Aggregated Cultural Positioning: Why AI Cannot Serve as a Sociocultural Mediator in Education

Zia Hassan

Abstract

As artificial intelligence tools proliferate in educational settings, proponents argue they will democratize access to personalized learning and narrow achievement gaps. This paper argues the opposite. Drawing on Vygotsky, Freire, and Gee, it contends that AI is structurally incapable of serving as a sociocultural mediator in education because large language models produce what this paper terms aggregated cultural positioning – statistically weighted outputs that reflect dominant cultural norms while presenting them as neutral. Deploying AI as a stand-in for culturally positioned human others does not expand access to genuine learning; it removes the mechanism by which genuine learning occurs, at greatest cost to the students the technology claims to serve. A framework centered on culturally situated dialogue, productive friction, and positioned feedback is proposed as a response.

Introduction

Many schools, districts, and higher education institutions are rushing to implement Artificial Intelligence (AI) tools into all facets of their work. Prominent technology companies such as OpenAI, Microsoft, and Anthropic have communicated to students, teachers, and administrators that AI not only has a place in their workflows, but that it is essential for their productivity, and for student learning (Anthropic Education Report, n.d.; Team, 2025). Proponents argue that students must be "exposed" to this technology because they will one day need to use AI in their workplace, and that these tools will achieve a type of ubiquity similar to that of the Microsoft Office suite.

Some extend this argument to say that not only will AI tools help all students, they will help the most needy students the most (The Potential Impact of Artificial Intelligence on Equity and Inclusion in Education, 2024). The claim is that AI tools can determine the exact level of mastery a student demonstrates, and then appropriately "scaffold" to the next level, acting as a personal tutor, despite these effect sizes shrinking or vanishing when they are held up against human tutoring (Létourneau et al., 2025). Still, the argument goes that these tools can scale and are cheaper than human resources. However, while these tools could indeed tutor students effectively, it will not close the achievement gap; instead, it will widen the inequities in the student population. The tools being deployed with the promise of solving inequity issues are structurally blind to the cultural dynamics that produce inequity in the first place.

Much of the current literature on AI in education focuses cognitive offloading (Fan et al., 2025), academic integrity (Ngo Cong-Lem et al., 2024), and intelligent tutoring (Lin et al., 2023). This article is unique due to its sociocultural lens. Even a perfectly accurate, unbiased AI system cannot do what sociocultural learning requires, because the problem goes deeper than a faulty algorithm. While generative AI may become faster and more sophisticated with time, it will never overcome its biggest sociocultural hurdle: it occupies no genuine cultural position from which to be challenged, questioned, or held accountable by a learner. Further, it cannot provide feedback from a cultural positioned being, since it isn't one.

Instead, AI chatbots are pre-trained on an enormous corpus of text that contains culturally relevant information; however, most of it is from the dominant culture in the western world (Fenech-Borg et al., 2025). It cannot take a sociocultural position; it can only imitate having one. I will refer to this imitation as an aggregated cultural position. The most common patterns that exist in the pre-training corpus of all the major Large Language Models (LLM) are the ones most likely to appear in its output (Bender et al., 2021).

This paper argues that AI's aggregated cultural positioning makes it structurally incapable of serving as a sociocultural mediator in education, and that deploying it as such does not reduce inequity but reproduces it. I will first explain what sociocultural learning requires. Then, I will discuss why AI cannot fill that role. Finally, I will propose a framework as a response to this problem that can be used to ensure that students in a school or classroom are truly engaging with culturally situated beings rather than a system that cannot be held accountable.

What Sociocultural Learning Requires

A common view of learning is that a teacher transmits their knowledge to a student in form of assessment and instruction. This view, sometimes called the "banking model" of education, was critiqued by Paulo Freire (1970) as a fundamentally dehumanizing conception of what happens between a teacher and a student. In the banking model, students are merely passive receptacles. They receive, store, and reproduce information deposited by the teacher. Freire argued that this model does not just fail pedagogically; it fails morally, because it does not consider the life experience that students bring to a classroom and it is selective about what is considered knowledge (exclusively limited to what the teacher knows). For Freire, genuine education requires two subjects in dialogue, each capable of being changed by the encounter. A teacher must have the propensity to be genuinely challenged by their students, and they must be able to be changed by a student's perspective. This is what separates sociocultural learning from knowledge transmission.

