



The Human Mind in Recursion: Ethical AI Policy in the K–12 Classroom

Ava DiGiuseppe

I. EXECUTIVE SUMMARY

Artificial intelligence has rapidly entered K–12 classrooms, promising personalized learning, enhanced accessibility, and real-time feedback. From adaptive and AI-generated writing support to predictive attendance models and automated grading tools, schools across the United States are increasingly integrating AI into the fabric of daily instruction. For under-resourced schools—those facing systemic funding gaps, staffing shortages, outdated materials, or digital infrastructure deficits—AI offers potential relief for students and administrators. Likewise, students historically underserved by conventional education models, including English language learners, students with learning differences, and those in high-poverty districts, may benefit from more immediately responsive instruction. Yet the promise of AI cannot be separated from its evolving ethical implications. This brief adopts a definition of ethical AI usage that centers student autonomy, developmental integrity, and cognitive growth. Ethical usage in the education sector does not seek to replace autonomous thought, existing epistemological frameworks, automate creativity, or diminish pedagogical relationships, but operates as an assistive tool, reinforcing rather than outsourcing critical processes of learning. Because AI is a rapidly

developing field, this definition will evolve in tandem, but at present, it must be grounded in student-centered values and an educational mission of intellectual formation

Without proactive policy, the use of AI in the education sector risks being shaped by opaque systems, unregulated data usage, and disparities in access. This brief outlines pathways for ethical, equitable, and educationally sound implementation across U.S. public schools. Drawing from recent federal frameworks and institutional case studies, it recommends policy options that promote algorithmic transparency, human oversight, inclusive design, data privacy, and AI literacy.

II. OVERVIEW

AI that is operationalized for the education sector refers to a broad range of technologies that simulate cognitive functions such as prediction, personalization, and feedback. These include adaptive learning platforms that adjust to student performance in real time, AI tutors offering individualized instruction, predictive analytics systems used to forecast academic outcomes or flag students for intervention, and programs used to provide quality assurance for pedagogical instruction. While many of these tools fall under the term *narrow AI*—systems designed for specific tasks—*generative AI* models, such as large language models (LLMs), are now entering the classroom,

raising new questions about authorship, cognitive development, and intellectual autonomy.

AI holds considerable promise: it can enhance instructional efficiency, support overburdened teachers, and tailor educational experiences to the needs of the learner; yet, when deployed without sufficient oversight, AI systems risk eroding administrative integrity, amplifying bias, and undermining student development. This has been starkly reflected in cases like Houston's teacher evaluation algorithm, which moved from incentivizing performance to punishing educators for factors beyond their control, a modern echo of the punitive outcomes of No Child Left Behind [1].

Recognizing such risks, the Biden administration released the Blueprint for an AI Bill of Rights to establish civil protections in the age of automation [2]. Its principles—privacy, explainability, safety, and equity—are directly applicable to K-12 academic settings and affirm that ethical AI governance is a bipartisan imperative. Complementing this is the OECD's AI Capability Indicators, which compare machine performance to nine core human abilities, offering policymakers a lens through which to assess how AI can aid, rather than replace, human learning and cognition [3].

III. RELEVANCE

In 2023, the rapid deployment of AI in K-12 classrooms has transitioned from experimentation to expectation. School districts are adopting tools like Khanmingo, Google Translate, and AI writing assistant with little delay, often under pressure to compete with digitally forward peers or respond to shrinking teacher workforces. A December 2023 EdWeek Research Center survey

found that one-third of K-12 teachers are already using generative AI at their own discretion to build rubrics, draft communications, or design lessons, often without system-wide guidance [4]

The acceleration has provoked concern among privacy advocates and IT Leaders. At the 2023 CITE Conference, education technologists emphasized that few districts have robust review protocols or equity-centered policies in place. In response, districts like California's Palmdale School District now require liability insurance and collective IT oversight for ed-tech procurement—guardrails often absent in under-resourced districts [5]

Additionally, AI adoption intersects sharply with civil rights law, particularly as policymakers navigate the existence and usage of AI under federal protocols that protect student data privacy like FERPA and the ADA, which were not designed to address behavioral prediction, 24/7 surveillance, or opaque algorithmic profiling. As Stanford Law Review author Danielle Citron outlines in "The Surveilled Student," monitoring platforms embedded in school-issued devices can track everything from keystrokes to location data, often outside school hours, which undermines trust, discourages expressive learning, and disproportionately harms marginalized students [6].

