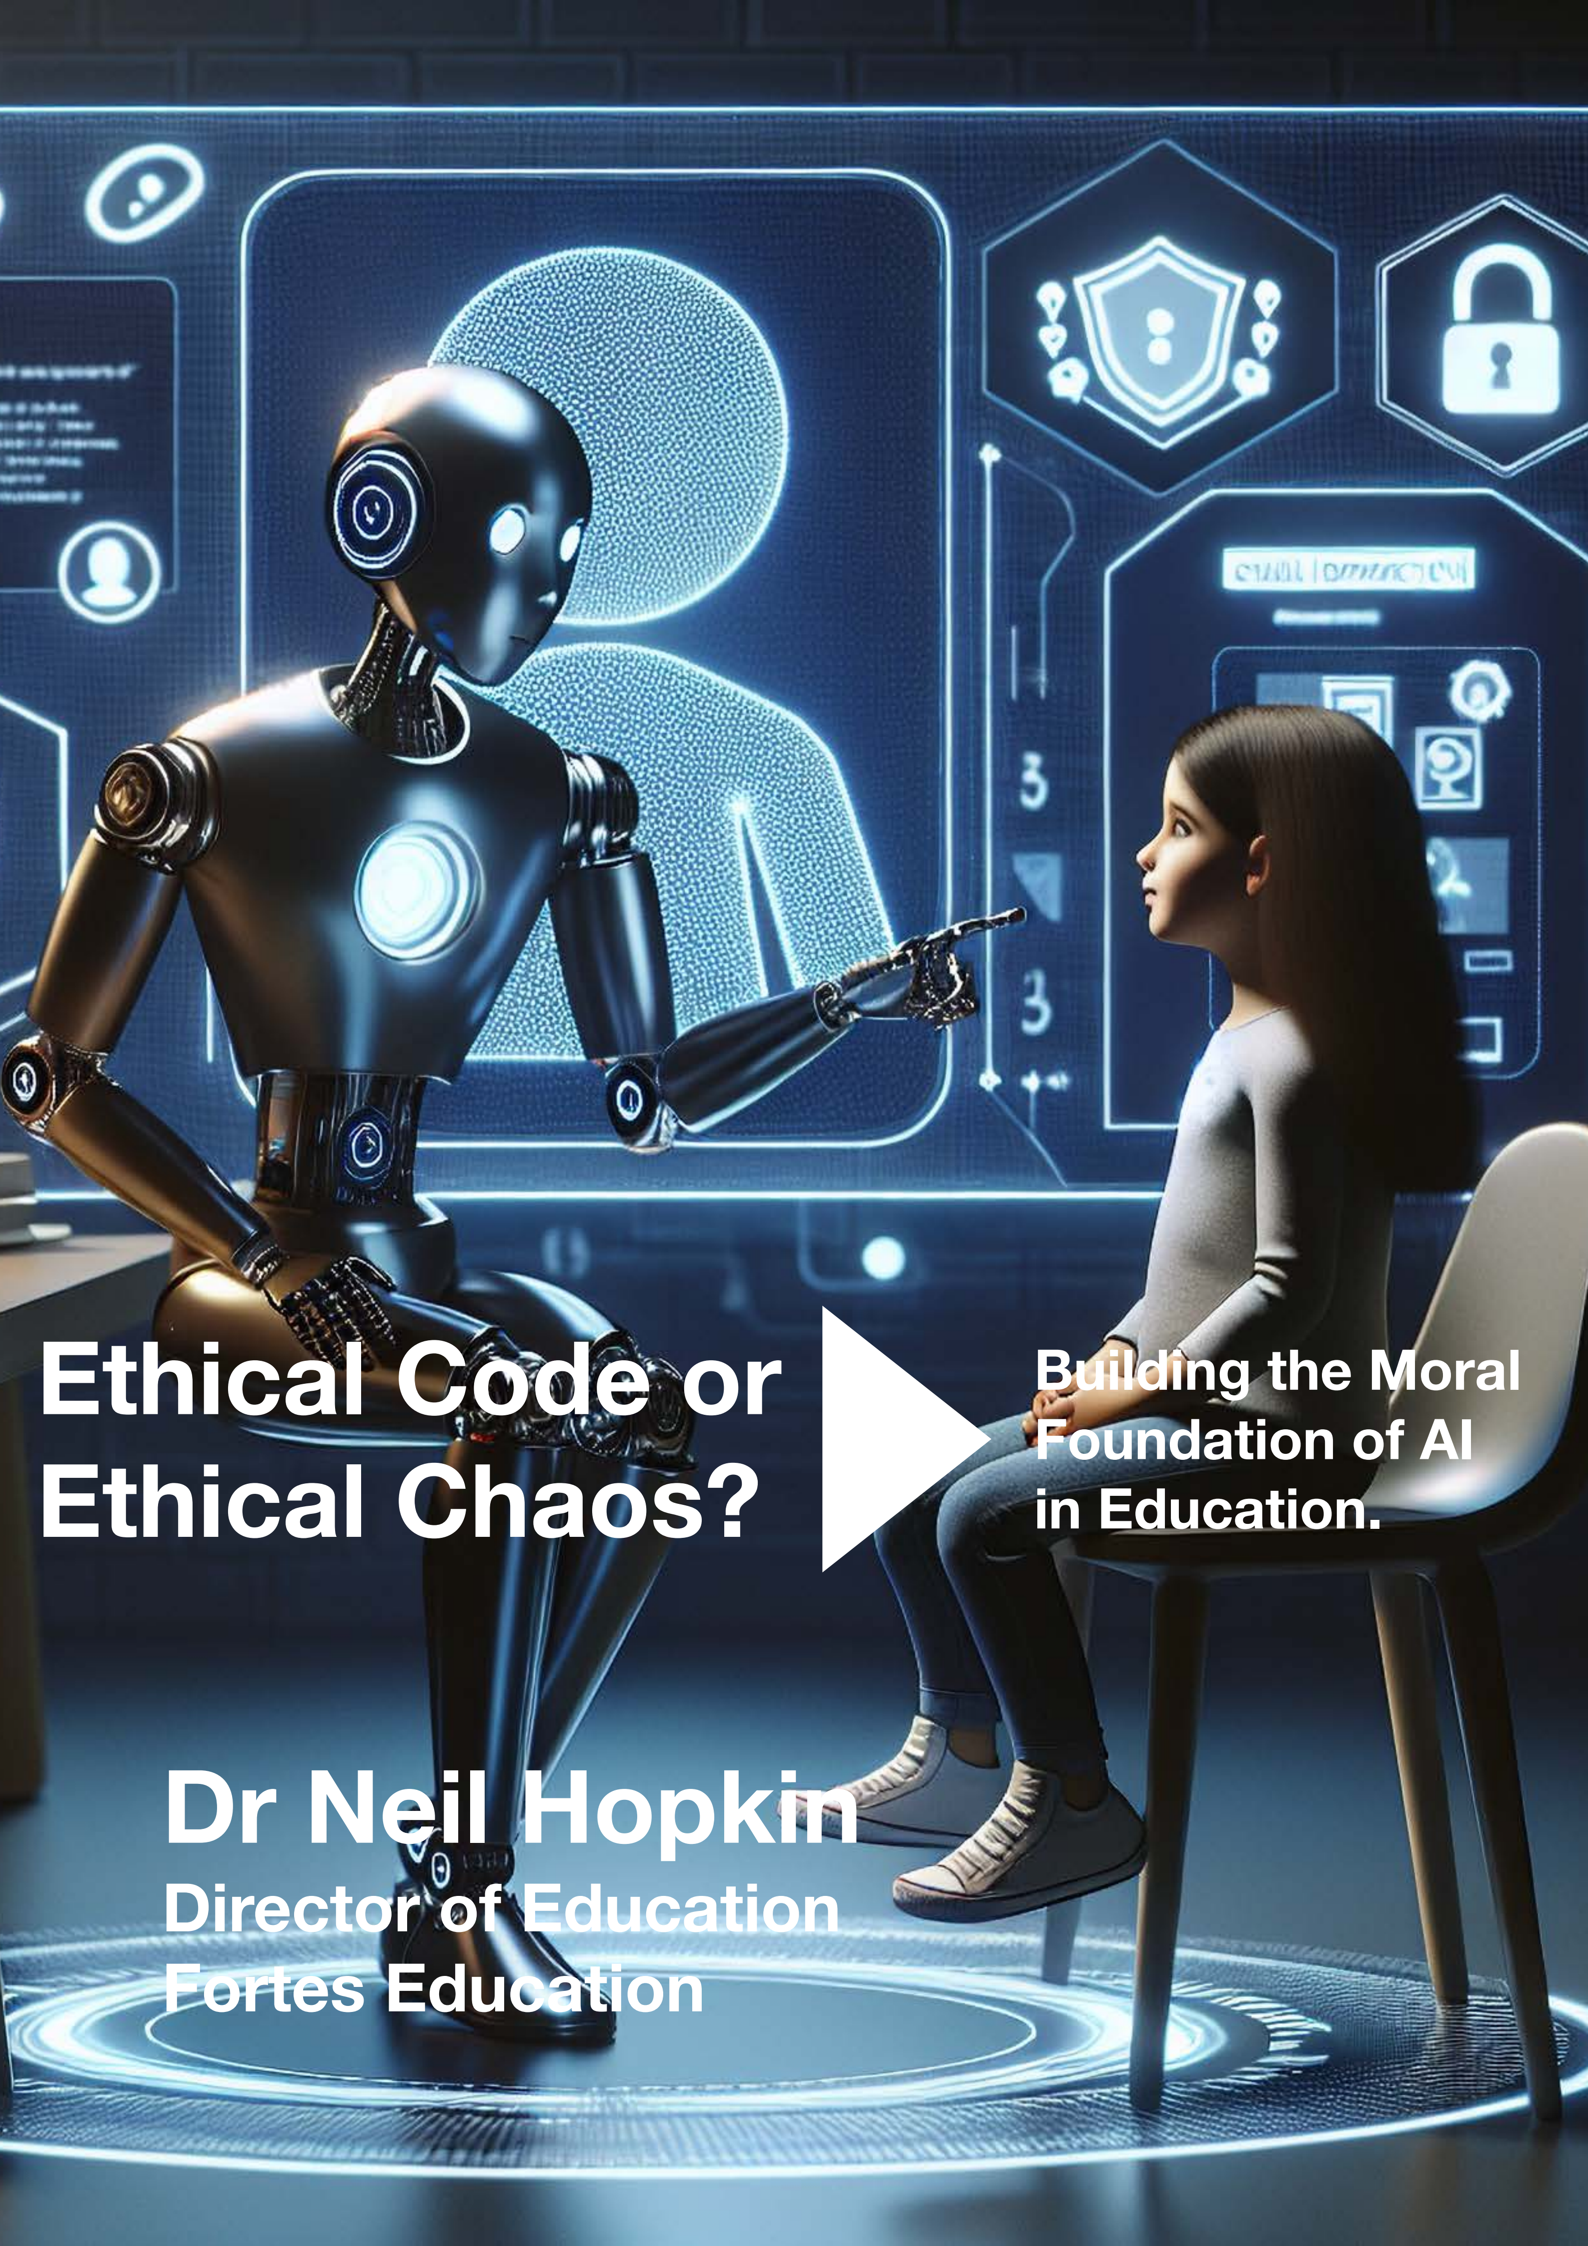

Ethical Code or Ethical Chaos?

Building the Moral Foundation of AI in Education.

Dr Neil Hopkin
Director of Education
Fortes Education

Introduction: Ethical Code or Ethical Chaos? Building the Moral Foundation of AI in Education

A few years ago, automated essay scoring tools began quietly filtering into schools, promising faster feedback and "fairer" marking. At first glance, it seemed a logical step: if artificial intelligence could streamline the slog of grading, why not embrace it wholeheartedly? Yet, behind those tidy claims lay unsettling questions. What if the model consistently underestimated creative or culturally grounded writing styles? What if it disfavoured a student's phrasing because of dialect or second-language usage? Who, if anyone, had verified that these automated judgments aligned with the values of a diverse learning community? As soon as educators pressed for specifics, they discovered corporate secrecy: developers guarded their algorithms, claiming proprietary advantage. That tension—between convenient automation and the messy moral realities of education—captures the core dilemma this article seeks to unravel.

AI's expansion in education was, from the beginning, propelled by rosy visions. By harnessing data analytics, personalisation, and real-time insights, proponents imagined we could tailor lessons precisely, lighten teachers' workloads, and guide each learner along an optimal path. Yet as Ruha Benjamin (2019) has noted, technology introduced into social systems seldom leaves power structures untouched. In the realm of schooling—where young minds develop literacy, social awareness, and moral compasses—AI can shape more than reading levels or arithmetic drills: it can influence self-esteem, identity, and even the latent expectations of who gets to succeed. Recognising this broader impact, some educational authorities rushed to adopt "ethical guidelines," but found them vague or riddled with corporate disclaimers. The question remained: where does the ultimate responsibility lie? Who ensures that data is gathered justly, who decides the parameters of personalisation, and what moral frameworks govern these decisions?

For decades, ethical technology discourse often revolved around Nick Bostrom (2014) and the existential risks of superintelligent AI. But in schools, the immediate concern is not a rogue AI taking over the planet—it's an automated system that might quietly disadvantage certain groups, or commodify personal data for profit. Timnit Gebru (2020) and Joy Buolamwini (2018) have both demonstrated how algorithmic bias can infiltrate face recognition tools, reinforcing stereotypes and misclassifying individuals from certain ethnic backgrounds. Translate that scenario into an educational setting: an AI might misread dialectical English or label certain cultural expressions as "incorrect," compounding inequalities. The moral stakes—especially when these systems become de facto gatekeepers to scholarships, advanced courses, or critical interventions—are enormous.

