intentional design of systems to affirm and represent these traditions.

Ethical imperialism: the imposition of one cultural or philosophical tradition's values onto global systems, marginalizing others.

Digital colonization: the embedding of Western assumptions into digital infrastructures in ways that perpetuate inequity.

Global Inequities in AI

AI does not exist in isolation; it is a product of its development environment. Today, its integration into governance, education, and media often amplifies Western norms at the expense of others (Bender et al., 2021). According to Roche et al. (2023), “the ethical principles and values being proposed mirror and replicate the context in which they are developed” (p. 1098). This results in a global distribution of AI benefits that is far from equitable, with wealthier regions reaping most of the rewards. Under-resourced communities, which are often the most vulnerable, bear a disproportionate share of the risk (3DBear, 2023).

These inequities are tangible. In education, algorithmic personalization may narrow rather than expand cultural horizons, privileging efficiency over critical engagement (Neuroscience News, 2023). In healthcare, diagnostic models trained on English-language, Western-centric datasets often produce skewed or harmful results for non-Western populations (Hutchinson et al., 2020). UNESCO (2023) estimates that mainstream AI tools represent less than 1% of the world's language, which is an exclusion that scholars call ethical imperialism (Floridi et al., 2018). Even UNESCO's own guidelines, while influential, risk universalizing ethics in ways that erase diversity (Jobin et al., 2019; Green, 2021).

Asymmetries persist in the global information economy because the dominance of large technology companies, which are highly concentrated in the Global North, determines which languages, values, and cultural perspectives are prioritized in AI development. Encouraging diversity and intersectionality among development teams helps AI developers identify and address blind spots, thereby filling gaps in representation and creating more inclusive technologies. Research has demonstrated that diversifying teams and deliberately addressing biased language and feedback processes not only enhances equity in hiring and promotion but also leads to better outcomes and a more equitable AI landscape for end-users of all backgrounds. By centering diversity, equity, and intersectionality, the industry can move toward an AI environment that

both reflects and serves the breadth of humanity rather than reinforcing the silences and erasures of the past.

Cultural Risks of AI: Bias, Loss, and Digital Colonization

The risks of AI to cultural diversity are multifaceted. Generative systems often replicate stereotypes and steer users toward prevailing linguistic and cultural norms (Azgin & Kıralp, 2024). In doing so, they usually reinforce biases related to gender, profession, or ethnicity (Noble, 2018).

The erosion of cultural narratives represents a deeper threat. AI-generated content diminishes the richness and nuance of traditions by reducing them to generalizations (Wakunuma & Eke, 2024). With fewer than one percent of global languages digitally supported (UNESCO, 2023), low-resource communities are at risk of digital extinction. These inequities exemplify digital colonization, whereby algorithmic infrastructures marginalize and silence minority voices while consolidating the power of dominant groups (Mohamed et al., 2020).

Different cultures interact with AI in distinct ways, which can introduce new forms of inequity. Cultural tightness also shapes interactions with AI. High-tightness cultures that emphasize strict norms often value anthropomorphic AI, whereas loose cultures prioritize functionality (Sui et al., 2024). Without a culturally responsive design, AI risks alienating entire populations. The concentration of decision-making power in a few corporate and geopolitical centers amplifies these risks. Without intersectional representation and structural accountability, AI will continue to reflect dominant worldviews while marginalizing others.

Cultural Frameworks in AI Ethics

Cultural diversity shapes not only daily AI interactions but also the broader ethical questions we ask of these technologies. This complexity came into sharp focus with the MIT Media Lab's Moral Machine experiment. The experiment involved the collection of millions of responses from around the world. The results revealed how collectivist societies, valuing group welfare, often reach strikingly different conclusions than individualist cultures do, which places personal rights at the forefront (Awad et al., 2018). These global findings remind us that the cultural context is not just a backdrop; it actively informs how people judge the morality of AI in real-life scenarios.