It is notable that recent journalism on Congress's 2025 One Big Beautiful Bill Act reveals an intraparty fracture within the GOP on the issue of AI regulation, with Trumpian ideologues favoring unregulated use and moderate conservatives calling for guardrails to prevent potentially existential outcomes of unrestricted AI usage [7]. This divide reflects a volatile national

policy climate, one in which both partisan gridlock and unified, but ideologically extreme, governments delay meaningful protections even as the education sector absorbs the consequences [8].

IV. HISTORY

A. Current Stances

Over the past half-decade, the role of AI in K-12 education has undergone a paradigmatic shift. Early applications were bounded by specificity: individualized education program (IEP) optimization, adaptive math platforms, and predictive attendance tools. These interventions functioned largely as assistive mechanisms supported by human oversight; yet, the proliferation of generative AI from 2022 to 2023 destabilized this model, introducing systems that simulate human cognition and creativity. This expansion has proceeded largely in the absence of robust governance frameworks.

A February 2024 EdWeek Research Center survey reveals that 79% of educators report no formal district policy guiding the use of generative AI, while 56% anticipate greater reliance on such tools in the coming year. Notably, educators retain wide latitude in use, but students remain subject to restrictions, which reflects a reactive, often punitive posture that privileges institutional control over technological literacy or student learning development [9]. These findings align with academic research revealing that high schools are significantly less likely than higher education institutions to draft comprehensive policies on AI usage, particularly on issues of algorithmic transparency, data governance, and student privacy [10].

This policy vacuum is shaped by ideological discord at the federal level. The proposed AI regulation provision in the original House version of the One Big Beautiful Bill Act would have barred states from regulating AI for a decade, effectively subordinating local autonomy to national laissez-faire ambitions. Moderate Republicans such as Susan Collins, alongside Democrats, voiced concerns that this bill subverts core tenets of states' rights and risks entrenching disparities between states in educational AI access and oversight [11]. Legislative stagnation has thus delegated responsibility to the states, where hundreds of AI-related education bills have emerged in uneven waves.

Indeed, at the state level, AI governance increasingly reflects the contours of American political polarization. Regulatory inertia is no longer merely a consequence of federal gridlock, but stems from epistemological divisions over the risks and purposes of AI, entrenched policy monopolies in technology sectors, and disparities in stakeholder influence. These dynamics have begun to shape education-specific AI legislation, indicating that the sector may become a battleground for broader ideological realignment in science and technology policy.

V. POLICY PROBLEM

A. Stakeholders

The regulation of AI in K-12 classrooms implicates both direct stakeholders—students, teachers, administrators—and indirect ones, such as policymakers, politicians, and ed-tech developers, each with competing sociopolitical interests and constitutional concerns. From a

prosocial standpoint, it must be emphasized that students—particularly those from historically underserved populations, including low-income, neurodiverse, and multilingual learners—are most vulnerable to algorithmic bias and systemic disparities. Parents increasingly express concern over ambiguous surveillance practices, data retention, and the erosion of expressive freedoms, raising First and Fourteenth Amendment questions about students’ rights to intellectual privacy and equal protection under the law.

Teachers, meanwhile, are expected to integrate emerging technologies without standardized training or sustained professional development, resulting in inconsistent implementation. School districts—especially those lacking sufficient funding or technical infrastructure—often adopt AI platforms under vendor pressure, in response to staffing shortages, or in pursuit of rapid gains in performance metrics. At both federal and state levels, education departments have issued technically dense guidance, but legislative action remains limited due to a lack of legal precedent and jurisdictional clarity. In this regulatory vacuum, private tech companies—largely unaccountable to pedagogical standards—wield disproportionate influence over classroom cognition. Without coherent policy frameworks and collaborative development, AI governance in education risks being shaped more by commercial imperatives and administrative expedience than by constitutional values or student-centered educational goals.

