

Education in the Age of AI: Adaptive Systems, Assessment, and Responsible Governance

Gabrijela Dimić¹ , Žaklina Spalević^{2*} , Milutin Nešić¹ , Ratko Ivković³ , David Sotiroski⁴ ,
Čedomir Vasić³  & Dragan Vučković⁵ 

¹Academy of Technical and Art Applied Studies, Belgrade, Republic of Serbia,
e-mail: gabriela.dimic@viser.edu.rs, milutin.nesic@viser.edu.rs

²Singidunum University Belgrade, Faculty of Tourism and Hospitality Management, Republic of Serbia,
e-mail: zspalevic@singidunum.ac.rs

³MB University, Faculty of Business and Law, Belgrade, Serbia, email: rivkovic@mbuniverzitet.edu.rs, cedomir.vasic@ppf.edu.rs

⁴Qinshift, Skopje, Republic of North Macedonia, e-mail: sotiroski_d@yahoo.com

⁵Academy of Applied Studies Polytechnic, Beograd, Serbia, e-mail: dvuckovic@politehnika.edu.rs

Abstract: This paper examines how Artificial Intelligence in Education (AIED) is reshaping teaching and learning, drawing on a systematic literature review alongside policy analysis to explore practical applications, the theories behind them, and their governance consequences. Adopting the EU AI Act's risk-based lens, we investigate the ways in which regulatory demands—ranging from transparency and data stewardship to human oversight and provider accountability—influence how AIED tools are built and taken up in practice. We group current uses into four areas—adaptive learning, intelligent assessment, learner profiling, and emerging tools—and read them through the prism of well-known learning theories such as constructivism. The analysis underscores that while these technologies hold real promise, several prominent use cases—automated grading and learner profiling, for instance—fall squarely within the EU AI Act's higher-obligation categories, which means equity, explainability, and genuine human control are not optional but essential for public trust. On the basis of these findings, we put forward concrete, compliance-oriented recommendations aimed at helping educators, institutions, and policymakers deploy AI responsibly across varied educational settings.

Keywords: *Artificial Intelligence, Education, Adaptive Learning, EU Artificial Intelligence Act, Transparency, Policy and Governance.*


Introduction

In the last ten years or so, advances in artificial intelligence (AI) have profoundly changed the way educators and researchers think about what educational technology can achieve. Whether through personalized tutoring platforms or systems that grade student work automatically, AI now sits at the heart of education's digital shift. Yet this growing ubiquity also brings thorny questions about ethical boundaries, built-in biases, learner privacy, and fairness (U.S. Department of Education, 2023). The U.S. Department of Education (2023) characterizes AI as a form of pattern-driven automation: machines go beyond mere data collection to recognize regularities and act on them within educational settings.

When AI takes on tasks that were once the sole province of educators—deciding what a student should study next, or flagging someone as at risk—it introduces responsibilities that schools have not previously had to manage. What makes AI significant in education is not just what it can do technically, but how it is beginning to reshape the very acts of learning, teaching, and assessment. “AI in education can only grow at the speed of trust,” the report highlights (U.S. Department of Education, 2023).

For that reason, any deployment needs to be anchored in transparency and shared accountability,

*Corresponding author: zspalevic@singidunum.ac.rs



with teachers, students, families, and policymakers all having a seat at the table. We ground our discussion in the European Union's AI Act and its risk-based logic, which treats applications like automated assessment and learner profiling as areas where stronger safeguards—greater transparency, tighter data governance, and real human oversight—are non-negotiable.

Our review of the literature identifies four broad clusters of AI use in education: adaptive learning and personalized tutoring; intelligent assessment and management; learner profiling and outcome prediction; and novel AI-driven products (Wang et al., 2024).

The bulk of existing research concentrates on higher education, leaning on theoretical lenses like constructivism, learning-styles frameworks, and cognitive models (Wang et al., 2024). Scholarly attention to AI in education surged after 2017, a trend closely tied to the rise of generative AI tools and the rapid pivot to remote learning prompted by COVID-19. By 2022, researchers had scrutinized over 2,200 publications in an effort to chart the intellectual terrain of this field—work that reveals an increasingly multidisciplinary landscape bridging computer science, education studies, and information-systems research (Wang et al., 2024). At the same time, the U.S. Department of Education (2023) flags several pressing dangers: biased training data, erosion of student privacy, automation that outpaces human oversight, and AI models built without adequate theoretical foundations (U.S. Department of Education, 2023).

