The Algorithmic accent: investigating the homogenization of linguistic identity in AI-mediated classrooms

ISSN 2657-9774; https://doi.org/10.36534/erlj.2026.05.03

Tehreem Fatima*, Muhammad Zain ul Hassan**, Muhammad Nawaz***
*Government College University, Pakistan; tehreemfatima2191@gmail.com
**Air University, Pakistan; zainhadri5@gmail.com
***Northern University, Pakistan; Nawaz@northern.edu.pk

Abstract
AI has become ubiquitous in schools, and personalised learning carries a cost that is rarely discussed: the flattening of student identity. This paper examines the algorithmic accent - the way Large Language Models (LLMs) often default to a standardized, Western-centric linguistic style that can overshadow diverse student voices. The technology is quick. But that speed comes with a quieter effect: it eliminates local dialects and cultural speaking style along the way. Most research so far asks only whether students achieve improved grades using AI; how this impacts a student's sense of self, and their ability to negotiate multiple cultures, remains an enormous gap. We report on a two-month, teacher-led observation of 15 university students (semesters 1-5) enrolled in an academic writing course at an academy in Pakistan, all of whom gave informed consent to participate. Drawing on translanguaging (Wei, 2018) and the concept of Global English standardisation (Agarwal et al., 2025), we suggest these tools function as a filter rather than a mirror. Students are led to exchange their own voice for the words of the algorithm-more fluent, perhaps, but missing the social texture that makes a language human. We recommend a pedagogical approach in which AI becomes a tool for exploring language, not a template for conforming to it.

Keywords*: algorithmic accent, linguistic identity, linguistic homogenization, translanguaging, ai-mediated classrooms*

Introduction

The red pen is mostly gone. Instead, there is a glowing fix button that the majority of the students press without giving a second thought. A student comes in with a brilliant idea, but very messy, ideas laden with local slang, something weird and beautiful, or a home sounding rhythm. They in turn submit a paper ten minutes later that resembles a corporate press release. It is grammatically flawless, but the human is omitted. This is the algorithmic accent. We discuss Large Language Models as though they are unbiased assistants when in fact, they have a massive Western bias - and that trade-off is exactly what this paper sets out to examine.

We are seeing the invisible prompt usurp in our classrooms. Students would keep on struggling with a sentence until it was right. They now let the machine determine what is right. This does not mean only spelling. It is a sociolinguistic crisis. Agarwal et al. (2025) show that AI suggestions push multilingual writers toward Western styles, which we refer to as a Global English narrowing. The software tends to erase the distinct traits that characterise a student with a multilingual background when he or she is using AI. It is not merely correcting a typo, but something to do with identity. This gives the illusion of proficiency wherein the student appears to be a master over the language yet has lost authority over the language. Bastani et al. (2025) found that students who relied on unrestricted GPT support performed better during practice but worse once it was removed, which suggests that the appearance of mastery can outrun real learning; in writing, the result can be syntactically correct but culturally empty text.

Linguistic identity is something that we cannot simply classify as a box. It is a process in which one breathes life. Wei (2018) revealed that high intelligence is manifested in translanguaging, which is the ability of individuals to incorporate their languages to communicate complex ideas. It is the way the brain really functions in a globalised world. We therefore adopt translanguaging as our theoretical lens: a student's mixed or hybrid language use is not treated as error here, but as evidence of active, cross-linguistic meaning-making. However, according to AI models, this is noise. They are socialised on a particular part of the internet that is largely Western and affluent. This training data is an agent of filtering. When the natural manner of thinking of a student does not correspond to that filter, the AI points to it as a mistake. With time, the student ceases to fight. They begin to write in a manner which the AI will not redline. This is the research gap with which we are most concerned. What does a student think when they are repeatedly informed via the machine that the dialect they use is incorrect due to being their heritage to begin with?

Digital hidden curriculum is something that we as educators can observe being created in real time. Giroux and Penna (1979) described the hidden curriculum as the unstated norms and values that schooling teaches beyond what is formally tested, and AI writing tools can be read as a digital extension of it. Holmes and Tuomi (2022) note that AI does not have the compassion and human imagination to be truly close to humans - and it is in the struggle of choosing words, not the speed of producing them, that learning occurs.

These models carry a problem in their very data: the majority are trained on high-resource, Western languages. Joshi et al. (2020) and Bender et al. (2021) both show that this reflects how data are collected and prioritised, a design choice, not a technical accident.