This is not merely a political argument. It reflects something structural about how human beings develop knowledge. Lev Vygotsky's sociocultural theory offers a developmental account of why the positioned other is not optional but essential. For Vygotsky, learning occurs in the space between what a learner can do independently and what they can do with the support of a more capable other, or what teachers call the Zone of Proximal Development (ZPD) (Vygotsky et al., 2012a). Critically, the "more capable other" in Vygotsky's framework is not a knowledge transmission device. Rather, they are a culturally situated being whose language, values, and ways of knowing are embedded in the interaction itself. Therefore, in a sociocultural learning environment, meaning is co-constructed. The scaffolding that moves a learner through the ZPD, then, is a relational process, one between two culturally situated beings.

What this means in practice is that the quality of the ZPD interaction is inseparable from the cultural positioning of the person providing support. When a teacher helps a student work through a problem, they are not simply providing a hint or reducing cognitive load. They are modeling an entire orientation toward problem-solving. They are communicating what counts as a good question, what counts as sufficient evidence, what kind of uncertainty is tolerable, and what kind of precision is required. These examples represent culturally specific ways of knowing, and they are transmitted not through explicit instruction but through sustained contact with someone who embodies them. Vygotsky was clear that development is always mediated by tools, by language, and importantly, by other people. Remove the cultural specificity of the mediating other, and you have removed the core of what makes it learning.

Gee (Gee, 1992) extends this argument into the domain of identity and membership. For Gee, what schools are really teaching when they teach anything is a "discourse," which he considers a socially recognized way of being in the world that includes values, ways of talking, ways of acting, and ways of interpreting experience. Discourses are not acquired through instruction but through immersion in communities of practice where a learner is apprenticed to someone who already inhabits that discourse. A student does not learn to think like a historian, a scientist, a critical reader, or a mathematician simply by receiving correct information about how those people think. They learn it by being in sustained, dialogic contact with someone for whom that way of thinking is lived, embodied, and consequential.

An important distinction is between what Gee refers to as primary discourses — those acquired in early life through family and community — and secondary discourses, which are acquired through participation in institutions like schools, workplaces, and civic organizations. Secondary Discourses are always learned in relation to primary ones, and the degree to which a learner's primary Discourse aligns with the secondary Discourse being taught has profound implications for how easy or difficult it will be absorb the secondary discourse. A student whose home culture already mirrors the epistemic norms of academic discourse may find that they are rewarded when they engage in the discourse at school. On the other hand, a student whose primary dscourse is more distant from that norm faces a more arduous and often more alienating process of acquisition. But the presence of a mentor who can bridge those discourses, one who can meet the learner where they are culturally and walk with them toward where they would like to be, can be path-changing.

And, this is precisely why the identity dimension of learning cannot be separated from its cognitive dimension. When a student learns, they are not just acquiring skills or knowledge. They are negotiating membership in a community of practice (Lave and Wenger-Trayner, 2020). They are trying on ways of being, testing whether those ways of being are available to someone like them, and looking to the people around them for evidence about whether they belong. A teacher who shares cultural markers with a student (or who has navigated the distance between cultures themselves) can provide something that no amount of correctly sequenced content can replicate: evidence that the discourse is is worth the engagement, and the learner will indeed by changed by their efforts to engage with it.

Taken together, Freire, Vygotsky, and Gee converge on a single structural requirement for genuine sociocultural learning: it requires a positioned other. The requirement goes beyond mere knowledge. The teacher, tutor, or mentor must occupy a genuine cultural position, one that can be observed, questioned, pushed back against, and ultimately internalized or rejected by the learner. It is this mechanism by which learning, in the full sociocultural sense, actually happens.

This requirement has consequences that are easy to overlook in discussions focused narrowly on knowledge transfer and performance outcomes. For example, a positioned other can be wrong in culturally revealing ways, and they can also change their mind. They can be challenged and forced to defend their perspective. They can demonstrate what it looks like to hold a view with conviction, to change a view in response to evidence, or to maintain a position under social pressure. They bring with them a history, a community, and a set of genuine stakes in the conversation. The learner gets to absorb this intuitively. They are, therefore, learning how to be.

It is precisely the requirement for a positioned, accountable, culturally embedded other that AI cannot meet. Not because of current technical limitations that future development might resolve. But for structural reasons that are independent of sophistication, speed, or scale.

Why AI Cannot Fill the Sociocultural Role

The previous section established that genuine sociocultural learning requires a positioned other: a culturally situated being whose perspective can be questioned, challenged, and held accountable. This section argues that AI cannot fulfill this role, not because of current technical limitations, but for structural reasons rooted in how large language models are built and what they are capable of representing. Further, it argues that deploying AI as though it can fill this role actively reproduces the inequities it claims to address.