If left unchecked, AI risks deepening the inequalities already present in education, which is precisely why current guidelines stress keeping humans in the loop—making sure that people, not algorithms, hold the final say in decisions that affect learners (U.S. Department of Education, 2023). Among the most exciting possibilities is formative assessment: AI can deliver feedback as students work, giving them guidance that would be impossible for a single teacher to provide at scale.

Automated essay scoring is a good case in point—it can lighten a teacher's workload considerably, yet if deployed carelessly it can also undermine the pedagogical goals it is supposed to serve. There is also a growing consensus that AI should not merely be something students encounter passively; they need to understand how it works, where it falls short, and how to engage with it critically (U.S. Department of Education, 2023). This makes it essential that educational AI systems be transparent in their reasoning, open to being overridden, and subject to audit—because a wrong algorithmic decision in a school setting can leave marks that last well beyond a single semester (U.S. Department of Education, 2023).

Getting there will require researchers to build AI tools that are rooted in educational theory and adapted to local contexts, while educators, policymakers, technology firms, and academics work together on shared standards that keep AI use safe, fair, and genuinely helpful. Put simply, education needs its own regulatory guardrails—ones that protect both students and teachers without stifling the innovation that makes AI worth pursuing in the first place (U.S. Department of Education, 2023).

To situate this analysis within current regulatory developments, we briefly outline the EU Artificial Intelligence Act (Regulation (EU) 2024/1689) (European Union, Regulation 1689/2024). The Act establishes a risk-based framework, prohibiting certain practices—including the use of systems to infer emotions in educational institutions—and classifying several educational uses as high-risk (e.g., AI for admissions/access decisions, evaluation of learning outcomes, placement/streaming, and test proctoring/monitoring).

High-risk systems must implement risk management, data and data-governance measures, transparency to deployers, human oversight, and accuracy/robustness/cybersecurity; the Act also introduces transparency duties for AI interactions and synthetic content.

For AIED, this implies that learner profiling, automated assessment, and exam proctoring typically fall under high-risk obligations and must be designed and deployed with compliance-by-design and meaningful human control.

This paper integrates insights from a systematic review of over 2,200 peer-reviewed publications on AI in education, drawing primarily from indexed databases (e.g., Scopus, Web of Science, ERIC). Inclusion criteria focused on works published post-2017 addressing empirical applications, theoretical models, and governance frameworks. Thematic coding enabled the extraction of four main domains of application. The review is complemented by an interpretive policy analysis grounded in legal and ethical instruments (e.g., EU AI Act, UNESCO Guidance).

Given the rapid pace of technological advancement and the growing adoption of AI tools in educational contexts, this manuscript integrates theoretical foundations, current applications, and policy recommendations aimed at ensuring the responsible, equitable, and pedagogically grounded use of AI in education.

By synthesizing findings from recent research and international policy guidelines, it offers a solid

foundation for future decision-making in educational technology and serves as a resource for teachers, researchers, and policymakers seeking to integrate AI meaningfully, ethically, and effectively into practice.

Foundations and conceptual framing of AI in education

AI is now woven into almost every corner of digital life, yet its importance for education runs deeper than hardware and code—it raises fundamental questions about how we construct knowledge, how we ought to teach, and how we can fairly judge what students have learned.

This section lays out the theoretical groundwork for AIED, sketches the field's intellectual architecture, and traces the main lines of thinking that inform how practitioners and scholars make sense of AI in schools and universities today.

What Counts as AI in Educational Practice?

In the scholarly record, AI is typically portrayed as a broad family of computational systems that exhibit capacities we associate with human cognition—learning from experience, reasoning through problems, perceiving their environment, and processing language. Alongside these conceptual descriptions sit more operational definitions stressing that such systems pursue predefined objectives either through rule-based logic or by learning patterns from data (U.S. Department of Education, 2023), (IEEE-USA Board of Directors, 2017). For educators, these are not merely academic distinctions: AI now routinely steps into roles that teachers used to fill on their own—selecting readings, adjusting the difficulty of a test, or even assigning marks (Friedman et al., 2021).

Narrow vs. General Intelligence: Implications for Teaching and Learning

The AI tools that actually show up in classrooms today are overwhelmingly “narrow”—each one built for a specific job, whether that means an expert system that scores quizzes, a neural network that recommends reading material, or a learning-analytics agent that tracks how students are progressing (Gartner, 2025). The idea of “general” AI—a system that could handle any intellectual task a human can—opens up even weightier ethical territory for education, particularly around empathy and care. Today's AI cannot forge the emotional connections with students that lie at the heart of good, relationship-driven teaching (The White House, 2022).

