

# AI-Driven Personalized Learning Systems for K-12 Education: Enhancing Educational Equity and Outcomes in the United States

Jason Miller

Department of Computer Science and Education Technology, University of Texas  
at Arlington  
jason.miller@uta.edu

Mary Johnson

College of Education, California State University, Long Beach  
mary751@csulb.edu

## Abstract

AI-driven personalized learning systems are increasingly positioned as a transformative infrastructure for K-12 education in the United States, promising to tailor instruction, accelerate learning, strengthen teacher decision-making, and reduce long-standing disparities in educational opportunity. Yet the practical significance of these systems lies less in their adaptive interface features than in the institutional, technical, and policy arrangements through which they are designed, deployed, governed, and sustained. This paper develops a system-level analysis of AI-driven personalized learning for U.S. elementary and secondary education, arguing that the core challenge is not simply whether personalization technologies can improve student performance in isolated settings, but whether they can be integrated into public education in ways that are educationally meaningful, operationally robust, fiscally sustainable, and normatively just. The analysis synthesizes scholarship from intelligent tutoring systems, learning analytics, educational data infrastructure, algorithmic fairness, public-sector technology governance, and U.S. education policy. It examines architectural design choices, interoperability constraints, teacher-facing workflows, procurement incentives, district capacity variation, data governance regimes, and the political economy of platform adoption. The paper contends that personalized learning systems can contribute to equity only when they are treated as socio-technical infrastructures rather than standalone software products. This requires deliberate attention to curriculum alignment, human oversight, public accountability, broadband and device access, representational fairness, model transparency, and the preservation of pedagogical professionalism. The paper concludes by proposing a governance-oriented framework for the next generation of K-12 personalization systems, one that balances innovation with democratic accountability and situates AI within a broader educational mission concerned with opportunity, inclusion, and durable institutional trust.

## **Keywords**

AI in education; personalized learning; K-12 education; educational equity; learning analytics; intelligent tutoring systems; algorithmic fairness; education policy; socio-technical systems; digital infrastructure

## **1. Introduction**

The growing interest in AI-driven personalized learning systems reflects a broader historical pattern in U.S. education reform in which technological innovation is repeatedly presented as a solution to entrenched structural inequities. From teaching machines and programmed instruction to computer-assisted learning, adaptive tutoring, and contemporary platform-based analytics, personalization has long been associated with the aspiration to provide every learner with instruction calibrated to individual needs, pace, prior knowledge, and motivation. In its contemporary AI-inflected form, this aspiration has intensified. Schools and districts are now presented with systems that claim to diagnose misconceptions in real time, recommend differentiated content pathways, generate formative assessments, predict learning risk, automate aspects of feedback, and provide teachers with dashboards that promise fine-grained insight into student progress. In a national educational landscape marked by uneven academic recovery, persistent achievement gaps, workforce concerns, and accelerating digitalization, such promises carry considerable political and institutional appeal.

Yet the educational and social implications of these systems cannot be understood through a narrow focus on instructional efficacy alone. The practical question facing U.S. K-12 education is not whether machine learning, recommendation engines, or adaptive content sequencing can produce statistically significant gains under selected conditions. Rather, the central issue is whether AI-driven personalization can function as a reliable and equitable component of public education systems that differ widely in resources, governance structures, teacher capacity, curriculum coherence, broadband access, data maturity, and public trust. This paper adopts that broader frame. It treats personalized learning systems not as isolated pedagogical applications but as socio-technical infrastructures embedded in public institutions, shaped by procurement markets, accountability policy, interoperability standards, labor relations, and community expectations.

This shift in analytical perspective is essential because many of the strongest claims about AI in education obscure the distinction between local technical optimization and system-level educational improvement. An algorithm may successfully recommend the next exercise for a learner, but that does not ensure that the resulting system aligns with state standards, supports teachers' professional judgment, protects children's privacy, avoids amplifying bias, or remains usable under conditions of unstable staffing and limited technical support. Similarly, a platform may demonstrate benefits in well-resourced pilot environments while failing to deliver comparable results in districts serving higher concentrations of low-income students, emergent bilingual learners, rural communities, or students with disabilities. A genuinely equity-oriented evaluation of AI-driven personalization must therefore examine not only

learner-model accuracy or effect sizes, but also infrastructure distribution, implementation fidelity, governance capacity, and the institutional conditions of scale.

The U.S. policy environment makes this inquiry especially urgent. The country's K-12 system is decentralized, locally administered, and heavily unequal in fiscal and technical capacity. School districts differ dramatically in vendor evaluation expertise, legal resources, data architecture, curricular standardization, and professional development capability. These variations matter because AI-driven personalized learning systems depend on dense and continuous flows of data, stable classroom routines, and coherent implementation support. In such an environment, the same platform can be empowering in one context and extractive or disruptive in another. Moreover, as national and international organizations increasingly emphasize the need for human-centered and ethically governed AI in education, the tension between innovation and accountability has become a defining feature of the field. UNESCO has stressed that AI in education must be aligned with human-centered values and governance, while U.S. educational guidance has increasingly highlighted both the potential of personalized digital resources and the need to bridge unequal access to technology.

This paper argues that AI-driven personalized learning systems can contribute to improved educational outcomes and greater equity in the United States only when they are governed as public-interest infrastructures rather than adopted as market-ready efficiency tools. The argument proceeds by situating personalized learning within the historical development of adaptive educational technologies, examining the technical architectures that support personalization, analyzing equity claims and their structural limitations, investigating governance and accountability dilemmas, and outlining a forward-looking framework for robust and just deployment. The paper maintains a system-level orientation throughout. Its central premise is that the future of AI in K-12 education depends less on increasingly sophisticated personalization algorithms than on whether public institutions can shape those systems in ways that sustain pedagogical integrity, democratic legitimacy, and distributive fairness.

## **2. From Individualization to AI-Driven Personalization: Historical and Conceptual Foundations**

The contemporary discourse on AI-driven personalized learning is best understood as the latest phase in a longer intellectual and institutional history of educational individualization. Earlier visions of individualized instruction emerged from behaviorist approaches that decomposed knowledge into sequenced units and sought to optimize reinforcement through tightly structured learner interactions. Computer-assisted instruction expanded these efforts by enabling branching logic, automated scoring, and limited forms of adaptive response. Later, intelligent tutoring systems introduced richer learner modeling, domain representation, and adaptive scaffolding, often informed by cognitive science and human-computer interaction research. The current wave differs not because it invents personalization from nothing, but because it combines earlier aspirations with large-scale data collection, cloud-based delivery, predictive analytics, and increasingly powerful machine learning techniques.