To understand why, it is necessary to examine what AI actually does when it appears to take a cultural position. Large language models are trained on vast amount of text drawn mostly from digitized, English-language, western sources (Bender et al., 2021). The outputs these models produce are not the result of cultural experience, community membership, or embodied participation in a discourse. They are the result of statistical patterns. What presents as output are simply the most common ways that certain ideas, values, and social norms have been expressed in the training data. When an AI system responds to a student's question, offers feedback on their writing, or models a social skill, it is not drawing on a genuine cultural position. It is producing an aggregated cultural position: a weighted average of the dominant patterns in its training corpus, presented as though it were a perspective.

This distinction matters. A genuine cultural position can be interrogated. A student can ask a teacher why they value the kind of argument structure they are rewarding, and the teacher must answer (or refuse to answer), which is itself revealing. A positioned other can misrepresent themselves in culturally specific ways that expose the assumptions underlying their feedback. They can be pushed, surprised, and changed by the encounter. An aggregated cultural position has none of these properties. It cannot be held accountable (Flenady and Sparrow, 2025) because it has no stake in the conversation. It cannot be genuinely challenged because there is no mini-human inside the machine. Its position, therefore, is not so much wrong or inaccruate; it is the result of a categorically different process. It carries no weight and no actual meaning, other than what the user interprets. While a user interpretation may lead them to insight, it cannot be said that this is the same as dialogue between two culturally situated beings.

This gap between culturally situated dialogue and AI output is larger than it might seem on first blush. One may argue, for example, that using an AI chatbot like ChatGPT is akin to looking up information in a library. Libraries, they might argue, are not culturally situated beings, but the texts within them might have cultural value. However, this view is misguided. The aggregated cultural position that AI produces is not a balanced synthesis of all human cultural expression. It is weighted heavily toward the dominant patterns in its training data. These patterns reflect, disproportionately, white, western, middle-class, English-language ways of knowing, communicating, and being (Fenech-Borg et al., 2025). Pierre Bourdieu's concept of cultural capital helps clarify the consequences. For Bourdieu, cultural capital refers to the knowledge, skills, and dispositions that are valued by dominant institutions — the kind of capital that schools reward and that is transmitted primarily through families with access to those same institutions (Bourdieu, 2018). Students who arrive at school already in possession of this capital and whose home cultures mirror the epistemic and social norms that school rewards move through educational institutions with relative ease. Students whose cultural capital diverges from institutional norms, on the other hand, face the additional burden of translation, assimilation, or alienation.

When AI is introduced into this already unequal field, it risks accelerating inequities. A student who submits writing to an AI writing assistant receives feedback calibrated to the dominant norms of academic prose. The feedback nudges them, deceptively and insidiously, toward assimilation. There is no person to ask about the reasoning for the chatbot's feedback. There is no teacher who has navigated the distance between cultures themselves and can name the cultural preference embedded in the feedback, or provide alternate feedback that may not be purposefully culturally situated, but will be nonetheless. The reproduction happens in the background, disguised by a confident and the AI's anthropomorphic character. Students whose primary discourse already approximates the dominant norm receive confirmation, while students whose primary discourse diverges receive correction. While the chatbot can explain its reasoning, there is no verification, no accountability, and no recourse.

The same dynamic operates in domains far beyond writing. Consider the teaching of social and emotional skills, an area where AI is increasingly being deployed as a personalized development tool or as a way for teachers to plan SEL lessons (Henriksen et al., 2025). A student seeking to improve their communication skills might receive guidance from an AI system trained on frameworks like CASEL's widely-used SEL model (Domitrovich et al., 2025). The guidance will likely include advice to help teach relationship skills and social awareness, such as maintaining eye contact, nodding while listening, and mirroring body language (Akechi et al., 2013). These behavioral signals that are coded as markers of attentiveness and respect in dominant western cultural contexts. What this guidance cannot do is teach the student the underlying relational capacity these signals are meant to represent. Eye contact is not listening. Nodding is not understanding. These are <u>performances</u> of connection, not connection itself. The difference is precisely what sociocultural learning is supposed to develop.