Within what Shoshana Zuboff (2019) has termed "surveillance capitalism," data on students' learning habits, emotional responses, and domestic backgrounds can be turned into commodities. Such data isn't always tracked in an overt, malicious sense. It might be as simple as capturing how often a student visits a reading platform, how quickly they answer a question, or which cultural references they respond to. Yet the potential for commodification—selling analytics to textbook publishers or commercial sponsors—lurks. Virginia Eubanks (2017) has documented how automated systems in welfare and public services can deepen exclusion for marginalised populations; a parallel risk is that AI in education could replicate a "digital poorhouse," allocating fewer resources or more rigid guidance to underperforming students. That outcome would be antithetical to any egalitarian vision of schooling.

Hence the impetus for ethical codes. If we can articulate guidelines—much like codes of conduct for doctors or lawyers—then perhaps teachers, administrators, and developers can all push back against exploitative or harmful AI uses. But codes are only as strong as their implementation. As Kate Crawford (2021) highlights, lofty statements about "accountability" can ring hollow if the underlying business model remains opaque, or if third-party audits are refused. The tension, then, is whether the drive for an ethical foundation can outpace the profit and data-harvesting imperatives of a booming ed-tech industry. Meredith Broussard (2018) calls this a mismatch between real human needs—personal growth, cultural respect, teacher autonomy—and the slick marketing narratives of "artificial intelligence."

The moral questions don't stop at bias or privacy. If generative AI can supply instant answers, or craft entire essays, how do we preserve the role of struggle, curiosity, and genuine discovery? The tension is reminiscent of what N. Katherine Hayles (2012) describes as a posthuman moment: where cognitive processes shift from human minds to networked code. In education, that shift surfaces daily. Students might rely on AI tutors for homework, never quite learning to reason from first principles. Teachers might offload planning to algorithmic recommendations. This "ease" could undermine the deeper moral growth that emerges when teachers guide students through adversity and reflection. A code that ensures humans remain in the loop is essential, but whose responsibility is it to enforce that?

Simultaneously, there's an undercurrent of techno-optimism. Jaron Lanier (2010) has long argued that, shaped conscientiously, digital platforms can empower creative expression and cultural exchange. In the educational arena, this might look like harnessing VR to let students role-play historical events, or employing natural language processing to give immediate feedback on creative writing. If the moral foundation is well constructed, such experiences might be truly inclusive—pulling in local knowledge from indigenous communities, recognising dialect diversity, celebrating different learning paces. But it demands concerted vigilance against one-size-fits-all solutions that flatten cultural nuances or marginalise unique voices.

Audrey Watters (2015) has offered a stark critique of "the history of the future of education," emphasising how profit motives and simplistic narratives of "disruption" overshadow the real complexities of teaching and learning. Her caution extends to AI: if deployed without rigorous ethics, these systems can commercialise classroom relationships, standardise creativity, and render teachers' expertise secondary to algorithmic logic. Add in Paulo Freire's (1970) notion of dialogue-based liberation, and we see a clash: can a structured AI truly accommodate the back-and-forth, teacher-student co-creation of meaning? Or will it inevitably channel learners into preset tracks that might rob them of autonomy?

To navigate these concerns, some educators and theorists invoke bell hooks (1994) on engaged pedagogy—emphasising love, mutual respect, and cultural affirmation. Her perspective suggests we need more than technical fairness: we need an AI environment that fosters empathy, critical consciousness, and communal ties. Meanwhile, Henry Giroux (2011) frames education as a cultural battleground, contending that if AI is left unregulated, corporate or political agendas could shape children's worldviews unnoticed. The moral foundation must thus address not only data privacy or bias, but also the broader cultural politics of who controls knowledge production.

Activists like Joan Donovan (2018) point out that civil society must be part of technology governance, rather than relegating oversight to tech executives or bureaucrats alone. For AI in education, that means parents, students, teachers' unions, and local communities need a voice in how these tools are designed, tested, and implemented. Without such a grassroots dimension, ethical frameworks can become mere window dressing, powerless against entrenched economic interests or political power plays. The last thing schools need is another top-down mandate that fails to reflect real classroom dynamics.

So how do we build a moral foundation that is neither vague nor oppressive? Helen Nissenbaum (2010) proposed "contextual integrity," arguing that privacy and ethics must be tailored to the specifics of a social domain. In education, that context includes not just test scores and teacher evaluations, but also intangible values like trust, emotional safety, communal identity, and the personal transformations that occur in adolescent years. Thus, an ethical code might require that AI tools be tested in real learning environments, with teachers and students actively evaluating whether the system supports or hinders the relationships that define good education.

Sandra Harding (1991) would add that standpoint theory can help: historically marginalised communities should define the priorities of such a moral code, ensuring it doesn't just mirror the worldview of Silicon Valley developers or global corporations. This approach resonates with critiques from indigenous scholars who see AI-based ed-tech as yet another wave of cultural erasure if local pedagogies are ignored. The universal code, in other words, must have a built-in mechanism for local adaptation. If a code states, for instance, that data collection must be minimal, how is "minimal" defined in a remote indigenous context lacking robust digital infrastructure? If a code demands culturally sensitive content, who is designated to judge content appropriateness?

At the same time, a purely relativistic approach—where every community crafts its own code—could degrade into ethical chaos, especially for global ed-tech vendors. Luciano Floridi (2013) insists that robust information ethics can balance universal principles (respect, transparency, accountability) with context-specific implementations. We might see frameworks akin to UNESCO guidelines or the EU AI Act, adapted for schooling. The code might articulate baseline values: no exploitative data gathering, no manipulative "nudges" that override learners' autonomy, and a guaranteed teacher oversight. Then local bodies can augment these guidelines with culturally specific rules.

Langdon Winner (1980) famously argued that "artefacts have politics," meaning the architecture of technology can embed social hierarchies. AI in classrooms is no exception. If the moral foundation is not specified, subtle design choices—like default language or recommended content—can shape cultural attitudes. For instance, an AI that predominantly sources Western authors might marginalise local literature, reinforcing a monolingual or monocultural notion of excellence. Conversely, a well-crafted ethical code could mandate robust representation of local languages, ensuring that a child in Nigeria or Malaysia sees their linguistic heritage validated in AI-driven lessons.

Meanwhile, Neil Selwyn (2019) critiques the hype in educational technology, urging a more measured approach that highlights teachers' professional expertise. He reminds us that no AI is a silver bullet: the intangible human element—those flashes of empathy and improvisation when a teacher spots a student's frustration—remains vital. So the moral foundation might emphasise the teacher as a decision-maker, not just a passive user of AI analytics. This includes the right to override algorithmic suggestions, the right to query the code's logic, and the duty to inform students about how AI shapes their learning experience.

One could imagine an iterative cycle: a school receives a new AI platform, teachers run it in a pilot setting, a local ethics board (including parents and possibly older students) reviews logs for questionable patterns, and the developer is required to refine the system. Over time, such cyclical feedback fosters a "living" ethical code—never static, always adjusting to fresh realities. That process points to the messy but essential domain of Joan Donovan (2018) again—public engagement. If the moral foundation remains the domain of experts alone, the social dimension that truly defines schooling is lost.

Thus, the stage is set for a multi-layered analysis in the chapters to come. We will investigate how algorithmic bias intersects with educational equity, how privacy controversies might erode trust, and how data governance frameworks might—or might not—defend student autonomy. We will trace the tension between big ed-tech's desire for profit and the public good demanded by teachers and communities. Through it all, we will circle back to the same core question: can there be a robust "ethical code" that is practically enforced, or are we careening toward "ethical chaos" under the weight of conflicting interests?

No doubt, the stakes have never been higher. Education shapes the next generation's worldview; if AI tips that scale in hidden ways, entire societies might inherit those distortions. Ruha Benjamin (2019) warns about the illusions of neutrality in technology, while Timnit Gebru (2020) reminds us that diverse, inclusive teams must lead AI's development. The articles ahead aim to forge a broad consensus on how to ensure that an "AI in education" revolution is anchored in moral depth rather than ephemeral commercial logic. If we succeed, we might set a precedent for other industries, demonstrating that advanced technology can operate under a transparent, equitable framework that honours the dignity and potential of every learner. But achieving that will require steady collaboration, critical reflection, and a refusal to succumb to the chaos that can arise when ambition outstrips our ethical imagination.

The Emerging Ethical Tensions in AI-based Education

When an education official in a large urban district first proposed using AI-driven "dropout predictors," many teachers were relieved: it sounded as though they might finally catch struggling students early. Yet, the moment someone enquired how the system arrived at its risk scores—on which data it relied, and whether it might inadvertently label students from certain neighbourhoods as "high risk"—administrators grew vague (Reich, 2020). On closer inspection, this dual promise and peril of AI-based education reveals a profound set of ethical tensions. Tools meant to lighten teachers' burdens and personalise student paths simultaneously raise deeper questions about who holds power, whose values underpin the algorithms, and how young minds may be shaped by invisible digital processes.

Schools are historically overloaded with tasks. Teachers must manage large class sizes, standardised tests, administrative paperwork, and the varying paces of student learning. It is no wonder that a suite of AI-based solutions has made such an appealing entrance: from automated homework grading (Kohn, 2015) to dynamic timetabling (Collins, 2018). Advocates claim that if we let algorithms handle routine tasks, teachers can focus on creative or pastoral roles. This argument draws on broader debates about how automation can release humans from drudgery (Bryson, 2019). But the moral dimension emerges when we examine what counts as "routine." Are we certain that personal feedback or subtle cues—like a student's anxiety in class—can be cleanly separated from mere "marking tasks"? Once a system automates chunk after chunk of teaching labour, do we risk diminishing the relational bonds and intangible teacher judgement that many see as essential to education?

Moreover, the impetus for AI solutions often arises not from educators themselves but from policy goals shaped by data-centric visions of accountability. Since the late 1990s, standardised testing regimes have popularised the notion that educational performance can be measured quantitatively, spurring a climate where tech solutions promise "objective" analytics (Chomsky, 2017). Schools under pressure to raise test scores or reduce absentee rates might leap at AI dashboards. Yet critics like Selwyn (2019) warn that such dashboards can reframe student life as a collection of metrics. Rather than a space for moral growth or discovery, the classroom becomes a site where each child's "learning velocity" is tracked and compared, potentially overshadowing intangible qualities like empathy or resilience.

One of the clearest ethical dilemmas emerges around decision-making authority. Teachers who adopt AI-based recommendations for interventions—say, deciding which pupils receive extra tutoring—may inadvertently cede professional judgement to an opaque model. The concept of "teacher empowerment" can invert if the teacher now feels compelled to follow algorithmic advice. This phenomenon parallels insights from Bryson (2018), who notes that even well-intentioned automation can erode human agency if people begin trusting machines as impartial arbiters. Instead of educators personally knowing each pupil, relationships may be filtered through a system's lens, which classifies some as "gifted," others as "at risk," occasionally reinforcing stereotypes (Chowdhury, 2020).

The question is not simply whether AI is good or bad, but how teachers and administrators navigate these tools. Tensions arise if teachers cannot interrogate the underlying assumptions. Suppose an algorithm identifies a correlation between single-parent households and lower reading scores, then lumps all such students into a "high-risk" bracket (Ben Green, 2019). The teacher might intuit that this one-size-fits-all approach is misguided, yet if the software is integrated district-wide, ignoring its advice may become a career risk. At the macro level, the tension is: does adopting an AI-based model strengthen educational equity—or replicate systemic biases under the banner of objectivity?

Another crucial tension involves the data AI systems rely on. Modern ed-tech platforms collect not only test scores but also keystroke dynamics, eye-movement data in VR settings, emotional expression analyses, and more (boyd, 2014). On the surface, such data could refine personalisation: a system might sense boredom and interject a new approach. However, the more personal the data, the greater the intrusion (McKee, 2017). Are we comfortable letting an AI track subtle emotional cues for each pupil? Who stores that data, and for how long? Does the child (or parent) have a right to expunge it after a certain time?

In many contexts, the commercial nature of these platforms raises alarm. If a private vendor obtains real-time emotional analytics from students, it may glean advanced marketing insights—like how best to shape content for engagement. Eubanks (2017) warns that in vulnerable communities, data-driven systems can amplify disadvantage, especially if monetised or shared with third-party companies. The tension is between innovative personalisation and the spectre of "surveillance capitalism" (Zuboff, 2019). Without robust ethical guidelines and public oversight, these platforms could become data goldmines, with children as unwitting test subjects.

Education inevitably contains a paternalistic dimension: teachers guide minors who may not yet grasp long-term consequences. AI can intensify this paternalism if it "nudges" students toward certain subjects or reading materials, presumably for their benefit (Warschauer, 2016). Yet what if those nudges conflict with a learner's genuine interests or cultural identity? At times, paternalism may be appropriate—like steering a child away from harmful content. But friction arises if the system's logic classifies certain cultural expressions as deviant or less academically relevant. Consider a scenario where a pupil's creative writing is overshadowed by standard language norms in the AI's model, which flags particular dialectical expressions as "erroneous" (Noble, 2018). The student might internalise that judgement, losing confidence in their home language or community identity. Freed from human nuance, paternalism can morph into a silent assimilation tool.

The notion that AI-driven decisions are more "objective" than human ones also draws attention. A wealth of scholarship—Gebru (2020) and Buolamwini (2018) among others—shows how training data can embed structural biases. In education, the dataset might overrepresent certain socio-economic or linguistic groups, shaping the AI's assumptions of what "normal progress" looks like. Or the algorithm may rely on historical data that reflect discriminatory practices. The net effect is an AI that appears neutral while replicating biases (O'Neil, 2016). Some, like Kohn (2015), question whether this chase for objectivity overrides more vital subjective dimensions in education, such as moral growth or creative exploration.

On a practical level, these systems typically come from ed-tech firms seeking profit. That raises the risk that business considerations overshadow deeper ethical reflection (Morozov, 2013). Districts strapped for funds or lured by philanthropic grants may adopt AI tools rapidly, sidelining teachers' concerns or local cultural context. Arora (2019) calls this dynamic "digital colonialism," where universal solutions overshadow local pedagogies. The rhetorical tension is whether schools see themselves primarily as consumers of a product or as co-creators of an educational environment with public accountability.