This historical continuity matters because it reveals a recurring pattern: educational technology repeatedly promises to solve the heterogeneity problem of classroom teaching by making instruction more responsive to individual variation. The appeal is understandable. K-12 classrooms in the United States are characterized by wide differences in prior knowledge, language background, socio-emotional readiness, attendance stability, and access to out-of-school learning support. Personalized systems appear to offer a means of addressing this complexity more systematically than traditional one-size-fits-all instruction. However, the history of educational technology also shows that claims of transformation often exceed the institutional capacity required for durable implementation. Many technologies have performed well in controlled settings but underperformed when diffused into diverse school environments with competing demands, fragmented data systems, and limited teacher time. This historical pattern cautions against viewing AI as a straightforward solution to instructional variation.

Conceptually, personalized learning itself is not a singular construct. It encompasses multiple educational logics that are often conflated in public discourse. One meaning refers to adaptive instructional sequencing, in which software modifies content difficulty, pacing, or support based on student interactions. Another refers to learner-centered pedagogical design, emphasizing student agency, relevance, and differentiated pathways. A third concerns data-informed decision support for teachers, including dashboards, alerts, and grouping recommendations. A fourth involves administrative prediction, such as early warning systems for academic risk or disengagement. These forms of personalization differ in purpose, evidence base, ethical profile, and infrastructure requirements. A system that adapts mathematics practice items based on mastery is not equivalent to a platform that predicts dropout risk or one that generates individualized narrative feedback. Treating them as interchangeable obscures important design and governance distinctions.

AI intensifies this conceptual ambiguity because it is invoked across a wide range of technologies, some of which rely on advanced predictive modeling and some of which use relatively conventional rule-based adaptation. In the K-12 context, AI typically refers to systems that infer learner states from interaction data, optimize recommendations, automate content generation or assessment, or support classification and prediction tasks. But the presence of AI does not in itself determine educational quality. What matters is how the inferential capabilities of AI are integrated into teaching and learning processes. A high-performing predictive model can be pedagogically trivial if it recommends content poorly aligned with classroom instruction. Conversely, a comparatively simple adaptive mechanism can be educationally valuable when embedded in coherent curriculum, strong teacher mediation, and meaningful feedback loops.

The distinction between personalization as technical adaptation and personalization as institutional arrangement is particularly important. Technical adaptation concerns how the system modifies inputs, tasks, or support in response to student data. Institutional arrangement concerns how schools organize curriculum, staffing, assessment, procurement, and accountability around the use of such systems. In many cases, the success or failure of personalization depends less on the sophistication of the algorithm than on the quality of the

institutional arrangement. If teachers are not given time to interpret dashboard data, if students encounter fragmented and incoherent digital content across platforms, if district leaders cannot audit vendor claims, or if the system undermines rather than complements classroom relationships, the promise of personalization collapses into operational friction.

A system-level perspective therefore requires that personalization be treated as a layered construct. At the instructional layer, it concerns differentiated experiences and supports. At the technical layer, it concerns learner models, content recommendation engines, and interoperable data flows. At the organizational layer, it concerns school routines, professional development, and support structures. At the policy layer, it concerns standards, privacy, procurement, and accountability. At the normative layer, it concerns what kinds of educational differences should be recognized, amplified, or mitigated. Personalized learning systems do not simply respond to learners; they also encode assumptions about what counts as progress, what forms of variation are educationally salient, and whose goals matter in the design of educational pathways.

This layered view is essential in the U.S. context, where personalization is often framed as a means of advancing equity. That claim can be partially valid, but only under specific conditions. Personalization can make previously invisible forms of learner variation more legible and can help tailor supports for students who might otherwise be poorly served by standardized pacing. At the same time, data-driven personalization can also harden labels, reproduce historical inequities through biased training data, and fragment common educational experiences in ways that undermine democratic schooling. The key analytical challenge is not whether personalization is inherently equitable or inequitable, but how its educational logic interacts with the structural conditions of public education.

### **3. Architectural Foundations of AI-Driven Personalized Learning Systems**

AI-driven personalized learning systems in K-12 education are composed of multiple interacting technical layers, each of which has important implications for educational function and institutional governance. At the most basic level, these systems depend on data capture, learner modeling, decision logic, content delivery, and feedback presentation. The data layer typically includes clickstream records, response accuracy, time on task, completion patterns, item-level performance, attendance-linked data, behavioral logs, and occasionally writing samples or speech data. These streams feed into learner models that attempt to estimate mastery, engagement, misconception patterns, or future performance risk. The decision layer uses these estimates to recommend activities, generate interventions, assign pathways, or alert teachers. The content layer delivers exercises, explanations, assessments, multimedia materials, or generative feedback. Finally, the interface layer presents outputs to students, teachers, or administrators through dashboards, adaptive lesson sequences, or notification systems.

The architecture of these systems is not merely a technical concern; it shapes the institutional distribution of authority. A platform that centralizes decision-making within opaque recommendation engines gives educators less visibility into how learning pathways are

determined. A system that allows teachers to inspect mastery assumptions, override recommendations, and contextualize model outputs preserves greater human agency. Similarly, a system built on interoperable standards and modular components creates more public leverage than one that locks districts into proprietary data schemas and bundled content ecosystems. Architectural decisions therefore influence not only performance but also contestability, auditability, and long-term dependence.

One of the defining tensions in contemporary personalized learning architecture concerns the balance between prediction and interpretation. Machine learning models may improve the predictive accuracy of recommendations, risk flags, or next-step sequencing, but increased accuracy can come at the cost of reduced interpretability. In K-12 settings, where system outputs affect children and interact with teacher judgment, this trade-off is especially consequential. A district may be persuaded by strong vendor claims regarding predictive precision, yet precision alone does not resolve the practical question of whether teachers can understand, trust, and appropriately use the recommendations produced. If systems generate outputs that appear authoritative but are poorly explainable, schools may either over-rely on them or ignore them entirely. Both outcomes weaken educational value.

Another architectural issue concerns the granularity of personalization. Some systems personalize at the level of individual items, adjusting difficulty or hints within a narrow skill domain. Others personalize at the level of broader learning trajectories, recommending units, pacing shifts, or intervention categories. More expansive systems attempt to integrate academic, behavioral, and socio-emotional data into holistic learner profiles. Each level of granularity brings different risks and opportunities. Fine-grained item adaptation can support mastery but may overemphasize micro-skills at the expense of conceptual coherence. Broader trajectory recommendations may better align with curriculum but require stronger model validity and more sophisticated educator oversight. Holistic systems promise comprehensive support but raise acute privacy, consent, and fairness concerns, especially when inferences extend beyond observable academic performance.