More troubling still is that the framework being reproduced here carries its own acknowledged limitations. Researchers conducting content analyses of CASEL-derived SEL standards have cautioned that the very notion of educational standards is a culturally embedded construct that may not be applicable across cultures, calling for more culturally informed approaches to SEL. A CASEL-affiliated equity analysis has similarly acknowledged that without explicit attention to cultural diversity, prevalent SEL frameworks risk reflecting a bias toward white institutional norms (Williams and Jagers, 2022). An AI system that defaults to CASEL as the authoritative framework for social and emotional development may indeed be reproducing a framework that has internally flagged its own cultural limitations, while presenting those limitations as universal best practice to students who have no way of knowing otherwise.

The research on SEL and socioeconomic status makes the consequences of this concrete. Using PISA data across 74 countries, Gruijters et al. (2024) found that socioeconomically advantaged students consistently demonstrate higher levels of socio-emotional skills than their disadvantaged peers, with these skill differences accounting for a meaningful share of the overall achievement gap. The authors acknowledge cultural factors as a plausible mechanism without fully investigating them. What remains underdeveloped in that article is precisely what this paper argues: the SEL skills being measured are not culturally neutral competencies. They are culturally specific performances that students from dominant cultural backgrounds are more likely to have already acquired through their primary discourses. When AI delivers SEL instruction calibrated to those performances it teaches students who lack dominant cultural capital to perform its signals, a fundamentally different, and more insidious, thing.

This is what the progressive veneer problem means in practice. The promise of AI in education is falsely believed to be student learning. The myth is that it will help the students who need it most, by providing personalized, scalable access to the kind of support that has historically been available only to those with resources (The Potential Impact of Artificial Intelligence on Equity and Inclusion in Education, 2024). But the personalization being offered is personalization within a culturally fixed frame. The more an educational system relies on AI to deliver instruction, feedback, and social-emotional guidance, the more it removes the culturally positioned human other from the learning relationship. And, it is precisely that human other who is capable of bridging cultural distances, naming cultural assumptions, and providing the kind of accountable, positioned engagement that Vygotsky, Freire, and Gee identify as the engine of genuine learning. The students who lose the most when that human is replaced are those whose primary discourses are furthest from the dominant norm, which happen to be the students the technology claims to serve.

Bourdieu's concept of symbolic violence applies to this scenario. Symbolic violence refers to the process by which dominated groups come to accept the terms of their own domination as natural, inevitable, or even beneficial (Bourdieu and Wacquant, 2004). The story with AI in classrooms is more nuanced but even more destructive than Bourdieu could have ever imagined. For example, the student who receives AI feedback nudging their writing toward dominant prose norms is not being coerced. In fact, they think they are getting support. The student who learns to perform eye contact and nodding from an AI social skills teacher is receiving conditioning. The progressive framing of AI as an equity tool is not a cynical deception. It is something more structurally concerning: a genuine belief that access to a culturally dominant aggregated position is the same as access to a culturally positioned other. It is not. And the consequences of that confusion fall, as they always do, most heavily on those who were already furthest from the cultural center.

Viewed through Dewey's framework (Dewey, 1986a), AI systems fail to satisfy even the most basic conditions of genuine learning experience, as Table 1 illustrates.

Table 1: Dewey's Qualitative Thought (1931): Human vs. AI Learning

Dimension	Human Learning (Dewey)	AI "Learning"
Experience	Qualitative, felt, immediate — "total seizure" of wholeness	Statistical pattern recognition — no felt quality
Modification	Each experience modifies the learner — continuous growth	Fixed weights after training — no modification through use
Embodiment	Body is site of experience — sensing, feeling, doing	No body, no sensory experience
Social Context	Learning through interaction with positioned others	Aggregates patterns from text about interactions
Purpose	Learner has stakes, desires, intentions	No purposes — responds to prompts without intention
Integration	Experience forms unified whole — "an experience"	Outputs are stitched-together patterns — no unity
What Changes	Habits, dispositions, ways of being — self transforms	Weights, parameters — no self to transform

A Framework for Sociocultural Reintegration

The previous two sections have established a pointed critique: that AI's aggregated cultural positioning makes it structurally incapable of serving as a sociocultural mediator in education, and that deploying it as though it can do so reproduces the very inequities it claims to address. This section proposes a framework that allows educational institutions to understand what genuine sociocultural learning environments must preserve.

The framework is organized around a single central mechanism: the culturally situated human relationship. This mechanism has three necessary expressions:

1. Culturally situated dialogue — the ongoing process of exchange between a learner and a positioned other.
2. Productive friction — the necessary resistance that genuine dialogue creates and that propels learning forward.
3. Positioned feedback — the application of cultural situatedness to the specific work a student produces. These three are not independent features to be checked off a list. They are expressions of the same underlying requirement, each incomplete without the others.

