

GOVERNING THE ALGORITHMIC CLASSROOM: A COMPARATIVE POLICY ANALYSIS OF NATIONAL AI STRATEGIES IN PUBLIC EDUCATION SYSTEMS

Debopom Ghosh

Department of Public Administration, University of Chittagong, Chittagong, Bangladesh

debopomghosh@std.cu.ac.bd

RESEARCH ARTICLE | DOI: <https://doi.org/10.37792/8s7zjb81>

Received: May 27, 2026 | Revised: July 9, 2026 | Accepted: July 16, 2026 | Published: September 2026

License: Creative Commons Attribution 4.0 International (CC BY 4.0)

ABSTRACT

The rapid adoption of artificial intelligence (AI) within public education systems represents a major shift in state governance, regulatory oversight, and public administration. This paper is positioned primarily as a comparative policy analysis: it examines the structural, regulatory, and pedagogical dimensions of national AI strategies across seven core geopolitical archetypes — rights-based restrictive regimes (the European Union), decentralized market-driven systems (the United States), an emerging federal-democratic regime (India), state-centric optimization regimes (China and Singapore), and sovereign digital-development regimes (Bangladesh and Nigeria’s Edo State) — with Brazil, Chile, and the United Arab Emirates analyzed as supplementary, illustrative cases. Using a hybrid comparative design combining Most Different Systems Design (MDSD) and Most Similar Systems Design (MSSD), the paper’s central original contribution is the Educational AI Policy Rigor Index (EAPRI), a five-domain, 0–4 scoring rubric that operationalizes the “rhetoric-reality gap” in national policy as a measurable construct rather than a rhetorical one. The analysis further examines high-stakes algorithmic failures, including the 2020 Ofqual examination scandal in the UK, and recent experimental evidence on the “Performance-Learning Gap,” to characterize the systemic risks of “metacognitive laziness” and corporate control capture in public education. Building on these findings, the paper operationalizes the Sustainable Public Innovation Ecosystem (SPIE) model as a governance blueprint, offering a capital-sequenced set of policy recommendations and a non-algorithmic alternative assessment architecture intended to preserve pedagogical agency and democratic equity in the algorithmic age.

Keywords : *Algorithmic Governance, Comparative Policy Analysis, Sociotechnical Imaginaries, Data Colonialism, Pedagogical Agency, Sustainable Public Innovation Ecosystems (SPIE).*

I. INTRODUCTION

The global use of artificial intelligence (AI) in public education systems is a radical paradigm shift in state governance, regulatory authority, and public administration. In 2025 and 2026, the international dialogue on EdTech definitively transitioned from micro-level classroom interventions and pilot pedagogical experiments to macro-level concerns about systemic regulation, national sovereignty, and democratic accountability. Public education is no longer simply being technically upgraded, but structurally consolidated, where the basic mechanisms of teaching, student assessment, admissions and institutional management are being mediated by automated decision-making systems.

The global administrative change can be characterized as a structure tension between two competing visions of public sector reform: the **Optimization Paradigm** and the **Restructuring Paradigm**. The Optimization Paradigm seeks to automate already existing bureaucratic and pedagogical processes in order to increase the system's efficiency. This paradigm employs automated tools for low-stakes tasks such as automated grading, reporting, and standardized scheduling under doctrines of new public management. States are using technology to combat severe teacher shortages and stretched public school budgets by reducing administrative friction and automating resource allocation. However, this approach assumes that public education is a static operational environment, where historical inequalities and systemic biases are routinely hard-coded into predictive models in the name of technocratic neutrality.

Conversely, the Restructuring Paradigm imagines AI as a lever for shifting the power balance in the classroom. Democratic equity is the underlying principle of this paradigm, which seeks to increase student data sovereignty, safeguard teacher professional autonomy, and leverage technology to advance critical inquiry, collaborative problem-solving, and culturally responsive learning environments. The Restructuring Paradigm eschews the image of the machine as an infallible evaluator or replacement for human teaching, preferring to see AI as an adaptive Socratic partner who preserves the “productive friction” necessary for genuine cognitive growth. The progressive potential of this paradigm is rarely scaled up due to persistent bureaucratic inertia, political path dependency, and standardization tendencies in legacy public education systems.

This tension is compounded by a serious “rhetoric-reality gap” in national education policies. National strategies and ministerial guidelines tend to use aspirational language, with commitments to human-centered, inclusive, and responsible AI, but are often lacking in binding statutory mechanisms, technical audit standards, and budgetary resources to make good on these ethical aspirations. At the same time, the unregulated use of generative AI has caused a serious pedagogical crisis. The OECD Digital Education Outlook 2026 reports that the outsourcing of cognitive tasks to general purpose commercial chatbots has led to a large “Performance-Learning Gap”. Although students utilize generative tools to enhance the performance of tasks in the short

term, their real cognitive retention, critical thinking, and performance in the examination deteriorate resulting in a phenomenon called “metacognitive laziness”.

We conduct a comparative policy analysis to critically evaluate the structural, regulatory and ethical aspects of national AI strategies in terms of the limits of **rights-based restrictive models** (European Union), **decentralized market-based systems** (United States), **state-centric optimization regimes** (China and Singapore), **emerging democratic models** (India) and **sovereign digital development models** (Bangladesh and Sub-Saharan Africa). This analysis maps these various geopolitical archetypes, revealing systemic loopholes that jeopardize public education and proposes a way forward through a Dynamic Governance Framework based on Sustainable Public Innovation Ecosystems (SPIE).

This analysis is positioned primarily as a comparative policy analysis. Its central original contribution is the Educational AI Policy Rigor Index (EAPRI), introduced in Section III, which operationalizes the rhetoric-reality gap as a standardized, five-domain measurement applied consistently across cases rather than as a descriptive label. The theoretical apparatus developed in Section II and the Sustainable Public Innovation Ecosystem (SPIE) framework proposed in Section VI support this comparative contribution; they are not advanced as freestanding theoretical or design contributions in their own right.

Three research questions guide this investigation. **RQ1:** To what extent does the formal statutory design of national educational AI policy diverge from its operational enforceability, and can this divergence be measured systematically across jurisdictions? **RQ2:** Which structural and political-economic mechanisms — platformisation, principal-agent misalignment, and bureaucratic path dependency — best explain the persistence of the Optimization Paradigm despite widespread rhetorical commitment to rights-based, human-centered AI governance? **RQ3:** What sequence of institutional reforms would allow a resource-constrained state to transition from rhetorical to operational governance without capital investments it cannot presently afford?

Seven jurisdictions constitute the paper's core, EAPRI-scored comparative set: the European Union, the United States, India, China, Singapore, Bangladesh, and the sub-national case of Nigeria's Edo State. Brazil, Chile, and the United Arab Emirates are examined as supplementary, illustrative cases that stretch the boundaries of these archetypes — through multilateral digital-sovereignty partnership, constitutional neurorights, and mandatory state-directed curriculum design, respectively — without being treated as core comparators. The full case-selection logic is set out in Section III.

II. LITERATURE REVIEW

This analysis employs a theoretical framework that combines science and technology studies (STS), public choice theory and critical data studies to systematically decode the legislative and administrative design of national AI strategies in education.

a. Sociotechnical Imaginaries in Education Policy

Sociotechnical imaginaries are collectively held, institutionally stabilized and publicly performed visions of desirable futures, animated by shared understandings of forms of social life and social order attainable through, and supportive of, advances in science and technology. In the world of public education policy, states write their geopolitical anxieties, economic aspirations and civilizational narratives into their digital education blueprints. An analysis of national guidelines reveals that policy integration is actively pursued by three dominant competing sociotechnical imaginaries:

- **The Imaginary of Inevitable Disruptiveness:** This paradigm positions AI as an external, deterministic force of nature that public education systems need to rapidly adapt to, or risk immediate economic obsolescence. This kind of imaginary imagines human actors as passive subjects of technological advance, provoking rapid, unvetted procurement cycles under the banner of national competitiveness.
- **The Imaginary of Transformative Potential:** This is a utopian techno-solutionist vision that posits AI will perpetually address systemic educational inequities and institutional deficits through hyper-personalized software. It imagines a future in which the human teacher is supplemented or ultimately replaced by a perfect, omnipresent digital tutor customized to the exact cognitive makeup of each learner.
- **The Imaginary of the Surveillance Society:** This framework argues that the algorithmic classroom, as a dystopian counter-narrative, will generate a panoptic regime of biometric monitoring, behavioral profiling, and continuous extraction of student data. It points to the brutal erosion of student privacy and intellectual autonomy under the weight of automated tracking and behavioral classification systems.

Table 1: Typology of Sociotechnical Imaginaries in National Educational AI Policy

Sociotechnical Imaginary	Stated Policy Objectives	Prominent Geopolitical Adopters	Key Strategic/Ministerial Policy Documents	Core Institutional Actors
Inevitable Disruptiveness	Rapid technical adoption; global economic competitiveness; workforce computer science readiness.	United States (Georgia) ; United Arab Emirates.	Georgia Senate Bill 179 ; UAE "From Skill to Leadership" Campaign.	State chambers of commerce; Ministry of Advanced Technology (UAE).
Transformative	Hyper-personalized student tutoring;	Singapore ; Nigeria (Edo)	Singapore National AI Strategy 2.0 ; Edo State Universal Basic	Ministry of Education (Singapore) ;

Potential	automated administrative workload relief.	State).	Education (SUBEB) Reforms.	Board	World Bank Group.
Surveillance Society	Fundamental human rights protection; biometric data limits; cognitive and mental privacy shielding.	European Union ; Chile.	European Union ; Artificial Intelligence Act (2024) ; Constitutional Article 19.	Chilean Reform	European Parliament ; Future Challenges Commission (Chile) ; Supreme Court of Chile.

b. Metric Power, Platformisation, and Principal-Agent Dynamics

The administrative enactment of these imaginaries is sustained by the ‘metric power’ of the systematic datafication of human behaviour to sort, classify and evaluate individuals. In education, the power of the metric produces ‘student data subjects’ from durable digital traces of learning activity. When proprietary algorithms are used to analyze these data doubles to make high-stakes determinations, they risk locking students into predictive trajectories that foreclose future academic and economic opportunities.

In public administration, the Principal-Agent problem is the misalignment of the goals of the principal (the state, public school systems, and the democratic public) and the goals of the agent (private EdTech corporations) that are charged with delivering public services. The more public schools rely on proprietary black-box AI platforms from a monopolistic tech oligarchy, the more decision-making power is ceded from democratically elected school boards and public administrators to corporate product managers.