The effectiveness of personalization also depends on content architecture. Adaptive systems are only as educationally meaningful as the content libraries and pedagogical assumptions they draw upon. If the underlying content corpus is shallow, culturally narrow, poorly sequenced, or weakly aligned with state standards, no amount of algorithmic adaptation can compensate. Conversely, a robust content architecture can magnify the benefits of even modest adaptive logic. This is why many of the most durable systems historically have been domain-specific, especially in mathematics, where knowledge structures lend themselves more readily to procedural decomposition and mastery progression. Extending personalization across broader disciplinary areas such as writing, social studies, or inquiry-based science is more complex because these domains rely heavily on interpretation, context, discussion, and open-ended reasoning.

Infrastructure architecture further shapes the scalability of personalized learning systems. Districts must contend with device availability, bandwidth stability, identity management, authentication, system integration, and data governance. A platform that assumes

uninterrupted connectivity, one-to-one device access, and well-maintained student information systems may function adequately in affluent suburban districts while failing in under-resourced schools with shared devices, unstable Wi-Fi, or fragmented vendor ecosystems. This infrastructural dependency is one reason digital innovation often reproduces rather than reduces educational inequality. The promise of personalization presupposes an underlying baseline of digital functionality that remains unevenly distributed across the United States. U.S. educational guidance has explicitly linked effective personalization with access to varied digital resources and the need to address the digital divide, underscoring that personalization cannot be separated from material conditions of access.

A final architectural challenge concerns the growing incorporation of generative AI. Unlike earlier adaptive systems that primarily selected from predefined content or rule-based hints, generative models can produce explanations, examples, questions, summaries, and feedback on demand. This expands flexibility but also introduces instability. Generated content may be inaccurate, misaligned with curricular intent, developmentally inappropriate, or inconsistent across students. In K-12 settings, such variability creates serious quality assurance problems. The use of generative AI therefore heightens the need for constrained design, human review pathways, age-appropriate safeguards, and explicit governance rules. UNESCO's recent guidance emphasizes that generative AI in education should be implemented through human-centered and policy-guided approaches rather than unconstrained technological enthusiasm.

#### **4. Educational Equity: Promise, Ambiguity, and Structural Limits**

The equity case for AI-driven personalized learning rests on a compelling intuition: students differ in readiness, exposure, support, language background, and learning pace, and educational systems that ignore these differences tend to privilege those already aligned with dominant instructional norms. Personalization, in theory, can respond to learners who need more time, alternative explanations, scaffolded practice, or advanced challenge. It can help identify hidden struggles earlier, reduce idle time for students working below or above grade level, and provide teachers with more precise information about where support is needed. For historically underserved students, such responsiveness seems especially valuable. In districts confronting large class sizes, teacher shortages, and uneven access to specialists, personalized systems are often framed as mechanisms for extending differentiated support more consistently than human labor alone can sustain.

However, the equity promise of personalization is double-edged because it addresses symptoms of instructional mismatch without necessarily changing the structural conditions that produce educational inequality. Many of the most consequential disparities in K-12 education arise from school funding inequities, segregated housing patterns, differential access to experienced teachers, varied course offerings, exposure to environmental stressors, and disparities in health and family support systems. Personalized learning platforms can modulate the instructional experience within schools, but they do not eliminate these broader structural conditions. There is therefore a risk that personalization becomes a compensatory

technological strategy that is rhetorically linked to equity while leaving deeper inequalities intact.

Moreover, personalization can itself become a vector of inequity when systems are trained on data shaped by prior disadvantage. If models infer aptitude, engagement, or risk from historical performance patterns, they may reproduce the effects of unequal opportunity rather than identify latent capacity. Students who have had inconsistent access to high-quality instruction, unstable attendance due to transportation or caregiving burdens, or limited broadband at home may generate data traces that systems interpret as low persistence or limited readiness. When such inferences drive recommendations, personalization can quietly convert structural disadvantage into individualized educational pathways. The result is not overt discrimination but algorithmically mediated stratification.

The risk becomes even greater when personalization systems operate as de facto sorting mechanisms. In some contexts, adaptive software can provide additional support that helps students close gaps. In other contexts, it can channel lower-performing students into narrow remedial loops while their more advantaged peers gain access to richer, faster, or more exploratory content. This dynamic resembles earlier forms of tracking, except that the classification process is embedded in software and often less visible to families and teachers. The rhetoric of individualized support can therefore mask the emergence of differentiated learning ecologies that are difficult to monitor and contest.

Educational equity also involves more than equalizing measured academic performance. It concerns participation, belonging, access to intellectually rich curricula, recognition of cultural and linguistic diversity, and the capacity to engage in shared civic learning. Personalized systems optimized narrowly for test-aligned progression may improve some short-term outcomes while weakening exposure to common curricular experiences and collaborative learning. This is particularly salient in K-12 schooling, where education is not only about the efficient transmission of skills but also about socialization, democratic formation, and collective participation. An equity framework that focuses exclusively on individual optimization risks sidelining these broader public purposes.

The digital divide further complicates equity claims. Although device access improved substantially during and after the COVID-19 pandemic, meaningful digital inclusion still depends on home connectivity quality, technical support, quiet learning environments, language-accessible interfaces, accessibility features, and family digital literacy. Rural communities, low-income households, students in temporary housing, and some Indigenous and migrant populations remain especially vulnerable to infrastructural exclusion. Personalized learning systems that assume stable access can widen opportunity gaps even when nominally available to all students. This is why evidence from digital learning more broadly has often cautioned that large-scale implementation does not automatically improve outcomes for disadvantaged students without deliberate attention to context and support. Brookings noted that digital learning has frequently been oversold relative to the complexity of implementation and the challenge of improving outcomes at scale for disadvantaged learners.

There are also important issues of representation and disability justice. Students with disabilities may benefit significantly from adaptive pacing, multimodal presentation, speech-based interfaces, or feedback personalization. At the same time, systems that are insufficiently designed for accessibility can create new barriers. Automated attention metrics, behavioral flags, or normative engagement thresholds may misread neurodivergent patterns of participation. Speech and writing models may perform inconsistently for students with assistive devices or atypical language production. Equity-oriented deployment therefore requires a universal design perspective in which accessibility is foundational rather than retrofitted.

For emergent bilingual learners, personalization presents a parallel tension. Adaptive systems can provide translation supports, targeted vocabulary scaffolds, and differentiated language exposure. Yet many commercially available platforms remain Anglocentric in both interface design and pedagogical assumptions. When multilingual support is shallow or inaccurate, personalization can place additional interpretive burdens on students and teachers. More fundamentally, systems designed around monolingual progression metrics may under-recognize the assets and cognitive complexity associated with bilingual development. Equity in this domain requires that personalization systems be linguistically and culturally responsive at the level of data, content, and evaluation.