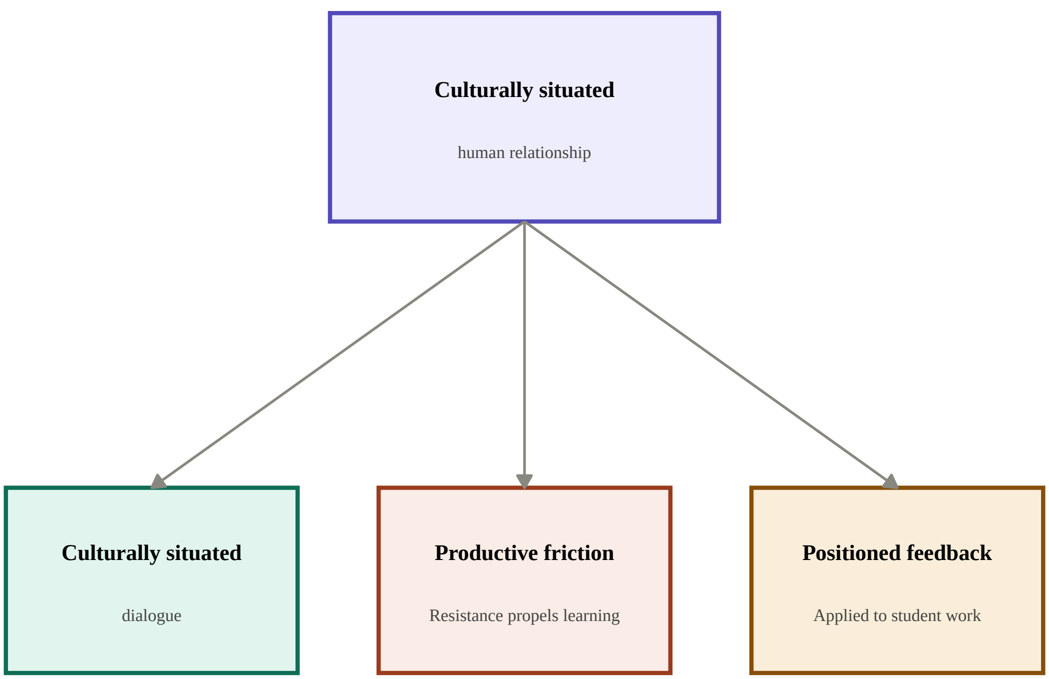

Figure 1: The culturally situated human relationship framework

Culturally Situated Dialogue

Dialogue, in the Freirean sense, is not just conversation. Dialogue requires two subjects with genuine cultural positions, each capable of being changed by the encounter (Freire, 1970). What makes dialogue educationally generative is not its pleasantness or its efficiency. It is the fact that it puts two different ways of being in the world into productive contact with each other.

This has direct implications for how schools and classrooms should be structured in an era of increasing AI integration. AI chatbots can simulate dialogue, but this simulated dialogue cannot do what genuine dialogue can, for reasons previously made clear. What follows from this is that schools must treat genuine dialogic exchange between students and culturally positioned humans not as a luxury or an enrichment activity, but as a structural requirement of the learning environment itself.

In practice, this supports protecting and prioritizing the conditions under which real dialogue can occur. Students need to have sustained, repeated contact with teachers and mentors who occupy genuine cultural positions, including positions that may differ from the student's own. And, we must think twice before replacing teachers with AI chatbots, as a growing number of schools are doing (Parents Fell in Love With Alpha School's Promise. Then They Wanted Out | WIRED, n.d.). Even more importantly, students must be clear on what dialogue is and why it matters. Overlooking this means that they may not care or be motivated to engage in the humanistic situations we facilitate for the good of their learning. Over the course of their schooling, they will come to understand that the discomfort of genuine exchange with a positioned other is what makes education worthwhile.

Culturally situated dialogue also means diversifying whose cultural positions are represented in the learning environment. A classroom in which all the positioned others share the same cultural background, even a non-dominant one, is impoverished in a different way than one dominated by a single cultural norm. The goal is not to replace dominant cultural positioning with a different dominant positioning. It is to ensure that students encounter multiple genuine positions, develop the capacity to engage with them, and in doing so, develop a richer and more flexible sense of their own.

Productive Friction

AI systems are optimized for user satisfaction. They are designed to be responsive, agreeable, and frictionless. This is intentional. This design choice reflects the commercial imperatives of the companies that build them. Frictionless may be a worthwhile aim in industry, but it is the antithesis of what makes education worthwhile.