This change in control is accelerated by the “platformisation” of public education. In this model, public schools are not merely licensing software but becoming part of corporate ecosystems designed for the continuous extraction and monetization of data. As the infrastructure of schooling – from learning management systems to communication portals to assessment engines – is outsourced to private tech conglomerates. This vertical integration leads to deep “control capture”. Private EdTech companies earn monopoly rents on educational services and set the standards of curriculum, pedagogy and behaviour. Democratic accountability and the public good of education are under major threat as the state (as the principal) has little leverage left to audit or replace these proprietary systems.

c. Decoloniality and Its Discontents: Critiquing Data Colonialism

The platformisation of educational technology is a modern extension of colonial logics via “Data Colonialism” from the perspective of critical data studies. As Nick Couldry and Ulises Mejias put it, data colonialism is the “contemporary process by which global technology appropriates human life as a direct, free source of raw data for capital accumulation.” This extraction is very specific in the form of global education:

Technocratic dependency from the adoption of generic, Western-centric Large Language Models (LLMs) and predictive analytics suites in public education systems (especially in the Global South). Epistemic injustice and epistemicide are the result of this – the systematic marginalization of local languages, indigenous worldviews, and regional pedagogical values to make way for a homogenized, Silicon Valley-approved curriculum, optimized solely for Northern market demands.

But a rigorous approach to global educational AI governance must grapple with the contested assumptions embedded in the data colonialism framework. The Couldry-Mejias paradigm has been criticized for its conceptualization of technology adoption in developing nations as passive extraction, which obscures actual agency, regional resistance and strategic navigation of Global South governments. It risks reducing developing states to homogeneous, helpless grounds, ignoring how sovereign governments employ state power to negotiate technology transfer that is mutually beneficial, or to demand data localisation or build indigenous alternatives.



Moreover, a blanket dismissal of EdTech as necessarily extractive overlooks documented examples of targeted culturally adapted digital interventions that have produced measurable learning gains and closed historical achievement gaps. Thus, a nuanced analysis is required to balance the uncovering of the very real structures of Big Tech monopolization with recognition of the agency of local actors to co-design, adapt and assert digital sovereignty over their educational futures.

d. The Political Economy of Bureaucratic Inertia

The Persistence of the Optimization Paradigm is not a mere technical accident, but is structurally sustained by the political economy of public administration. At their core, bureaucratic organizations are path dependent, budget maximizing, and risk averse.

First, the procurement cycles in public school systems are extremely complex and resource intensive, creating large high-cost barriers to exit; When a school district adds a corporate learning management system (eg, Google Classroom or Microsoft Teams) to its administrative infrastructure, the institutional switching costs are prohibitive. This structural lock-in allows private vendors to extract monopoly rents and capture long-term public funding.

Second, school administration and public policy are structured to reward short-term, quantitative performance measures at the expense of long-term qualitative educational development. Automated grading software, standardized scheduling algorithms, predictive attendance trackers all deliver immediate administrative efficiency gains easily reportable to oversight boards and political superiors while replicating systemic class biases. It is a powerful partnership of risk-averse public administrators and profit-oriented technology vendors that harnesses educational policy to the Optimization Paradigm and effectively sidelines the messier, more resource-intensive reforms called for by the Restructuring Paradigm.

Read together, these four theoretical lenses perform distinct and complementary analytical functions in the sections that follow. Sociotechnical imaginaries (Section II.a) supply the interpretive vocabulary used to classify each jurisdiction's national policy narrative in Table 1 and in the case analyses of Section IV. The principal-agent and platformisation literature (Section II.b) supplies the causal mechanism behind “control capture,” used throughout Section IV to explain why states retain nominal regulatory authority while ceding effective decision-making power to EdTech vendors. Critical data studies and the data colonialism framework (Section II.c) — applied with the caveats discussed above — inform the Child Data & Sovereign Privacy Guardrails (CDP) and Equitable Resource & Infrastructure Allocation (ER) domains of the EAPRI index constructed in Section III. Finally, the political economy of bureaucratic inertia (Section II.d) explains the persistence of the Optimization Paradigm identified empirically across nearly every jurisdiction scored in this analysis, and motivates the phased, capital-sequenced implementation roadmap proposed in Section VI.c. This sequencing — imaginaries to narrative, principal-agent theory to mechanism, critical data studies to measurement, bureaucratic inertia to persistence and remedy — is what allows the four theories to function as a single analytical apparatus rather than four independent literatures. Table 2 below fixes the definitions of the recurring analytical terms used throughout the remainder of the paper.

Table X-DEF: Conceptual Definitions of Key Analytical Terms

Term	Definition
Rhetoric-Reality Gap	The measurable divergence between a national policy's aspirational, rights-based language and its binding statutory enforceability; operationalized in this paper via the EAPRI index (Section III).
Control Capture	The transfer of effective decision-making authority over curriculum, pedagogy, and behavioral standards from democratically accountable public administrators to private EdTech vendors, arising from

		platformisation and vendor lock-in (Section II.b).
Metacognitive Laziness		A form of cognitive offloading in which learners bypass active self-regulatory processes (goal-setting, error diagnosis, iteration) by accepting AI-generated outputs uncritically; documented empirically in Section V.b–c.
Productive Friction / Productive Struggle		The cognitively necessary difficulty of unassisted problem-solving that drives schema formation and long-term retention — the pedagogical property that frictionless, general-purpose chatbots remove.
Sovereign Infrastructure	AI	State-owned, state-funded, or state-commissioned AI systems hosted on domestically controlled cloud infrastructure, used in place of commercial platforms to avoid data extraction and enable pedagogical ring-fencing (Section VI.a–b).
Optimization Paradigm		A technocratic governance orientation that automates existing bureaucratic and pedagogical processes to improve efficiency without altering underlying power relations or the historical biases embedded in the system.
Restructuring Paradigm		A governance orientation that uses AI to redistribute power toward students and teachers, through data sovereignty, pedagogical autonomy, and Socratic rather than directive AI design.

III. RESEARCH METHODOLOGY AND FRAMEWORK DESIGN

The analysis uses a robust comparative policy framework to systematically unpack legislative and administrative design of national AI strategies in education. This study employs a hybrid comparative design, the **Most Different Systems Design** (MDSD) and the **Most Similar Systems Design** (MSSD) to provide a solid methodological basis. The country cases are not illustrative examples.

The main methodological tool is the Most Different Systems Design (MDSD). This method involves the selection of cases that are very different on many systemic, cultural, economic and political-institutional variables but have the same outcome or are confronted with the same basic phenomenon. A common outcome of this research is the systemic shift back to the Optimization Paradigm and the growing “rhetoric-reality gap” in educational AI governance. The analysis examines jurisdictions with significant structural variation – the European Union, the United States, India, China, Singapore, Bangladesh and the sub-national structures of Nigeria – to isolate whether this technocratic drift is an inevitable consequence of commercial platformisation or a function of particular state-level regulatory choices.

In addition to this macro MDSD, the study employs a Most Similar Systems Design (MSSD) to compare cases with similar key explanatory factors and institutional structures but different regulatory outcomes. This MSSD is used only to disaggregate Singapore and China.

Both are highly centralized, high-performing East Asian systems, driven by national human capital development and technological optimization. But Singapore is under a technocratic-democratic framework under the rule of law. China is under single-party authoritarian optimization. Likewise, the framework of MSSD is used to compare Bangladesh and India. Both are densely populated, fast-digitizing South Asian states with acute rural-urban digital divides, but they employ different structural legal mechanisms to protect data privacy and guide educational development.

To ensure full transparency of case selection, this study analyzes ten jurisdictions in total, divided into two tiers. Seven constitute the core comparative set subjected to the MDSD/MSSD logic described above and scored on the full EAPRI index in Table 4: the European Union, the United States, India, China, Singapore, Bangladesh, and the sub-national case of Nigeria's Edo State. Three further jurisdictions — Brazil, Chile, and the United Arab Emirates — are analyzed as supplementary, illustrative cases in Section IV.e. These cases were selected because each stretches the paper's core archetypes in an analytically useful way that the seven core cases do not capture: Brazil illustrates multilateral, G20-coordinated digital-sovereignty partnership; the United Arab Emirates illustrates mandatory, state-directed K-12 curriculum design without the authoritarian character of the Chinese case; and Chile illustrates constitutional-level regulation of a narrower but adjacent policy domain — cognitive and neural privacy — that the EAPRI index's five domains do not fully capture. For this reason, Chile is discussed narratively but is not assigned an EAPRI score in Table 4: its governing instrument (the constitutional neurorights reform) speaks directly to only two of the five EAPRI domains (Child Data & Sovereign Privacy Guardrails and Pedagogical Autonomy & Agency Protections) and does not address algorithmic transparency, statutory enforceability of EdTech procurement, or infrastructure equity in a way that would make a composite score meaningful. Brazil and the United Arab Emirates, by contrast, are scored in Table 4, since their governing instruments speak to all five domains.

Table 3: MDSD/MSSD Comparative Design Logic

Comparison Type	Cases Compared	Shared / Divergent Factor	Analytical Purpose
MDSD (primary)	European Union, United States, India, China, Singapore, Bangladesh, Nigeria (Edo State)	Maximally different political systems, wealth levels, and legal traditions	Isolate whether drift toward the Optimization Paradigm is structurally inevitable or a function of specific state choices
MSSD (secondary pair 1)	Singapore vs. China	Both centralized, high-performing, East Asian, technology-driven human-capital states	Isolate the effect of rule-of-law / democratic accountability, holding centralization and human-capital strategy constant
MSSD	India vs. Bangladesh	Both densely	Isolate the effect of differing

(secondary pair 2)	populated, digitizing South Asian states with acute rural- urban divides	fast- architecture, holding regional and developmental context constant
-----------------------	---	--

Evidence for EAPRI coding was drawn from primary legislative and regulatory texts (e.g., the EU AI Act and its annexes, India's DPDP Act and Rules, California's proposed AB 1159 alongside its enacted KOPIPA and ELPIPA predecessor statutes, and Bangladesh's National AI Policy Draft V2.0), supplemented by official government strategy documents, international assessments (the OECD, UNESCO, and UNDP reports cited throughout), and, where primary legislative text was unavailable, reputable secondary legal analysis. Coding covers the period from January 2024 through May 2026. Scores were assigned by a single coder (the author); this is a limitation acknowledged here rather than presented as inter-rater-validated, and independent replication by a second coder is identified as a priority for future extensions of this index. Within each domain, whole-number scores (0, 1, 2, 3, 4) correspond to the anchor descriptions given below; half-point scores (e.g., 1.5, 2.5, 3.5) are assigned where a domain contains a binding legal mechanism whose coverage is partial — for instance, enforceable at the national level but not harmonized across sub-national jurisdictions (as in the United States and India), or covering some but not all AI use-cases identified in Annex III of the EU AI Act. The detailed documentary basis for each domain score, jurisdiction by jurisdiction, is embedded in the corresponding narrative case analysis in Section IV; Table 4 should be read alongside that discussion rather than as a standalone quantification.

In this paper, we move beyond purely descriptive narratives and introduce and apply a standardized scoring rubric, the Educational AI Policy Rigor Index (EAPRI). This index assesses the gap between rhetoric and reality across five core statutory areas on an 0-4 scale.