B. Risks of Indifference

Without proactive regulation, the integration of AI into K-12 classrooms risks formalizing

inequities under the guise of innovation. For example, AI-powered early warning systems designed to flag students at risk of academic failure often rely on incomplete or biased datasets (e.g. attendance, disciplinary records, test scores), disproportionately labeling students from low-income or racially marginalized communities as underperforming. These systems can unintentionally replicate historic patterns of educational tracking, precluding students from advanced coursework or enrichment opportunities based on algorithmic predictions rather than educator judgment or demonstrated growth.

Meanwhile, the absence of modernized student data privacy protections—particularly the lack of FERPA updates—allows AI vendors to retain, analyze, and potentially commercialize sensitive student information. Companies that offer so-called “free” ed-tech platforms may operate under business models that harvest data to train future models or target users with tailored educational content.

In pedagogical terms, automation threatens to deskill teachers by reducing lesson planning, assessment, and feedback to preprogrammed templates, diminishing professional agency and responsiveness to student needs. The pressure to adopt AI tools without sufficient training or instructional alignment also creates inconsistency across classrooms, often reinforcing disparities between affluent districts with tech-literate faculty and under-resourced schools reliant on default vendor packages.

Most critically, unregulated AI may reshape how students develop foundational skills: auto-generated reading summaries weaken text

engagement; AI writing assistants can displace authentic drafting; and adaptive learning platforms may oversimplify complex problem-solving. Left unchecked, this risks eroding intellectual autonomy and compromising the developmental integrity of the learning process itself

C. Nonpartisan Reasoning

The ethical governance of AI in K-12 education must transcend partisan agendas and instead be rooted in enduring educational and democratic principles such as intellectual autonomy, developmental integrity, and equity of access in order to support beneficial usage. Rather than reducing AI regulation to a question of political control, a nonpartisan framework demands careful attention to what a robust, American education ought to preserve—human judgment and the cultivation of the mind

- 1) *Equity as an ethical imperative:* A just education system cannot allow advanced technologies to reinforce structural inequality. Without coordinated policy, AI will remain concentrated in affluent districts, exacerbating disparity in cognitive development and educational outcomes. Founded in distributive justice, equitable access to ethically sound tools is not only a matter of fairness, but a condition for sustaining the public purpose of education
- 2) *Federal infrastructure, local oversight:* While education in the U.S. has long favored local control, ethical regulation requires standardization in areas that implicate student rights and human development. National standards on data use,

algorithmic transparency, and pedagogical oversight can coexist with local discretion on curricular matters; these conditions honor federal subsidiarity without sacrificing normative guardrails.

- 3) *Supporting cognitive and moral development:* If the primary purpose of public education in the U.S. is to support the cognitive development of self-governing individuals, AI tools must support, not supplant, the teacher-student relationship, which is foundational to moral development and epistemic trust. Policies should prioritize professional development and philosophical literacy to ensure educators can guide students through AI-mediated environments while preserving the necessity of human oversight.
- 4) *Proportional regulation:* Excessive surveillance, predictive profiling, and behavioral datafication risk violating students' intellectual privacy and expressive freedoms. At the same time, a laissez-faire approach invites market capture and ethical erosion. Ethically proportionate governance limits coercion while securing students' developmental rights under the First and Fourteenth Amendments.
- 5) *Democratic safeguards against algorithmic reductionism:* Educational technology must not reduce learning to measurable outputs alone. Under these ethical conditions, an AI-supported framework resists technocratic determinism and affirms that the learner is not a data point, but a moral

agent capable of choice, resistance, reasoning, and conscience development.

- 6) *Preserving the telos of education:* In a democratic society, education must cultivate minds that are not merely efficient, but intellectually independent. AI, if left unregulated, risks prioritizing result-driven metrics and optimization over rudimentary learning skills, which inhibits not only scholastic ability, but also personal and professional self-actualization. Policies developed under this ethical framework ensure that technological innovation remains subordinate to pedagogical purpose.

VI. TRIED POLICY

Though comprehensive AI regulation remains elusive, recent federal, state, and district-level initiatives illustrate the emerging contours of governance in K-12 education. The White House's 2023 Blueprint for an AI Bill of Rights, though nonbinding, articulates five foundational principles—safe and effective systems, algorithmic transparency, privacy, notice and explanation, and human alternatives—that have shaped federal discourse and informed public-sector innovation. In California, Senate Bill 1288 (2024) requires school districts to inventory and report their AI tools to the California Department of Education, perform risk audits, and disclose third-party data-sharing practices—a first-of-its-kind accountability framework designed to address algorithmic opacity in public schools [12]. California's Assembly Bill 2876 (2024) complements this effort by embedding digital and media literacy, including instruction on algorithmic influence and AI-generated content,

into the state's core curriculum for K-12 students [13].