Another dimension of educational equity extends beyond race or class to cultural sovereignty. Many indigenous communities, for instance, transmit knowledge orally and weave it with land-based spiritual practices (Smith, 1999). A standard AI might interpret knowledge in text-based, compartmentalised ways, losing the holistic or relational aspects crucial to indigenous epistemologies. Without careful guidance, generative systems might caricature local lore into simplistic mini-games, raising ethical questions about appropriation and misrepresentation.

At first glance, teachers appear overshadowed by sleek AI dashboards that track progress, generate quizzes, and even answer students' queries. Yet some argue that these dashboards can free teachers to become mentors, focusing on emotional and conceptual development (Collins, 2018). The tension arises if an under-resourced system expects fewer certified teachers, leaning on AI for "efficiency," thereby reducing teacher autonomy. Nissenbaum's (2010) principle of contextual integrity suggests teachers need domain-specific norms for how AI data is used. A teacher might want detailed analytics only for formative feedback, not for final grading. If the system's default is to feed data into a standardised scoreboard, the teacher's professional insight is undermined.

As AI grows more sophisticated, it might propose "nudges" to keep learners on task or suggest new challenges. While behavioural economics frames nudging as helpful (Sunstein, 2014), in an educational setting, moral complexities abound. Could these nudges quietly shepherd all students into prioritising STEM subjects over the arts? That choice might be beneficial for future careers but inadvertently erode cultural diversity in skill sets. The ethical question is who defines "optimal" learning paths, and whether students or communities have space to contest algorithmic suggestions.

In recent years, attempts to draft AI-ethics guidelines for schools resemble the general-purpose statements from tech companies promising fairness, transparency, and accountability (Crawford, 2021). Critics point out that these assurances can mean little without real enforcement or local adaptation. If teacher unions or parent associations lack technical literacy, they might accept a superficial compliance statement. Alternatively, a robust public framework—similar to institutional review boards in healthcare—could require that any AI introduced in classrooms undergo a thorough ethical review, with parents and teachers empowered to veto designs that conflict with communal values.

Ultimately, the tension swirling around AI-based education is not purely technical but deeply human. Teachers, parents, and students find themselves at an intersection: they crave the promise of expanded possibilities and reduced workload, yet fear the potential for standardisation, data exploitation, and cultural disregard. Without an overt moral foundation to guide design and deployment, a silent shift may occur where corporate aims or hidden biases reshape schooling beyond recognition. If ethical chaos is to be avoided, it will require forging a genuinely collaborative process, one that marries local knowledge, teacher autonomy, and inclusive technologies in a collectively defined code of AI conduct. The following question is whether such a coherent moral framework can truly be developed and enforced in our rapidly evolving digital landscape, or if these tensions will merely deepen as schools rush headlong into an AI-driven future.

Algorithmic Bias – Historical Legacies and Technical Realities

Algorithmic bias in educational AI emerges from a tangled web of historical injustices, technical oversights, and the uncritical allure of "neutral" computation. Many of today's advanced systems rely on machine-learning algorithms that interpret student data to generate predictions—everything from reading comprehension scores to behavioural risk flags. But these predictions, however sophisticated, can reproduce the very social divides they purport to address. The contrast is striking: on the one hand, AI-driven analytics promise to expand personalisation and equity, yet the technical pipelines often reflect cultural assumptions and historical imbalances that have dogged education for centuries (Green, 2019).

One cause is the reliance on datasets shaped by biased institutions. If past academic records systematically undervalued certain dialects, penalised immigrants, or segregated educational opportunities, the patterns embedded in that data might unfairly "teach" the AI to regard particular linguistic forms or cultural backgrounds as lower performing (Ben Green, 2019). When an algorithm infers risk based on historically biased data, it can unintentionally push teachers to concentrate resources on an already privileged cohort, or label entire communities as underachieving. This phenomenon parallels insights from Joy Buolamwini (2018), who found that face-recognition tools trained on predominantly light-skinned datasets failed to detect darker-skinned faces accurately. In an educational environment, the consequences run deeper: a system might misinterpret a child's subtle reaction or creative style, entrenching harmful feedback loops and effectively gatekeeping advanced classes or scholarships.

Timnit Gebru (2020) has emphasised the importance of inclusive AI development teams to forestall such biases. If only one demographic steers data curation and algorithm design, they might overlook how certain traits manifest in different cultural contexts. For instance, a neural network that monitors "engagement" by measuring how often a student speaks up might not accommodate culturally diverse norms about classroom participation (Gandy Jr., 1993). It might even classify respectful silence—common in some communities valuing quiet introspection—as disinterest, resulting in skewed recommendations. The tension here isn't about ill intent, but about the mismatch between local cultural realities and the assumptions coded into the model's training phase.

Historical legacies also linger in subtle forms. Leonard Wantchekon's (2018) studies on colonial education show how entire syllabi once legitimised certain languages while erasing others. If an AI's language module regards standard English as the only valid reference point, it may mark expressions from Jamaican patois or African American Vernacular English as errors. Students end up internalising that judgement, believing their home language lacks academic legitimacy (Noble, 2018). Under the veneer of efficiency, the system perpetuates a hierarchy of linguistic capital. Cathy O'Neil (2016) notes that such "weapons of math destruction" can produce cyclical harm: each time the AI signals a "deficit," the student is funnelled into remedial tracks. Over time, the algorithm becomes a self-fulfilling prophecy, reinforcing the disenfranchisement it purports to address.

Moreover, biases rarely appear in a singular dimension. Rumman Chowdhury (2020) argues for intersectional awareness: a system might treat a low-income female student of colour differently than a low-income male student of the same background, because it unintentionally merges multiple social signals that intensify stereotypes. If the AI lumps households with certain postcodes into high-risk categories, also factoring in free-lunch status and grammar patterns, the result can be an outsized alarm that leads to disproportionate scrutiny (Arvind Narayanan, 2018). Teachers, if not equipped to question those red flags, may second-guess the potential of pupils from these communities—harking back to older forms of discrimination that we hoped data-driven analytics would dissolve.

At a technical level, the concept of "fairness" is often oversimplified into numeric definitions—like ensuring equal false-positive rates across groups. While such metrics can help, Dorothy E. Roberts (2011) reminds us that deeper structural issues require more than statistical parity. Historical disenfranchisement cannot be undone by merely balancing error rates; it demands transformative approaches that question the very categories the AI uses. Kim TallBear (2013) similarly critiques how Western data classification can reduce indigenous identities to checkboxes, ignoring more relational or land-based ways of knowing. So even with good intentions, an educational AI might treat a student's background as a risk factor, never delving into the cultural wealth that background signifies.

The belief in "neutral technology" also feeds the illusion that an AI's verdicts are somehow above social bias (D'Ignazio and Klein, 2020). Teachers who see a refined interface or pseudo-scientific label can fall into the trap of regarding it as infallible. Michael Macy (2017) explores how group perceptions shift when a technology is framed as scientifically objective. In practice, many educational AI tools do not reveal their data sources or assumptions. The result is that frontline educators might uncritically follow the system's suggestions on course placement or discipline—invoking a potent form of algorithmic hegemony. In such scenarios, the teacher cedes professional judgement, and the pupil's unique context is overshadowed by generic predictions that might be anchored in historical prejudice (Wachter, 2018).

Archaic forms of standardised testing, originally designed with eugenicist beliefs in mind (Tom F. Green, 1970), exemplify how biases can become institutional norms. As new AI systems ingest historical exam results, they incorporate these same embedded norms unless forcibly corrected. Annie Jean-Baptiste (2020) advocates that tech creators explicitly track how training data is formed, ensuring that test-based or performance-based records are not decontextualised from their socio-political moorings. For instance, if a school district historically underfunded certain neighbourhoods, pupils' test data from those areas might reflect under-resourced teaching rather than personal capacity. Yet an AI oblivious to that difference might codify it as "low ability." Over time, the system funnels extra resources to those it considers promising, ironically exacerbating the imbalance.

Some attempt to mitigate bias by adding "awareness modules" or post-hoc audits. Yet Zeynep Tufekci (2018) points out that once technology is deployed at scale, it becomes entrenched in institutional habits. Educators rely on it for daily tasks, and district administrators seldom revisit fundamental design decisions. A more proactive remedy might be found in the approaches of Virginia Dignum (2019), who calls for "responsible AI" from the outset. Instead of grafting fairness onto a finished model, responsible AI demands inclusive co-design, diverse user testing, and continuous iteration where flagged biases lead to immediate model retraining. That approach parallels what Mireille Hildebrandt (2020) describes as due process in algorithmic systems: an enforceable standard requiring transparency, contestability, and redress mechanisms for students or teachers.

Yet these solutions are not easy to implement. Budgets, timelines, and the complexity of AI architecture can deter thorough bias checks (Ben Green, 2019). Additionally, AI vendors might guard proprietary details, claiming that revealing them would expose trade secrets. Teachers, especially in under-resourced areas, cannot easily mount legal or technical challenges. Oscar H. Gandy Jr. (1993) discusses the "panoptic sort" in which data classification quietly shapes individuals' life chances, absent robust oversight. Transposing that logic to schools illuminates a sobering reality: if a pupil's "academic identity" is determined by an opaque system trained on historically skewed data, we risk perpetuating a hidden caste-like system of digital sorting.

In the quest to address these dilemmas, various proposals emerge. Some call for open-sourcing educational AI models—ensuring that local communities can peer into the code, edit weighting factors, or tailor the features so they reflect local knowledge rather than universal norms (Chowdhury, 2020). Others propose that each major AI deployment in schools must pass an ethical review board comprising parents, educators, civil rights experts, and data scientists. Such a board might veto an algorithm that lumps special education students into sweeping risk categories or that equates certain language forms with deviance. This recall echoes older arguments from Tom F. Green (1970) that moral judgement in education cannot be outsourced to raw metrics.

Though these proposals seem promising, they will succeed only if we acknowledge and dismantle the centuries-old prejudice that shaped educational data in the first place. It is not enough to tweak an algorithm's fairness function if the entire knowledge pipeline—test design, curriculum selection, teacher training, resource allocation—already exhibits bias. Algorithmic bias in schools thus becomes a microcosm of larger socio-political struggles. When an ed-tech firm touts its "AI for equality," the real question is whether they have systematically engaged with historically marginalised voices, whether they have re-examined how training data was generated, and whether they allow meaningful oversight into the system's daily operation.

In the final analysis, algorithmic bias in education underscores how new technologies often magnify old inequities unless carefully reined in. The tension is not about halting AI altogether but guiding it within a moral framework that foregrounds diversity, historical awareness, and local authority. Rather than wholly rejecting predictive analytics, many suggest we couple them with context—a teacher's or community's lived knowledge—so that no numeric label is treated as destiny. That synergy might pave a path where generative AI becomes a tool for unveiling injustice rather than perpetuating it: shining a bright light on the hidden predispositions in historical records, and prompting educators to address them intentionally. But if we continue in a naive hope that data is destiny, we risk enthroning the biases we hoped to eradicate, letting them hide behind the persuasive veneer of an AI's computations.

Privacy, Autonomy, and Data Governance in Ed AI

When schools first adopted automated proctoring systems to prevent cheating, proponents hailed it as a solution to teacher burnout. The software would scan a student's gaze, posture, and audio to detect suspicious behaviour, freeing educators from hours of manual invigilation. Yet beneath the convenience lay a privacy quandary: the system's data collection reached far beyond capturing test infractions, documenting every micro-expression and background noise (boyd, 2014). This kind of monitoring raises profound questions about the meaning of educational autonomy—whether students, often minors, have a voice in how their private data is harvested, stored, and possibly shared.

Debates over the balance between personalisation and privacy have a long history in education. Even before AI, administrators debated whether it was appropriate to gather extensive personal details (social background, health records) to shape "bespoke" learning plans (McKee, 2017). The promise of AI magnifies such tensions: data-driven personalisation implies a continuous gathering of metrics, from keystrokes to emotional signals gleaned by advanced cameras or wearable devices (Livingstone, 2014). Teachers might appreciate the new insights—seeing which topics truly captivate or frustrate a pupil—yet a world where every fidget is tracked can easily slip into an environment of hyper-surveillance.

The line between beneficial adaptation and intrusive paternalism is notoriously blurry. If an algorithm perceives that a student is bored, it might immediately switch content, scaffolding tasks in a fresh way (Downes, 2012). But do we want the system to interpret bodily cues at all times? What if it logs students' blinking rates or posture for extended analysis? Parents might be unaware of how intimately the platform monitors their child. In many jurisdictions, minors do not consent to data collection themselves; forms are sometimes signed by a parent who lacks the time or expertise to parse dense terms of service (Nissenbaum, 2010). Coupled with the commercial impetus—vendors keen to monetise data flows or prove "engagement metrics" to district buyers—this environment can erode genuine autonomy.

Even beyond bodily or emotional data, the domain of self-directed learning is at stake. Digital tutoring systems typically shape content, pace, and difficulty based on real-time analytics (Warschauer, 2016). That approach can be liberating if it spares a pupil from boredom and ensures timely interventions. Yet some critics argue that it disempowers the learner, subtly funnelling them into predefined tracks, absent the chance to wander, self-reflect, or follow tangential curiosity (Cohen, 2012). If a student's identity—shy or extroverted, methodical or impulsive—becomes codified into an algorithm's permanent profile, we risk reinforcing a static idea of who they are. Freedoms to reinvent or explore outside the system's recommended path shrink when teachers, pressed for time, defer to the platform's suggestions.

This tension resonates with philosophical queries about moral agency. In classical education theory, autonomy is nurtured by allowing learners to make mistakes, question authority, and push beyond their comfort zone (Freire, 1970). If an AI system "nudges" them relentlessly towards certain correct answers or curated content, does that hamper the reflective dimension of learning—where slow, meandering exploration and even strategic failures lead to deeper understanding? Autonomy demands space for doubt. An AI focused on optimising performance can unintentionally remove that space, favouring streamlined progress over messy intellectual wandering (Selinger, 2018).