- **Domain 1: Statutory Enforceability (SE):** Assesses whether the policy depends on voluntary, non-binding guidelines (0) or includes legally binding, codified mandates with designated enforcement bodies and punitive sanctions (4).
- **Domain 2: Algorithmic Transparency & Audit Standards (AT):** Evaluates if proprietary vendor systems are protected as trade secrets with no external verification (0) or are subject to mandatory, independent pre-deployment conformities and continuous bias testing (4).
- **Domain 3: Child Data & Sovereign Privacy Guardrails (CDP):** Assesses if child data is obtained without age-gating or parental consent (0), or if it is subject to absolute prohibitions on behavioral profiling, targeted advertising, and cross-border commercial extraction (4).
- **Domain 4: Pedagogical Autonomy & Agency Protections (PA):** Analyzes whether automated systems can make high-stakes, unreviewable decisions regarding student grading, placement, or teacher employment (0) or if human-in-the-loop overrides and Socratic pedagogical designs are legally mandated (4).

- **Domain 5: Equitable Resource & Infrastructure Allocation (ER):** Evaluates the use of technology reinforcing existing rural-urban or socioeconomic disparities with unfunded mandates (0) or is the state actively funding, subsidizing, and building sovereign, multilingual, and localized open source digital infrastructure (4).

The composite EAPRI score is calculated as the unweighted arithmetic mean of the five domains:

$$EAPRI = (SE + AT + CDP + PA + ER) / 5$$

A high score (≥ 3.0) indicates a highly operationalized, rights-based governance framework aligning with the Restructuring Paradigm, whereas a low score (< 2.0) exposes a highly rhetorical, optimization-focused regime.

Table 4: Educational AI Policy Rigor Index (EAPRI) Classifications by Geopolitical Archetype

Geopolitical Archetype	Statutory Enforcement (SE)	Algorithmic Transparency (AT)	Child Data Privacy (CDP)	Pedagogical Autonomy (PA)	Infrastructure Equity (ER)	Composite EAPRI Score	Classification
European Union	4.0	4.0	3.5	3.5	2.0	3.40	Restructuring Paradigm (Rights-Based)
Singapore	3.5	3.0	3.0	3.0	4.0	3.30	Restructuring Paradigm (Technocratic)
United States	2.0	1.0	2.5	1.5	1.5	1.70	Optimization Paradigm (Decentralized)
India	3.0	2.0	3.5	2.0	2.0	2.50	Transitional (Hybrid-Democratic)
China	3.5	2.0	2.0	1.5	3.5	2.50	Optimization Paradigm (Centralized)

Brazil	1.5	1.5	2.0	1.5	2.0	1.70	Optimization Paradigm (Decentralized)
United Arab Emirates	4.0	2.5	2.5	3.0	4.0	3.20	Restructuring Paradigm (Technocratic)
Bangladesh	1.0	1.0	1.5	1.0	1.0	1.10	Rhetorical Optimization
Nigeria (Edo State sub-case)	3.0	3.0	2.0	3.5	3.0	2.90	Restructuring Paradigm (Sub-National)

IV. RESULTS: COMPARATIVE ANALYSIS OF GEOPOLITICAL ARCHETYPES

a. The Rights-Based & Restrictive Model (European Union)

The European Union leads the world in “hard governance” of educational AI, with sweeping supranational legislation designed to contain algorithmic overreach. The EU decisively departed from non-binding ethical pronouncements to a strict, enforceable regulatory regime with the phased implementation of the pioneering EUAI Act.

i. Annex III and the High-Risk Taxonomy

At the core of the EU AI Act's educational taxonomy is the classification of specific AI use cases as **High-Risk** under Annex III (Point 3). These high-risk educational systems include :

- AI systems intended to be used for the purposes of access or admission decisions or to allocate natural persons to educational and vocational training institutions at all levels.
- AI systems used to assess learning outcomes, including where the learning outcomes are used to influence and guide the learning of individuals.
- AI systems designed to assess the appropriate level of education an individual will receive or access.
- AI systems for monitoring and detecting students’ prohibited behavior during exams (automated proctoring).

Such a high risk classification triggers a full set of legally binding compliance requirements before any EdTech tool can be legally placed on the European market or used in public schools. Providers must implement a continuous, documented risk-management process, use training and testing datasets that are highly representative and free of discriminatory bias, provide extensive technical documentation showing that the system is robust, allow automatic

logging of events and provide clear instruction sets to deployers. The schools are needed critically as “deployers” of these high-risk systems to keep human oversight of them, so that the automated recommendations do not replace the final human pedagogical decision.

ii. Article 5 and the Biometric Emotion Prohibition

In addition to regulating high-risk systems, the EU AI Act introduces a de-facto zero-tolerance threshold for applications deemed to be an ‘Unacceptable Risk’. *Article 5(1) (f)* - the Act bans the use of AI systems that infer the emotions or intentions of a natural person in the workplace and educational institutions on the basis of biometric data in its entirety.

The European legislature implemented this ban acknowledging the structural power imbalances and asymmetric relationships that are intrinsic to educational settings where students are in a position of extreme vulnerability. A strong scientific consensus supporting this ban on biometric emotion recognition points to the scientific invalidity, cultural specificity and unreliability of inferring emotional states from facial, vocal or bodily cues. The EU’s ban on these systems is a pre-emptive strike against the commercialization of panoptic “smart classrooms” that seek to measure student engagement and compliance through biometric surveillance. Paired with this broad rights framework, the Product Liability Directive (EU) 2024/2853 updates liability rules for defective digital goods, providing students harmed by algorithmic grading errors or discriminatory profiling a simple path to claim monetary redress from tech providers.

This rights-based architecture should not be read as fully resolved in practice. The AI Act is enforced through each member state’s own designated national market surveillance authority, and these authorities vary considerably in technical auditing capacity, staffing, and enforcement experience; meaningful compliance gaps between better-resourced and less-resourced member states are therefore likely, despite uniform statutory text. The EAPRI score assigned to the European Union in Table 4 should accordingly be read as a ceiling on best-case, well-resourced enforcement rather than a description of uniform practice across all twenty-seven member states.

b. The Innovation-Driven & Decentralized Model (United States)

The European model is one of accepted centralized supranational mandates. The United States is not. It is a fractured and highly decentralized model. The federal government has only issued non-binding blueprints and voluntary frameworks. Individual states and local school districts have been left to do the heavy lifting of regulation, creating high compliance friction and a wide gap between rhetoric and reality.

The US approach is based on the idea of inevitability of disruptiveness and heavily relies on the Optimization Paradigm to maintain international tech competitiveness and ease administrative processes. However, by 2026, the legislative landscape began to shift with aggressive state-level actions to address critical data privacy vulnerabilities. For instance, California’s proposed Assembly Bill 1159 (AB 1159) would, if enacted, establish the Higher

Education Student Information Protection Act (HESIPA). As of mid-2026 the bill remains pending in the California Legislature and has not been signed into law. AB 1159 would expand the state's current K-12 student privacy laws – the K-12 Pupil Online Personal Information Protection Act (KOPIPA) and the Early Learning Personal Information Protection Act (ELPIPA) – to higher education students.

If enacted, the primary structural interventions of California's AB 1159 would include:

- **The Prohibited AI-Training Mandate:** It explicitly prohibits EdTech providers and third-party vendors from using covered student personal information to train, develop, or improve artificial intelligence and generative AI models, unless the data is strictly deidentified by strict statistical standards.
- **Sensitive Category Shielding:** The bill prohibits the collection, use, or disclosure of highly sensitive student data to protect vulnerable populations from automated profiling, including reproductive and sexual health records, immigration status, sexual orientation, and gender identity.
- **Data Minimization Requirements:** It limits the amount of student data a company can keep to only that which is “reasonably necessary” to provide the specific educational service, and requires timely deletion of data after a student leaves a school.
- **The Private Right of Action:** Most importantly, the bill creates a private right of action, enabling students and parents to directly sue technology providers for privacy violations, rather than waiting for the traditionally slow process of state attorneys general enforcement.

This fight is happening now in state legislatures. Big tech firms have fought hard against AB 1159. Business coalitions warn that private rights of action and strict training bans will choke off capital investment in next-generation personalized learning tools.

Other states are moving to pass laws to accelerate AI adoption, while California is obsessed with restrictive data boundaries. Georgia's Senate Bill 179 (SB 179) is the workforce-aligned vector of the decentralized model, mandating computer science and AI literacy instruction as a core high school graduation requirement starting with the 2031-2032 school year. The California Public Schools AI Working Group, established by Senate Bill 1288 (SB 1288), meets regularly to develop statewide guidance, striving to balance AI literacy with high standards for academic integrity. And this patchwork of education has produced a very uneven landscape. Students in well-resourced, progressive California districts are protected by robust data protections and public oversight, while students in states without privacy regulation are subjected to aggressive corporate data harvesting. Such divergence in regulations creates significant compliance friction for EdTech providers navigating competing state mandates and denies the nation's public education system a common standard for algorithmic fairness.

c. The Federal Democratic Model (India)

India is the world's second largest EdTech market after the United States and China and offers a highly complex federal democracy with stark rural-urban digital divides and a robust domestic technology sector. In response to the rapid and unregulated commercial expansion, the Indian state passed the Digital Personal Data Protection (DPDP) Act 2023, the operational backbone of which was codified in the DPDP Rules 2025.

India's legal framework considers the data of children to be a category of heightened regulatory concern and defines children's data as the data of natural persons below the age of 18. India has a strict consent driven architecture under Section-9 of the DPDP Act and Rule 10 of the DPDP Rules 2025 :

- **Verifiable Parental Consent:** Data Fiduciaries (such as schools, universities and EdTech platforms) are not allowed to process a minor's personal data without verifiable parental consent. Simple and deceptive UI mechanisms like generic pre-ticked checkboxes or passive "I agree" clicks are not legally valid. Rule 10 operationalises verification by mandating secure technical integrations such as government-backed DigiLocker-verified credentials, eKYC or OTP verification mapped directly to an adult's registered identity.
- **The Absolute Tracking Prohibition:** Under clause 9(3), it is clear that any processing of children's data which includes behavioural tracking, profiling, targeted advertising or any design practices likely to have a negative impact on the well-being of the child is not permitted. Learning analytics platforms that utilize behavioral data to label students as "low potential" or to predict behavioral infractions fall under harmful profiling, making platforms subject to statutory penalties.
- **Exemptions for Public Education and Welfare:** Absolute consent is not feasible in public administration, so the state has reserved the power to make exceptions. Rule 10 also provides an exception to the verifiable consent requirement for critical services such as public health, real-time safety, and core public education operations undertaken by the state. This allows public schools to collect and use the basic data needed for enrollment without being mired in bureaucratic red tape.