Models, Data, and Decisions: Why the Architecture Matters

At the technical core of every AI application is a model: a mathematical abstraction that takes in data and generates outputs—a usage of the word “model” quite different from how educators speak of pedagogical or institutional models (Ivković, 2025). Large language models like GPT embody this principle, working by predicting each successive word (or token) within enormous, high-dimensional parameter spaces. A key open question for the field is whether these statistical constructs capture enough of what actually matters in real classrooms. If a model's training data skew toward certain populations or contexts, the resulting system can discriminate in ways that hurt the very students it is meant to help. A robust framework for responsible AIED therefore requires explicit understanding of model assumptions and limitations.

Learning Theories as Design Anchors for AIED

AI-powered learning environments tend to be most effective—and most defensible—when they are rooted in well-tested learning theories. Constructivism is the framework cited most often, though researchers also draw on cognitive-load theory, item-response theory, and learning-styles accounts when designing personalized content and assessments. Intelligent Tutoring Systems are a good illustration: they diagnose individual learner needs and adjust pacing, feedback, and materials accordingly, going well beyond rote content delivery to offer support that is genuinely informed by pedagogical theory.

Augmented Intelligence: Partnering Humans and Machines

Current thinking in the field does not cast AI as a replacement for teachers but rather as a collaborator—a form of augmented intelligence that amplifies what human educators can do. In day-to-day terms,

this looks like handing off repetitive chores—sending reminders, marking low-stakes quizzes—to AI, freeing teachers to spend more time on the creative, relational, and interpretive parts of their work. Critically, the teacher must always be able to inspect, question, and overrule whatever the system produces. The overarching idea is straightforward: AI ought to function as an ally in the learning process, not as an opaque arbiter whose verdicts cannot be challenged.

Mapping AIED in practice: A four-domain typology

AI has moved well past the stage of being an occasional classroom novelty; it is now built into the routines and decision structures of education systems around the globe. Recent scholarship tends to organize these applications into four clusters: adaptive learning and individualized tutoring; intelligent assessment and classroom management; student profiling paired with predictive analytics; and newer AI-driven products that push the boundaries of what educational technology can look like. This typology is useful not only for comparing what each tool does, but also for thinking about the distinct pedagogical possibilities and trade-offs each category brings.

Personalized Adaptation: Intelligent Tutors & Hypermedia

At its core, AI-driven tutoring is about meeting each student where they are—adjusting what they see, how fast they move, and what feedback they receive. Intelligent Tutoring Systems (ITS) function much like attentive mentors: they watch what a student does, infer what help is needed, and steer the learning path accordingly. ZOSMAT, a mathematics-focused system that pairs learning analytics with individualized navigation, is one well-documented example. Common features of such systems include diagnostic quizzes that pinpoint gaps, practice exercises calibrated to each learner's level, and reading or activity suggestions that match demonstrated competence. A complementary line of work, Adaptive Hypermedia Learning Systems (AHLS), adapts presentation and sequencing to students' learning preferences using methods like neural networks and Bayesian classifiers—advancing the ideal of “education to the learner,” with self-paced progression matched to ability and interest.

Assessment and Orchestration: AI-Enabled Evaluation & LMS

Providing personalized attention to every student is hard enough in a single classroom; at institutional scale, it becomes all but impossible without automated support—hence the growing role of AI in assessment and classroom management. Intelligent assessment platforms handle tasks like scoring, competency mapping, and ongoing formative tracking—MI Write, for instance, evaluates student essays—while specialized pronunciation trainers offer detailed, immediate feedback on spoken language. Alongside these tools, learning management systems enhanced with AI consolidate student data, take over repetitive administrative work, and flag issues that instructors might otherwise miss. Typical capabilities comprise collaborative learning support, exam and evaluation workflows, and resource management with content recommendations; AI's role within these platforms is increasingly pivotal for observing progress and orchestrating timely interventions.

Learner Trajectories: Profiling & Predictive Modeling

A third major area of AI application brings together educational data mining and learning analytics to build profiles of individual learners and forecast their academic trajectories. Algorithms—Naive Bayes classifiers, neural networks, decision trees, support vector machines, among others—sift through behavioral data to estimate which students are academically at risk and how they are likely to perform. In practice, these systems generate early alerts about potential dropouts, predict grades and completion rates, model patterns of student behavior, and help advisors guide course selection. These predictive pipelines are already used to anticipate course selection or satisfaction, enabling administrators to align programs with learner needs and to advance precision education through proactive support.