Thus, the central equity insight is not that personalization is futile, but that its benefits are conditional and institutionally mediated. Personalized learning can support equity when it expands access to high-quality instruction, enhances teacher responsiveness, and provides flexible supports without narrowing opportunity. It undermines equity when it operates as a low-cost substitute for human investment, reproduces historical bias through predictive inference, or fragments educational quality along existing lines of advantage. A serious equity agenda for AI in K-12 education must therefore move beyond access and efficacy to include governance, transparency, representation, and the preservation of broad educational opportunity.

## **5. Teachers, Professional Judgment, and the Reorganization of Instructional Labor**

AI-driven personalized learning systems are often marketed through a dual narrative of support and efficiency. On the one hand, they are said to relieve teachers of burdensome administrative tasks, automate routine feedback, and make differentiation more manageable. On the other hand, they are presented as augmenting teacher insight through analytics dashboards, risk alerts, and fine-grained mastery data. Both claims contain elements of truth, but both also obscure the deeper organizational transformation of instructional labor that accompanies platform adoption.

Teachers are not merely end users of personalized learning systems. They are interpreters, mediators, and local governors of those systems. In classrooms, algorithmic recommendations do not act directly on students; they are incorporated into pedagogical routines shaped by teacher judgment, curriculum plans, disciplinary norms, and relational knowledge. This means that the educational value of personalization depends significantly on whether teachers

can make sense of system outputs in context. A dashboard that indicates a skill deficit has limited value unless the teacher understands how that deficit relates to current instruction, student motivation, language proficiency, and classroom dynamics. The more decontextualized the system output, the more interpretive labor it imposes on educators.

This interpretive burden is often underestimated. Personalized platforms can generate an abundance of data while providing limited guidance about actionability. Teachers may receive item-level accuracy patterns, growth indicators, risk categories, or attention alerts, yet have insufficient time during the school day to convert those signals into meaningful instructional adjustments. In such cases, the system does not reduce workload so much as displace it into new forms of cognitive and organizational labor. Teachers must decide which alerts to trust, when to override automated recommendations, how to communicate platform assignments to families, and how to reconcile software pathways with district-mandated curricula. Without sustained professional development and collaborative planning structures, personalization can become another source of fragmentation.

There is also a governance issue concerning professional autonomy. When AI systems are positioned as authoritative decision aids, they can subtly shift the locus of pedagogical judgment away from teachers and toward vendors' embedded models and content structures. This does not necessarily occur through explicit coercion. It can happen through interface design, performance dashboards, district monitoring practices, or the apparent objectivity of data visualizations. Teachers may feel pressure to conform to recommended pathways even when those recommendations conflict with their contextual understanding of students. Over time, such dynamics can produce what might be called platform pedagogies, in which the software's internal logic becomes the practical curriculum regardless of local educational intentions.

Yet the relationship between AI and teacher professionalism need not be adversarial. Personalized learning systems can strengthen teaching when they are designed as decision-support infrastructures rather than pedagogical replacements. This requires systems that are inspectable, overrideable, curriculum-aligned, and responsive to teacher feedback. It also requires recognizing that the most consequential educational work often involves motivational support, classroom culture, culturally responsive explanation, and relational attunement, all of which remain difficult to encode within adaptive systems. The challenge is therefore not how to replace teacher judgment with algorithmic optimization, but how to design collaborative human-AI arrangements that preserve teacher agency while extending instructional responsiveness.

The organizational implications are especially significant in districts facing staffing shortages. In such contexts, leaders may view AI-driven personalization as a way to compensate for limited specialist support, large class sizes, or uneven content expertise among novice teachers. There may be genuine short-term benefits in using adaptive systems to provide structured practice or targeted remediation. However, there is a long-term risk that reliance on technology becomes a substitute for investment in human capacity. If districts deploy personalized systems primarily because they cannot recruit reading specialists, mathematics

coaches, bilingual staff, or special educators, the technology may stabilize scarcity rather than solve it. Equity-oriented use of AI must therefore be distinguished from austerity-oriented use. The former augments professional capacity; the latter normalizes underinvestment.

Teacher education and ongoing professional learning are central to this distinction. Educators need more than technical training on platform features. They need conceptual grounding in model limitations, data interpretation, fairness concerns, and the pedagogical implications of adaptive sequencing. They also need opportunities to examine how personalization interacts with classroom discourse, student identity, and assessment practices. Where such capacity-building exists, teachers are more likely to use systems critically and productively. Where it does not, platforms are often either underused or over-trusted, neither of which supports robust educational outcomes.

The politics of teacher trust also matters. In many U.S. districts, technological reform initiatives are layered onto environments already shaped by accountability pressure, evaluation systems, and rapid policy turnover. Under such conditions, educators may reasonably interpret AI-driven systems as additional surveillance tools rather than instructional supports. If teacher-facing analytics are also used by administrators for performance monitoring, the boundary between pedagogical support and managerial oversight becomes blurred. This can erode trust and reduce the likelihood that teachers will engage openly with system data. Sustainable deployment therefore requires institutional arrangements that separate formative support functions from punitive accountability mechanisms wherever possible.

Ultimately, personalized learning systems should be judged partly by the kinds of professional roles they create. Systems that reduce teachers to platform supervisors narrow the intellectual and relational dimensions of teaching. Systems that provide meaningful insight, flexible support, and time-saving assistance without displacing professional reasoning can strengthen instructional practice. The difference lies not only in interface design but in governance choices about implementation, evaluation, and labor.

## **6. Data Governance, Privacy, and Public Accountability**

AI-driven personalized learning in K-12 settings depends on continuous data extraction and interpretation. This makes data governance a central concern rather than a peripheral compliance issue. Personalized systems rely on longitudinal records of student interaction, performance, pacing, error patterns, and often behavioral metadata. In some cases, they also integrate attendance, demographics, special education indicators, language status, and assessment results from student information systems or learning management systems. Such integration can increase instructional responsiveness, but it also creates complex questions about privacy, consent, purpose limitation, retention, secondary use, and institutional accountability.

Children constitute a particularly sensitive population for data-intensive technologies. Unlike adults using commercial learning platforms voluntarily, K-12 students often interact with

educational software under compulsory institutional arrangements, with limited capacity to consent meaningfully or avoid participation. Parents may receive privacy notices, but such notices rarely provide actionable clarity regarding model training, third-party access, retention policies, or inferential profiling. In public education, where the state has heightened responsibilities toward minors, this asymmetry raises serious ethical and legal concerns.