India's federal structure creates a lot of friction in coordination. The federal DPDP Act is implemented by the central government but education is implemented at the state level. Some creative ideas are being tested out by state governments. For example, state governments of Maharashtra, Chhattisgarh, Assam and Punjab and the Union Territory of Ladakh have bypassed central procurement and collaborated directly with Google India to launch the Google AI Educator Series (GES). This initiative breaks away from the standard desktop infrastructures and instead takes a mobile first training architecture in regional languages (including Assamese, Hindi, Marathi, Telugu, Odia and Punjabi) to train teachers on responsible AI classrooms. This state level variation provides local relevance but results in highly unequal digital safeguards across the country and high regional variation in data protection enforcement.

d. Disaggregating China and Singapore

i. Singapore: The Faith of the Sovereign Technocrat

From an analytic perspective, the earlier conflation of China and Singapore as one ‘State-Centric Optimization’ archetype is untenable, obscuring differences of political governance, judicial structure and administrative rule of law.

Singapore has a highly centralized, technocratic public administration that works under the rule of law. The city state’s National AI Strategy 2.0 (NAIS 2.0) has been systematically embedding AI literacy in all levels of the national curriculum, coupling technological adoption with massive state-funded teacher training initiatives.

Importantly, corporate procurement in Singapore is not fragmented. MOE designs, owns and manages Singapore Student Learning Space (SLS), its national learning portal. In this sovereign digital infrastructure, student data is not given to foreign corporate monopolies, but securely hosted in the state’s sovereign highly encrypted cloud infrastructure.

An elite teacher upskilling pipeline supports the state-directed technocratic trust. Singapore is the world’s most AI-ready country for teachers, according to the OECD’s Teaching and Learning International Survey (TALIS) 2024, with some 75% of lower secondary teachers using AI in their professional work and providing high rates of formal, state-provided professional development. Singapore demonstrates the ability of a centralized state to optimize technology delivery while maintaining public accountability, student data sovereignty, and teacher professional dignity.

ii. China: Authoritarian Optimization by State Direction

The Chinese state-centric optimization paradigm is a one-party authoritarian uniform top down project. The Chinese central government sees technological literacy as an important vector of national security and geopolitical competition and thus requires direct integration of AI in the national school curriculum (Beijing Consensus on Artificial intelligence in Education).

This model effectively erases the boundaries between the state and the private sector. The national standardized smart classroom platforms are built by private domestic tech conglomerates, and funded and implemented by the central government. They employ state-of-the-art computer vision , automated grading pipelines , and predictive analytics on a scale unprecedented . They have turned administrative efficiency and ideological conformity into the primary measures of educational success .

While the Chinese model sometimes tempers the most extreme forms of real-time biometric behavioral surveillance in primary classrooms in the wake of public anxiety over student stress, it is still firmly rooted in the optimization paradigm, where automated standardization ensures political conformity and creates a highly disciplined workforce.

e. Latin America and the Middle East

Brazil, Chile and the United Arab Emirates are very different sociotechnical experiments, with emerging governance models that stretch conventional binaries in global policy.

i. Brazil: Digital Autonomy and Global Partnerships

The Brazilian approach to the algorithmic classroom features a high stakes transition from a fragmented, market driven adoption to an organized fight for digital sovereignty. During Brazil's G20 Presidency in 2024, the Ministry of Science, Technology and Innovation (MCTI) coordinated the development of the G20 Toolkit for AI Readiness and Capacity Assessment, framing "digital and AI autonomy" as a fundamental national right. This framework views autonomy as the ability of a country to develop, control and manage its digital future with its own localised data, computing power and sovereign cloud infrastructure, protecting its people from digital dependency and Northern data extraction.

In order to put this vision into practice, Brazil has not imported strictly commercial platforms, but has developed strategic public-private partnerships. The 2026 Google.org-funded global partnership between Google, UNICEF and the Brazilian government is a prime example. The initiative pairs Google's cutting-edge models (including Gemini and NotebookLM) with UNICEF's systemic knowledge in education to co-develop localized, culturally adapted digital learning systems. Most critically, the partnership is built on rigorous public accountability metrics, a commitment to upskill educators, bridge the rural-urban digital divide and adapt all algorithmic tools to the regional languages and curricular standards of Brazilian public schools.

ii. Chile: The Constitutional Protection of Neurorights

Chile is ahead of the world, taking the ambitious path of regulating digital rights upstream, and addressing the physical and cognitive boundary of human-machine interaction. In 2021, the Chilean Senate ratified a historic constitutional reform of Article 19 of the Magna Carta, unanimously establishing a protected class of human rights called neurorights. It changes the legal status of personal brain data and neural activity to the same ontological status of a physical human organ, saying that brain data cannot be bought, sold, trafficked or manipulated.

This constitutional order was challenged in August 2023 when the Supreme Court of Chile unanimously decided that the US-based neurotechnology company Emotiv must delete all brain data collected from a former Chilean Senator using a consumer electroencephalogram (EEG) headset. This constitutional protection has important implications for educational AI. With adaptive learning systems and consumer EEG headsets increasingly common in mainstream education to monitor relaxation, cognitive load and attention, Chile's constitutional protections ensure that students maintain explicit, revocable control over their mental privacy and cognitive identity. It bans "neuro-cookies" and algorithmic behavioral tools intended to decode or manipulate student decision-making, establishing a global precedent for cognitive liberty in the algorithmic classroom.

iii. The United Arab Emirates: Mandatory K-12 AI Literacy

The United Arab Emirates is leading the way in policy-driven adoption of AI in K-12 education. The UAE Cabinet has approved a compulsory national AI curriculum that was launched at the beginning of the 2025-2026 academic year, skipping over pilot projects.

The initiative, led by Sarah Al Amiri, Minister of State for Public Education and Advanced Technology, is a mandatory K-12 course and will reach an estimated 400,000 learners in public schools and private schools following the Ministry of Education (MOE) national curriculum. The curriculum is based on 20 lessons per grade level per year, in a highly structured and age-differentiated three-cycle progression :

- **Cycle 1 (Kindergarten to Grade 4):** Focuses on visual, play-based activities to develop computational thinking without active screen time. Children compare the abilities of humans and machines and investigate basic algorithmic concepts. No screen time before Grade 5, so an explicit limitation on passive screen time.
- **Cycle 2 (Grades 5 to 8):** Provides students with hands-on experience with AI design, data literacy, and critical thinking about algorithmic bias.
- **Cycle 3 (Grades 9 to 12):** Emphasis on prompt engineering, supervised AI simulations, and real-world Capstone projects.

To underpin this national drive, the UAE trained a specific cohort of 1,000 teachers in preparation for the academic launch, utilising allocated class time in the Computing, Creative Design and Innovation subject blocks to prevent timetable congestion. Most significantly, the UAE curriculum dedicates of all lessons to AI ethics and social responsibility, ensuring technological mastery is structurally coupled with critical evaluation and ethical reflection.

f. Emerging Sovereignty in the Global South: Bangladesh and Sub-Saharan Africa

Bangladesh and Sub-Saharan Africa Initiatives in emerging governance and structured pedagogy challenge the historic marginalization of Global South states in EdTech literature that has tended to disallow technological agency and regional sovereignty.

i. Bangladesh: Structural Deficits and Sovereign Language Infrastructure

Infrastructure Bangladesh is a case of the Emerging Sovereignty archetype, a rapidly digitizing developing country exercising post-colonial data sovereignty and defensive regulation. Following the student-led July Revolution of 2024, the interim government of Bangladesh under Nobel laureate Muhammad Yunus sought to modernize the country technologically within the framework of inclusive democratic governance and the SDGs. This socio-political transition resulted in the formulation of Bangladesh National AI Policy 2026-2030 (Draft V2.0).

The National AI Policy articulates five core strategic pillars to transform the nation from a passive “AI consumer” to an active “AI producer” of sovereign technologies. Meanwhile, Bangladesh has adopted a hands-on, competency-based learning of robotics and STEAM in its secondary school curriculum, moving away from rote memorization of the past. National

programs, supported by the Smart Bangladesh 2041 Master Plan, provide subsidies for classroom robotics kits (including locally developed Dbot and MONU) and technology labs in rural districts to improve problem-solving skills, resulting in a 20% improvement in STEM test scores.

However, despite such advanced policy objectives, Bangladesh suffers from a huge implementation gap. The country is placed 75th in the world in the Oxford Insights Government AI Readiness Index. The UNESCO and UNDP Joint AI Readiness Assessment Report 2025 noted critical structural deficits in fragmented national data systems, acute scarcity of GPUs, rural-urban digital divides (less than 38% internet access in rural areas as against 68% in cities) and a “nearly absent” AI ethics training pipeline. In addition, the domestic language infrastructure is severely underdeveloped as only one entity (Hishab) is producing sovereign Bengali LLMs. Hishab’s voice-over-GSM platform and Series A capitalization address deep rural financial and digital exclusion by translating natural speech over a basic GSM call into structured records, but its use in formal classrooms is limited by basic infrastructural issues. The proposed regulatory institutions, including the National AI Ethics Board, remain on paper, waiting for formal legislative acts and capacity benchmarks.

ii. Sub-Saharan Africa: Rapid EdTech Expansion and the EdoBEST Model

EdTech is growing fast in sub-Saharan Africa, frequently poorly regulated, funded by multilateral development banks and Silicon Valley philanthropy, with over 400 million school-age children in its sights. In 2022, the World Bank and Mastercard Foundation established a strategic partnership to support African countries in developing policies for digital education, investing in affordable connectivity and upskilling youth to address the global learning crisis.

Launched in Edo State in April 2018 under the leadership of Governor Godwin Obaseki, Nigeria’s flagship model for sub-national educational reform in Africa is the Edo Basic Education Sector Transformation (EdoBEST) program. EdoBEST is a technology-enabled structured pedagogy program supported by the World Bank’s Accelerator scheme and a \$75 million funding approval.

The EdoBEST model’s key structural components are:

- **Standardized Lesson Plans via Tablets:** Instead of unleashing unguided generative tools, the program arms more than 20,000 public school teachers with handheld tablets loaded with highly structured, daily lesson plans derived from the cognitive science of learning.
- **Real-Time Administrative Monitoring:** The tablets capture real-time data on student and teacher attendance, lesson completion rates, and test scores, allowing the State Universal Basic Education Board (SUBEB) to immediately identify and intervene in failing classrooms.
- **Ongoing Teacher Coaching:** Standardized pedagogy is combined with ongoing, in-service classroom observations and skills-based coaching, restoring dignity to the teacher and consistency of instruction for urban and rural districts alike.