Vygotsky's ZPD depends on resistance. The space between what a learner can do independently and what they can do with support is not a comfortable space, and it is not meant to be. It is, by design, a space of productive struggle. What makes the struggle productive is the presence of a more capable other who will do more than validate what the learner already knows (Vygotsky et al., 2012b). Further, it requires a human that can gauge in an intuitive way what the student might need to grow (Dewey, 1986b). If the learning process becomes too "smooth" (Hassan, 2025), the risk is that students will come to rely on the sycophancy or warmth over the chatbot (Ibrahim et al., 2026) in order to engage with the work. Staring at a blank page or working so hard at a math problem that the paper tears will become a rare event, and yet that is exactly the type of friction needed for good learning. This is what Bjork & Bjork (2020) refer to as "desirable difficulties." Students engage with a learning process when they find the challenge harder; the result is stronger retention and deeper processing.

Productive friction is the resistance that a genuinely positioned other provides. It is the specific kind of pushback that comes from someone who sees the problem differently and will not waver in their expectations. It is the teacher who asks the student to extend their thinking, not because an algorithm told them to, but because they care. It is the mentor who reads a student's work and says this argument doesn't hold, and here is why. It is the senior colleague who disagrees in a meeting and is willing to defend their disagreement. What makes friction productive rather than merely frustrating is precisely its cultural situatedness. The pushback comes from the position, from a real set of values and ways of knowing. Importantly, the student can push back against it, as well.

This is why productive friction is inseparable from culturally situated dialogue. Friction without cultural situatedness is merely a way to induce frustration needlessly. Cultural situatedness without friction is affirmation. Together, they constitute the dynamic that Freire describes as genuine education: two subjects, each capable of changing the other, in sustained and honest exchange.

For schools navigating AI integration, this means actively designing for friction rather than against it (Hassan, 2025). Schools can increase the amount of this kind of friction by prioritizing learning and deprioritizing efficiency and student comfort. This has implications for teacher training programs, too, since we can teach new teachers to recognize these moments of productive struggle and understand their immense value.

Positioned Feedback

The third expression of the framework concerns the feedback that students receive on the work they produce. This is the domain where AI has made its most aggressive inroads. Teachers have used AI to create virtual writing assistants, automated grading tools, AI tutors that provide instant feedback on student responses (Liu et al., n.d.). Therefore, it is the domain where the consequences of aggregated cultural positioning are most immediately visible.

Feedback is never neutral. When a teacher reads a student's essay and responds to it, they are not applying a culturally independent algorithm for quality assessment. They are bringing their entire cultural history to the encounter, including their sense of what an argument should look like, what evidence is compelling, what voice is appropriate, what matters. This cultural specificity is what makes the feedback educationally meaningful rather than merely corrective.

Positioned feedback does two things that AI feedback cannot. First, it can be interrogated. A student who receives feedback from a human teacher can ask why. They can push back, can request justification, and can disagree. In doing so, the student engages in exactly the kind of dialogue that the first expression of this framework describes. The feedback becomes an occasion for learning not just about the work but about the cultural norms that shape how the work is evaluated. Second, positioned feedback can be explicitly culturally aware. A teacher who understands that a student's rhetorical choices reflect a different cultural tradition can name that. This kind of explicit cultural metacognition is not possible from a system that does not know it has a cultural position. It is important to note that AI systems can, once again, imitate these processes. A student can indeed ask an AI chatbot why it gave the feedback it did, and perhaps even disagree with the chatbot. However, disagreement with a chatbot often prompts the chatbot to begin to agree (Ibrahim et al., 2026), at least partially, with the student which renders the AI useless as a situated being.

This does not mean that AI tools have no role in the feedback process. It means that their role must be carefully circumscribed. AI can flag grammatical errors, identify structural inconsistencies, and provide information. What it cannot do is provide the culturally positioned, accountable, dialogic feedback that moves a learner through the ZPD and helps them negotiate membership in a community of practice.

Structural Implications

Table 2 operationalizes each expression as observable indicators for educators and researchers.