Novel Interfaces: Robots, XR, and Generative Systems

Some of the most visible innovation in AIED comes from robotics, extended reality (VR and AR), and generative AI. Educational robots—many of them chatbot-driven—take on roles as varied as co-

teacher, study buddy, emotional-support companion, and telepresence proxy for students who cannot be physically present. Published examples range from toy robots that help children practice Chinese idioms to conversational bots that facilitate book-club discussions. Virtual and augmented reality add another dimension, placing students inside immersive simulations—virtual labs, historical reconstructions—that can boost both engagement and experiential understanding. Meanwhile, generative tools such as ChatGPT open up possibilities for essay drafting, conversational tutoring, and the collaborative creation of learning materials—pushing AIED beyond its traditional focus on content delivery and assessment feedback.

Table 1 maps the principal AIED application domains onto the risk categories set out in the [EU Artificial Intelligence Act \(2024\)](#), offering a bird’s-eye view of where each type of tool sits in the regulatory landscape. Because the impact on learners and educational outcomes differs substantially from one domain to another—adaptive learning environments carry different risks than, say, generative AI chatbots—so too do the levels of regulatory scrutiny that apply. The high-risk tier is especially pertinent for automated grading, student profiling, and decisions about who gets access to educational opportunities; here the Act requires stringent safeguards—transparency, auditability, human oversight, and verified data quality. Reading the table from left to right gives practitioners a way to weigh both the pedagogical promise and the compliance burden of different AI deployments in their own institutional context.

Table 1. Mapping AIED Domains Against EU AI Act Risk Categories and Regulatory Obligations

AIED Application Domain	Examples of Tools / Functions	Theoretical Foundation	Risk Classification (EU AI Act)	Regulatory Obligations
1. Adaptive Learning and Personalized Tutoring	Intelligent Tutoring Systems (ITS), Adaptive Hypermedia Systems (e.g., ZOSMAT)	Constructivism, Cognitive Learning Theories	Limited / High Risk (depending on their impact on learner evaluation or progression)	Transparency to users- Human oversight- Risk assessment (if affecting learner access or progression)
2. Intelligent Assessment and Learning Manag.	Automated essay scoring (e.g., MI Write), AI- enhanced LMS platforms, E-proctoring tools	Assessment theory, Constructivist Pedagogy	High Risk (especially for automated grading, test proctoring, and evaluation systems)	Risk management systems- High-quality datasets- Meaningful human control- Incident reporting mechanisms
3. Learner Profiling and Predictive Modeling	Dropout prediction, performance forecasting, behavioral analytics	Learning Analytics, Data-Driven Instruction	High Risk (when influencing educational decisions through profiling or prediction)	Bias testing and documentation- Auditability- Explainability of decisions- Right to human review
4. Innovative AI Tools: Robotics, XR, Gen. AI	ChatGPT, educational robots, VR/AR for immersive learning	Constructivism, Experiential Learning	Limited / Minimal Risk (if not involved in decision-making or profiling)	Disclosure of AI-generated content- Notification of AI interaction- Integration of AI literacy in curricula

Governance and action agenda for trustworthy aied

As AI spreads through every level of education, the policy conversation has shifted from whether it should be governed to how, precisely, to go about doing so. These are no longer speculative technologies; they are tools that teachers and students encounter every day—automated grading, adaptive tutoring, predictive analytics—and they carry genuine opportunities for better learning alongside real worries about fairness, privacy, accountability, and the preservation of human agency. Building a coherent policy architecture that speaks directly to education’s needs is therefore urgent if we want AI to be deployed safely, equitably, and to genuine effect.

Human-in-the-Loop and Pedagogical Sovereignty

The starting point for any sensible policy is that AI should strengthen human judgment, not supplant it. Regulations should insist that AI systems used in schools be explainable, auditable, and capable of being overridden—so that teachers, students, and administrators can grasp what the system is recommending, check whether that recommendation makes sense, and set it aside when it does not. Preserving this

kind of control is what keeps pedagogical sovereignty intact: educators must be free to step in whenever an AI output clashes with what they know about their students or the realities of their classroom (U.S. Department of Education, 2023).

Equity by Design: Bias, Fairness, Inclusion

Equity in digital education is not a nice-to-have; it is a baseline requirement. Models trained on data that underrepresent certain groups risk replicating—and even amplifying—existing structural injustices. Policy should therefore require concrete safeguards: routine checks on data quality, independent fairness testing, and the active involvement of diverse stakeholders at every stage of design and evaluation, all aimed at preventing AI from disproportionately harming the very students who most need support (The White House, 2022).