The empirical evaluations of the EdoBEST program—aptly titled *The EdoBEST Effect*—demonstrate massive, rapid learning gains. Within a single term, primary school students in EdoBEST classrooms outperformed their peers in traditional public schools, scoring 2% higher in mathematics and 5% higher in science. A rigorous, two-year randomized controlled trial co-authored by Nobel laureate Michael Kremer alongside Gray-Lobe, Keats, Mbiti, and Ozier evaluated Bridge International Academies (now operating as NewGlobe) — the same structured-pedagogy provider whose model underlies EdoBEST, EKOEXCEL, and KwaraLEARN — in Kenyan primary and pre-primary schools, rather than evaluating EdoBEST directly. The study found that the intervention increased pre-primary learning by 1.35 standard deviations and primary-school learning by 0.81 standard deviations, placing these results among the largest ever rigorously documented in the international education literature. Learning gains were statistically indistinguishable between girls and boys — evidence that the model closed the gender gap in learning outcomes within the program, a narrower and better-supported claim than closing the broader historical gender gap across Sub-Saharan African education systems. EdoBEST represents a powerful model of the Restructuring Paradigm in the Global South. Rather than allowing unguided, corporate-driven generative AI to induce cognitive offloading and data extraction, the state utilizes highly structured, public-interest digital tools to augment teacher capacity, enforce pedagogical consistency, and deliver equitable, proven learning outcomes at scale.

g. Cross-Case Synthesis: Findings, Interpretation, and Governance Implications

Three patterns emerge when the ten jurisdictions analyzed above are read side by side. Descriptively, EAPRI scores in Table 4 show an uneven profile across domains rather than uniform weakness or strength: Statutory Enforceability (SE) is, on average, the strongest-scoring domain across the nine EAPRI-scored jurisdictions (mean 2.83 of 4), while Algorithmic Transparency (AT) and Pedagogical Autonomy (PA) are the weakest (means of 2.22 and 2.28, respectively). This pattern indicates that most states have been more willing to legislate the existence of a binding rule than to build the technical audit capacity and human-oversight mechanisms needed to make that rule operative — the rhetoric-reality gap, in other words, is concentrated less in whether a law exists than in whether it can be independently verified and whether it preserves human judgment in practice.

Interpreting this pattern through the theoretical apparatus developed in Section II, the persistent weakness of Algorithmic Transparency is consistent with the principal-agent and platformisation mechanisms described in Section II.b: EdTech vendors have structural incentives to resist independent auditing regardless of a jurisdiction's statutory environment, since proprietary opacity is itself a source of market power. The weakness of Pedagogical Autonomy tracks the political economy of bureaucratic inertia described in Section II.d: automated, high-stakes sorting is administratively cheaper to implement and easier to defend to oversight bodies than human-in-the-loop review, which requires sustained investment in teacher capacity. Infrastructure Equity (ER), by contrast, scores comparatively well in centralized systems that retain direct control over cloud and curriculum infrastructure (Singapore, China, and the United

Arab Emirates, at 4.0, 3.5, and 4.0 respectively) — though, as Section II.c cautions against reading centralized control as costless, China's high ER score coexists with one of the sample's lowest Pedagogical Autonomy scores (1.5, tied with the United States and Brazil), underscoring that infrastructure sovereignty and rights protection are analytically distinct achievements that do not automatically travel together.

For governance, the practical implication is that Algorithmic Transparency and Pedagogical Autonomy — not infrastructure spending — are the binding constraints on closing the rhetoric-reality gap for most states, including resource-constrained ones, since AT and PA reforms (independent audit mandates, human-in-the-loop requirements) are primarily legislative and administrative rather than capital-intensive. This directly motivates the phased sequencing proposed in Section VI.c, which places defensive statutory and transparency reforms in Years 1–2, ahead of any capital-intensive infrastructure investment.

V. DISCUSSION: ALGORITHMIC FAILURES, COGNITIVE OFFLOADING, AND INSTITUTIONAL DYNAMICS

a. The Pedagogy-Performance Crisis: Cognitive Offloading and the Socratic Safeguard

The swift penetration of generative AI into public classrooms has spawned a deep pedagogical crisis, revealing a structural tension between the demand to accomplish urgent tasks and the demand for deep cognitive development.

b. The Turkey Mathematics Randomized Controlled Trial

The principal empirical anchor for this pedagogical crisis is a landmark, large-scale randomized controlled trial (RCT) by Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı, and Rei Mariman, published in the Proceedings of the National Academy of Sciences (PNAS). In Fall 2023, the researchers evaluated almost 1,000 high school mathematics students in around 50 classrooms (9th, 10th, and 11th grades) in Turkey to investigate the differential impact of software design on learning outcomes. This study was designed as a three-arm randomized comparison to:

- **The Control Group:** Students with standard learning resources (course notes, textbooks) and no access to generative AI.
- **The "GPT Base" Group:** Students with full access to a standard ChatGPT-4 chat interface, similar to general-purpose commercial chatbots.
- **The "GPT Tutor" Group:** students with access to a custom-secured version of GPT-4, which used prompt engineering to act as a Socratic guide (offering hints, strategy prompts, and refusing to directly give the final answer).

The AI-enabled groups had explosive gains in performance in the assisted self-study practice sessions. The open-ended GPT Base group improved performance on math exercises by

48% over the control group, while the guided GPT Tutor group showed an astounding 127% increase in performance.

c. Explaining "Metacognitive Laziness" and Cognitive Offloading

The researchers conducted a detailed qualitative and quantitative analysis of the chat logs to reveal the causal mechanisms behind this learning collapse. The logs revealed that students in the unrestricted GPT Base group used the chatbot as an intellectual “crutch”. The vast majority of these students did not engage in the active cognitive work of solving the problems, they pasted the math problems into the chat, and copied the generated answers directly onto their practice sheets. Crucially, GPT-4 only gave the correct answer in math 51% of the time, hallucinating in arithmetic and logic. But what the students did was to blindly copy these flawed calculations without checking them because they had completely delegated their thinking and had no real comprehension of the material.

This empirical pattern supports the notion of metacognitive laziness, a form of cognitive offloading in which the learner avoids critical self-regulatory and diagnostic processes (e.g., setting goals, assessing information, brainstorming alternatives, iterating on errors) by blindly accepting and executing automated outputs.

This evidentiary base warrants a scope condition: the Turkish RCT's findings pertain specifically to unrestricted, general-purpose, answer-providing chatbot use — the “GPT Base” condition. The same study's “GPT Tutor” condition, which used identical underlying model capability but restricted output to Socratic scaffolding, produced a 127% performance gain with no corresponding evidence of the exam-performance collapse observed in the GPT Base condition. The claim that generative AI harms learning is therefore not a claim about generative AI as such, but about a specific, common, and currently dominant mode of deploying it — a distinction that motivates the Pedagogical Ring-Fencing proposal in Section VI.b.

“Productive struggle” and “productive friction” drive cognitive development in a natural, unmediated learning cycle. When the student is faced with a difficult concept, the brain is forced to go through the active retrieval, schema rebuilding, and error diagnostic loops to embed knowledge in long term memory. General purpose commercial chatbots are explicitly optimized to reduce user friction, delivering instant, very fluent and very confident outputs. By removing the productive struggle of learning, the machine induces cognitive atrophy, reducing student attention spans, critical reading capacities and cognitive stamina.

Then there is the added cognitive degradation from a profound psychological mismatch. In the Turkish RCT, students in the unrestricted GPT Base group failed the independent examination at higher rates, but believed that they had learned just as much as their peers. This results in an “illusion of competence” in which the student internalizes the machine’s fluency and speed as personal intellectual mastery, leaving them completely unprepared for unassisted performance.

The Turkish RCT results and the broader OECD Digital Education Outlook 2026 call for a clear distinction between two radically different types of technological deployment:

Table 5: Architectural and Pedagogical Divergence: General-Purpose Chatbots vs. Purpose-Built Educational AI

Architectural Attribute	General-Purpose Chatbots (e.g., ChatGPT-4, Gemini)	Purpose-Built Educational AI (e.g., MATHia, Socratic Tutors)
Optimization Target	Maximizing prompt completion speed, transactional fluency, and immediate output generation.	Maximizing long-term knowledge retention, conceptual understanding, and critical evaluation.
Pedagogical Mechanism	Direct answer provision; cognitive offloading of thinking and writing tasks.	Real-time diagnostic mapping; Socratic questioning; adaptive hints and strategic nudging.
Friction Level	Zero-friction; designed to minimize cognitive or operational resistance.	Controlled, productive friction; designed to enforce active cognitive retrieval loops.
Data Governance	Opaque; student inputs are extracted to train commercial models; minimal child-safety auditing.	Transparent; bound by strict student data localization and non-commercial erasure mandates.
Learning Outcome	Performance-Learning Gap: immediate task success followed by unassisted exam score collapse.	Sustained, durable learning gains; independent skill acquisition matches or exceeds traditional tutoring.

d. Multimodal AI and Biometric Surveillance

Multimodal AI and Biometric Surveillance While a lot of the 2023 and 2024 discussion was on text-based LLMs, the technology landscape of 2025 and 2026 is increasingly dominated by multimodal AI architectures. These advanced systems integrate text and audio speech processing, computer vision and real time physiological tracking. This change in the mode of delivery has been reflected in the classroom through remote examinations that usually involve

surveillance through webcams, recording of the screen and audio, and analysis of behaviour. That kind of surveillance is very intrusive and has important civil rights implications. If a student's eyes stray from the screen, or if background noise is detected, the algorithm flags the behavior, generating an automatic "suspicion score" that can result in automatic disqualification or grade penalties. Some schools have experimented with biometric cameras that use models for emotion-recognition. The systems track Head orientation to determine student attention and compliance. The physical bodies of children are converted into streams of metric data. This is the central structural risk of the algorithmic classroom.

e. The Political Economy of Algorithmic Failures: The Ofqual Standardisation Scandal

The dangers of the technocratic Optimization Paradigm are not abstract policy projections; they are mathematically hardcoded into global administrative failures. The definitive historical warning remains the United Kingdom's **Ofqual Exam Standardisation Scandal of August 2020**.

f. The Standardisation Algorithm: Mathematical Modeling of Socioeconomic Bias

The UK's Department for Education, led by Education Secretary Gavin Williamson, ordered the exams regulator, Ofqual, to come up with and deploy an automated statistical system after the COVID-19 pandemic led to the sudden cancellation of GCSE and A-level exams. A ministerial decree set out the overt political goal: to preserve rigorous historical standardization and combat "grade inflation".

To do this Ofqual developed the Direct Centre Performance (DCP) model. Instead of predicting the performance of individual students based on their own academic trajectory, the algorithm predicted grade distributions at the level of the institution, treating the school's past performance as a hard cap on the achievement of its current students.

The mathematical standardization formula used by Ofqual to calculate the calculated grade distribution can be systematically modelled as below :

To model this mathematically, let j denote a specific educational center (school or college) and k denote a specific subject grade (e.g., A*, A, B, C). Five further quantities are defined for grade k at school j :

- $P(k,j)$: the calculated final grade distribution assigned to grade k at school j .
- $C(k,j)$: the historical grade distribution of grade k at school j over the previous three years (2017–2019).
- $q(k,j)$: the predicted grade distribution of grade k at school j based on the current cohort's prior academic attainment (their GCSE scores).
- $p(k,j)$: the predicted grade distribution of grade k at school j based on the prior academic attainment of historical cohorts.
- $r(j)$: the data-completeness fraction ; the proportion of students in the current class at school j for whom historical prior-attainment data (GCSE scores) is fully trackable.