Table 2: Framework implementation indicators

Dimension	Indicators of presence	Indicators of absence
Culturally situated dialogue	Educator draws on shared cultural knowledge to reframe concepts; learner's background shapes the exchange; dialogue reflects awareness of both parties' positioning	Generic explanations with no reference to learner context; exchange could apply to any learner in any setting; cultural background treated as irrelevant or invisible
Productive friction	Learner pushes back or expresses confusion; educator holds the tension rather than resolving it prematurely; discomfort is treated as productive rather than corrected away	Smooth redirection at signs of resistance; consensus sought quickly; learner's uncertainty is treated as a problem to be resolved rather than a site of learning

Dimension	Indicators of presence	Indicators of absence
Positioned feedback	Feedback references the learner's specific cultural or social positioning; educator acknowledges how their own positioning shapes their reading of the work	Feedback addresses only technical correctness; no acknowledgment of how context shapes meaning; same response would be given regardless of who produced the work

The framework proposed here makes demands that run against the current direction of educational technology adoption. It insists on the irreplaceability of the culturally positioned human other at precisely the moment when institutional pressures are pushing toward replacement. It argues that the sociocultural features of human learning relationships that are most difficult to scale are, in fact, what makes the learning work.

This does not mean rejecting AI tools in educational settings wholesale, but it does mean being precise about what those tools can and cannot do, and designing learning environments accordingly. AI can handle information retrieval, some peripheral teaching tasks, and administrative burden. But it cannot do the irreducibly human work of culturally situated dialogue, productive friction, and positioned feedback. No amount of sophistication in AI development will change that.

Conclusion

The promise at the center of AI's educational turn is a genuinely appealing one. Every student, regardless of zip code or family income, with access to a patient, knowledgeable, always-available tutor. This paper argues that school leaders make the wrong assumption about what disadvantaged students lack. It is not access to information, to correct answers, to well-sequenced content. If that were true, the library would have closed the achievement gap long ago.

The concept of aggregated cultural positioning developed in this paper is an attempt to name something that is easy to overlook in the excitement of a new technology. AI arrives in the classroom already weighted toward dominant language patterns, dominant epistemic norms, dominant ways of performing competence and connection. It arrives having already decided, statistically, what good writing looks like, what attentive listening looks like, what a well-reasoned argument looks like. And it delivers these weightings without accountability, without cultural self-awareness, and without any possibility of being genuinely challenged by the student on the receiving end.

The consequence, as Bourdieu would recognize immediately, is reproduction dressed as innovation. The students who benefit most from AI's aggregated cultural position are those whose primary discourses already approximate it. The students for whom the technology was ostensibly designed receive, at best, efficient assimilation into a cultural frame they had no hand in constructing, and at worst, the systematic replacement of the human relationships that might have helped them navigate that frame on their own terms. The progressive veneer does not merely fail to close the gap. It closes off the conversation about why the gap exists.

This is not argument for technological abstinence. It is an argument for theoretical clarity. AI tools will continue to proliferate in educational settings, and there are genuine peripheral tasks they can perform. But the framework proposed here insists on a distinction that institutions are currently at risk of losing: the distinction between what AI can do and what education requires.

References

Akechi, H., Senju, A., Uibo, H., Kikuchi, Y., Hasegawa, T., and Hietanen, J. K. (2013). Attention to Eye Contact in the West and East: Autonomic Responses and Evaluative Ratings. PLoS ONE, 8(3), e59312. https://doi.org/10.1371/journal.pone.0059312

Anthropic Education Report: The AI Fluency Index. (n.d.). Retrieved May 6, 2026, from https://www.anthropic.com/research/AI-fluency-index

Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? . Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610–623. https://doi.org/10.1145/3442188.3445922

Bjork, R. A., and Bjork, E. L. (2020). Desirable difficulties in theory and practice. Journal of Applied Research in Memory and Cognition, 9(4), 475–479. https://doi.org/10.1016/j.jarmac.2020.09.003

Bourdieu, P. (2018). Cultural Reproduction and Social Reproduction. In R. Brown (Ed.), Knowledge, Education, and Cultural Change (1st ed., pp. 71–112). Routledge. https://doi.org/10.4324/9781351018142-3

Bourdieu, P., and Wacquant, L. (Eds.). (2004). An invitation to reflexive sociology (Nachdr.). University of Chicago Press.