However, as powerful as the Moral Machine experiment has been, its influence is not without controversy. Critics contend that by boiling ethical dilemmas down into simple, trolley-problem-style choices and by sampling disproportionately from the Global North, the experiment risks undermining the rich diversity of global moral perspectives (Greene et al., 2019; Noothigattu et al., 2018). Acknowledging these limitations, scholars

argue that there is no one-size-fits-all set of principles to guide AI development globally.

Despite this awareness, most current systems still echo Western technocratic ideals—championing efficiency, speed, and precision above all. In the process, they often disregard values such as empathy, stewardship, or community that are deeply rooted in other traditions (Dignum et al., 2022; Roche et al., 2023). Suppose AI reflects the world's true ethical spectrum. In that case, it will need frameworks as plural as its users, such as Ubuntu, Confucianism, Indigenous knowledge, and beyond, to expand the potential definition of AI ethics (Cambridge, 2023; UNESCO, 2021).

Opportunities for Cultural Preservation and Empowerment

While much of the global discourse on AI focuses on risks and limitations, recognizing its potential for cultural affirmation and empowerment is equally essential. When carefully designed in partnership with local communities, AI can help revitalize traditions, preserve languages, and advance equity. For example, generative models aid in the reconstruction of ancient texts, restore fragile artifacts, and support the survival of endangered scripts such as Nüshu. Nüshu was once a crucial way for women to express themselves in China (Dartmouth, 2023; UNESCO, 2023).

AI benefits also extend to the marketplace, where algorithmic insights are increasingly used to identify and address gaps in diversity initiatives (Mariyono & Hidayatullah, 2024; Gupta & Kaul, 2024). However, these opportunities must be balanced with constant caution: algorithms designed to promote inclusivity can inadvertently reinforce stereotypes if trained on biased data. Studies have shown that practices such as targeted advertising risk reducing complex identities to simplistic caricatures, reinforcing the very biases that AI is intended to disrupt (Noble, 2018).

In education, culturally responsive AI tutors help reduce bias and increase engagement among marginalized student groups, whereas in healthcare, the adoption of diverse datasets improves diagnostic accuracy and outcomes for underserved populations (Birhane et al., 2022; MDPI, 2023; Samuel et al., 2023). These real-life examples demonstrate how AI, when guided by the values of equity, foresight, and active community engagement, can promote diversity and make meaningful contributions to cultural empowerment.

Toward Ethical Frameworks and Development Strategies

Beyond critique, inclusive AI requires reorienting design priorities. Compassion, empathy, relational accountability, and patience must be balanced with efficiency and precision (Roche et al., 2023). Cultural frameworks offer practical paths.

Table 1. Cultural Frameworks and Their Influence on AI Design

Cultural framework	Key principles	Impact on AI design
Ubuntu	Community, interdependence	Optimizes outcomes for collective well-being
Confucianism	Harmony, moral obligations	Promotes context-sensitive and socially stable decisions
Indigenous thought	Relationality, stewardship	Embeds sustainability and intergenerational accountability

These framework challenges narrow Western paradigms, expanding the ethical imagination of AI toward pluralism.

SIX STRATEGIES FOR INCLUSIVE AI DEVELOPMENT

1. Ethical Pluralism

Ethical pluralism challenges the notion that a single framework can adequately address the complexities of AI systems deployed across culturally diverse societies (Youvan, 2024). Rather than relying on only Western philosophies, such as Kantian or utilitarian philosophies, and non-Western philosophies, including Ubuntu, which emphasizes communal harmony, Confucianism has centered on relational ethics and various indigenous frameworks in the normative models guiding AI decision-making (Cambridge, 2023; UNESCO, 2021). This approach affirms that ethical values are inherently context dependent and that genuine inclusivity requires moving beyond surface-level appropriation of non-Western principles.