The legal framework in the United States offers only partial guidance. Federal statutes such as FERPA and COPPA establish important protections, but they were not designed for contemporary AI-enabled analytics ecosystems characterized by cloud vendors, interoperable APIs, behavioral inference, and large-scale machine learning pipelines. As a result, much of the practical governance burden falls on districts and states, which vary considerably in legal sophistication and procurement capacity. Wealthier districts may negotiate stronger data terms, conduct security reviews, and require deletion guarantees. Under-resourced districts may accept standard vendor contracts with limited leverage. This creates a paradox in which the districts serving the most vulnerable students may also face the weakest governance conditions.

Data governance is not only about preventing breaches or unauthorized disclosure. It is also about governing inference. Personalized learning systems do not simply store student data; they generate probabilistic claims about learners, such as mastery status, motivation, future risk, or recommended pathways. These inferences can have pedagogical and psychological consequences even when the underlying raw data never leave the system. A student categorized as chronically disengaged or low readiness may receive different learning opportunities than one inferred to be advanced or self-directed. If such classifications are inaccurate or insufficiently contestable, the governance problem lies not in access control but in the social effects of automated categorization.

The rise of generative AI magnifies these concerns because it expands both data inputs and output forms. Systems may use student writing to generate feedback, summarize progress, or create adaptive prompts. They may log conversational interactions and retain them for model improvement or quality assurance. Unless carefully constrained, these practices can produce opaque data flows that are difficult for districts to monitor. International guidance has emphasized that AI in education should be governed in ways that protect privacy, support human oversight, and ensure that technological capacity does not outpace ethical and policy safeguards.

Public accountability requires several concrete governance capacities. First, districts need procurement frameworks that distinguish instructional value from data extraction value. Second, they need technical and legal mechanisms to audit what data are collected, how models are trained or updated, and what downstream decisions are influenced by system outputs. Third, they need communication practices that make system use legible to teachers, families, and governing boards. Fourth, they need incident response procedures for model failures, biased outputs, or security problems. Fifth, they need governance structures that incorporate stakeholder voice, including that of teachers, parents, and where appropriate, students themselves.

Transparency alone is insufficient, but opacity is corrosive. Many AI systems rely on proprietary claims that make meaningful external scrutiny difficult. Vendors may disclose broad functional descriptions while withholding model documentation, feature importance, or validation practices on commercial grounds. This creates a public-sector governance dilemma: school systems are asked to trust systems they cannot fully inspect. In high-stakes domains involving children, such arrangements are difficult to justify without stronger public-interest standards. At minimum, districts and states should require documentation regarding data sources, intended use, known limitations, subgroup performance, update cycles, and human oversight expectations.

There is also a need to distinguish between educationally justified data use and opportunistic data accumulation. Not every technically available signal should be used for personalization. The collection of affective indicators, webcam-derived attention measures, keystroke dynamics, or inferred emotional states is especially contentious, given questionable validity and substantial privacy risk. In K-12 settings, governance should be guided by data minimization and pedagogical relevance. Systems should collect what is necessary to support legitimate educational functions, not whatever data might improve model performance marginally or expand future commercial uses.

For public education systems, trust is itself an infrastructure. Once families or educators come to view personalized learning platforms as opaque extraction tools, implementation legitimacy erodes. Robust data governance is therefore not an external constraint on innovation but a condition of sustainable adoption. Systems that cannot justify their data practices in public-interest terms are poorly suited for educational institutions charged with serving children fairly and transparently.

## **7. Evidence, Evaluation, and the Problem of Scale**

The research literature on personalized and AI-supported learning presents a mixed but instructive picture. Across intelligent tutoring systems, adaptive practice platforms, learning analytics tools, and decision-support systems, there is evidence that well-designed technologies can improve selected learning outcomes, particularly in constrained domains and under supportive implementation conditions. However, the distribution of effects is highly uneven across grade levels, subjects, student populations, and institutional contexts. Moreover, much of the strongest evidence comes from settings where the technology is tightly aligned with instructional goals, teachers receive substantial support, and outcome measures are closely matched to the platform's design. This should not be dismissed, but it should temper generalization.

The gap between efficacy and scale is especially important in K-12 education. A system may show gains in a pilot involving motivated teachers, stable technical support, and focused use cases, yet yield weaker or no effects when districtwide deployment introduces heterogeneous classrooms, varying curricula, inconsistent implementation, and competing digital tools. This is a recurring challenge in educational innovation. The problem is not that early findings are necessarily wrong, but that institutional complexity increases dramatically with scale.

Learning analytics research at the PK-12 level has highlighted both opportunities and persistent design and implementation challenges, underscoring that adoption outcomes depend on more than technical capability alone.

A major evaluation limitation concerns what outcomes are measured. Many studies focus on proximal indicators such as item performance, short-term test gains, time on task, or platform completion. These metrics are useful but incomplete. System-level evaluation should also examine transfer of learning, retention, teacher workload, student motivation, subgroup effects, curricular coherence, implementation burden, and the extent to which platform use alters instructional patterns. In equity-oriented settings, it is essential to assess not only average gains but also differential access, differential benefit, and whether the system narrows or widens opportunity across student groups.

Causal evaluation is further complicated by the recursive nature of personalized systems. Adaptive platforms change the learning environment based on learner behavior, meaning the treatment itself evolves over time. Traditional evaluation designs may struggle to capture these dynamic interactions. In addition, model updates, content library changes, and vendor-side modifications can alter the intervention during the study period. This creates a moving-target problem in which the system being evaluated is not static. For public education decision-making, this implies that one-time evidence reports are insufficient. Ongoing monitoring and local validation are necessary.

There is also a representational problem in the evidence base. Technologies are often evaluated with populations that do not fully reflect the diversity of U.S. public education, or subgroup analyses are underpowered. As a result, claims about broad equity benefits may rest on evidence that is thin for students with disabilities, emergent bilingual learners, rural populations, Native students, or highly mobile students. Similarly, systems that perform well in mathematics may not generalize to literacy-rich or discussion-based subject areas. Districts therefore need evaluative frameworks that distinguish domain-specific strengths from inflated general-purpose claims.

The What Works Clearinghouse and related U.S. evidence-building efforts provide an important institutional model for disciplined scrutiny, emphasizing transparency in study quality and intervention review. At the same time, the pace of AI platform iteration now exceeds the tempo of conventional education evidence cycles. This mismatch creates pressure to adopt tools based on vendor claims, selective pilot findings, or peer district enthusiasm rather than mature evidence. It suggests the need for hybrid evaluation strategies that combine rigorous external research with district-level continuous improvement methods, implementation analytics, and stakeholder feedback.