New York State's Assembly Bill A6972 (2023) sought to implement procurement regulations and ethical guardrails for AI tools used in classrooms, but failed in committee amid concerns about enforcement feasibility and technological neutrality [14]. Meanwhile, New York City's Department of Education entered into a high-profile partnership with Microsoft to deploy an AI-powered teaching assistant to support lesson planning and student feedback, effectively reinstating AI tools after an earlier moratorium [15].

Elsewhere, Senate Bill 382 (2024) and Florida Senate Bill 702 (2025) both proposed regulating AI procurement and ethical implementation in schools, but neither advanced to law, revealing legislative hesitation and ideological fault lines over the federal and local governments' role in AI oversight and [16] [17].

VII. POLICY OPTIONS

Mandate Transparency and Human Oversight: AI tools used in K-12 education must be subject to clear, enforceable transparency mandates. All school districts should be required to conduct regular audits of AI systems for algorithmic bias, data accuracy, and pedagogical efficacy. These audits should be made publicly available to foster community trust and accountability. Most critically, teachers must retain ultimate decision-making authority over instructional design and student evaluation, with AI functioning solely in a support role. Public disclosure of which algorithms are in use, along with documentation of their intended educational

outcomes are known limitations, can further demystify ethical AI usage and prevent its misuse

Enforce Privacy and Data Protection

Standards: To protect students rights in the age of machine learning, federal legislation must expand the Family Educational Rights and Privacy ACT (FERPA) to explicitly regulate AI-generated data, such as behavioral predictions or writing samples produced through student interaction with chatbots. Schools should be prohibited from selling student data or using it to train commercial models. High-risk AI tools—those that track emotional states, predict academic outcomes, or generate content—should require opt-in consent from families, particularly when deployed with minors. Clear limits on retention, access, and third-party sharing must be codified in law.

Promote Inclusive and Equitable AI Design:

Federal and state governments should provide competitive grants to ed-tech developers who adhere to universal design principles, ensuring AI tools are accessible to students with disabilities, multilingual learners, and others who face systemic barriers. Independent review panels composed of educators, ethicists, special education professionals, and technologists should vet proposed tools before classroom implementation. To reduce inequality in technological access, the federal government can pilot AI platforms in Title I schools under strict oversight protocols, accompanied by research partnerships that evaluate equity outcomes in real time.

Support AI Literacy for Teachers and Students:

States should launch professional development programs, particularly in regions where rural and underfunded schools often lag in digital readiness.

These programs must offer training in AI ethics, data interpretation, and classroom integration, emphasizing the importance of teacher agency. Concurrently, AI literacy should be embedded in middle and high school curricula, focusing on civic reasoning, authorship, and algorithmic accountability. Strategic partnerships between public schools, universities, and nonprofits can provide scalable resources for equitable AI education and prepare students for a future in which algorithmic systems shape economic and civic life.

VIII. CONCLUSIONS

To realize AI's full potential in public education, ethical design must serve as both foundation and filter as opposed to a reactive fix to structural vulnerabilities. This brief proposes a regulatory framework rooted in developmental theory, and constitutional protections to guide technology policy implementation. Federal and state interventions must not only encourage inclusive and transparent design, but also cut exploitative and pedagogically reductive applications. Absent such a vision, classrooms risk becoming experimental zones for unregulated technologies, undermining cognitive development, eroding public trust, and displacing human relationships at the heart of a rigorous and robust education.

ACKNOWLEDGMENT

The Institute for Youth in Policy wishes to acknowledge Mason Carlisle, Lilly Kurtz, Asher Cohen, Paul Kramer, and other contributors for developing and maintaining the Fellowship Program within the Institute.