What remains with us is the long-term effect. A majority of studies will only focus on a single semester. Change does not happen in a period that short. Doshi and Hauser (2024) address a related problem: AI-assisted writing raises individual quality but makes outputs more alike across writers, so a tool that suits everyone tends to flatten what is distinctive. So, what will become of the sense of self in any person once they have three years of the algorithm refining every thought? This is the longitudinal gap which we are bridging. If the AI is the one constantly refining, who does the thinking? We are in danger of producing a generation of brilliant writers who all happen to have the same robotic, neutral accent. This echoes Fan et al. (2024), who warn that reliance on ChatGPT can foster technology dependence and "metacognitive laziness".

This study asks a simple question: how does sustained, AI-mediated writing support affect multilingual students' linguistic identity and their sense of authorial voice? Over two months, we observed 15 university students (semesters 1-5) in an academic writing course at an academy in Pakistan, recording the precise moments where a student had to choose between their own voice and the algorithmic ideal. We wanted the friction. We wanted the silences - the moment a student decides a local idiom simply is not good enough for the machine. This paper bridges the research gap identified above with an account of that sociolinguistic change, the kind that data points alone cannot capture.

Literature review

AI bias and the "global English" effect

There is already some alarm in present studies. Agarwal et al. (2025) show that AI writing suggestions pull writers toward a standardised, Western-leaning form of Global English - a business, refined form of the language which is no one's and everyone's simultaneously. When a student in a city such as Manila or Nairobi asks an AI to improve their work, the software tends to remove the local taste, polishing away the rhythms that make that student sound like themselves. Rosa and Flores (2017) argue that standards of correct language are shaped by power rather than being neutral, so this is not only a technological problem but also a power struggle over the authority to determine what counts as "correct" speech. Most existing studies remain at the level of describing Global English narrowing. We push further, into the space of individual ownership: does a student eventually stop attempting their own expressions, simply because

they know the machine will erase them anyway? That shift - from language pluralism to a kind of mute acquiescence - is what interests us here.

Translanguaging as a theoretical lens

Cenoz and Gorter (2021) argue that pedagogical translanguaging is crucial to successful learning, yet AI tends to interfere with this natural process. By flagging hybrid, code-switched phrasing as error, these tools implicitly teach students that their natural way of thinking is a bug in the system - even a tool given a bilingual dictionary would still lack an explanation for the code-switch itself. Fan et al. (2024) caution that reliance on generative AI can foster dependency and metacognitive laziness, which may weaken learners' engagement with their own linguistic judgement.

The proficiency illusion

The gap in question is hazardous when it comes to the so-called illusion of proficiency. Bastani et al. (2025) found that students with unrestricted AI support performed better while it was available but worse once it was removed: a gap between apparent and actual competence. In our setting this appears as text that is eloquent on paper but lacking in social flavour. This is because AI is not pragmatically nuanced - it is not aware when a word is too serious or too crude for a particular local culture. Bender and Koller (2020) argue that systems trained on linguistic form alone do not learn the communicative meaning behind it, which helps explain why AI can work as a dictionary but a poor cultural guide. We are concerned with the gap between technical fluency and authenticity: when a student speaks like a CEO but thinks like a child because the machine has done all the heavy work, the student has been failed rather than served. According to Holmes and Tuomi (2022), AI does not possess the human ability to empathise or create a true connection with others.

Decolonizing AI and data bias

The problem itself is in the data. Bender et al. (2021) note that training data that overrepresents dominant groups encodes their viewpoints into model output; the majority of models are trained on high-resource languages, leaving what Joshi et al. (2020) show to be a severe data gap for everyone else. This is not merely a technical bug but a design decision that reproduces existing structures of power. Birhane (2020) warns that technologies built in the Global North can reproduce colonial relations, which is why these tools need to be decolonized; personalization risks becoming a novel way of imposing old norms. Our response is to propose a Resistance Pedagogy, a culturally responsive approach in which the teacher becomes a mediator. Instead of letting the AI silently rewrite a student's text, we stop and ask: why did it change that word? Does the result still sound like you? This lines up with what Mohamed et al. (2020) call unrepresentativeness in the AI development pipeline.