A key challenge with putting ethical pluralism in practice in AI development is avoiding ethical tokenism, which is the superficial adoption of values from diverse cultures without genuine engagement and choosing values out of context. To address this challenge, ethical guidelines and practices should be cocreated through ongoing dialog, interdisciplinary collaboration, and participatory governance, ensuring that diverse stakeholders are meaningfully engaged at every stage of the AI lifecycle (Birhane, 2021; Costanza-Chock, 2020). The African Union's Continental AI Strategy is a good example of this approach. It exemplifies how non-Western values can be intentionally built into large-scale AI

governance, paving the way for more equitable, accountable, and culturally aware technological development (African Union, 2022; UNESCO, 2021).

2. Inclusive Design and Development

Diversity within design teams is a critical driver for reducing algorithmic bias and building equitable AI systems (Amodei et al., 2016; Hall & Ellis, 2023; Partnership on AI, 2022). However, simply having diverse representations is not enough; meaningful inclusion depends on shifting power so that voices traditionally left out can shape agendas, influence decisions, and ensure that AI products meet a broad spectrum of needs (Ahmed, 2012). For inclusion to translate into real change, diverse stakeholders must be actively involved in every stage of AI development, from the initial conception to deployment and auditing.

Ethical pluralism calls for an ongoing commitment to dialogue, integrating multiple moral traditions and inclusive governance structures that involve a wide range of perspectives (Birhane, 2021; Costanza-Chock, 2020; UNESCO, 2024). Establishing participatory mechanisms, such as codesign frameworks, open forums, advisory boards composed of various cultural and ideological representatives, and partnerships with international organizations, can advance distributed ethical oversight, ensuring that guidelines are both globally informed and locally relevant (Youvan, 2024;

Practices such as codesign, participatory auditing, and robust accountability measures further strengthen the connection between diversity and positive outcomes (Whittaker et al., 2018). These practices not only mitigate bias but also foster trust and innovation, recognizing that people must design and govern equitable AI with, not just for, historically excluded communities (UNESCO, 2024).

3. Linguistic and Cultural Equity through Participatory Design

Participatory design is a collaborative method that actively engages stakeholders, particularly end-users and affected communities, in every step of the AI design process. In this way, their knowledge and lived experience are integrated with the expertise of system designers to ensure that outcomes are democratic and contextually relevant (Yuvan, 2024). This framework emphasizes reciprocal learning, cocreation, and collaborative decision-making to integrate user and community values directly into technical solutions (Simonsen & Robertson, 2012).

In AI, participatory design has been crucial for promoting fairness and language and cultural equity. Projects such as NūshuRescue illustrate how direct involvement from native communities maintains cultural authenticity and reduces the risk of misrepresentation (Dartmouth, 2023). While participatory methods may be resource-intensive, they offer long-

term benefits by avoiding costly errors, reputational damage, and even legal consequences tied to cultural insensitivity or exclusion (Costanza-Chock, 2020; Green, 2021).

Implementing effective participatory design in AI requires organizations to allocate resources, establish interdisciplinary teams, and partner with NGOs, universities, or community groups to share costs and strengths (Simonsen & Robertson, 2012). This commitment to inclusivity and shared ownership not only enhances project legitimacy but also enables technology to be adopted in an equitable and sustainable manner.

4. Culturally grounded governance and language preservation

With fewer than 1% of the world's languages supported by current AI tools (UNESCO, 2023), linguistic diversity remains starkly underrepresented in artificial intelligence. Koch (2025) recommended the development of codesigned frameworks that integrate language communities as key partners throughout the AI development lifecycle, from data collection and model training to application design and deployment. This approach encompasses creating tools and platforms that empower communities to manage and contribute directly to their own linguistic data and AI initiatives, strengthening both equity and authenticity in language preservation efforts. This approach of engaging language communities as core partners in the development of their own linguistic data and AI initiatives is not just a procedural step but a fundamental requirement for ensuring accuracy, authenticity, and ethical stewardship in language preservation efforts. (Birhane, 2021; Mohamed et al., 2020).

5. Education and capacity building

In the context of AI and cultural pluralism, education is about more than just teaching codes or technical skills. It's about helping people—whether they are future technologists, policymakers, or everyday users—understand how culture shapes the way we build and use AI. This means learning to recognize and respect the many languages, histories, and values that exist in our interconnected world.