Dewey, J. (1986b). Experience and Education. The Educational Forum, 50(3), 241–252. https://doi.org/10.1080/00131728609335764

Dewey, J. (1986a). Experience and Education. The Educational Forum, 50(3), 241–252. https://doi.org/10.1080/00131728609335764

Domitrovich, C. E., Li, Y., Kendziora, K., and Greenberg, M. T. (2025). Impact of schoolwide support for social and emotional learning: The effect of the CASEL school guide on organizational capacity for SEL and systemic SEL. Social and Emotional Learning: Research, Practice, and Policy, 6, 100156. https://doi.org/10.1016/j.sel.2025.100156

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., and 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

Fenech-Borg, E. Z., Meznaric-Kos, T. P., Lekovic-Bojovic, M. D., and Hentze-Djurhuus, A. J. (2025, December 29). The Cultural Gene of Large Language Models: A Study on the Impact of Cross-Corpus Training on Model Values and Biases. https://doi.org/10.48550/arXiv.2508.12411

Flenady, G., and Sparrow, R. (2025). Cut the bullshit: Why GenAI systems are neither collaborators nor tutors. Teaching in Higher Education, 1–10. https://doi.org/10.1080/13562517.2025.2497263

Freire, P. (1970). Freire, P. (1970). Pedagogy of the Oppressed. New York Seabury Press. - References - Scientific Research Publishing. https://www.scirp.org/reference/referencespapers?referenceid=1415936

Gee, J. P. (1992). A Situated Sociocultural Approach to Literacy and Technology.

Hassan, Z. (2025). Friction and Coordination in Schooling:

Henriksen, D., Creely, E., Gruber, N., and Leahy, S. (2025). Social-Emotional Learning and Generative AI: A Critical Literature Review and Framework for Teacher Education. Journal of Teacher Education, 76(3), 312–328. https://doi.org/10.1177/00224871251325058

Ibrahim, L., Hafner, F. S., and Rocher, L. (2026). Training language models to be warm can reduce accuracy and increase sycophancy. Nature, 652(8112), 1159–1165. https://doi.org/10.1038/s41586-026-10410-0

Lave, J., and Wenger-Trayner, É. (2020). Situated learning: Legitimate peripheral participation (31st printing). Cambridge Univ. Press.

Létourneau, A., Deslandes Martineau, M., Charland, P., Karran, J. A., Boasen, J., and Léger, P. M. (2025). A systematic review of AI-driven intelligent tutoring systems (ITS) in K-12 education. NPJ Science of Learning, 10, 29. https://doi.org/10.1038/s41539-025-00320-7

Lin, C.-C., Huang, A. Y. Q., and Lu, O. H. T. (2023). Artificial intelligence in intelligent tutoring systems toward sustainable education: A systematic review. Smart Learning Environments, 10(1), 41. https://://doi.org/10.1186/s40561-023-00260-y

Liu, A., Esbenshade, L., Sarkar, S., Sun, M., Zhang, Z., Han, T., Lapicus, Y., and He, K. (n.d.). 1 University of Washington, 2 Colleague AI.

Ngo Cong-Lem, Tin Nghi Tran, and Tat Thang Nguyen. (2024). Academic Integrity in the Age of Generative AI: Perceptions and Responses of Vietnamese EFL Teachers. Teaching English with Technology, 24(1), 28–48. https://search.ebscohost.com/login.aspx?direct=true&AuthType=ip,shib&db=eric&AN=EJ1417822&site=ehost-live&scope=site&authtype=ip,shib&custid=s3555202

Parents Fell in Love With Alpha School's Promise. Then They Wanted Out | WIRED. (n.d.). Retrieved October 28, 2025, from https://web.archive.org/web/20251027180145/https://www.wired.com/story/ai-teacher-inside-alpha-school/

Team, M. E. (2025, August 20). AI in Education Report: Insights to support teaching and learning. Microsoft Education Blog. https://www.microsoft.com/en-us/education/blog/2025/08/ai-in-education-report-insights-to-support-teaching-and-learning/

The potential impact of Artificial Intelligence on equity and inclusion in education (OECD Artificial Intelligence Papers 23; Vol. 23). (2024). (Vol. 23). https://doi.org/10.1787/15df715b-en

Vygotsky, L. S., Cole, M., John-Steiner, V., Scribner, S., and Souberman, E. (2012b). Mind in society: Development of higher psychological processes. Harvard University Press. https://books.google.com/books?id=u2PP6b0ddtoC

Vygotsky, L. S., Cole, M., John-Steiner, V., Scribner, S., and Souberman, E. (2012a). Mind in society: Development of higher psychological processes. Harvard University Press. https://books.google.com/books?id=u2PP6b0ddtoC

Williams, B. V., and Jagers, R. J. (2022). Transformative Social and Emotional Learning: Work Notes on an Action Research Agenda. Urban Education, 57(2), 191–197. https://doi.org/10.1177/0042085920933340