Methodology

The approach to this inquiry is teacher-led rather than a formal experiment, using an autoethnographic approach combined with classroom observation. Fifteen university students across semesters 1 to 5 took part, all enrolled in an academic writing course at an academy in Pakistan. They came from classes the first author was already teaching, and joined on a voluntary basis. This is a useful way of addressing the research gap, since it captures the tension of lived experience between the student, the teacher, and the screen. This study uses a teacher-led observation approach over a two-month period. We are not looking at these students as distant researchers; we are watching them work in real time. This allows us to see the "micro-moments" of linguistic surrender that a one-off survey would miss. Most studies on AI in education are done in labs or through data mining; we instead intervene in the gap between the screen and the student's desk. As the class teacher delivering instruction, the first author's observations are centred on her own reflexive position in the room - the moments she, too, had to decide whether to

accept, question, or override an AI-generated correction to a student's work. That is what makes this autoethnography rather than a simple field study: the teacher-researcher's own evolving judgement counts as data, not just the students' writing. We recognise it also borrows from participant observation, and we treat the two as complementary rather than mutually exclusive. Informed consent was obtained from every participating student before observation began.

We keep a "linguistic log" of our classroom, kept over two months, with roughly 40 entries recorded in total - some students came up repeatedly, others only once or twice. Every time a student shows us an AI-generated draft, we compare it to their natural speaking voice. This builds on the work of Fan et al. (2024) on how generative AI support changes the way learners work on tasks. We look for the "erasure" of local phrasing. We log an instance whenever the AI swaps out a culturally specific word, a code-mixed phrase, or a metaphor - a personification of nature, say, or a line drawn from a student's first language - for something semantically equivalent but culturally flat. When a student from a multilingual background turns in a "perfect" essay, we look for the missing translanguaging markers that Wei (2018) identifies as signs of deep thought.

Our data collection is simple and direct: we collect "before and after" samples. The "before" is the student's raw handwritten notes or initial brainstorm. The "after" is the final AI-polished submission. Over the two months we collected roughly 15 such pairs; ChatGPT was the AI writing assistant students used. By comparing these, we can see where the algorithm "corrected" a student's cultural framing, addressing the cultural bias concerns raised by Bender et al. (2021). We are not looking only for grammar fixes; we are looking for changes in tone, metaphor, and social rhythm.

We also focus on "intervention dialogues." When we see that an AI has changed a student's unique phrasing into a generic, corporate one, we discuss it with the student. These are not formal interviews but "pedagogical conversations" - well over 80 of these conversations took place over the two months - some students we spoke with more than once, others only briefly - each one written up in notes during or right after the exchange. We want to see whether the student even noticed the change; this helps us track the "proficiency illusion" that Bastani et al. (2025) point to. If a student believes the AI's version is "better" simply because it sounds more "Western," we treat this as a recorded moment of linguistic assimilation.

Our analysis uses the "decolonizing" framework suggested by Birhane (2020) and Mohamed et al. (2020): we examine whether the AI's suggestions consistently push students toward a high-resource language centre. The second author, Muhammad Zain ul Hassan, coded the bulk of the log and samples. The other two authors then cross-checked a subset independently, and where our coding disagreed, we talked it through until we agreed. This is a qualitative audit of the "algorithmic accent" in our own classroom; we are the primary instrument of research here, positioned as mediators between the student and the machine, able to locate the point where personalization stops helping and starts erasing.

This begins to address the gap in "long-term identity" that short-term studies such as Agarwal et al. (2025) leave open - though at two months, it is a first window, not a full longitudinal account. We come back to that limitation in the Conclusion. Even so, it is an honest look at what happens to the educational role of language when the teacher is the only one left fighting for the student's authentic voice.

Results

Our first major observation was the "passive delete." We watched students write a sentence rich with local metaphor, something that tied their heritage to the assignment. One example: a student writing about family responsibilities used the phrase "*meri maa ke haath ka khana*" (an Urdu language sentence meaning “food prepared by my mother's hands”) to describe the emotional importance of home; ChatGPT changed it to "home-cooked meals prepared by my mother." When we discussed the change with her, she explained that the original phrase felt more personal because it reflected a Pakistani understanding of care and family bonding. More broadly, when students ran their writing through a standard LLM, the AI would almost always flag such metaphors as "unclear" or "informal," and instead of arguing with the

machine, the students usually just clicked "accept." This matches the findings of Agarwal et al. (2025) about the narrowing of Global English. The AI isn't just suggesting a better word; it is acting as a gatekeeper for what counts as "academic." We saw that the more a student used the tool, the less they even tried to use their own idioms in the first place. They started "pre-filtering" their thoughts to please the algorithm.