A useful way to frame the evidence question is to distinguish between three levels of claim. The first level concerns technical functionality: can the system accurately estimate mastery, recommend content, or generate coherent feedback? The second concerns instructional effectiveness: does using the system improve learning outcomes under particular conditions? The third concerns public-system value: does the system improve educational opportunity,

teacher capacity, and institutional functioning in ways that justify sustained adoption? Many platforms can demonstrate the first, some the second, and far fewer the third. Public procurement and policy decisions often collapse these levels, treating technical capability as a proxy for public value. That shortcut is especially risky in K-12 settings.

The strongest future evaluation agenda will therefore be socio-technical. It will examine how platform design interacts with curriculum, labor, infrastructure, and governance. It will require subgroup validity reporting, implementation-cost analysis, and public transparency regarding where evidence is strong, weak, or absent. It will also resist the assumption that faster iteration is inherently better in educational contexts where stability, trust, and developmental appropriateness matter. Schools are not beta-testing environments in the commercial sense; they are public institutions responsible for children's learning and well-being.

## **8. Policy, Procurement, and the Political Economy of Platformization**

The spread of AI-driven personalized learning in U.S. K-12 education is shaped as much by political economy as by pedagogy. Vendors, philanthropic actors, state agencies, district leaders, and federal guidance all influence which systems are adopted, how they are framed, and what forms of evidence or accountability are treated as sufficient. Personalized learning platforms do not enter neutral institutional spaces. They enter markets structured by budget constraints, administrative overload, uneven expertise, and recurring reform pressure. As a result, adoption decisions often reflect a combination of educational aspiration, vendor persuasion, compliance needs, and operational pragmatism.

Procurement is a pivotal but underanalyzed site of governance. Districts frequently acquire platforms through processes that prioritize price, functionality claims, compatibility, and implementation timelines over long-term questions of transparency, auditability, and educational alignment. Because educational technology procurement is decentralized, thousands of districts negotiate with vendors from highly unequal positions of leverage. Larger districts may have dedicated legal counsel, data privacy officers, and instructional technology teams. Smaller districts may rely on a handful of administrators managing multiple portfolios. This asymmetry affects not only contract quality but the capacity to define terms under which AI systems can be meaningfully governed.

Platformization intensifies dependency. As vendors integrate assessment, content delivery, analytics, communications, and administrative features into unified ecosystems, districts may gain convenience at the cost of flexibility. Once student data, teacher workflows, curriculum materials, and parent communication channels are routed through a single ecosystem, switching costs rise substantially. The district becomes dependent not merely on software but on an evolving vendor logic that may reshape priorities over time. This is especially consequential when vendors introduce AI features incrementally into existing platforms. Districts may find themselves using new predictive or generative tools without having made an explicit governance decision about those functions.

The policy challenge is therefore not only to regulate AI as a category but to govern the infrastructures through which AI becomes normalized in schooling. This includes interoperability standards, documentation requirements, audit rights, and procurement templates that prioritize public-interest safeguards. States can play a stronger role by creating approved-vendor frameworks tied to transparency and equity criteria, supporting shared technical expertise across districts, and funding independent evaluation capacity. Federal guidance can also help clarify norms for human oversight, student data protection, and the use of AI in instructional rather than purely administrative domains.

There is an additional risk that personalization aligns too comfortably with performance-management cultures in education. Systems that generate dashboards, risk scores, and growth metrics can be attractive not only for instructional reasons but because they fit accountability logics centered on measurable outputs. In such environments, AI may be valued less for enriching learning than for making performance more legible to administrators and policymakers. This can distort adoption incentives, favoring systems that produce visible metrics over those that support deeper but less easily quantified educational work.

Austerity politics further complicate matters. Districts facing fiscal stress may be drawn to personalized platforms as scalable substitutes for human-intensive supports. Vendors may explicitly or implicitly market AI as a way to deliver more with fewer staff, especially in intervention, tutoring, or feedback-heavy domains. While some automation can be beneficial, cost-saving frames can erode the public mission of schooling by encouraging technological replacement where relational and professional capacity are indispensable. A system-level policy approach must therefore ask not only what AI can do, but what public education ought not delegate to AI, especially where judgment, care, and democratic accountability are central.

International and national guidance increasingly supports a more cautious approach. UNESCO has emphasized governance, human agency, and policy readiness in the educational use of AI, while U.S. evidence-building bodies continue to stress transparency, evidence, and implementation realism. The challenge is translating these high-level principles into operational policy instruments. This likely requires standard contractual clauses for AI documentation, mandated impact assessments for sensitive educational uses, public reporting on district AI deployments, and stronger support for collaborative procurement consortia that reduce informational asymmetry between vendors and districts.

Political economy also influences whose educational values are encoded in personalization systems. Content sequencing, mastery thresholds, feedback styles, and behavioral assumptions are not neutral. They reflect curricular priorities, cultural expectations, and theories of learning. When such design choices are made predominantly by private firms with limited democratic accountability, public education risks outsourcing not only technology provision but elements of pedagogical governance. This is one of the deepest stakes in the expansion of AI-driven personalized learning. The issue is not merely software adoption, but the gradual reconfiguration of who defines the practical architecture of teaching and learning.

## 9. Robustness, Resilience, and Sustainability in Diverse School Systems

A system-level assessment of AI-driven personalized learning must also consider robustness and sustainability. Educational technologies are often introduced through pilot enthusiasm and short-term funding, but long-term value depends on whether systems remain functional, trusted, and educationally coherent under ordinary and adverse conditions. Robustness in this context includes technical reliability, pedagogical stability, institutional adaptability, and resilience to environmental stressors such as staffing turnover, cyber incidents, policy shifts, and infrastructure disruptions.

Technical robustness begins with uptime, latency, security, and integration quality, but it extends further. Personalized systems should behave predictably across different devices, bandwidth conditions, and user populations. They should fail gracefully when data are incomplete or connectivity is unstable. They should not generate volatile recommendations in response to minor fluctuations in input behavior. In K-12 settings, where teachers operate under tight time constraints, fragile systems quickly lose legitimacy. A platform that requires extensive troubleshooting or produces inconsistent outputs imposes costs that often outweigh theoretical instructional benefits.

Pedagogical robustness concerns the stability of educational meaning. Systems should not merely function technically; they should support coherent learning over time. This requires alignment with standards, curriculum pacing, assessment structures, and developmental appropriateness. It also requires avoiding excessive dependence on narrow metrics that can be gamed or misinterpreted. A robust system should produce benefits that persist even when teachers use it in varied but reasonable ways, rather than requiring idealized implementation conditions.