Such hazards multiply when data-laden software intersects with private companies hoping to monetise usage patterns (Zuboff, 2019). If the platform logs not just academic performance but also times of peak activity, emotional states, or personal interests, corporations might package that intelligence—selling it to marketing firms or adding it to broader user profiles. Pupils effectively become data sources. In certain school systems, administrators justify it under budget constraints, finding that a free AI tool partially funded by user data is better than no resource at all (Eubanks, 2017). Others protest that the trade-off is morally unacceptable: a child's privacy is priceless, and educational institutions should remain sanctuaries of trust, not commercial data pipelines.

Even if all the data stays within a district, internal privacy concerns remain. Teachers might access a dashboard revealing a student's emotional volatility index or micro-attention profile (Calo, 2017). The teacher might then treat that student differently, whether out of empathy or suspicion. Equally, the school psychologist or a district official could override teacher insights with algorithmic interventions. The problem extends to how decisions are made: if parents or guardians do not fully grasp what these metrics imply, they cannot challenge or contextualise them. The result can be a chilling effect on free expression. Pupils may fear that every emotional slip or personal quirk is being logged, creating a climate of self-censorship.

In many countries, educational data is subject to regulations like FERPA in the United States or the GDPR in Europe. Yet these frameworks, as influential as they are, were not designed for the continuous, granular data streams generative AI collects (Gasser, 2016). They often revolve around notions of personally identifiable information, overshadowing subtler forms of inference-based profiling. If an algorithm never stores a direct name, but retains a highly detailed behavioural signature, does that sidestep confidentiality laws? Many legal experts say no—behavioural data can re-identify individuals. But enforcement lags behind the speed of AI deployments, and ed-tech companies may claim compliance as long as they anonymise superficial details, even if re-identification risks persist (Hartzog, 2018).

The consequences for autonomy become tangible when these platforms suggest or even assign tasks, well beyond the teacher's immediate view (Sonia Livingstone, 2014). If a student consistently struggles with algebra, the system might keep them in remedial modules, convinced it's "for their own good." Over time, the pupil is locked into a narrow set of experiences, seldom seeing advanced topics or extracurricular challenges. They miss the chance to discover latent interests. Teachers might not intervene because the AI's dashboards reassure them that the student is "progressing at an appropriate level." The child's voice—wanting to try something different—fades under algorithmic paternalism. For older students, ironically, the greater the personal data the system collects, the more it might sculpt a learner's future path without their informed input.

In search of solutions, some propose data minimisation: collecting only what's strictly necessary for essential personalisation (Solove, 2010). Instead of real-time emotional tracking, a platform could rely on periodic self-reports or optional teacher observations. Another approach is building robust data governance councils—where teachers, parents, and students can veto or refine data practices (Nissenbaum, 2010). This method aims to keep local agency at the core, so families have a say in how data is used. A third proposal is to separate the roles of personalisation and performance monitoring, ensuring that the AI that offers learning recommendations is not the same tool used for academic records or discipline decisions (Chomsky, 2017). Such a firewall might reduce the risk of each micro-behaviour becoming part of official transcripts.

Still, cultural contexts vary. In some Asian systems with strong central planning, data-driven paternalism might be seen as beneficial, reflecting communal values or the duty of educational authorities to guide youths (Arnett, 2021). By contrast, in Western contexts, the emphasis on personal privacy and individual autonomy is more pronounced, fuelling intense backlash against any perceived "surveillance" of minors (Cohen, 2012). This variation complicates the idea of universal ethical codes. If we treat data-mining as a transgression in one cultural setting, it might be routine or even welcomed in another that prioritises collective well-being over personal privacy. Achieving truly global frameworks demands sensitivity to these differences.

Underlying all of this is the evolving question: who ultimately shapes a pupil's identity? Education has historically been a collaborative dance—teacher, child, family, community—forming a moral and intellectual trajectory. AI's intrusion can tilt the dance if an external system wields disproportionate influence. A teenage learner might see the AI's judgments as final, especially if teachers appear to defer to it. Freed from human nuance, the technology's interventions can become normative: certain emotional expressions are flagged as signs of "risk," while certain linguistic patterns are deemed "unprofessional." Over time, the student may adapt their behaviour and worldview to conform.

To ensure that tools do not compromise learners' rights or self-determination, policymakers and educators must demand more than superficial disclaimers. They need explicit design choices—like local data storage under public oversight, default data minimisation, and consent processes that respect children's developing autonomy. If these systems remain black boxes, teachers risk losing professional agency, and students lose the promise of exploration. The conversation must also reach beyond institutions, inviting civic voices, ethicists, and young people themselves to articulate how technology can serve rather than shape them (Freire, 1970).

In an ideal scenario, a school might adopt an AI reading assistant that respects personal boundaries. It collects minimal analytics, enabling personalisation without micromanaging emotional or behavioural cues. It defers to teacher judgement for major decisions, offering suggestions rather than imposing. The data it gathers is locked down locally, with community oversight committees reviewing any extended usage or potential commercial partnerships. Meanwhile, families can opt in or out of advanced monitoring features, and the system's design ensures no single user profile is used for punitive or high-stakes classification. This vision attempts to marry innovation with conscientious governance, letting children benefit from adaptive support without surrendering their personal lives to corporate or bureaucratic watchfulness.

Achieving that balance is no small task. Yet if education is to remain a domain where moral agency and personal growth thrive, the deployment of AI cannot be left to market forces alone. Maintaining student autonomy and safeguarding privacy require deliberate guardrails, shaped by those who understand that schooling goes beyond producing test scores. The next step involves forging clear, enforceable rules—built from democratic input—that champion a child's right to shape their own learning journey, free from unwarranted intrusions or manipulative algorithms. Only then can personalisation truly reflect the best of what education can be, rather than a Faustian bargain between comfort and control.

The Global Perspective – Equity, Access, and Digital Colonialism

Global discussions on AI in education often revolve around sophisticated data analytics or personalisation features, yet the question of equitable access underpins every one of those technological promises. When advanced AI platforms demand stable high-speed internet, robust hardware, and continuous software updates, the result can exacerbate existing educational divides across regions, socioeconomic brackets, and cultural contexts. Many rural schools or low-income urban districts already struggle to maintain a basic digital infrastructure. In such places, adopting sophisticated AI tools is hardly realistic, a point repeatedly emphasised by Kentaro Toyama (2015), whose "amplification principle" argues that technology magnifies underlying social inequalities more often than it resolves them. Even a well-intentioned AI initiative, introduced to improve reading scores, may entrench privilege if only well-resourced schools can deploy it effectively.

The situation becomes more complex when we examine the subtle forms of "digital colonialism" that can accompany AI-based learning solutions. Commercial platforms often assume a universal set of cultural references, linguistic norms, or pedagogical methods, which can overshadow the knowledge traditions of local communities (Narayan, 1997). Indigenous learners, for instance, might find the system's definitions of "progress" or "success" at odds with relational, land-based epistemologies that emphasise communal well-being over individual performance. Similarly, in certain African contexts, as Sylvia Tamale (2020) argues, standard Western metrics around academic rigour may disregard communal responsibilities or intergenerational teaching approaches. An AI might interpret collective tasks as signs of cheating or measure-based inconsistency, instead of recognising them as culturally rich collaborative practices.

The tension is compounded by how such technology enters these regions: typically, major ed-tech vendors partner with governments or philanthropic bodies, offering free or discounted services in exchange for data. This model can echo what Payal Arora (2019) calls a "leisure commons" phenomenon: technology arrives under the banner of bridging digital gaps, yet the real exchange involves capturing user data that becomes valuable for corporate analytics. In a resource-poor environment, administrators and parents may welcome the arrival of cutting-edge AI, unaware of the long-term consequences of data extraction and cultural homogenisation (Zuckerman, 2014). The core dilemma is whether these solutions truly serve local needs or primarily reinforce external commercial or policy agendas.

Equity in access also goes beyond hardware. Teachers need training and pedagogical support to integrate AI tools into existing curricula. Without robust professional development, advanced platforms become an underutilised novelty at best, or a confusing burden at worst (Arul Chib, 2010). Many schools in low-income areas rotate inexperienced teachers who seldom receive stable contracts, let alone AI-related training. If the system's interface and analytics are tailored for well-staffed institutions with dedicated IT support, those educators may lack the bandwidth to engage with it meaningfully. This disparity harks back to Amartya Sen's (1999) capabilities approach, which argues that real freedom or empowerment requires more than surface-level resources; it demands the ability to utilise them effectively within one's social context.

Some NGOs strive to deliver regionally adapted AI solutions, ensuring multiple languages or offline modes. Firoz Lalji (2020) has documented community-based programmes in African settings that adapt open-source tutoring systems for local usage, embedding culturally relevant examples. While promising, such initiatives often operate on small scales, overshadowed by large ed-tech deployments guided by central authorities. John Willinsky (2006) underscores the tension between open educational resources that emphasise local appropriation and the proprietary nature of mainstream AI software. If a multinational vendor provides a platform with locked-down code, teachers in remote settings cannot localise content or fix issues on their own. The friction intensifies when updates from headquarters disrupt local customisations, disempowering the very communities the platform pledged to uplift.

Whether in Asia, Africa, or Latin America, the notion of "local knowledge" often collides with monolithic data structures. Pranav Mistry (2012) once demonstrated how innovative hardware prototypes could be adapted for distinct contexts—like simple gesture-based inputs for communities less accustomed to QWERTY keyboards. Yet large-scale AI systems typically assume standard user interfaces. The result might impose a cognitive barrier: learners who handle phones or tablets differently, or rely on oral transmission of knowledge, find themselves trying to conform to an interface logic not designed for them. Chika Ezeanya-Esiobu (2019) calls this a forced assimilation that treats local ways of learning as anomalies rather than valid methods. Meanwhile, global tech leaders brandish the term "digital transformation," skirting the question of how diverse cultures define transformation on their own terms.

A parallel equity concern lies in the mismatch of content. AI tutors often rely on corpora that reflect certain historical narratives or language patterns, ignoring local histories or marginalised viewpoints (Bhatia, 2015). If a rural school in Southeast Asia wants the AI to teach English using examples from its own folklore, the system might lack any data on local legends, resorting to references from Western pop culture. This disparity fosters a sense of cultural alienation: the child sees their environment mirrored nowhere in the lessons, possibly reinforcing a colonial mindset that true knowledge resides only in foreign sources (Fricker, 2007). Uma Narayan (1997) has long argued that educational reliance on external narratives can perpetuate hidden hierarchies, stalling genuine cultural pride and self-determination.

Mariana Mazzucato (2018) emphasises the role of public investment in steering technological innovation for the common good. If governments commit resources to creating open platforms that communities can shape, equity might improve. They could demand that all AI solutions for schools be open-licensed, or at least allow local tailoring. The challenge is whether government budgets permit such bold moves, especially when philanthropic or corporate donors dangle "free" solutions. Often, those free packages involve data extraction or subscription lock-ins, a short-term fix with long-term costs (Katz, 2012). Ethan Zuckerman (2014) sees this as a stark choice between locally governed solutions and the convenience of globally scaled corporate platforms—a dilemma that must be navigated if we want to avoid a new wave of digital dependency.

Another aspect is the interplay of language. Much cutting-edge natural language processing thrives on English-language datasets. Where local language corpora exist, they may be too small or under-curated for advanced AI training. Arora (2019) highlights the phenomenon of "linguistic overshadowing," where less dominant languages are relegated to second-class status in AI's worldview. If the system rarely "sees" them, it cannot deliver the same quality of personalisation to those learners. The result is not just a hardware or bandwidth gap but a data gap, further fracturing opportunities for inclusive education.

Meanwhile, the concept of digital colonialism intensifies when we see a pattern: wealthy regions produce the AI tools, define the ethics frameworks, and then export them to less wealthy regions, where they "solve" local problems. Roman R. Williams (2013) emphasises that technology is never a neutral import; it reshapes local social structures. If teachers in a remote area must adjust their methodology to fit corporate standards for AI usage, we risk undermining their cultural autonomy. Over time, such subtle shifts can supplant local educational philosophies, eroding intangible heritages that once thrived. For instance, a system that rewards individual test success might conflict with communal approaches to learning, eventually nudging the community to adopt more individualistic, exam-centric values.

Payal Arora (2019) extends the argument by asking who truly benefits from these deployments. If the data gleaned from rural schools helps refine an AI model that corporate entities sell at premium rates in richer countries, are the original communities merely data suppliers with minimal reciprocal gain? This phenomenon, reminiscent of historical resource extraction, fosters resentment and cynicism, eroding trust in ed-tech as a liberating force. As with mineral resources, local communities witness external powers profiting from raw materials—here, their students' learning footprints—without enjoying improved local infrastructure or educational autonomy.

Various initiatives attempt to break this cycle. Pranav Mistry's low-cost computing prototypes, if integrated with open AI software, can empower local innovators to craft their own versions of generative tutors or self-assessment modules (Mistry, 2012). In a parallel domain, Arul Chib (2010) has documented success stories where small-scale educational tech was co-designed with teachers, ensuring culturally relevant content. With enough institutional support, these grassroots models can scale. The obstacle is a global ed-tech market that privileges big brand solutions, overshadowing smaller local efforts.

A deeper shift might emerge through a "network of networks" approach, championed by some philanthropic organisations, in which communities share best practices, local data sets, and indigenous knowledge bases. This horizontal collaboration contrasts with top-down, one-size-fits-all solutions. Combined with open licensing, it can combat digital colonialism by giving each region the means to customise AI or develop micro-updates that reflect local tradition and language (Willinsky, 2006). The dream is a patchwork of globally connected but locally attuned AI platforms, bridging the digital gap without imposing uniform cultural metrics.