In our intervention dialogues we attempted to interrupt this process. When a student presented us with a polished piece of paper that read like a corporate memo, we would go back to their original notes and ask why they had allowed the AI to modify a certain phrase. In the majority of cases, the student was not even aware that they had lost their voice. They were living in the illusion of proficiency suggested by Bastani et al. (2025): so pleased that the paper had a professional look that they did not notice it no longer sounded like them. This is a dangerous trade-off. It suggests that students are starting to perceive the machine's output as more neutral than their own lived experience. Rosa and Flores (2017) are right to read this as a power struggle: the software is currently winning, since it offers the path of least resistance.

We also followed the translanguaging silence. Wei (2018) claims that the ability to mix languages is an indication of a high-functioning brain, yet the AI we studied approached it as a broken system. The students who were most heavily corrected by the AI in our classroom were often the most fluent in dialect switching, which brought out a real level of frustration. The "*meri maa ke haath ka khana*" (Urdu language) example above is one such case: the phrase was not corrected but replaced outright, leaving no trace of the translanguaging move the student had made. This reinforces the fears of Fan et al. (2024) regarding dependence on AI: when a tool that is meant to assist you only ends up making you feel inferior, it is not really a tool anymore. It is an assimilation filter.

Among the most common discoveries was so-called idiom erasure. One clear case: in a writing task about climate change, a student described the flooding in her village by writing that "the river became angry and entered our homes." ChatGPT proposed "excessive rainfall caused severe flooding in residential areas." Though technically correct, this took away the geographic and cultural soul of the sentence - the personification of nature as an active force that is common in local communities. It is the standardisation of Global English which Agarwal et al. (2025) cautioned about: not necessarily a tutor, but a filter that eliminates anything that does not conform to a mid-Atlantic, corporate norm. Over the course of the observation, students had largely ceased even trying to use local metaphors; they started writing to the machine and hoping that the AI would approve them.

To push back, we began a resistance log, the practical core of what we call Resistance Pedagogy. We asked all 15 students to identify at least three AI suggestions per paper that they would deliberately reject, and to justify why their own phrasing was not an error; roughly half of the rejected suggestions were ultimately kept in the students' final drafts. This proved harder than we expected, given the pressure to sound standard. Birhane (2020) calls for decolonizing these tools, and we understood why it is such an emergency: without a teacher present to say it was okay to sound like that, the algorithm becomes a new form of colonial power - one that needs no whip, only a rephrase button. As our observations suggest, when students were granted the right to be imperfect, they appeared to engage more with the subject matter at hand, although this is an impression drawn from our classroom notes rather than a measured outcome. They had stopped playing to the machine and were writing again.

Discussion

Another gap that we explored was the empathy gap. With creative writing, the AI-polished pieces were technically flawless and emotionally one-dimensional; they lacked the grounding in communicative intent that Bender and Koller (2020) argue meaning requires. When read aloud in class, other students were not able to relate to them - they sounded as if written by a ghost. This confirms what Holmes and Tuomi (2022) argue: that AI does not have the heart to connect with people in an authentic sense.

Another issue, reported by Joshi et al. (2020) and confirmed in our own observations, is the problem of the data desert. Because the models are trained on Western datasets (Bender et al., 2021), the AI was considerably more accommodating toward students who already used a standard dialect, and functioned as an incessant critic for linguistically disadvantaged and marginalised students. This supports the point Doshi and Hauser (2024) make about homogenisation: in attempting to create a tool that everyone could use, developers have succeeded only in making everyone sound like a particular, privileged someone.

This brings out a greater moral debate. According to Birhane (2020) and Mohamed et al. (2020), there is a need to decolonize AI ethics. In our classroom, in the absence of teacher intervention, the AI functioned as a kind of digital colonial agent, imposing a "correct" way of speaking rooted in the Global North. We proposed Resistance Pedagogy to push back against this: students were encouraged to reject AI suggestions and defend their original wording. We found that students who justified their non-standard choices this way became more invested in their own writing - they were no longer simply producing text, but reasserting their authority over it.

At the end of this discussion, one conclusion stands out: the algorithmic accent is not a necessity, but a choice. When we leave the machine as the ultimate judge of good writing, we are agreeing to a homogenized linguistic future. But when we treat the machine as a mirror instead - a device that helps students recognise the gap between their own voice and the so-called standard one - this same tension can become an Intercultural Communicative Competence (ICC) lesson. As Bender and Koller (2020) imply, a system trained on form alone can supply the vocabulary, but only a human supplies the meaning. A student's "accent," literal or metaphorical, should not be refined away, the machine can never really imitate everything that it represents.