Institutional resilience is equally important. Schools are high-turnover environments. Principals change, technology coordinators leave, teacher teams reconfigure, and district priorities shift. Personalized learning systems that depend on a few local champions may show early promise but prove unsustainable when personnel change. Durable adoption requires documentation, onboarding processes, role clarity, and support structures that outlast individual enthusiasts. It also requires local capacity to manage vendor relationships, evaluate updates, and revise usage policies as technologies evolve.

Sustainability includes financial sustainability. Many personalized learning initiatives begin with federal stimulus funds, philanthropic grants, or short-term innovation budgets. Once those funds end, districts must decide whether subscription costs, implementation support, device replacement, and professional development remain affordable. In under-resourced systems, this can lead to abrupt discontinuation or reduced support quality. Sustainable deployment therefore requires realistic total-cost planning, including hidden costs related to integration, training, accessibility compliance, and governance oversight.

Environmental sustainability, while less frequently discussed, is becoming relevant as AI systems grow more computationally intensive. Large-scale generative models and

cloud-based analytics consume substantial energy and rely on data center infrastructures far removed from school decision-making. Although individual districts have limited influence over these broader infrastructures, public procurement can still incorporate questions about computational efficiency, model proportionality, and responsible use. Not every educational task requires the most resource-intensive form of AI.

Resilience also has a normative dimension. Personalized learning systems should be resilient to misuse, including overreliance, inappropriate delegation, or coercive accountability applications. A system originally designed to support formative instruction can be repurposed for high-stakes tracking or teacher monitoring if governance boundaries are weak. Designing for resilience therefore involves not only technical safeguards but institutional guardrails regarding acceptable use.

The COVID-19 pandemic made visible how educational resilience depends on digital infrastructure, but it also exposed the limitations of assuming technology alone can stabilize learning under crisis. Device distribution and online platforms were necessary, yet learning continuity depended equally on family support, teacher adaptability, local trust, and public coordination. This lesson applies directly to AI-driven personalization. Technology can strengthen educational resilience when embedded in supportive systems, but it cannot substitute for the institutional and relational capacities that make schooling viable under stress.

Robust and sustainable AI adoption in K-12 education thus requires a move away from novelty-driven implementation and toward infrastructure thinking. Districts and policymakers must ask whether systems are maintainable, governable, interoperable, and pedagogically durable over multiple years. Without that shift, personalized learning risks becoming another cycle of uneven adoption, promising pilots, and quiet abandonment.

## **10. Toward a Public-Interest Framework for AI-Driven Personalized Learning**

If AI-driven personalized learning is to contribute meaningfully to educational equity and outcomes in the United States, it must be governed through a public-interest framework that recognizes the distinct institutional character of K-12 schooling. Such a framework begins by rejecting two inadequate positions: first, the technologically deterministic view that better models will automatically produce better schooling; and second, the purely precautionary view that AI is too risky to play any constructive educational role. The appropriate path lies between these extremes, grounded in democratic governance, pedagogical integrity, and differentiated institutional capacity.

A public-interest framework would treat personalized learning systems as accountable socio-technical infrastructures. This means that system design and adoption should be evaluated across multiple dimensions simultaneously: educational validity, equity impact, teacher usability, privacy protection, accessibility, fiscal sustainability, and contestability. It would also require that districts and states develop the institutional means to govern these dimensions rather than outsourcing judgment to vendors. In practice, this implies stronger

procurement standards, shared technical expertise, impact assessments for sensitive uses, and transparent documentation obligations.

Curriculum alignment should be a first-order principle. Personalization is educationally valuable only when adaptive pathways remain anchored to coherent learning goals and intellectually substantive content. Systems that optimize engagement or completion without robust curricular grounding can distract from rather than support meaningful learning. Public-interest governance should therefore require evidence that personalization logic aligns with standards, disciplinary practices, and developmental expectations, not merely that users spend more time on the platform.

Human oversight should also be built into both interface and policy design. Teachers must be able to inspect, interpret, and when necessary override system recommendations. Families should have understandable explanations of system purpose and limitations. Administrators should avoid converting formative data into punitive accountability metrics without clear justification and public deliberation. Human oversight does not mean symbolic involvement; it means preserving meaningful discretion and review pathways where automated outputs shape educational experiences.

Equity auditing must become routine. Districts and vendors should examine subgroup access, usage patterns, recommendation differences, outcome gaps, and error distributions across race, language status, disability status, socioeconomic background, and geography. Such auditing should not be reduced to a technical fairness metric. It should include qualitative investigation into how students and teachers experience the system, whether pathways differ in educational richness, and whether personalization expands or narrows access to opportunity. Ethical reviews of AI in education have repeatedly identified privacy, bias, transparency, and informed consent as core concerns, reinforcing the need for systematic rather than ad hoc oversight.

Capacity building is another essential component. A public-interest framework assumes that equitable AI adoption requires investment in people, not only products. Teachers need time and learning opportunities to use systems critically. District leaders need support in procurement, evaluation, and governance. States can help by offering shared review resources, model contract language, and technical assistance. Without such capacity supports, governance ideals remain aspirational and unequal implementation persists.

There is also a need for proportionality. Not every educational challenge should be addressed with AI, and not every task benefits from maximal personalization. In some domains, simpler tools may be more transparent, reliable, and pedagogically appropriate than complex predictive systems. Public-interest governance should therefore encourage fit-for-purpose design, matching technological complexity to educational need rather than rewarding novelty for its own sake.

Finally, the framework must preserve the public purposes of schooling. Education in a democratic society involves common experiences, civic formation, social interaction, and the

cultivation of judgment, not just individualized skill acquisition. Personalized learning systems can be useful components of this mission, but they should remain components, not governing logics. When personalization is treated as the master principle of schooling, the risk is that education becomes a sequence of privately optimized pathways rather than a public practice of shared intellectual and civic development.

The future of AI-driven personalized learning in the United States will be determined less by computational breakthroughs than by whether public institutions can establish the norms, capacities, and safeguards needed to align these systems with educational justice. That is the central systems problem, and it is ultimately a governance problem.

## **11. Conclusion**

AI-driven personalized learning systems occupy an increasingly influential place in the imagination and practice of U.S. K-12 education. Their appeal rests on a powerful and partially justified claim: that digital systems can help schools respond more effectively to learner variability, support teachers with better information, and improve outcomes in contexts where one-size-fits-all instruction has long proved inadequate. Yet the significance of these systems cannot be assessed through technological capability alone. Their real educational value emerges only through the socio-technical arrangements that shape how they are designed, integrated, governed, and sustained in public institutions.