Nevertheless, critics warn that such efforts require continuous funding and technical capacity. Small communities might face challenges sustaining or updating AI systems over time. Mazzucato (2018) sees the solution in strong public investment frameworks that treat educational AI as a public good, akin to roads or healthcare. By guaranteeing ongoing support and open governance, states can set conditions that align with local empowerment instead of relying on precarious philanthropic cycles or corporate freebies. The pragmatic question is whether political will exists. Advocates must battle entrenched interests benefiting from data extraction or from the sales of proprietary solutions (Ranjit Singh, 2021).

In the end, bridging equity, access, and local sovereignty calls for more than rhetorical nods. It demands local co-design, open data governance, teacher training, and robust policies that limit the extractive side of ed-tech. A universal moral code on AI in education cannot succeed unless it explicitly addresses the threat of digital colonialism, the uneven distribution of connectivity, and the complexities of sustaining regionally grounded innovations. Without such recognition, the rush toward advanced AI classrooms will likely replicate the divides we have seen with simpler ed-tech: a privileged minority leaps ahead, while the rest remain bystanders or unwitting data providers.

Yet global activism and community resilience offer hope. From community mesh networks in underserved areas to local hacking clubs that adapt open AI modules, a groundswell of initiatives suggests the possibility of inclusive modernisation (Chika Ezeanya-Esiobu, 2019). These projects reveal that educators, parents, and youth themselves can be powerful architects when not bypassed by monolithic solutions. Ultimately, forging ethical AI in education will require consistent attention to infrastructural realities, linguistic justice, and respect for cultural difference—refusing to assume that a single piece of software can or should conquer the world of learning. By weaving local insight into global connections, we stand a chance of realising a more balanced and dignified landscape for AI-augmented classrooms everywhere.

Moral Complexity – Ownership, Commercialisation, and Student Agency

One of the most striking aspects of AI adoption in schools is how it introduces a corporate dimension into relationships that were historically defined by public values and community trust. For many decades, educational resources—textbooks, facilities, even teacher training—were funded and managed by public entities or non-profit institutions. Yet the surge in AI-driven tools has come largely from private ed-tech firms, some of them global giants, others backed by venture capital. Their solutions are licensed rather than purchased outright, and their updates and data-handling protocols lie outside local control (Zuboff, 2019). This dynamic raises fundamental questions about who truly owns the educational experiences of children and what that means for student agency.

In older classrooms, ownership was a simpler affair: a school might buy textbooks, which teachers and pupils used freely. If a teacher wanted to annotate or adapt a lesson, they could do so without seeking permission. Today, generative AI systems often arrive with strict terms of service, limiting how teachers can modify content or export data. Because the software's processes are proprietary, educators lack the freedom to see how recommendations or materials are generated (Giroux, 2011). The moral significance of ownership emerges here: if an algorithm suggests particular reading sets or "auto-creates" quizzes, does the teacher retain ultimate authority to question or reshape them? Or has the vendor effectively become a remote orchestrator of each child's learning path?

Ownership also pertains to data. When an AI platform logs student assignments, tracks their progress, or uses personal queries to refine its generative capabilities, it transforms intangible educational moments into a resource with commercial value (Rushkoff, 2019). These data points can feed research and development pipelines, helping the vendor improve its product, which it then sells elsewhere. In many cases, the terms of service disclaim any kind of shared ownership—schools and families are merely service users. The tension is that these private companies benefit from collectively amassed knowledge about how children learn, yet do not necessarily reinvest that benefit back into local communities or educational ecosystems (Morozov, 2013). Instead, new features might be locked behind paywalls, or data insights might be used to develop a premium-tier product. Effectively, the intangible interactions of pupils become the firm's intellectual property.

This dynamic can undermine student agency, particularly when we consider that educational processes are deeply personal. A pupil's difficulty in algebra, their sudden flair for creative writing, or their emotional reaction to historical content each reflect unique facets of who they are. If these become fodder for commercial analytics and algorithmic improvements, there is a risk of commodifying students' intellectual growth (Noble, 2018). By default, we might accept it as a trade-off, trusting that personalisation gains justify the partial loss of data sovereignty. But a moral foundation for AI in education would require a clear stance on data usage and commercial exploitation. Are children implicitly consenting to their learning journeys being re-sold or integrated into private R&D?

High-profile philanthropic ventures can also blur these boundaries. For instance, a wealthy foundation might fund the rollout of advanced AI tutors for an entire district, proclaiming it a philanthropic gesture (Williamson, 2017). Yet behind the scenes, the ed-tech partner collects vast troves of user data, refining proprietary algorithms. The philanthropic narrative glosses over how local autonomy might be undermined: teachers are trained to rely on a single platform, and any adjustments to suit cultural nuances require the vendor's approval or additional fees. On paper, this looks like a charitable gift, but in practice it can lock schools into a commercial ecosystem that erodes local ownership (hooks, 1994).

Paulo Freire (1970) argued that education at its best is a dialogue, empowering learners to question the world and co-create knowledge. If an AI system's licensing terms forbid users from delving into its logic, or it shapes content by referencing a single worldview, that dialogue is stifled. Pupils become consumers of pre-packaged narratives. Equally, teachers who want to blend indigenous stories or local archives into the AI's generative prompts might find the software lacks an open interface, or charges extra. The effect is to centralise intellectual production in the vendor's domain, leaving marginalised cultures voiceless or forced to reformat knowledge to fit the platform (Smith, 1999).

Another dimension emerges when we consider who profits from these systems. Shoshana Zuboff (2019) described how surveillance capitalism can monetise user data and behaviour patterns. In educational AI, such patterns might extend to biometric or emotional metrics, as well as a child's trajectory of conceptual mastery. If the vendor sells or trades these behavioural profiles for targeted marketing—whether for commercial goods or tutoring services—students unknowingly become the product. Some might argue that personal data from children should be off-limits for any commercial exploitation, akin to how certain jurisdictions prohibit child labour. Yet, lacking a robust ethical or legal framework, schools often sign contracts that permit broad usage of data, out of cost-saving necessity or ignorance of the deeper implications (Crouch, 2004).

bell hooks (1994) insisted that genuine educational freedom requires safe spaces for self-definition, uncolonised by market rationales. But if an AI's main objectives revolve around collecting "engagement data" to please ed-tech investors, the software design might promote addictive or superficial tasks that artificially boost usage stats. Pupils could be enticed to chase extrinsic rewards, overshadowing the reflective, socially engaged aspects of learning (Kohn, 2015). The moral question is: does a business-centric AI platform have any incentive to cultivate deeper social responsibility or critical consciousness if these do not translate into profitable data metrics?

Teachers might see themselves compromised too. Some ed-tech companies plan or already offer "AI coaching" for teachers, analysing their classroom interactions and giving performance feedback. That feedback loop can be valuable, but it also centralises control: if teachers' promotions or salaries hinge on meeting metrics the platform deems key, they may become subservient to the software's logic rather than their professional expertise (Giroux, 2011). Their agency to experiment with alternative methods, slow the pace for vulnerable pupils, or incorporate local narratives could wane if those decisions conflict with the system's data-driven benchmarks.

In certain contexts, philanthropic or corporate entities openly speak of "disruption" in education—implying that old teacher-centric models are obsolete, and that AI solutions can produce better outcomes with fewer staff (Morozov, 2013). This approach can be seductive in cash-strapped districts, but it risks commodifying the teacher's role, framing them as mere facilitators behind the scenes. The ethics of ownership here concern not just data but also the soul of teaching. If a private solution claims ownership of the "learning content pipeline," teachers effectively become service intermediaries (Langdon Winner, 1980). Students, for their part, may struggle to see the human, relational side of education as anything but an extension of the AI's instructions.

Another subtle but crucial area relates to intellectual property in AI-generated outputs. Suppose a platform helps students draft essays, quizzes them, or co-authors creative works. Who owns the final piece? If a pupil invests genuine creativity but the AI handles the structure, does the platform claim partial ownership? Some terms of service stipulate that the company retains rights to any data or content produced within the system (Boyle, 2008). This can be jarring: the student might wish to share or adapt that content, but faces legal hurdles. Meanwhile, the notion of a teacher collaboratively refining an AI lesson for their class might breach proprietary guidelines, as they are effectively "modifying" or "copying" the vendor's intellectual property. The very concept of open educational resources—once seen as a democratising force—runs contrary to these closed commercial licences (DeRosa, 2016). The result is a paradox: advanced generative AI could greatly facilitate knowledge sharing, yet commercial constraints lock it down.

Critics like Douglas Rushkoff (2019) argue that we should reimagine digital tools as public utilities, freed from a profit motive that stifles genuine collaboration. In the educational sphere, that could mean open-source generative AI that fosters local content creation, respects user privacy, and treats student data as a collectively owned resource, not a commodity. Teachers, local historians, or cultural groups could co-develop modules that reflect their traditions. But realising this vision faces stiff resistance from business models predicated on licensing fees, data-based services, and intellectual property claims (Mulligan, 2019). A moral code for AI in education would thus need robust stipulations on open licensing or at least local content autonomy, alongside a ban on profit schemes reliant on the extraction of children's data.

The deeper worry is that if we allow commercial priorities to overshadow the public interest, students grow up in an environment where everyday learning is monitored, shaped, and monetised, normalising the idea that knowledge belongs to external entities (hooks, 1994). They might accept a permanent sense of indebtedness to corporate platforms or philanthropic brands, as though their educational opportunities hinge on private goodwill rather than being a public right. Even the subtle design choices—for instance, an AI brand's logo plastered across every lesson—can imprint a commodified worldview, implying that advanced learning and personal growth are gifts from corporations, rather than a natural societal function.

One glimmer of optimism is found in teacher unions and activist networks that push for "ethical ed-tech procurement." They demand transparent contracts, data sovereignty, and the freedom to adapt resources for local contexts (Giroux, 2011). A parallel impetus comes from open educational resource communities that champion collaborative licensing, ensuring that if AI is used to generate content, that content can be freely modified and shared. In such a scenario, vendors might still sell support services, but do not hold absolute ownership over the educational experience.

In short, the moral complexities around ownership, commercialisation, and student agency intertwine deeply. AI can broaden horizons, but if it does so by rendering pupils into data assets, or by curtailing teachers' autonomy, it undermines core educational values. A child's intellectual journey should not be subject to hidden economic calculations that place corporate gain above human development (Freire, 1970). Nor should teachers be reduced to onlookers, overshadowed by algorithms. Rather, the moral compass would direct us to treat knowledge creation and sharing as communal endeavours, belonging to learners, educators, and the broader public interest. Without explicit protections, we risk entrenching a digital feudalism in which large ed-tech players hold decisive sway over the future of education. The next step, then, is to articulate clearer frameworks —policies that embed teacher empowerment, restrict exploitative licensing terms, and preserve the principle that each learner's creative spark remains theirs, not a commodity to be bartered or controlled by external gatekeepers.

Transparency, Accountability, and the Teacher's Role

The spread of advanced AI in education has raised new uncertainties about how teachers might retain—or even enhance—their roles as ethical stewards of the classroom. Historically, teachers have acted as mediators, tailoring lessons to local contexts and challenging students to grow beyond prescriptive boundaries. With sophisticated software now suggesting lesson plans, tracking student performance, and even providing discipline alerts, the teacher's function can appear overshadowed by algorithmic authority. Yet the same technology offers a chance for educators to reclaim professionalism, provided they have the transparency, training, and institutional support to question and refine what AI systems propose.

Teachers often carry a nuanced understanding of cultural values, interpersonal dynamics, and the subtle cues that signal a student's emotional needs (Wiliam, 2011). A data dashboard can display test scores or engagement metrics, but it cannot replicate the compassion or on-the-fly adaptation a teacher exercises when a pupil is distressed or lost in thought. In many schools, AI solutions are marketed as time-savers that free staff from repetitive tasks. That promise resonates, yet it can become a trap. If the system intrudes too deeply, teachers might find themselves deferring to software for decisions once grounded in their expertise—like when to move a student to advanced materials or how to handle disruptive behaviour (Turkle, 2015). Over time, a climate of technological paternalism may emerge: data-driven dashboards commanding the flow of instruction while teachers simply facilitate.

One of the crucial protective measures is interpretability. If an AI flags a pupil as "at risk," the teacher should see clearly on which features or patterns the system based that conclusion (Shneiderman, 2020). When a tool simply doles out final labels—low engagement, high potential, early dropout risk—educators can feel pressured to act without fully understanding the rationale. If they disagree, they may need to justify why they overrode a machine's "objective" assessment. That scenario reverses normal accountability: the teacher, trained in pedagogy and local knowledge, ends up subordinate to a black box. The remedy is to insist on explanatory systems that share reasoning in plain language or visual form. Teachers then gauge whether an AI's logic fits their contextual understanding before adopting any recommendation.

The concept of accountability runs both ways. If data reveals that a teacher's class consistently struggles with certain concepts, the system might suggest the teacher refine methods or incorporate new materials. This can empower staff, offering timely insight into patterns that might otherwise remain hidden (Downes, 2012). Some educators appreciate how AI analytics bring awareness to subtle issues, like reading-speed mismatches or overlooked skill gaps. The teacher's role becomes that of a discerning interpreter, selecting which insights to follow. Yet tensions persist when official policies or performance reviews hinge on AI metrics. Teachers can end up adjusting their practice to satisfy an algorithm's preferences—perhaps emphasising quickly scored tasks over deeper explorations that the system struggles to measure (Collins, 2018).

A parallel challenge involves the risk of data overreach. If the software collects emotional cues via webcams or micro-behaviours in VR, teachers might face a flood of sensitive details about their pupils. In principle, that could help identify anxiety issues or social isolation early. But in practice, it can also breed a culture of surveillance. Teachers, already juggling multiple responsibilities, might feel uneasy about monitoring each micro-expression, not wishing to reduce the classroom to a data-tracking zone (Nissenbaum, 2010). They may prefer the relational approach—an intuitive, empathetic bond—over focusing on dashboards. This choice, however, relies on the institution's stance: do administrators expect teachers to use every new AI feature, or are they granted autonomy to pick and choose?