Effective AI education enables the integration of cultural awareness and critical thinking into every aspect of the learning process. It encourages students and professionals to examine their own perspectives, spot biases in the technology they use, and design AI that works for everyone, not just a select few. For example, programs such as CATE-AI show how teaching frameworks can help learners connect local culture and global technology, leading to more intelligent and fairer decisions.

Universities and organizations that incorporate cultural context and lived experience into their lessons equip graduates with the tools they need to challenge stereotypes, address bias, and create technology that

brings people together rather than leaving some behind. Equally important is ensuring that voices from underrepresented communities—especially in the Global South—have access to funding, training, and leadership roles in AI. When these perspectives are included, AI becomes a more meaningful and equitable force for change (UNESCO, 2025).

Education, understood in this way, is not a one-time event but an ongoing, collective journey. It is about constantly learning from one another and working together to ensure that AI uplifts everyone and truly reflects the world’s diversity.

6. Cultural Auditing and Accountability

A cultural audit is a thorough examination of whether a technology, policy, or organization aligns with the values and daily realities of the communities it is intended to serve (The Oxford Review, 2024; Li et al., 2025). Concerning AI, a cultural audit is not just about spotting data bias or technical errors—it means looking beneath the surface to see if an AI system fits with what matters to real people (Li et al., 2025).

The most meaningful audits encompass diverse viewpoints, including those of outside experts, internal stakeholders, and representatives from the broader community. This mix helps ensure that we are not just listening to one group’s perspective (Raji et al., 2020; Costanza-Chock, 2020). Audits are also not a “one-and-done” process. They need to occur early, during, and after an AI system is rolled out so that the technology can keep up with how communities grow and change (Whittaker et al., 2018).

For these reviews to work well, it helps to use more than one set of rules—setting global guidelines but also customizing standards for different regions or communities. To keep information safe while remaining open to the process, organizations can utilize approaches such as data trust, which protects sensitive materials while still holding companies accountable (Rahwan et al., 2019). At their best, cultural audits help AI systems earn trust—and truly connect with the people they are built to help.

CONCLUSION

Cultural diversity is not a nice extra—it is the heart of ethical AI. Whenever technology overlooks differences between people, it risks not only making mistakes but also repeating centuries-old errors, deepening divides, and erasing the richness that makes communities unique (Birhane, 2021; Mohamed et al., 2020; UNESCO, 2024). However, when we make pluralism a goal, AI can become a much more hopeful force: a catalyst that helps cultures thrive, values grow, and languages survive.

Throughout this paper, I have explored what it takes to build inclusive AI. It is not just about mixing up design teams or ticking boxes for ethics reviews. This means embedding pluralism right from the beginning, giving people real power to shape technology, inviting communities to codesign solutions, and ensuring that every step—from governance to education—considers lived experience and local wisdom. Supporting linguistic diversity, funding leaders from under-resourced regions, and making space for genuine audits are practical ways to keep AI on track.

If we commit to sharing power, listening across traditions, and giving collaboration a real seat at the table, a world of pluralistic AI does not have to be a distant dream. We stand at a crossroads: technology can reinforce old systems of dominance, or it can help build a world where everyone's story counts. Making AI culturally inclusive is not just an ethical imperative—it is necessary for creating a future where people and cultures can thrive together.