This paper has argued that the central question is not whether AI can personalize learning in a narrow technical sense, but whether personalized learning systems can serve the public goals of education under conditions of structural inequality, institutional diversity, and democratic accountability. That question requires attention to architecture, curriculum alignment, professional labor, infrastructure distribution, data governance, evidence quality, procurement incentives, and long-term sustainability. Across each of these domains, the same conclusion emerges: AI-driven personalization can contribute to equity and improved outcomes, but only conditionally. It is most promising when it augments teacher judgment, expands access to high-quality instruction, remains transparent and contestable, and is supported by robust institutional capacity. It is most dangerous when it functions as a low-visibility sorting mechanism, a substitute for human investment, or an opaque infrastructure of behavioral inference embedded in unequal school systems.

For the United States, the stakes are especially high because K-12 education is both decentralized and deeply unequal. In such a context, leaving AI-driven personalization to market diffusion and local improvisation is likely to reproduce familiar patterns of uneven benefit, weak oversight, and reform churn. A more responsible path requires treating personalized learning systems as public-interest infrastructures subject to deliberate governance. This includes stronger procurement standards, equity auditing, interoperability requirements, teacher-centered design, privacy protections, and sustained investment in implementation capacity. It also requires a broader understanding of educational success, one that values common learning, inclusion, trust, and civic purpose alongside individualized academic progression.

AI will almost certainly remain part of the future educational landscape. The decisive issue is whether it will deepen platform dependency and stratification, or whether it can be shaped into a more accountable and equitable support for public education. Achieving the latter outcome will depend not on the inevitability of technological progress, but on the quality of collective institutional judgment brought to bear on its development and use.

## References

1. Akgun, S., & Greenhow, C. (2022). Artificial intelligence in education: Addressing ethical challenges in K-12 settings. *AI and Ethics*, 2(3), 431–440.
2. Baker, R. S. (2016). Stupid tutoring systems, intelligent humans. *International Journal of Artificial Intelligence in Education*, 26(2), 600–614.
3. Benjamin, R. (2019). Race after technology: Abolitionist tools for the new Jim code. *Polity*.
4. Bienkowski, M., Feng, M., & Means, B. (2012). Enhancing teaching and learning through educational data mining and learning analytics: An issue brief. U.S. Department of Education, Office of Educational Technology.
5. Bowers, A. J., & Zhou, X. (2019). Receiver operating characteristic predictions of student success in school: A systematic review and future directions. *Review of Educational Research*, 89(5), 689–732.
6. Coburn, C. E., Hill, H. C., & Spillane, J. P. (2016). Alignment and accountability in policy design and implementation. In D. H. Gitomer & C. A. Bell (Eds.), *Handbook of research on teaching* (5th ed., pp. 243–273). American Educational Research Association.
7. Darling-Hammond, L., Flook, L., Cook-Harvey, C., Barron, B., & Osher, D. (2020). Implications for educational practice of the science of learning and development. *Applied Developmental Science*, 24(2), 97–140.
8. Dee, T. S. (2011). *The importance of skills: Hard and soft*. Brookings Institution.
9. Dede, C. (2019). The role of AI in promoting personalized learning. *Educational Technology*, 59(4), 3–6.
10. Fletcher, J. D. (2018). *Teaching machines: Learning from the intersection of education and technology*. Johns Hopkins University Press.
11. Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial intelligence in education: Promises and implications for teaching and learning*. Center for Curriculum Redesign.
12. Holmes, W., Miao, F., & Akwasi, A. (2023). *Guidance for generative AI in education and research*. UNESCO.
13. Itani, A., Holroyd-Leduc, J., & Verma, A. A. (2025). Ethical considerations for teaching with artificial intelligence: A scoping review. *Humanities and Social Sciences Communications*, 12, Article 1042.
14. Khalil, M., Prinsloo, P., & Slade, S. (2024). Learning analytics in school education: Opportunities, limitations, and governance challenges. *British Journal of Educational Technology*, 55(2), 487–506.
15. Klimova, B., Pikhart, M., Kacetl, J., & Cierniak-Emerych, A. (2023). Ethical issues of the use of AI-driven mobile applications for education. *Education and Information Technologies*, 28(1), 613–632.

16. Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8, 30–43.
17. Means, B., Toyama, Y., Murphy, R., Bakia, M., & Jones, K. (2010). *Evaluation of evidence-based practices in online learning: A meta-analysis and review of online learning studies*. U.S. Department of Education.
18. Miao, F., Holmes, W., Huang, R., & Zhang, H. (2021). *AI and education: Guidance for policy-makers*. UNESCO.
19. Pane, J. F., Steiner, E. D., Baird, M. D., & Hamilton, L. S. (2015). *Continued progress: Promising evidence on personalized learning*. RAND Corporation.
20. Pane, J. F., Steiner, E. D., Baird, M. D., Hamilton, L. S., & Pane, J. D. (2017). *Informing progress: Insights on personalized learning implementation and effects*. RAND Corporation.
21. Paolucci, C., Börner, D., & El-Haddad, R. (2024). A review of learning analytics opportunities and challenges for K-12 contexts: A qualitative metasynthesis. *Education and Information Technologies*, 29(4), 4517–4548.
22. Reich, J. (2020). *Failure to disrupt: Why technology alone can't transform education*. Harvard University Press.
23. Roschelle, J., Feng, M., Murphy, R. F., & Mason, C. A. (2016). Online mathematics homework increases student achievement. *AERA Open*, 2(4), 1–12.
24. Selwyn, N. (2019). *Should robots replace teachers? AI and the future of education*. Polity.
25. Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist*, 57(10), 1510–1529.
26. UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO.
27. U.S. Department of Education, Office of Educational Technology. (2024). *National educational technology plan 2024: A call to action for closing the digital access, design, and use divides*. U.S. Department of Education.
28. van de Oudeweetering, K., & Decuyper, M. (2023). Understanding openness and opacity in AI in education. *Learning, Media and Technology*, 48(2), 175–189.
29. Williamson, B. (2017). *Big data in education: The digital future of learning, policy and practice*. Sage.
30. Williamson, B., Bayne, S., & Shay, S. (2020). The datafication of teaching in higher education: Critical issues and perspectives. *Teaching in Higher Education*, 25(4), 351–365.
31. Wise, A. F., & Jung, Y. (2019). Teaching with analytics: Towards a situated model of instructional decision-making. *Journal of Learning Analytics*, 6(2), 53–69.
32. Woolf, B. P. (2010). *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning*. Morgan Kaufmann.
33. Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education—Where are the educators? *International Journal of Educational Technology in Higher Education*, 16, Article 39.

34. Zhao, Y. (2018). *What works may hurt: Side effects in education*. Teachers College Press.
35. Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory Into Practice*, 41(2), 64–70.