REFERENCES

- [1] AI Now Institute, Center on Race, Inequality, and the Law, and Electronic Frontier Foundation. *Litigating Algorithms: Challenging Government Use of Algorithmic Decision Systems*. Litigating Algorithms. 2018. <https://ainowinstitute.org/wp-content/uploads/2023/04/litigatingalgorithms.pdf>.
- [2] Executive Office of the President White House Office of Science and Technology Policy. *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People*. By The White House Office of Science and Technology. 2022. <https://bidenwhitehouse.archives.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.
- [3] *Introducing the OECD AI Capability Indicators*. OECD, 2025. <https://doi.org/10.1787/be745f04-en>.
- [4] Langreo, Lauraine. "ChatGPT, Bard, Copilot, Khanmigo: What AI Tools Do Teachers Use?" EducationWeek. Editorial Projects in Education. Last modified February 28, 2024. Accessed June 29, 2025. <https://www.edweek.org/technology/chatgpt-bard-copilot-khanmigo-what-ai-tools-do-teachers-use/2024/02>.
- [5] Torchia, Rebecca. "CITE 2023: Protect Student Data Privacy with Ed Tech Policies." EdTech. Last modified November 30, 2023. Accessed June 29, 2025. <https://edtechmagazine.com/k12/article/2023/11/cite-2023-protect-student-data-privacy-ed-tech-policies-perfcon>.
- [6] Citron, Danielle Keats. "The Surveilled Student." *Stanford Law Review*, 2024, 1439-72. <https://review.law.stanford.edu/wp-content/uploads/sites/3/2024/10/Citron-76-Stan.-L.-Rev.-1439.pdf>.
- [7] One Big Beautiful Bill, H.R. H.R.1, 119th Cong., 2025-2026. (as engrossed in House, May 22, 2025). <https://www.congress.gov/bill/119th-congress/house-bill/1/text>.
- [8] Milman, Oliver, and Dharna Noor. "Trump's Tax Bill Seeks to Prevent AI Regulations." The Guardian. Last modified June 27, 2025. Accessed June 29, 2025. <https://www.theguardian.com/technology/2025/jun/27/trump-tax-bill-ai-climate-emissions>.
- [9] Klein, Alyson. "Schools Are Taking Too Long to Craft AI Policy. Why That's a Problem." EducationWeek. Editorial Projects in Education. Last modified February 19, 2024. Accessed June 30, 2025. <https://www.edweek.org/technology/schools-are-taking-too-long-to-craft-ai-policy-why-thats-a-problem/2024/02>.
- [10] Ghimire, Aashish, and John Edwards. "From Guidelines To Governance: A Study Of AI Policies In Education." *Communications in Computer and Information Science*, 2024, 299-307. https://doi.org/10.1007/978-3-031-64312-5_36.
- [11] Chow, Andrew R. "Why AI Regulation Has Become a 'States' Rights' Issue." Time. TIME USA. Last modified June 25, 2025. Accessed June 30, 2025. <https://time.com/7297580/ai-moratorium-senate-big-beautiful-bill/>.
- [12] Public Schools: Artificial Intelligence Working Group Group, S. 1288, 2024th Leg. (Cal. Sept. 28, 2024). Accessed July 16, 2025. https://calmatters.digitaldemocracy.org/bills/ca_202320240sb1288.
- [13] Pupil Instruction: Media Literacy: Artificial Intelligence Literacy: Curriculum Frameworks: Instructional Materials., A. 2876, 2024th Leg. (Cal. Sept. 29, 2024). Accessed July 16, 2025.
- [14] A. 6972, 2025th Leg. (N.Y.). Accessed July 16, 2025.

<https://www.nysenate.gov/legislation/bills/2025/A6972>

- [15] Donaldson, Sahalie. "After Initially Shunning Artificial Intelligence, NYC Schools Partner with Microsoft on AI Teaching Assistant." City and State New York. Last modified September 14, 2023. Accessed July 16, 2025. <https://www.cityandstateny.com/policy/2023/09/after-initially-shunning-artificial-intelligence-nyc-schools-partner-microsoft-ai-teaching-assistant/390292>.

- [16] S. 382, 2025th Leg. (Tex.). Accessed July 16, 2025.

<https://legiscan.com/TX/text/SB382/id/3028209>

- [17] Provenance of Digital Content, S. 702, 2025th Leg. (Fla.). Accessed July 16, 2025. <https://www.flsenate.gov/Session/Bill/2025/702>