For large cohorts (defined by Ofqual as classes where $n(j) \geq 15$ students), the standardized grade distribution is computed as:

$$P(k,j) = (1 - r(j)) \times C(k,j) + r(j) \times [C(k,j) + q(k,j) - p(k,j)]$$

Where prior-attainment data is perfectly complete ($r(j) = 1$), this simplifies to:

$$P(k,j) = C(k,j) + q(k,j) - p(k,j)$$

(This is a structural reconstruction of Ofqual's published Direct Centre Performance methodology, simplified for analytical clarity; it illustrates the model's mechanics rather than reproducing Ofqual's full technical specification.)

This mathematical architecture established a severe socioeconomic bias through the implementation of a separate, non-standardized track for small cohorts :

- For small subject cohorts (where $n_j < 5$ students), standardization was completely bypassed. The algorithm discarded the historical center limits, accepting the teacher's original **Center Assessed Grade (CAG)** entirely without downgrading.
- For intermediate subject cohorts (where $5 \leq n_j \leq 15$ students), the standardization model was only partially applied, placing the dominant mathematical weight on the CAG.

Table 6: Operational Variables and Policy Logic of the 2020 Ofqual Direct Centre Performance (DCP) Model

Mathematical Variable	Operational Definition	Policy standardisation Model	Role in Undergirding Administrative Logic
$p_{k,j}$	Calculated final grade distribution of grade k at school j .	The final automated grade quota allocated to the school; students are slotted into this distribution based on rank order.	Represents the "standardized" reality determined by the algorithm.
$C_{k,j}$	Historical average proportion of grade k at school j over baseline years (2017–2019).	Acts as a computational “hard ceiling” or anchor, linking current year grades to historical performance.	Enforces historical homologation, assuming a school cannot exceed its historical mean.

$q_{k,j}$	Predicted proportion of grade k at school j based on current cohort's prior GCSE prior attainment.	Nudges the historical baseline upward or downward depending on the incoming cohort's prior test scores.	Integrates "value-added" adjustments based on standardized prior cognitive metrics.
$p_{k,j}$	Predicted proportion of grade k based on prior GCSE attainment of historical cohorts at school j .	Corrects for any unique changes in prior historical cohort ability, preventing unfair historical penalties.	Standardizes the relationship between Key Stage 4 and A-level progression over time.
r_j	Prior attainment data matching/completeness fraction at school j .	Controls the statistical weight placed on historical center performance vs. prior GCSE cohort achievement.	If matched prior attainment is unavailable ($r_j \rightarrow 0$), the model defaults entirely to historical averaging ($P \rightarrow C$).
CAG	Centre Assessed Grade (unmoderated teacher grade prediction).	Completely bypassed for large cohorts; accepted directly as final grade ($P = \text{CAG}$) only for small classes ($n < 5$).	Treats unmoderated teacher judgment as an unreliable and systematically over-generous grade metric.

g. A Worked Numerical Example of Algorithmic Discrimination

To illustrate how this mathematical division directly translated into systemic class bias, consider two students of equal academic ability studying A-level Physics in 2020: Student A (at a large, state-funded Sixth-Form College in an underserved public district) and Student B (at an elite, expensive independent private boarding school).

Case 1: Student A (State Sixth-Form College)

- **Cohort Size (n_j):** 40 students (classified as a large cohort, triggering 100% standardization).
- **Data Completeness (r_j):** 1.0 (all students have trackable GCSE records).
- **Teacher's Center Assessed Grade (CAG):** A^* (Student A is ranked 1st in the class of 40).
- **School's Historical Grade Distribution ($C_{k,j}$):** Over the previous three years (2017-2019), the college's average percentage of A^* grades in Physics was exactly 0% ($C_{A^*,j}$).

- **Cohort Prior Attainment Adjustments ($q_{A*,j} - p_{A*,j}$):** The current cohort's prior GCSE scores are slightly higher than historical averages, calculating a minor upward adjustment of $+0.02$ (2%).

Applying the Ofqual DCP formula for the top grade ($k = A^*$):

$$P_{A*,j} = (1 - 1.0)(0.0) + 1.0(0.0 + 0.02) = 0.02$$

The algorithm dictates that this Physics class of 40 students is permitted to receive exactly $0.02 \times 40 = 0.8$ grades of A^* . Since grades must be integers, the distribution rounds down, meaning the college is awarded exactly **0 A^* grades** in Physics for 2020.

Despite Student A being ranked 1st and possessing an outstanding personal academic record, the algorithm overrules the teacher's CAG. Student A is automatically downgraded to an $\square$ or $\square$ solely because the school's historical average ($C_{A*,j}$) acted as an unbreakable computational ceiling, trapping Student A as a "prisoner of the past".

Case 2: Student B (Elite Private Boarding School)

- **Cohort Size (n_j):** 4 students (classified as a small cohort).
- **Teacher's Center Assessed Grade (CAG):** A^* (Student B is recommended for an A^* by the private tutor).

Because $n_j < 5$, the Ofqual algorithm completely suspends the DCP formula ($r_j = 0$), bypassing historical constraints and accepting the teacher's CAG directly :

$$P_{A*,j'} = CAG_{A*} = 1.0$$

Student B receives their A^* grade directly. They face 0% algorithmic downgrading, despite studying the exact same subject and possessing the exact same academic profile as Student A.

The implications for the real world of this mathematical model were disastrous. On August 13, 2020, the algorithm reduced 39.6% of all A-level results in England when the results were published. Almost all of the hits were experienced by state comprehensive schools and large public colleges, which experienced an instant windfall in high grades, and private independent schools. The resulting public outrage, student-led street protests and condemnation of "class warfare by algorithm" forced a complete, humiliating government U-turn within four days of the results release. The government scrapped the DCP model, withdrew the grades generated by the algorithm and re-issued all qualifications based on the unmoderated teacher-predicted CAGs. The political fallout included the immediate resignations of Ofqual CEO Sally Collier and the Department for Education's permanent secretary, Jonathan Slater.

h. The Aftermath: Regulatory Learning and Statutory Reforms

To assess the degree of real regulatory learning from the 2020 disaster, we need to monitor the structural reforms introduced by Ofqual and the UK Parliament from 2020 to 2026.

First, the House of Commons Education Select Committee held an enquiry and found that Ofqual's chief failing was its complete lack of procedural transparency and dogmatic determination to prefer grade inflation over individual fairness. In response, Parliament amended the Ofqual regulatory enforcement powers under the Apprenticeships, Skills, Children and Learning Act 2009 (ASCLA).

Secondly, Ofqual used the Chief Regulator's Rebuke as a public-facing enforcement tool to ensure market compliance in October 2025. The regulator used these new statutory powers to issue 26 formal enforcement actions and hefty monetary penalties, including a £750,000 fine on Pearson Education in December 2025 for systemic failures in internal quality assurance and delayed event notifications.

Third, after the final report of the state-sponsored Curriculum and Assessment Review in late 2025, the UK government began a fundamental reshaping of the secondary qualifications landscape. The key reform is not to try to automate academic grading, but rather to:

- **The Development of V-level Qualifications:** A new, federally standardised vocational qualification, designed to sit within the Level 3 technical streams, ensuring that non-academic pathways are informed by nationally set content and subjected to rigorous, human-led quality assurance.
- **Systemic Reductions in Examination Load:** * A statutory requirement to reduce the time spent in sitting exams by 2.5 to 3 hours on average for the average GCSE student, with the assessment model moving from the high-stakes exam halls on a single day to an ongoing, formative human assessment.

These reforms demonstrate that genuine regulatory learning occurred. The UK state recognized that the solution to algorithmic bias is not a more complex algorithm; rather, it requires a structural retreat from high-stakes mathematical sorting, the restoration of human teacher judgment, and the legislative anchoring of public accountability.

i. Integrating Student Voice

In the discourse of educational AI, students are mostly constructed as abstract subjects of policy or as passive objects of algorithmic harm, with their active lived experiences routinely silenced. Empirical evidence of the psychological and civil rights violations of students in the algorithmic classroom includes student testimony and legal action.

Research bodies collecting student testimony in the 2020 Ofqual standardisation crisis documented a profound sense of “existential erasure”. Gifted pupils from working-class estates said their individual work, hours revising and personal intelligence had been completely devalued by an algorithm that “judged them by their postcode, not their brains”. The sudden,

impersonal exclusion of hard-won university placements caused deep anxiety, leaving a generation to believe that their academic futures had been dictated by the historic averages of institutions and not personal merit.

At the same time, the rapid deployment of automated proctoring software has triggered aggressive legal resistance by students fighting to protect their constitutional rights. In August 2022, *Ogletree v. Cleveland State University* was a landmark civil rights case, chemistry student Aaron Ogletree sued Cleveland State University for requiring him to submit to a virtual 360-degree camera scan of his private bedroom workspace before a remotely proctored exam.

The US District Court for the Northern District of Ohio has ruled that the virtual room scans performed by public higher education institutions are an illegal search under the Fourth Amendment of the US Constitution. The university's administrative interest in academic integrity does not outweigh a student's subjective and objectively reasonable expectation of privacy in his or her home bedroom, ruled Judge Philip Calabrese, finding suspicionless room scans unconstitutional.

This landmark legal precedent shows a growing student backlash against panoptic AI and that the corporate platformisation of school infrastructure has serious constitutional and psychological costs that students are no longer willing to bear.

VI. POLICY RECOMMENDATIONS: OPERATIONALIZING SUSTAINABLE PUBLIC INNOVATION ECOSYSTEMS

a. Reconciling the Sovereign AI vs. Cognitive Atrophy Contradiction

A central tension sits at the heart of contemporary educational AI policy : if the cognitive harm of AI stems from the fundamental nature of automated task completion rather than from corporate ownership of the tools, then transferring sovereignty over AI to the public sector does not automatically solve the pedagogical problem.

This analysis resolves this tension by demonstrating that **corporate ownership directly dictates the pedagogical architecture of the software**. Under the market incentives of "Surveillance Capitalism," private technology conglomerates optimize their consumer chatbots (such as standard ChatGPT or Gemini) to maximize platform engagement, user throughput, and frictionless transactional speed. This commercial logic demands the absolute minimization of user effort, providing instant, direct, and complete essay drafts or mathematics solutions. This frictionless completion is what directly induces cognitive offloading, metacognitive laziness, and the subsequent collapse in unassisted human learning.

Conversely, public sovereignty over AI allows the state to insulate technological development from short-term commercial monetization and venture-capital exit pressures. Under the **Sustainable Public Innovation Ecosystem (SPIE)** model, the state leverages public funds

to build and host sovereign, open-source pedagogical models where **pedagogical ring-fencing is structurally coded into the system's core architecture**.