Conclusion

Over two months, we have been observing the smoothing of the jagged, beautiful lines of student identity by the redline of the algorithm. The question that started this study was simple: what happens to the human voice when the main editor is a machine? We have an indication that we are experiencing a silent, structural erosion. The algorithmic accent is not merely a technical anomaly; it is a pedagogical one. We found that students had swapped their cultural authority for digital speed, producing a proficiency illusion that conceals a deeper loss of voice. The AI presents the appearance of a standard, but in practice suppresses the translanguaging and local idioms that Wei (2018) and Agarwal et al. (2025) identify as markers of real cognitive and cultural richness. Two months is a first step, not a full longitudinal account, and we see what we found here as reason enough to follow it up over a longer stretch of time.

The research gap we outlined at the start of this paper remains the most pressing issue for educators. Most studies still quantify output and disregard outcome - they track how many words a student can produce, not how many of those words the student still feels are theirs. This gap is what our two-month observation begins to bridge, by pointing to the psychological and sociolinguistic cost of long-term AI reliance. According to Agarwal et al. (2025) and Doshi and Hauser (2024), the push of these models is toward a bland, Western-centric average - a new form of linguistic homogenisation. Unless we act, we risk raising a generation that understands their native tongue or local dialect only as something to be fixed with the press of a button.

We need a way out, which we call Resistance Pedagogy. It is neither realistic to prohibit these tools nor advisable to defer to them uncritically. Instead, following Birhane (2020) and Mohamed et al. (2020), we argue for actively decolonizing the way AI is used in the classroom. The teacher's role should shift from grammar-grader to voice-protector. Students should be encouraged to ask why the machine wants to turn their untidy metaphors into clean cliches. We should celebrate the fact that they write with an accent, rather than treat it as a 10-20% syntax error to be fixed. As O'Neil (2016) cautions, algorithms will not become equal-opportunity until we make deliberate choices to hold them to that standard.

After all, language is one of the clearest ways we demonstrate identity. When the algorithm makes the final decision about what is right, the pedagogical purpose of language teaching is lost. We should make sure that personalised education does not, in practice, lead to a world where everyone sounds alike. Our analysis suggests that the greatest skill we can teach in the age of AI is the willingness to be imperfect rather than polished. Education should not aim to produce a flawless product, but to help a student discover a voice that is unmistakably their own.

References

Agarwal, D., Naaman, M., & Vashistha, A. (2025). AI suggestions homogenize writing toward Western styles and diminish cultural nuances. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM. https://doi.org/10.1145/3706598.3713564

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences, 122*(26), e2422633122. https://doi.org/10.1073/pnas.2422633122

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM. https://doi.org/10.1145/3442188.3445922

Bender, E. M., & Koller, A. (2020). Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 5185–5198). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.463

Birhane, A. (2020). Algorithmic colonization of Africa. *SCRIPTed, 17*(2), 389–409. https://doi.org/10.2966/scrip.170220.389

Cenoz, J., & Gorter, D. (2021). *Pedagogical translanguaging*. Cambridge University Press. https://doi.org/10.1017/9781009029384

Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances, 10*(28), eadn5290. https://doi.org/10.1126/sciadv.adn5290

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2024). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*. https://doi.org/10.1111/bjet.13544

Giroux, H. A., & Penna, A. N. (1979). Social education in the classroom: The dynamics of the hidden curriculum. *Theory & Research in Social Education, 7*(1), 21–42. https://doi.org/10.1080/00933104.1979.10506048

Holmes, W., & Tuomi, I. (2022). State of the art and practice in AI in education. *European Journal of Education, 57*(4), 542-570. https://doi.org/10.1111/ejed.12533

Joshi, P., Santy, S., Budhiraja, A., Bali, K., & Choudhury, M. (2020). The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 6282–6293). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.560

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

O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown Publishing Group.

Rosa, J., & Flores, N. (2017). Unsettling race and language: Toward a raciolinguistic perspective. *Language in Society*. https://doi.org/10.1017/S0047404517000562

Wei, L. (2018). Translanguaging as a practical theory of language. *Applied Linguistics, 39*(1), 9-30. https://doi.org/10.1093/applin/amx039