References

- African Union. (2022). *Artificial intelligence continental strategy for Africa*. African Union Commission.
- Ahmed, S. (2012). *On being included: Racism and diversity in institutional life*. Duke University Press.
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. *arXiv Preprint*, arXiv:1606.06565.
- Azgin, B., & Kıralp, S. Z. (2024). Generative AI and linguistic homogenization. *Journal of Digital Cultures*, 15(2), 201–218.
- Bender, E. M., Gebu, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of FAccT '21*, 610–623.
- Birhane, A. (2021). Algorithmic injustice: A relational ethics approach. *Patterns*, 2(2), 100205. <https://doi.org/10.1016/j.patter.2021.100205>
- Birhane, A., Prabhu, V. U., & Kahembwe, E. (2022). Multicultural datasets in healthcare AI. *Health Informatics Journal*, 28(3), 1462–1478.
- Cambridge, E. (2023). *Ethics across cultures: Comparative traditions in AI*. Cambridge University Press.
- Costanza-Chock, S. (2020). *Design justice: Community-led practices to build the worlds we need*. MIT Press.
- Crenshaw, K. (1991). Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stanford Law Review*, 43(6), 1241–1299.
- Dartmouth. (2023). *NūshuRescue project*. Dartmouth College.

- Dignum, V., Corbett-Davies, S., & Zhou, Y. (2022). Plural values in AI ethics. *AI & Society*, 37(4), 1101–1113.
- Diversio. (2023). *Equity in global AI research: Funding disparities and opportunities*. Diversio Institute.
- Floridi, L., COWLS, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., ... Vayena, E. (2018). AI4People—An ethical framework for a good AI society. *Minds and Machines*, 28(4), 689–707.
- Green, B. (2021). The contestation of algorithms as regimes of justification. *Philosophy & Technology*, 34, 827–851. <https://doi.org/10.1007/s13347-020-00440-6>
- Greene, J., Hoffmann, A. L., & Stark, L. (2019). Better, nicer, clearer, fairer: A critical assessment of the Moral Machine experiment. *Ethics and Information Technology*, 21(2), 117–129.
- Gupta, R., & Kaul, M. (2024). AI for inclusive marketplaces: Cultural representation and equity. *International Journal of Business Ethics*, 19(1), 55–72.
- Hall, R., & Ellis, K. (2023). Diversity in design: Reducing algorithmic bias through inclusive practices. *AI & Society*, 38(2), 457–472.
- Hutchinson, B., Prabhakaran, V., Denton, E., Webster, K., Zhong, Y., & Denuyl, S. (2020). Social biases in NLP models. *Proceedings of FAccT '20*, 1–12.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.
- Li, X., Patel, A., & Gomez, R. (2025). Beyond bias: Cultural audits for inclusive AI. *AI & Society*, 40(1), 77–92.
- Li, X., Zhang, H., & Kumar, S. (2025). Assessing AI through cultural audits: A value-alignment approach. *AI & Society*.
- Lin, K. (2020). Defining cultural diversity in the digital era. *Journal of Global Ethics*, 16(3), 245–260.
- Mariyono, J., & Hidayatullah, A. (2024). Diversity and AI in global marketplaces. *Journal of Digital Business*, 12(1), 88–104.
- MDPI. (2023). *Culturally responsive AI curricula for higher education*. MDPI Education & Technology Report.
- Mohamed, S., Png, M. T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33(4), 659–684. <https://doi.org/10.1007/s13347-020-00405-8>
- Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press.
- Rahwan, I., Cebrian, M., Europa, L., Kovarik, J., & Zhao, W. (2019). Data trusts: Lessons from privacy, fairness, and accountability in AI. *Nature Machine Intelligence*, 1(12), 560–564.

- Raji, I. D., Smart, A., White, R. N., & Gebru, T. (2020). Closing the AI accountability gap: Defining an auditing standard for AI systems. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44.
- The Oxford Review. (2024). *Cultural audit – Definition and explanation*. <https://oxford-review.com/the-oxford-review-dei-diversity-equity-and-inclusion-dictionary/cultural-audit-definition-and-explanation/>
- Whittaker, M., Crawford, K., Dobbe, R., Fried, G., Kaziunas, E., Mathur, V., ... Schwartz, O. (2018). *AI Now Report 2018*. AI Now Institute, New York University.
- Youvan, D. C. (2024). *Ethical pluralism in AI: Challenging the monolithic values of red teams*. Retrieved from https://www.researchgate.net/publication/384014537_Ethical_Pluralism_in_AI_Challenging_the_Monolithic_Values_of_Red_Teams