Unions and professional bodies can step in to protect educators' rights and define boundaries around these tools. Drawing parallels from medical ethics—where hospital staff can question new procedures before adopting them—teacher associations might similarly demand a voice in decisions about AI deployment. They could insist that each system be trialled with teacher oversight, that data remain under local control, and that educators have final say in high-stakes determinations (Giroux, 2011). By championing the principle that no algorithm can supplant professional judgement, unions reinforce the teacher's moral responsibility: caring for the holistic growth of each pupil, not merely hitting numeric targets.

Another dimension is teacher training. Many educators have limited background in algorithmic thinking or data science (Warschauer, 2016). Faced with complex AI dashboards, they might rely on vendor demos or short workshops. If the training only covers operational usage—click here to see at-risk flags—teachers lose out on the critical literacy needed to interrogate bias or weigh contextual factors. That shortfall can be especially acute in under-resourced districts, where staff are already overstretched. Without a structured professional development strategy, AI becomes a "black box in the corner," its outputs largely accepted at face value. Over time, this can erode the culture of reflective practice historically vital to good teaching.

Those who do receive advanced training, conversely, can emerge as ethical gatekeepers. Skilled teachers might cross-reference AI insights with in-person observation or external data, reinforcing or challenging algorithmic claims. They could share anomalies with colleagues, collectively spotting patterns of bias. Some districts encourage teachers to keep a reflective log: if the AI suggests certain reading interventions, how did the student respond, and what local knowledge might refine future recommendations? This cyclical approach—machine suggestions, teacher reflection, iterative improvement—reflects a human-machine partnership that respects professional nuance. Teachers are not passively "managed" by the AI; they co-manage it, shaping its ongoing calibration (Scribner, 1984).

In light of cultural sensitivity, educators often see themselves as custodians of community identity. If an AI's recommended reading list includes few references to local history or indigenous practices, teachers can manually insert such content or override default suggestions. The system might then record teacher preferences and adapt its library, but only if it was designed to learn from educator input rather than from abstract data alone. Tech design must thus incorporate teacher feedback loops as first-class elements, not afterthoughts (Brown, 2002). Otherwise, the teacher's local knowledge—a vital asset for bridging culture and curriculum—might be sidelined by the system's universal approach.

Such teacher involvement grows more urgent when generative AI is used to create or adapt lesson materials. Educators should have the final edit, checking for cultural or factual appropriateness.

While a system might spontaneously craft a reading comprehension exercise, the teacher ensures it aligns with local moral norms and accurately represents historical complexities. The presence of an automatic lesson generator might tempt some staff to offload lesson planning. But in a robust model, staff remain central curators, weaving the AI's speed and creativity with their moral judgement and communal insights. This synergy fosters a new professional identity: part mentor, part data interpreter, part cultural guardian.

Yet pressures from top-down reforms can threaten this autonomy. If large-scale accountability frameworks push schools to adopt certain AI vendors for "efficiency," teachers might feel powerless. Their professional concerns—bias, privacy, or nuance—can be dismissed under the logic of "higher test scores" or "cost savings." Some school boards, constrained by budgets, may view the teacher's scepticism as Luddite resistance. The net effect is a precarious shift of moral accountability: teachers still get blamed if outcomes falter, but they have less control over the system that shapes daily practice. Instead of empowerment, teachers risk becoming scapegoats for decisions made by administrators, vendors, or algorithmic engineers (Downes, 2012).

An alternate future sees teachers forming local committees to evaluate ed-tech proposals before implementation, echoing the idea of "technology review boards" akin to ethical committees in research contexts (Edsger Dijkstra, 1982). Districts might require vendors to grant partial transparency into how AI suggestions are generated. Teachers, having read the summary of data handling and model assumptions, could ratify or reject the system. Over time, their professional insights refine guidelines: for instance, an AI used for reading practice must incorporate local dialect acceptance, or a predictive tool for dropouts must not rely purely on demographic data. This structure returns moral accountability to educators, bridging digital illusions with on-the-ground realities.

When teacher input weaves seamlessly with advanced AI, a new synergy emerges. The system might highlight students who overcame adversity, prompting teacher-led recognition or advanced opportunities. Alternatively, it might confirm a teacher's hunch that a student needs more challenge, letting the educator swiftly offer deeper tasks without labour-intensive searching. Freed from mechanical chores, teachers may dedicate more energy to one-on-one mentorship, Socratic debate, or community outreach, all while monitoring an AI dashboard that supplies hints or patterns. But crucially, the teacher retains final interpretive authority—no numeric ranking is definitive without their considered perspective. If this approach becomes the norm, we witness a genuine collaboration, not a hollow displacement of professional judgement by algorithmic scripts.

Thus, the debate over transparency, accountability, and teacher agency cuts to the heart of AI ethics in education. It reminds us that advanced analytics, while beneficial, must serve as assistants rather than rulers in the classroom (Langdon Winner, 1980). If technology is seen as a tool, not an oracle, teachers remain the moral bedrock, shaping how data-driven suggestions fit each pupil's story. To ensure that vision, policy frameworks might grant teachers the legal right to override AI-driven recommendations, or require that all high-stakes decisions undergo human verification. Unions, teacher associations, and local boards can stand guard against creeping algorithmic paternalism. Only then can schools harness the creative promise of AI without forfeiting the intimate, human dimension of education. At its best, the synergy realigns with a tradition of teacher-led inquiry and empathy, weaving advanced insights into pedagogical wisdom while ensuring the classroom stays a place of deep learning rather than robotic compliance.

Regulatory and Policy Frameworks – The Road to an Ethical Code

Educators, policymakers, and parents alike often voice discomfort at the thought of ceding crucial classroom decisions to opaque algorithms. Yet proposals to curb these excesses sometimes feel scattered or toothless. In the wake of high-profile controversies—like AI-based admissions tools that marginalise certain demographics (Broussard, 2018)—calls for more structured oversight have grown louder. The question is no longer whether we need regulation, but how robust and enforceable such frameworks should be, especially in contexts as varied as local primary schools or massive online platforms. A moral foundation, if it is to be more than an attractive slogan, must be concretised through policy instruments that strike a balance between encouraging innovation and safeguarding public values.

Many look to templates emerging from broader AI governance initiatives. Some echo what Pasi Sahlberg (2011) has argued about educational reforms: overly standardised global models can overlook local textures. In the domain of AI, a single, one-size-fits-all code might clash with cultural norms or legal peculiarities. Even so, there is a growing consensus that certain ethical pillars—like transparency, accountability, and fairness—deserve universal recognition. Andreas Schleicher (2018) at the OECD has hinted that educational AI could integrate standards akin to international testing frameworks but geared towards algorithmic ethics. If these standards remain purely voluntary, however, vendors may cherry-pick which guidelines they follow, a pattern that has haunted sustainability charters in other industries (Miller, 2018). The real question becomes how to enforce compliance in a domain dominated by proprietary software.

Sheila Jasanoff (2004) emphasises "co-production" in science and technology policy, reminding us that effective governance arises when developers, users, and affected communities actively shape rules together. In the classroom context, that implies an active role for teachers, parents, and students in vetting how AI solutions are integrated. Rather than school boards unilaterally signing deals with ed-tech vendors, local committees could examine potential bias, data handling, and adaptational flexibility before approving usage. This approach ensures that a proposed AI platform respects community-specific norms—say, honouring indigenous knowledge or avoiding invasive emotional tracking—while still adhering to broad ethical standards. The tension emerges when educators lack the resources or expertise to engage with technical details. If these decisions remain top-down, the ethical code becomes little more than a box-ticking exercise.

Another concern is that existing legal frameworks often revolve around the notion of privacy rights (Gasser, 2016). While privacy remains pivotal, it does not fully address questions like who sets the metrics for "academic success," or whether the AI can override teacher autonomy. Some argue for a structured approach akin to medical ethics committees: technology proposals for schools would undergo local "ethical review boards" that weigh the AI's potential benefits against cultural or pedagogical risks (Broussard, 2018). Ideally, these boards would involve child psychologists, data scientists, civil rights advocates, and teacher representatives, ensuring well-rounded scrutiny. Yet critics worry this might slow down AI adoption or hamper innovation. The moral question is whether such caution is an inconvenience or a necessary guardrail. Riel Miller (2018) of UNESCO emphasises futures literacy, implying that anticipating the social repercussions of AI is integral to policy, not an optional add-on.

Wendy Hall (2019), examining AI guidelines, notes that multiple industries have discovered codes with fine phrasing but minimal enforcement. In the ed-tech realm, a comparable risk is that vendors produce glossy brochures on ethics while continuing data extraction behind the scenes (Williamson, 2017). This gap between rhetoric and reality suggests that regulatory bodies may need the power to audit code and data usage. That means legally mandating that ed-tech companies open their algorithms or provide detailed model documentation for external inspection. Some developers resist, claiming intellectual property protections, yet educational institutions arguably have a higher moral claim: these are minors' data, after all. If the vendors refuse meaningful transparency, can they be entrusted to uphold the public good?

Pasi Sahlberg (2011) introduced the notion of "trust-based accountability," contrasting it with test-driven oversight. A parallel might be drawn for AI oversight: instead of punishing schools if the AI flags them for low compliance, we can emphasise collaborative improvement. Yet Sahlberg's concept relies on mutual respect among stakeholders, and the high commercial stakes in ed-tech complicate that relationship. If school districts perceive that refusing a major vendor's AI solution might brand them as regressive or cost them philanthropic funding, the dynamic shifts. Some boards yield, adopting systems despite teacher misgivings. The code's moral stance must empower educators to say no to intrusive or culturally insensitive designs without financial repercussions. This might require rethinking how public funds or philanthropic grants are distributed—perhaps awarding them only to AI deployments that pass rigorous, teacher-led reviews.

The push for universal frameworks also encounters the phenomenon of "local custom," which can be used to justify intrusive data collection in certain contexts or hamper essential reforms in others. A universal moral code might declare that no ed-tech vendor can gather biometric signals from students under a certain age, for instance. But in a region where a paternalistic model is widely accepted, stakeholders could argue they want that data to ensure pupils are "always on track." Conversely, some highly privacy-conscious communities might go further than the universal standard and ban any AI-driven personalisation that logs personal diaries or family details. The tension arises when universal codes meet local power structures, highlighting the messy interplay of policy, culture, and authority (Jasanoff, 2004).

Another thread addresses how these codes can evolve. AI innovation moves quickly: a system that only did adaptive quizzes last year might soon integrate generative content creation and emotional analytics. Relying on a static set of guidelines risks irrelevance. One idea is to create multi-stakeholder councils at regional or even national levels, regularly updating ethical rules or providing clarifications in response to new technological capabilities (Hall, 2019). Such councils might also listen to teacher-led investigations into how certain AI updates changed classroom dynamics. If the shift raises bias flags or undermines teacher autonomy, the council might revise policy or temporarily suspend usage. This nimble approach demands sustained funding and political will, which is never guaranteed.

Some experts look to parallels in bioethics, where genetic testing or stem-cell research confronted societies with new moral frontiers (Miller, 2018). In those domains, structured frameworks emerged—like the Helsinki Declaration in medical research—spelling out universal principles. Translating that to AI in education, we might see a "Global Declaration on Ethical AI for Learners," championed by bodies like UNESCO. It could state, for example, that data usage must be minimal and proportionate, no algorithm shall override teacher or parental judgement without recourse, and cultural context must be systematically respected. This declaration, if widely endorsed, would let communities pressure vendors to align or else forfeit access to big markets. Of course, signing onto a declaration can be performative if not tied to consequences, a challenge that has stymied many international charters.

Sheila Jasanoff's co-production lens insists that building trustful governance mechanisms requires local democratisation. Teachers, students, and civil society must shape the content of ethical codes, not just governments or ed-tech corporations. A code drawn exclusively by policy elites may reflect a narrow viewpoint. For instance, it might emphasise data encryption but gloss over teachers' rights to override predictive models. Or it might fixate on personal data consent while ignoring the broader corporate prerogatives in shaping learning content. By weaving local voices into code drafting, communities produce frameworks that address everyday tensions—like how to handle generative expansions of indigenous stories or how to discipline a child flagged as disruptive by an algorithm that lacks cultural nuance.

One might wonder whether such moral frameworks risk slowing technology's beneficial spread. Skeptics argue that burdensome checks stifle innovation and hamper the rapid iteration that AI thrives on (Crawford, 2021). Yet we can recall that educational integrity depends on trust. Rapid rollouts of questionable tools can erode public faith in ed-tech altogether. Pragmatic voices maintain that slow, deliberate policy fosters a stable environment where truly beneficial AI solutions can flourish without sparking backlash or causing harm. The notion of "ethical by design" (Dignum, 2019) extends beyond product rhetoric, insisting that from the earliest prototypes, developers incorporate teacher feedback and address local contexts. If robust policy demands it, developers adapt or risk losing the chance to operate in public schools.

One essential factor in forging an ethical code is the distribution of resources. Implementation requires not only guidelines but also staff who can interpret them. Smaller districts with minimal legal or technical capacity can be overwhelmed by complex compliance steps. Some propose an Ed-Tech Ombudsman model: a central resource offering free legal counsel and data science expertise to under-resourced schools (Sahlberg, 2011). This ensures that no community is left grappling alone with corporate terms of service. The impetus for such structural support underscores that policy is not an abstract statement but a real apparatus for empowerment.

Ultimately, building a moral foundation for AI in education demands parallel action at multiple levels. Grassroots teacher committees push for daily oversight, union negotiations set non-negotiables around data privacy or algorithmic override rights, district or national regulators define baseline rules, and international bodies craft aspirational declarations. The synergy of these layers can shape an environment where advanced AI can serve learners while respecting local culture, teacher judgement, and the child's own evolving autonomy (Freire, 1970). Without such synergy, codes remain symbolic, overshadowed by well-funded ed-tech expansions guided by profit or simplistic visions of digital "disruption."