Because the sovereign model is not designed to maximize commercial engagement, it is explicitly programmed to refuse to write the student's essay or solve their mathematics equation. Instead, it is calibrated to act strictly as a Socratic dialogue partner—interrogating student assumptions, suggesting reading paths, diagnosing logical errors, and enforcing the "productive struggle" necessary to build durable schema in long-term memory. Public sovereignty is therefore the only mechanism that can rescue the technology from frictionless commercial exploitation, converting AI from a cognitive crutch into a scaffolded, public-interest cognitive partner.

b. The SPIE Operational Architecture

Implementing the SPIE framework requires four structural interventions across the public education system :

- **Mandatory Algorithmic Impact Assessments (AIA):** School boards and public university administrators should refuse to buy any AI tool that is not subject to independent, third-party audits for algorithmic bias, data safety, and pedagogical validity.
- **The Funding of Sovereign Alternatives:** States should not surrender public education budgets to Silicon Valley platforms, but rather actively fund and develop open-source pedagogical AI tools that are culturally adapted, in local languages and indigenous epistemologies.
- **A Unified Federal Common Core:** Create strong national standards for student data protection to stop private tech companies from using public school systems as free data mines.
- **Pedagogical Ring-Fencing:** Imposing strict software and policy limits to stop general-purpose commercial chatbots from being used for core cognitive learning functions.

Table 7 makes this alignment explicit, mapping each EAPRI domain — including those identified as weakest on average in Section IV.g — onto the specific SPIE pillar designed to strengthen it.

Table 7: Mapping EAPRI Domains to SPIE Pillars

EAPRI Domain	SPIE Pillar		Mechanism of Alignment
Statutory Enforceability (SE)	Mandatory	Algorithmic Impact Assessments (AIA)	Converts the domain's binding-mandate requirement into a specific, auditable procurement instrument.
Algorithmic Transparency (AT)	Mandatory	Algorithmic Impact Assessments (AIA)	Requires documented bias testing and conformity assessment, directly targeting the “Soft Governance” trap responsible for

			AT's weak average score in Table 4.
Child Data & Sovereign Privacy (CDP)	Sovereign Data Guardrails & Public Infrastructures	Mandates Cloud	Legislates student data as a non-commercializable public good, addressing the data-colonialism mechanism identified in Section II.c.
Pedagogical Autonomy & Agency Protections (PA)	Participatory & Multi-Stakeholder Design; Pedagogical Ring-Fencing	Co-Creation	Embeds teacher and student judgment into tool design and mandates human-in-the-loop review, directly targeting PA's weak average score in Table 4.
Equitable Resource & Infrastructure Allocation (ER)	Sovereign Data & Public Infrastructures; sequencing (Section VI.c)	Mandates Cloud phased	Builds domestically controlled infrastructure on a sequenced timeline that avoids requiring upfront capital a resource-constrained state does not have.

c. Feasibility, Capital Allocation, and Phased Sequencing in Low-Resource Settings

Developing a production-grade sovereign LLM with multilingual capacity and pedagogical validity testing is extraordinarily capital-intensive. For an emerging country like Bangladesh—ranked 75th in AI readiness with severe GPU scarcity and near-absent ethics training infrastructure—a direct recommendation to build state-funded, massive foundation models is unfeasible at current resource levels.

To resolve this structural constraint, this analysis proposes a realistic, Phased Sequencing Roadmap designed to prioritize investments and build regulatory capacity over a five-year horizon.

YEAR 1-2: Defensive Regulation & Data Privacy

- Codify binding student data privacy laws (DPDPA 2023 / Rule 10 standards).
- Implement neutral age-gating and verifiable parental consent frameworks.
- Establish mandatory Algorithmic Impact Assessments (AIA) for all public procurement.

YEAR 3-4: Local Capacity Building & Strategic Partnerships

- Launch localized teacher upskilling pilots via mobile-first, regional language programs.
- Negotiate public-private partnerships bound by G20 Toolkit Autonomy standards.
- Fund national university consortia to develop specialized, low-parameter Socratic engines.

YEAR 5: Full Sovereign Infrastructure Deployment

- Establish Sovereign Data Commons and localized open-source model repositories.
- Deploy purpose-built educational AI across public school networks on state-owned clouds.
- Transition completely from frictionless automated tools to Socratic-guarded learning.

Under this phased capital allocation model, a developing nation spends zero capital on high-cost model training during the first two years. Instead, it utilizes its regulatory authority to enforce defensive privacy boundaries, protecting its child data from external commercial extraction.

In years three and four, the state leverages public-private partnerships (such as the Google-UNICEF education initiative) and regional university consortia to build local capacity.

Only in year five, once basic infrastructure, teacher literacy, and regulatory oversight bodies are fully operationalized, does the nation allocate capital to deploy sovereign, purpose-built educational AI tools within its public-interest cloud.

d. Designing an Alternative, Non-Algorithmic Assessment Architecture

If the critique of automated grading and standardization algorithms is correct, and the Ofqual disaster demonstrates that it is, then the solution requires not just better algorithms, but a complete rethinking of what is being assessed and how. To replace the failed paradigm of high-stakes, algorithm-friendly testing, this analysis proposes a concrete, **Alternative Assessment Architecture** designed to evaluate higher-order cognitive skills in an AI-ubiquitous world.

The architecture comprises three core, human-centered components :

- **The Process-Iterative Mapping (PIM) Framework:** Rather than grading the final product (e.g., an essay or a completed mathematics sheet), students are evaluated on their cognitive journey. Assessments are structured to require students to submit a multi-step design portfolio, including original brainstorming maps, error-correction logs, and documentation of how they critically evaluated, verified, and edited any AI-generated prompts. Grading rubrics are explicitly recalibrated to award up to 60% of total marks to these iterative, diagnostic, and reflective steps, rendering uncritical automated plagiarism mathematically obsolete.
- **The Portfolio-Based Oral Defense (PBOD):** The high-stakes written examination is, in turn, systematically supplemented by brief, in-person oral defense panels, firmly grounded in the classical Socratic tradition. Students submit their portfolios to a panel of human educators, who ask students to explain their logic, defend their citations of sources, and answer conceptual questions. This face-to-face human interaction completely closes the performance-learning gap and immediately reveals whether the student has achieved personal intellectual mastery or has merely outsourced his or her thinking to a commercial machine.

- **The Hybrid Intelligence (HI) Benchmarking Model:** Assessment tasks are carefully crafted to assess human-AI synergy, rather than to ask whether a student can perform on a task as well as an AI. Students are assessed on their capacity to use AI as a research partner to generate hypotheses, map complex literature or run code simulations, taking full responsibility for evaluating the validity of evidence, checking logical coherence and making final ethical and professional judgments.

These proposals carry real implementation costs that should not be understated. Portfolio-Based Oral Defense panels are labor-intensive relative to machine-graded assessment: at scale, they require additional teacher time for panel participation, training in oral-assessment moderation to ensure inter-rater consistency, and — in large-cohort public systems — either smaller class sizes or supplementary examining staff. In under-resourced systems such as Bangladesh's, where teacher-student ratios are already strained, a full transition to this architecture is unlikely to be feasible without the phased capacity-building described in Section VI.c; a realistic starting point is to pilot the PBOD component for final-year, high-stakes cohorts only, expanding coverage as teacher capacity allows, rather than mandating system-wide adoption immediately.

This alternative assessment paradigm shifts public education away from standardized, machine-like compliance. By prioritizing process over product, and human reflection over automated speed, this architecture ensures that technology serves human growth, critical inquiry, and democratic equity in the algorithmic age.

VII. CONCLUSION

In this comparative policy study, we show that the global embedding of artificial intelligence in public education is not primarily a technical matter, but rather a profound governance and political-economic struggle. While national strategies often embrace the progressive language of human-centered, ethical AI, the fact is that public administration has been mostly taken over by an optimization logic that emphasizes short-term bureaucratic efficiency, concentrated corporate market power, and high-stakes algorithmic sorting. The unregulated, frictionless deployment of this technology has created a deep pedagogical crisis of “metacognitive laziness,” eroding the productive struggle that is the basis of real cognitive development. Closing this rhetoric-reality gap and moving public education from technocratic optimization to democratic equity means states must be sovereign architects, not passive consumers, of their digital futures. The analysis offers a clear prioritization logic for policy-makers, curriculum designers and public administrators struggling with this transition.

“First, states should immediately implement Defensive Data Privacy Safeguards. This requires codifying binding, legally enforceable statutory frameworks (e.g., the EU AI Act or India’s DPDP Act) that require verifiable parental consent, impose absolute prohibitions on behavioral tracking and targeted advertising aimed at minors, and outlaw scientifically invalid biometric emotion-recognition surveillance in classrooms.

Second, Pedagogical Ring-Fencing should be mandated by governments. The public procurement codes need to be rewritten to explicitly ban the deployment of general-purpose, frictionless chatbots in classrooms. Public money should be spent on educational AI tools that are purpose built to be Socratic, to scaffold learning, and to maintain cognitive friction. We should be actively funding sovereign, localized, and open-source language infrastructures.

Third, public school systems must Empower the Human Element of Education. Put simply, it must be aligned with teacher unions through collective bargaining, it must be supported by robust and continuous state funding for professional upskilling of teachers, and it must be supported by an alternative assessment architecture that is non-algorithmic and that values process, reflection, and Socratic oral defense in lieu of high stakes algorithmic testing.

Operationalizing these strategic priorities in a Sustainable Public Innovation Ecosystem (SPIE) provides the global community with an opportunity to ensure that the digital revolution in education does not conclude in pervasive surveillance and cognitive decline but rather becomes a powerful beacon for intellectual autonomy, democratic accountability, and human flourishing.

REFERENCES

- Abbas, A., Azar, B., Mahrishi, M., Martín-Núñez, J., & Mishra, D. (2025). AI governance in higher education: A meta-analytic thematic review of current research trends, policy initiatives and knowledge gaps. *Equilibrium. Quarterly Journal of Economics and Economic Policy*. <https://doi.org/10.24136/eq.3551>.
- Ahmed, S. (2024). Beyond Human Teachers: Post-Humanist Perspectives on AI, Cultural Inequities, and Educational Transformation. *Journal of Posthumanism*. <https://doi.org/10.63332/joph.v4i3.1816>.
- Ali, O., Murray, P., Momin, M., Dwivedi, Y., & Malik, T. (2024). The effects of artificial intelligence applications in educational settings: Challenges and strategies. *Technological Forecasting and Social Change*. <https://doi.org/10.1016/j.techfore.2023.123076>.
- Alqahtani, N., & Wafula, Z. (2024). Artificial Intelligence Integration: Pedagogical Strategies and Policies at Leading Universities. *Innovative Higher Education*, 50, 665 - 684. <https://doi.org/10.1007/s10755-024-09749-x>.
- Alzahrani, A. (2024). Unveiling the shadows: Beyond the hype of AI in education. *Heliyon*, 10. <https://doi.org/10.1016/j.heliyon.2024.e30696>.
- Amiri, S., Islam, M., & Hossen, M. (2025). The Role of Artificial Intelligence in Shaping Future Education Policies. *Education Journal*. <https://doi.org/10.11648/j.edu.20251401.14>.
- An, Y., Yu, J., & James, S. (2025). Investigating the higher education institutions' guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. *International Journal of Educational Technology in Higher Education*, 22. <https://doi.org/10.1186/s41239-025-00507-3>.

- Anderson, G. (2026, May 25). Assemblymember Addis seeks to protect student privacy through new bill. Mustang News. <https://mustangnews.net/college-student-privacy-bill/>
- Baker, R., & Hawn, A. (2021). Algorithmic Bias in Education. *International Journal of Artificial Intelligence in Education*, 32, 1052 - 1092. <https://doi.org/10.1007/s40593-021-00285-9>.
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>.
- Berendt, B., Littlejohn, A., & Blakemore, M. (2020). AI in education: learner choice and fundamental rights. *Learning, Media and Technology*, 45, 312 - 324. <https://doi.org/10.1080/17439884.2020.1786399>.
- Boateng, O., & Boateng, B. (2025). Algorithmic bias in educational systems: Examining the impact of AI-driven decision making in modern education. *World Journal of Advanced Research and Reviews*. <https://doi.org/10.30574/wjarr.2025.25.1.0253>.
- Carney, S. (2022). Reimagining our futures together: a new social contract for education. *Comparative Education*, 58, 568 - 569. <https://doi.org/10.1080/03050068.2022.2102326>.
- Chan, C. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, 20. <https://doi.org/10.1186/s41239-023-00408-3>.
- Chen, L., Chen, P., & Lin, Z. (2020). Artificial Intelligence in Education: A Review. *IEEE Access*, 8, 75264-75278. <https://doi.org/10.1109/access.2020.2988510>.
- Daher, R. (2025). Integrating AI literacy into teacher education: a critical perspective paper. *Discover Artificial Intelligence*, 5. <https://doi.org/10.1007/s44163-025-00475-7>.
- Ferrante, P., Williams, F., Büchner, F., Kiesewetter, S., Muyambi, G., Uleanya, C., & Modén, M. (2023). Inequalities in digital education policy – sociotechnical imaginaries from three world regions. *Learning, Media and Technology*, 49, 122 - 132. <https://doi.org/10.1080/17439884.2023.2237870>.
- Filgueiras, F. (2023). Artificial intelligence and education governance. *Education, Citizenship and Social Justice*, 19, 349 - 361. <https://doi.org/10.1177/17461979231160674>.
- García-López, I., & Trujillo-Liñán, L. (2025). Ethical and regulatory challenges of Generative AI in education: a systematic review. *Frontiers in Education*. <https://doi.org/10.3389/feduc.2025.1565938>.
- Ghimire, A., & Edwards, J. (2024). From Guidelines to Governance: A Study of AI Policies in Education. 299-307. <https://doi.org/10.48550/arxiv.2403.15601>.
- Gray-Lobe, G., Keats, A., Kremer, M., Mbiti, I., & Ozier, O. (2022). Can education be standardized? Evidence from Kenya (BFI Working Paper No. 2022-68). Becker Friedman Institute for Economics, University of Chicago. https://bfi.uchicago.edu/wp-content/uploads/2022/06/BFI_WP_2022-68-1.pdf.

- Gulson, K., Sellar, S., & Webb, P. (2022). Algorithms of Education. <https://doi.org/10.5749/9781452968797>.
- Hoca, F., & Nuredin, A. (2025). ALGORITHMIC BIAS IN AI-ENHANCED EDUCATION: CULTURAL DIMENSIONS AND PEDAGOGICAL IMPACT. *Symposium Proceedings*. <https://doi.org/10.55843/isl2025symp163h>.
- Ifenthaler, D., Majumdar, R., Gorissen, P., Judge, M., Mishra, S., Raffaghelli, J., & Shimada, A. (2024). Artificial Intelligence in Education: Implications for Policymakers, Researchers, and Practitioners. *Technology, Knowledge and Learning*, 29, 1693 - 1710. <https://doi.org/10.1007/s10758-024-09747-0>.
- Ignjatović, G. (2024). AI technologies in education: Regulatory frameworks at the international, regional and national level. *Zbornik radova Pravnog fakulteta Nis*. <https://doi.org/10.5937/zrpf1-55374>.
- Jamaluddin, F., Jamaluddin, A., Jamaluddin, F., & Jamaluddin, F. (2025). Malaysia's AI-Driven Education Landscape: Policies, Applications, and Comparative Insights for a Digital Future. *ArXiv*, abs/2509.21858. <https://doi.org/10.48550/arxiv.2509.21858>.
- Jiang, Y., Xie, L., & Cao, X. (2025). Exploring the Effectiveness of Institutional Policies and Regulations for Generative AI Usage in Higher Education. *Higher Education Quarterly*. <https://doi.org/10.1111/hequ.70054>.
- Joshi, S. (2025). Strategic Integration of Artificial Intelligence in U.S. K–12 Education: A Comprehensive Review and Policy Roadmap. *International Journal of Computer Applications*. <https://doi.org/10.5120/ijca2025925428>.
- Kamalov, F., Calonge, D., & Gurrib, I. (2023). New Era of Artificial Intelligence in Education: Towards a Sustainable Multifaceted Revolution. *Sustainability*. <https://doi.org/10.3390/su151612451>.
- Karabatak, S., Ay, S., & Karabatak, M. (2025). Artificial Intelligence Applications in Education Within the Scope of the National Artificial Intelligence Strategy in Turkey. *2025 13th International Symposium on Digital Forensics and Security (ISDFS)*, 1-4. <https://doi.org/10.1109/isdfs65363.2025.11011998>.
- Karan, B., & Angadi, G. (2023). Artificial Intelligence Integration into School Education: A Review of Indian and Foreign Perspectives. *Millennial Asia*, 16, 173 - 199. <https://doi.org/10.1177/09763996231158229>.
- Kim, D., & Wu, J. (2024). Artificial Intelligence in Higher Education: Examining the AI Policy Landscape at U.S. Institutions. *Proceedings of the 25th Annual Conference on Information Technology Education*. <https://doi.org/10.1145/3686852.3687076>.
- Knox, J. (2020). Artificial intelligence and education in China. *Learning, Media and Technology*, 45, 298 - 311. <https://doi.org/10.1080/17439884.2020.1754236>.
- Kundu, A., & Bej, T. (2025). AI policies in school education: a comparative study on China, Singapore, Finland, and the US. *Journal of Science and Technology Policy Management*. <https://doi.org/10.1108/jstpm-06-2024-0218>.

- Kuziemski, M., & Misuraca, G. (2020). AI governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings. *Telecommunications Policy*, 44, 101976 - 101976. <https://doi.org/10.1016/j.telpol.2020.101976>.
- Lee, J., & Jeong, H. (2023). Keyword Analysis of Artificial Intelligence Education Policy in South Korea. *IEEE Access*, 11, 102408-102417. <https://doi.org/10.1109/access.2023.3317261>.
- Li, Z., & Zhang, W. (2024). Technology in education: Addressing legal and governance challenges in the digital era. *Education and Information Technologies*, 30, 8413 - 8443. <https://doi.org/10.1007/s10639-024-13036-9>.
- Lin, C., Huang, A., & Lu, O. (2023). Artificial intelligence in intelligent tutoring systems toward sustainable education: a systematic review. *Smart Learning Environments*, 10, 1-22. <https://doi.org/10.1186/s40561-023-00260-y>.
- Linderöth, C., Hultén, M., & Stenliden, L. (2024). Competing visions of artificial intelligence in education—A heuristic analysis on sociotechnical imaginaries and problematizations in policy guidelines. *Policy Futures in Education*, 22, 1662 - 1678. <https://doi.org/10.1177/14782103241228900>.
- Mahrishi, M., Abbas, A., & Siddiqui, M. (2024). Global Initiatives Towards Regulatory Frameworks for Artificial Intelligence (AI) in Higher Education. *Digital Government: Research and Practice*, 6, 1 - 9. <https://doi.org/10.1145/3672462>.
- Mallik, S., & Gangopadhyay, A. (2023). Proactive and reactive engagement of artificial intelligence methods for education: a review. *Frontiers in Artificial Intelligence*, 6. <https://doi.org/10.3389/frai.2023.1151391>.
- OECD. (2026). *OECD digital education outlook 2026: How artificial intelligence is transforming education*. OECD Publishing. https://www.oecd.org/en/publications/oecd-digital-education-outlook-2026_062a7394-en.html.
- Rahm, L. (2021). Educational imaginaries: governance at the intersection of technology and education. *Journal of Education Policy*, 38, 46 - 68. <https://doi.org/10.1080/02680939.2021.1970233>.
- Rincón, E., Jiménez, D., Aguilar, L., Flórez, J., Tapia, Á., & Peñuela, C. (2025). Mapping the use of artificial intelligence in medical education: a scoping review. *BMC Medical Education*, 25. <https://doi.org/10.1186/s12909-025-07089-8>.
- Sangwa, S. (2025). Navigating the Algorithmic Turn: A Dynamic Governance Framework for Ethical and Equitable AI Integration in Education. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.5281882>.
- Schiff, D. (2021). Education for AI, not AI for Education: The Role of Education and Ethics in National AI Policy Strategies. *International Journal of Artificial Intelligence in Education*, 32, 527 - 563. <https://doi.org/10.1007/s40593-021-00270-2>.
- Selwyn, N. (2024). On the Limits of Artificial Intelligence (AI) in Education. *Nordisk tidsskrift for pedagogikk og kritikk*. <https://doi.org/10.23865/ntpk.v10.6062>.

- Sposato, M. (2025). Artificial intelligence in educational leadership: a comprehensive taxonomy and future directions. *International Journal of Educational Technology in Higher Education*, 22. <https://doi.org/10.1186/s41239-025-00517-1>.
- Strielkowski, W., Grebennikova, V., Lisovski, A., Rakhimova, G., & Vasileva, T. (2024). AI-driven adaptive learning for sustainable educational transformation. *Sustainable Development*. <https://doi.org/10.1002/sd.3221>.
- Tran, M., Balasooriya, C., Jonnagaddala, J., Leung, G., Mahboobani, N., Ramani, S., Rhee, J., Schuwirth, L., Najafzadeh-Tabrizi, N., Semmler, C., & Wong, Z. (2025). Situating governance and regulatory concerns for generative artificial intelligence and large language models in medical education. *NPJ Digital Medicine*, 8. <https://doi.org/10.1038/s41746-025-01721-z>.
- Vargas-Salas, O., De Manchego, V., Cateriano-Chávez, T., & Molina-Rodríguez, F. (2025). Reform in University Education as a Consequence of Artificial Intelligence: A Systematic Review of the Literature. *Journal of Educational and Social Research*. <https://doi.org/10.36941/jesr-2025-0086>.
- Wang, Y. (2024). Algorithmic decisions in education governance: implications and challenges. *Discover Education*, 3. <https://doi.org/10.1007/s44217-024-00337-x>.
- Williamson, B. (2015). Governing software: networks, databases and algorithmic power in the digital governance of public education. *Learning, Media and Technology*, 40, 105 - 83. <https://doi.org/10.1080/17439884.2014.924527>.
- Xu, X. (2025). AI optimization algorithms enhance higher education management and personalized teaching through empirical analysis. *Scientific Reports*, 15. <https://doi.org/10.1038/s41598-025-94481-5>.