The road to a truly ethical code is thus an ongoing process rather than a single moment of policy enactment. From drafting to enforcement, from local pilot to global collaboration, each step must engage those who inhabit classrooms every day. As big AI moves further into schooling—revising curricula, generating tasks, or even grading emotions—no aspect of learning remains untouched. A code that sits passively in a file cannot safeguard the moral core of education, but a living framework shaped by teachers, learners, communities, and conscientious developers might. It is this evolving synergy that can keep AI's rapid progress anchored in humanity's broadest educational ideals, ensuring we do not unravel the very freedoms and communal bonds that make teaching a profoundly human endeavour.

Toward a Universal Framework – Cultural, Philosophical, and Community Input

A universal set of ethical guidelines for AI in education often sounds enticing, yet practitioners repeatedly uncover how challenging it is to translate broad principles into day-to-day practice. Cultural diversity, varied interpretations of fairness, and differences in local governance all suggest that any attempt at a single, rigid framework may lack the flexibility schools need. At the same time, minor or purely local codes can become toothless if powerful ed-tech vendors refuse to engage. Navigating this tension requires a multi-layered approach: specific enough to protect core values, yet open to adaptation by each community or region. Critics worry, however, that such flexibility might just let large companies circumvent inconvenient clauses, while local educators remain uncertain about their rights.

One crucial starting point is acknowledging how moral concepts such as equity, autonomy, or cultural respect mean different things across societies (Parekh, 2000). A child in a rural environment may define a successful education as tied to collective stewardship of land or oral traditions. Students in a densely populated city might look for advanced digital literacy and global competitiveness. A universal code that dwells only on privacy or algorithmic transparency might overlook these deeper contexts, prompting frustration on the ground. If teachers cannot reconcile the code's statements with their immediate concerns—like preserving local languages or ensuring that community elders' knowledge counts as legitimate content—the code remains a lofty document detached from reality. This discrepancy highlights the need for ongoing dialogue among teachers, community representatives, and AI designers, so that big-picture principles filter down to local specifics.

Amartya Sen (1999) argued that true development emerges when individuals possess the capabilities to shape their futures. Translating that logic to educational AI, we might see a moral framework emphasise each learner's right to co-determine how technology mediates their experiences. That would entail a baseline rule: no AI system in a school can override teacher or student choices without a tangible "consent" mechanism. If the system prompts a student towards a certain path, it would clearly state the rationale, allowing the user to accept or reject. Teachers would also maintain a right to intervene if they sense that the algorithm's logic conflicts with a pupil's unique needs. Enshrined in policy, such a rule could counterbalance the paternalistic tendencies that sometimes arise when algorithms interpret raw data without context.

Global bodies like UNESCO have started to explore how to encourage ethical AI through culturally inclusive standards (Miller, 2018). But the path from drafting a global declaration to seeing it honoured in local classrooms is fraught with administrative labyrinths, corporate lobbying, and funding pressures. Without something akin to an international coalition that ties compliance to real incentives—like public procurement rules or philanthropic grants—vendors may treat these declarations as optional branding. Meanwhile, local educators struggle to articulate demands for transparency, interpretability, or the freedom to adapt AI outputs. They may only discover the limitations after contracts are signed, at which point the negotiation leverage is minimal.

A more robust route might revolve around bottom-up synergy, reminiscent of how open educational resources gained traction. Teachers across different regions could pool experiences, highlighting how they overcame biased data sets or integrated local knowledge. If such collaboration is supported by structured networks—forums where teachers can share code modifications or strategies for bridging cultural contexts—then a universal code becomes dynamic, not just abstract. Each school that encounters a novel challenge might propose an amendment or best practice, which the global community refines. Over time, a living framework emerges, shaped by those who deal with AI tools daily (Willinsky, 2006). International guidelines become less about prescriptive rules and more about iterative design methods, oriented towards inclusive, context-sensitive usage.

Another stumbling block, though, is that many teachers do not have the time or resources to engage deeply with code or data sets. The pressure to deliver test results or handle large class sizes remains. If vendors claim that opening up systems or ceding local adaptation invites "chaos," schools might be swayed to accept locked-down solutions. That is where teacher associations or unions must come into the process—similar to how medical or legal professionals guard their domains. A shared universal code could require that, as part of any AI deployment, teachers be given professional development not just to operate the tool but to question its ethical underpinnings. This training might cover data sovereignty, interpretability, and methods for incorporating local culture. Only by embedding knowledge of ethical checks into everyday practice can educators move from passive recipients to active co-shapers.

The concept of "embedded moral reflection" also surfaces, echoing moral philosophers like MacIntyre (1984), who saw ethics as inseparable from the ongoing communal practice. If a universal code becomes an inert set of bullet points, it fails to integrate with the living experiences of teachers and students. Instead, schools might set aside regular reflection sessions, where staff discuss how well the AI has aligned with or violated the shared ethical commitments. Did it inadvertently treat certain dialects as errors, or classify socio-economically disadvantaged pupils as high risk without acknowledging their potential strengths? By collectively reviewing these misalignments, educators feed real-time corrections back to developers or district officials. The code's function then becomes a benchmark for continuous dialogue, not a static checklist.

Leanne Betasamosake Simpson (2014) highlights how indigenous knowledge traditions thrive in reciprocal interactions rather than top-down edicts. If a universal code aims to incorporate decolonial perspectives, it must invite local communities to articulate how AI should or should not engage with ceremonial knowledge, land-based teaching, or oral genealogies. The code might specify that any generative platform be restricted from trivialising sacred narratives into game-like tasks. Or that the platform's data retention policy never treat indigenous stories as commercial assets. Such specifics reveal the depth of cultural sensitivity required for universal guidelines to be more than broad statements about "respecting diversity."

Kwame Anthony Appiah (2006) argues for a "rooted cosmopolitanism," suggesting that people can uphold universal moral norms while remaining loyal to local contexts. Adapting that perspective, a universal code can champion broad ethical ends—equity, autonomy, transparency—while leaving the "how" to local professionals. So a district in Southeast Asia might adopt communal-living principles as part of the AI's approach to group tasks. Another region might emphasise the child's individual autonomy to override the platform's suggestions. Each expression resonates with the same moral foundation, but manifested in local idioms and pedagogical patterns. This approach can mitigate the tension that arises when universal frameworks appear to bulldoze local culture.

Walter Mignolo (2011) writes of "border thinking," where different epistemic traditions intersect. Global ed-tech efforts seldom harness this approach, defaulting to mainstream Western logic. But a universal code that fosters border thinking encourages cross-cultural co-design: bridging intangible knowledges, cross-linguistic references, and historically marginalised vantage points in the AI's design. Potentially, the code might demand that any large-scale tool incorporate local datasets or consult community elders, rather than rely on monolithic corpora. Enshrined in a policy requirement, this ensures that non-Western frameworks inform how AI interprets children's expressions of intelligence or creativity.

Some worry that universal codes risk being overly procedural. They might dwell on data usage or transparency while glossing over intangible relational aspects. Yet if done thoughtfully, the code could specify moral responsibilities in child-centred, teacher-guided language—reminding vendors that educational AI is not purely about performance metrics but also about nurturing social and emotional well-being. Gert Biesta (2010) underscores that education involves "being and becoming" as much as "knowing and doing." The code could thus forbid any design that subjects children to relentless data extraction or manipulative behavioural nudges. Even better, it could require that each new algorithmic feature be tested with real teachers and students for how it aligns with socio-emotional growth, not just skill acquisition.

Another dimension of universal frameworks concerns enforcement. While declarations can prompt moral reflection, actual compliance demands structural levers. That might involve linking compliance to public funding: if a vendor's platform fails an ethical review, they lose eligibility for government purchase or philanthropic grants. Alternatively, teacher associations might vow not to adopt solutions that do not meet key standards. In a robust scenario, multiple nations or district consortia unite, forming a bloc that collectively insists on meeting the code's guidelines—thereby compelling vendors to adjust. Critics might call it heavy-handed, but given the moral weight of shaping children's minds, such assertiveness might be warranted. After all, we accept stringent standards for child safety in toys and food; why not for advanced AI software that influences identity and learning?

In the end, a universal moral framework for AI in education appears both necessary and fraught with complications. Its success hinges on balancing top-level universals with local detail, ensuring not only that principles exist on paper but that everyday teachers and communities can adapt them. Historically, calls for ethical standards in tech have faltered when they encountered commercial realpolitik. Yet education, as an arena of future-shaping significance, has a unique claim to push beyond superficial self-regulation, establishing binding norms that protect children's data, creativity, and cultural identity. If local educators, parents, and civil society remain integral to defining and enforcing those norms, the code can become a lived reality, not just a set of abstract ideals. Perhaps then, AI can evolve as a genuinely emancipatory force in classrooms—fulfilling grand visions of personalised support while anchored in collective moral wisdom.

Conclusion

Long before algorithms entered the equation, schools were places of potential upheaval and hope—where diverse minds convened, wrestling with inherited injustices, sparking unforeseen aspirations. The promise of AI in education resurfaced this classic tension, heightening the stakes in every dimension: data, teaching, privacy, equity, and cultural identity. As we now look to consolidate the ethical imperatives that have emerged across the discussions around bias, privacy, data governance, teacher autonomy, and global inequities, it becomes clear that forging a moral foundation for AI in education is not a small task. It requires facing deep contradictions in how schooling is funded, how technology is developed, and how local voices can stand firm against global commercial expansions.

Supporters of AI highlight the new vistas it opens: personalisation at scale, timely feedback for teachers, or a reduction in menial tasks. They see pupils inspired by adaptive lessons that mirror their evolving curiosity, teachers relieved of repetitive grading chores, and administrators able to allocate resources based on real-time analytics. But in parallel, sceptical voices underline the pitfalls of corporate agendas, potential erosion of critical thinking, and overshadowing of intangible classroom bonds. One might observe that each new educational technology, from the radio to the internet, has sparked a similar cycle of hype and caution. The difference with AI is the level of autonomy these systems can wield. They do not merely transmit content: they intervene in decisions that shape a child's sense of self and academic destiny (Noble, 2018).

Some wonder whether the moral debate might fade if teachers and officials simply "use AI sensibly," balancing it against broader judgement. But the complexities run deeper, entwined with commercial IP claims, hidden training data, and algorithmic designs that privilege certain knowledge forms over others. If we treat advanced ed-tech like a benign tool, ignoring the structural asymmetries behind its deployment, we inadvertently legitimise an environment where children's data flows into profit-driven pipelines or standardised metrics, while teachers lose agency in the name of efficiency (Morozov, 2013). A genuine moral foundation would thus need explicit guardrails at multiple layers—legal, institutional, and cultural—rather than vague hopes of "responsible use."

At the heart of these debates are two overarching visions for education. One sees the classroom as a site of economic and cognitive optimisation, forging a future workforce adept at navigating data-saturated industries. The other conceives of schooling as a moral community, fostering holistic growth, empathy, and critical agency. AI can serve either vision, or straddle them in uneasy compromise. Ruha Benjamin (2019) insists that technology is not neutral: its design and usage reflect social priorities. If an AI is shaped by market imperatives, it might emphasise test performance and content mastery. If the moral blueprint instead centres on co-creation and cultural respect, the system might prioritise space for teacher oversight, local knowledge integration, and iterative reflection.

A robust moral foundation must first address **data sovereignty**. Shoshana Zuboff (2019) argues that we cannot champion children's dignity if we treat their emotional states, intellectual journeys, or behavioural patterns as commodities. A code that simply states "we respect privacy" is insufficient if vendors can gather granular data for indefinite retention. Instead, the foundation must specify strict boundaries: minimal data collection, local data storage, open logs available for teacher review, and a child's or parent's right to withdraw or delete records. Translating these principles into actual terms of service or procurement contracts is key. It clarifies that advanced analytics are not a pretext for boundless extraction.

Next is the matter of **bias and cultural nuance**. AI's capacity to replicate historical prejudices reminds us that technology can easily become a Trojan horse for systemic discrimination (Gebru, 2020). If we want a moral code, it must demand that each vendor present a verifiable method for identifying and mitigating bias—an ongoing process rather than a one-time test. In practice, it would require audit trails, third-party reviews, and teacher unions or local committees empowered to halt usage if significant biases surface. Such a stance resonates with what Joy Buolamwini (2018) calls "algorithmic accountability," ensuring that no pupil is automatically pigeonholed based on demographic correlations. The code would also define how to rectify harm: if a system flags pupils wrongly as "high risk," who corrects it and how swiftly? The moral imperative extends to acknowledging that the best defence against bias includes teachers' and families' local insights rather than purely numeric solutions.

Teacher autonomy remains another pivotal theme. For centuries, educators have shouldered the moral and intellectual guidance of children, forging relationships that transcend test scores. If advanced AI is to enrich rather than erode this heritage, educators must retain ultimate interpretive authority. We can imagine a code that stipulates no system can enforce a certain path or outcome without teacher validation and an opt-out for students where appropriate (Nissenbaum, 2010). Indeed, the concept of a teacher as a "human in the loop" is not just about operational correctness, but about upholding the relational, empathy-based dimension of schooling. Where an algorithm suggests an intervention, the teacher's role is to contextualise, verifying if the data truly aligns with the child's lived context or cultural background. The moral foundation, in effect, codifies the teacher's right—and duty—to override AI-driven decisions that clash with professional judgement.

Underpinning these concerns is **equity of access**, bridging divides that separate well-funded districts from rural or low-income ones. Kentaro Toyama (2015) notes that technology magnifies existing advantage unless specifically reoriented for inclusion. Thus, a moral foundation should push beyond rhetorical calls for "closing the digital gap" to ensure sustained public investment in infrastructure, open licensing that fosters local adaptation, and training so that teachers everywhere can harness AI meaningfully. We might specify that any publicly funded AI in schools must run effectively on low-bandwidth setups or integrate offline modes, guaranteeing that learners do not need always-on connectivity. In the same spirit, global authorities could bar ed-tech vendors from profiting off user data in impoverished regions, at least unless local councils explicitly endorse it for mutual benefit. Such measures, though ambitious, rectify the frequent pattern of digital colonialism, where external players gather data while local communities see few tangible improvements.

Tied to that is **cultural recognition**. Although universal codes talk about respecting diversity, the real test is whether AI includes the data sets or design logic that incorporate indigenous languages, diaspora histories, or communal values. If the moral blueprint says each ed-tech solution "must accommodate local epistemologies," someone must define the procedure for verifying that. Perhaps communities review the system's content library or run pilot sessions to see if the AI trivialises sacred ceremonies or interprets collaborative tasks as cheating. The code could set up an appeals process: if a teacher or elder finds content offensive or reductive, they can trigger an immediate reconfiguration or partial ban until the vendor addresses the issue. This level of detail demands local investment in oversight—a shift from paternalistic top-down approaches to real co-governance.

Global policy frameworks might then operate like a scaffolding, not a final blueprint. They would enshrine basic values—data minimisation, interpretability, teacher override, inclusive design—and require each region to embed them in culturally congruent ways. The synergy arises if philanthropic grants or government budgets only fund AI solutions that pass these thresholds, effectively using financial leverage to shape vendor behaviour. Even large firms respond to consistent global demand signals, adjusting offerings to meet ethical codes if they want to secure market share in public education. In an ideal scenario, teachers' unions, parent associations, and civil society collaborate to maintain that global stance, ensuring no company can sidestep accountability by pivoting to less regulated locales.

Of course, the fluid nature of AI calls for ongoing revision. A code written today might not anticipate future developments like emotional-linguistic hyper-personalisation or the integration of neural sensors that track advanced cognitive states. For the foundation to remain relevant, it must function as a "living charter," with regular updates triggered by new technologies or findings (Crawford, 2021). Each update would involve reflection: have we discovered new forms of bias, new modes of data exploitation, or creative expansions that require revised clauses? This iterative approach can mirror the concept of "living documents" in open-source communities, bridging global dialogue with local reflection sessions.

Some fear that these layered regulations and committees could stifle the impetus for bold AI experimentation in schools. They caution that teachers already bound by multiple policies might find more constraints unwelcome. Yet a well-designed moral foundation can create clarity and trust, not red tape. Freed from suspicion that the AI manipulates or profits off children, teachers can embrace generative possibilities with greater confidence. Pupils, too, can explore AI-based creativity or project-based learning, comfortable that their data usage remains respectful and not exploitative. Over time, the synergy of moral clarity and innovative pedagogy might yield deeper engagement rather than stifling it (Freire, 1970).

A more radical angle suggests we rethink the corporate presence altogether. Could advanced AI for schools be developed by public or nonprofit consortia, shaping open frameworks that local communities adapt? A moral blueprint might strongly favour open licensing and public ownership of core modules, so that teachers or local developers can tweak features to align with cultural or linguistic nuances. That approach resonates with the tradition of open educational resources, crossing from static textbooks to dynamic AI tutors. Critics retort that such an approach demands large-scale funding and technical capacity that governments often lack. They argue that partnerships with the private sector are inevitable. If so, the moral code could mandate a co-ownership model, ensuring data sets and generated materials remain in the public domain rather than locked behind vendor paywalls.

Another aspect is addressing the intangible moral growth schools aim to foster. If an AI system is primarily geared to track academic metrics, does it neglect deeper learning about empathy, civic engagement, or ethical reasoning? A moral foundation might specifically require that AI be tested against not just performance outcomes but also intangible socio-emotional measures, verified by teachers. Some might question how to standardise such intangible qualities, but the code can simply note that technology should not overshadow the social dimension of schooling. The vendor might be obliged to show that their AI fosters or at least does not hinder collaborative group tasks, cultural expression, and critical debate—elements crucial to the formation of democratic citizens. The code thereby extends beyond "technical fairness" into the realm of moral and civic values (hooks, 1994).

We also revisit the idea that moral frameworks can be a bulwark against creeping commodification. If ed-tech providers see children's learning purely as a marketplace, we risk commodifying intangible aspects like curiosity or cultural heritage. By contrast, an established ethical code underscores that schools are not typical consumer environments: the well-being and autonomy of children hold a special moral status. Corporate claims of proprietary data usage clash with that moral baseline, reinforcing the necessity for a thorough "informed consent" that includes children (as their capacity matures) and parents or guardians. Instead of a passive check-the-box form, this consent becomes a shared understanding of how the AI is used and how it might shape each student's learning trajectory.

In the end, forging a workable moral foundation for AI in education emerges as an ongoing, multi-actor endeavour. Philosophical abstractions about fairness or transparency must link with teacher training manuals, local oversight councils, global procurement guidelines, and open data practices. If the conversation remains theoretical, ed-tech vendors can continue business as usual, harnessing data under a veneer of social good. If, however, educators, parents, policymakers, and pupils mobilise around these guidelines—ensuring that each contract and deployment respects them—AI might truly evolve from a source of ethical anxiety into a conscientious partner for teaching and learning.

That future vision sees classrooms equipped with advanced tools, but never overshadowing the teacher's relational authority or the learner's capacity for wonder. Pupils enjoy dynamic personalisation, yet remain free to challenge or deviate from recommended paths. Teachers rely on analytics to glean deeper insights, but interpret the results within the broader tapestry of each student's life. Local cultures, languages, and traditions feed into the AI's generative logic, rather than the system imposing universal standards. Data is cherished as a means to refine educational support, not hoarded for monetisation or manipulative ends. By weaving these threads into law, funding models, and daily practice, we confirm that technology's place in schools is ultimately at the service of shared humanistic goals—equity, self-discovery, and communal flourishing.

The question, then, is whether society can muster the resolve. Ed-tech has become a lucrative frontier, and big players will resist constraints that limit revenue or demand open accountability. Teachers and families, often struggling with myriad responsibilities, may find it difficult to push for structural changes. Yet the moral stakes are too high to ignore. Education shapes how we see ourselves, how we treat one another, how we adapt to an ever-changing world. If we allow commercial or technocratic logics to dominate, the consequences will ripple for decades. Conversely, if local communities, teacher unions, civil rights advocates, and forward-thinking policymakers align around a robust moral code, the future holds promise. AI could evolve as a genuine ally—a flexible, context-aware resource that amplifies, rather than supplants, the magic of human teaching.

In that sense, the final image is not one of a sterile, data-driven classroom but a lively space where teachers and pupils harness advanced, ethically grounded AI to enrich their shared journey. The systems might adapt to shifting cultural nuances, integrate local narratives, and reflect teacher wisdom. Students remain at the centre, exploring new perspectives, safely guided by a code that safeguards their rights and upholds communal values. This equilibrium—of technology intertwined with moral clarity—represents education's best hope for upholding its democratic mandate and ensuring that the next generation inherits a world where human dignity and creativity remain paramount, even in the face of profound digital transformation.

About the Author

Dr. Neil Hopkin is a globally recognised thought leader in international K-12 education, and serves as the Director of Education at Fortes Education.

His extensive academic background includes advising UK government bodies and spearheading significant educational initiatives, particularly with the EdTech, Early Years, Higher Education and Teacher Professional Development fields, equipping him with invaluable insights and expertise.
As the head of Fortes' Academic Leadership Team, Dr. Hopkin is responsible for overseeing academic performance, operational efficiency, curriculum development, and staff professional development across Fortes Education institutions.

FORTES EDUCATION
1981 Enlighten • Empower • Excel

For more information contact Dr Neil Hopkin at:

www.sunmarke.com www.risdubai.com

Bibliography

Appiah, K.A. (2006). *Cosmopolitanism: Ethics in a World of Strangers*. New York: W.W. Norton & Company.

Arora, P. (2019). *Leisure Commons: A Spatial History of Web 2.0*. London: Routledge.

Ben Green, B. (2019). *The Smart Enough City: Putting Technology in Its Place to Reclaim Our Urban Future*. Cambridge, MA: MIT Press.

Benjamin, R. (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Cambridge: Polity.

Biesta, G. (2010). *Good Education in an Age of Measurement: Ethics, Politics, Democracy.* London: Routledge.

Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.

boyd, d. (2014). *It's Complicated: The Social Lives of Networked Teens*. New Haven, CT: Yale University Press.

Boyle, J. (2008). *The Public Domain: Enclosing the Commons of the Mind*. New Haven, CT: Yale University Press.

Broussard, M. (2018). *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge, MA: MIT Press.

Brown, J.S. (2002). "Learning in the Digital Age." In *The Internet and the University: Forum 2001*, edited by M. Devlin, R. Larson, and J. Meyerson, 65–91. Boulder, CO: EDUCAUSE.

Buolamwini, J. (2018). "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." *Proceedings of Machine Learning Research*, 81:1–15.

Crawford, K. (2021). *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven, CT: Yale University Press.

Crouch, C. (2004). *Post-Democracy*. Cambridge: Polity Press.

DeRosa, R. (2016). "OER: Bigger than Affordability." *Educause Review.* Available at: https://er.educause.edu/ [Accessed (date)].

D'Ignazio, C. and Klein, L. (2020). *Data Feminism*. Cambridge, MA: MIT Press.

Downes, S. (2012). "Connectivism and Connective Knowledge: Essays on Meaning and Learning Networks." *Self-published online monograph*. Available at: http://www.downes.ca/

Edsger Dijkstra (1982). Selected letter on "On the fact that the Atlantic Ocean has two sides," re: technological clarity. Available in EWD archive, University of Texas.

Eubanks, V. (2017). *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor.* New York: St. Martin's Press.

Floridi, L. (2013). *The Ethics of Information*. Oxford: Oxford University Press.

Freire, P. (1970). *Pedagogy of the Oppressed*. New York: Herder and Herder.

Gandy Jr., O.H. (1993). *The Panoptic Sort: A Political Economy of Personal Information*. Boulder, CO: Westview Press.

Gebru, T. (2020). "Reflection on AI Ethics and Inclusivity." *Medium*. Available at: https://medium.com/@timnitGebru [Accessed (date)].

Giroux, H.A. (2011). *On Critical Pedagogy*. New York: Bloomsbury.

Green, D. (2019). *The Smart Enough City: Putting Technology in Its Place to Reclaim Our Urban Future*. Cambridge, MA: MIT Press.

Hall, W. (2019). "Towards a New Ethical and Regulatory Framework for AI." In *Ethics in AI*, edited by M. Blaauw, 45–68. London: Springer.

Harding, S. (1991). *Whose Science? Whose Knowledge? Thinking from Women's Lives*. Ithaca, NY: Cornell University Press.

Hayles, N.K. (2012). *How We Think: Digital Media and Contemporary Technogenesis*. Chicago: University of Chicago Press.

hooks, b. (1994). *Teaching to Transgress: Education as the Practice of Freedom*. New York: Routledge.

Jasanoff, S. (2004). *States of Knowledge: The Co-production of Science and Social Order*. London: Routledge.

Kohn, A. (2015). *Schooling Beyond Measure and Other Unorthodox Essays about Education*. Portsmouth, NH: Heinemann.

Lanier, J. (2010). *You Are Not a Gadget: A Manifesto*. New York: Vintage.

Latour, B. (2005). *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford: Oxford University Press.

Mazzucato, M. (2018). *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*. London: Penguin.

McKee, H.A. (2017). "Policy, Purpose, Practice, and Privacy: Teacher Agency in the Age of Big Data." *Computers and Composition*, 44: 76–86.

Miller, R. (2018). "Futures Literacy: Transforming the Future to Transform the Present." *European Journal of Futures Research*, 6(25).

Morozov, E. (2013). *To Save Everything, Click Here: The Folly of Technological Solutionism*. New York: PublicAffairs.

Mulligan, D.K. (2019). "Imagining Governance for Privacy." *New Media & Society*, 21(2): 620–626.

Narayan, U. (1997). *Dislocating Cultures: Identities, Traditions, and Third World Feminism*. New York: Routledge.

Nissenbaum, H. (2010). *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford University Press.

Noble, S.U. (2018). *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: NYU Press.

O'Neil, C. (2016). *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown.

Parekh, B. (2000). *Rethinking Multiculturalism: Cultural Diversity and Political Theory*. London: Macmillan.

Postman, N. (1993). *Technopoly: The Surrender of Culture to Technology*. New York: Vintage.

Rushkoff, D. (2019). *Team Human*. New York: W.W. Norton & Company.

Sahlberg, P. (2011). *Finnish Lessons: What Can the World Learn from Educational Change in Finland?* New York: Teachers College Press.

Selwyn, N. (2019). *Should Robots Replace Teachers? AI and the Future of Education*. Cambridge: Polity Press.

Sen, A. (1999). *Development as Freedom*. Oxford: Oxford University Press.

Shneiderman, B. (2020). *Human-Centered AI*. Oxford: Oxford University Press.

Smith, L.T. (1999). *Decolonizing Methodologies: Research and Indigenous Peoples*. London: Zed Books.

Sunstein, C.R. (2014). *Why Nudge? The Politics of Libertarian Paternalism*. New Haven, CT: Yale University Press.

Toyama, K. (2015). *Geek Heresy: Rescuing Social Change from the Cult of Technology*. New York: PublicAffairs.

Turkle, S. (2015). *Reclaiming Conversation: The Power of Talk in a Digital Age*. New York: Penguin.

Tufekci, Z. (2018). *Twitter and Tear Gas: The Power and Fragility of Networked Protest*. New Haven, CT: Yale University Press.

Watters, A. (2015). "The History of the Future of Ed-Tech." *Hack Education*. Available at: http://hackeducation.com

Wiliam, D. (2011). *Embedded Formative Assessment*. Bloomington, IN: Solution Tree Press.

Williamson, B. (2017). *Big Data in Education: The Digital Future of Learning, Policy and Practice*. London: SAGE.

Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.