Pedagogy-First Alignment of AI Tools

Educational AI ought to be grounded in pedagogical theory and designed to foster active, experience-based, constructivist learning—not mere drill-and-recall. Close partnerships between tool developers, classroom teachers, and educational researchers are vital if the resulting products are to address genuine learning needs across different student populations. In other words, policy needs to keep AI firmly tethered to well-founded pedagogical ideas and established learning science (European Commission, DG EAC, 2022).

Data Governance and Radical Transparency

Educational AI runs on vast quantities of sensitive information—behavioral traces, interaction logs, grades, sometimes even socioeconomic background data—which means governance frameworks must enshrine strong privacy safeguards (like FERPA compliance), along with anonymization protocols, strict access controls, and disciplined rules for data sharing. Institutions must tell students and families plainly what data are being gathered, why, and how those data feed into decisions—and obtaining informed consent while shielding sensitive personal attributes is not optional (U.S. Department of Education, 2023).

Capacity Building and AI Literacy for Educators

Ironically, the people who deal with AI in classrooms most directly—teachers—are often the ones with the least training and support for doing so effectively. Mandatory professional development and AI-literacy initiatives should give teachers the skills to use these tools with confidence, read their outputs critically, and step in when something goes wrong. Policymakers, for their part, need to back up those programs with practical resources and everyday guidance on how to weave AI responsibly into regular teaching (IEEE-USA Board of Directors, 2017).

Participatory Policy and Co-Design with Stakeholders

Governance becomes more credible when the people affected by it have a hand in shaping it. That means bringing students, parents, teachers, school leaders, and technology companies together to co-develop guidelines—ensuring that the tools deployed in schools genuinely reflect the educational and human values of the communities they serve. When communities are involved, the result is more likely to respect local languages, cultures, and priorities—which in turn makes the technology more relevant and its governance more just.

National Frameworks and EU Guidance for Education

Individual institutional policies are not enough; countries need dedicated regulatory frameworks for education that set minimum technical and ethical standards, require impact assessments across different student groups, evaluate how AI affects both staff and learners, and mandate ongoing monitoring of live systems. As a reference point, the European Commission's educator-focused Ethical guidelines on the use of AI and data in teaching and learning offer a useful model for sectoral governance (Maslej et al., 2023). Legal scholarship further underscores alignment with constitutional guarantees and international norms on digital rights and the right to education.

Weaving these normative commitments directly into the fabric of policy does more than ensure

legal compliance—it lends the whole enterprise a measure of ethical credibility that purely technical standards cannot provide.

The EU AI Act (2024) and education

This chapter interprets the [EU Artificial Intelligence Act \(2024\)](#), arguing that the Act's risk-based architecture, graduated obligations, and governance innovations provide both a regulatory compass and a practical playbook for educational actors deploying adaptive systems and assessment technologies at scale; yet, translating that framework into classroom-ready practice will require sustained alignment with international pedagogical guidance, robust cybersecurity preparedness, and institution-level capacity building.

At the core of the AI Act is a four-tier risk taxonomy—unacceptable, high, limited, and minimal/no risk—that calibrates obligations according to anticipated harms and the salience of fundamental rights in context; crucially for education, the Act explicitly positions certain uses (e.g., scoring of exams, systems that shape access to education) within the high-risk category, thereby elevating educational assessment and learner-profiling beyond mere ed-tech functionality to the domain of regulated socio-technical systems with mandatory risk management, high-quality datasets, traceability/logging, human oversight, and demonstrable levels of robustness, cybersecurity, and accuracy ([European Commission, 2025](#)), ([European Parliament, 2025](#)).

In parallel, the Act enumerates a closed list of prohibited practices, including emotion recognition in educational institutions, social scoring, and untargeted scraping for facial recognition databases, which directly constrains a swath of speculative “affective computing” and invasive monitoring tools sometimes marketed to schools; for educators and administrators, this prohibition line clarifies not only what cannot be procured or piloted, but also how to frame vendor due diligence and contract clauses *ex ante* so that “innovation” initiatives do not drift into rights-eroding experimentation ([European Commission, 2025](#)).

The Act's transparency layer further touches classroom realities: interactions with AI must be disclosed, AI-generated content must be identifiable, and deepfakes or other synthetic media intended to inform the public require clear labeling—obligations that intersect with academic integrity policies, media literacy curricula, and institutional rules on the acceptable use of generative models in teaching and assessment ([European Commission, 2025](#)), ([European Parliament, 2025](#)). For general-purpose AI (GPAI) models—the engines increasingly powering adaptive learning platforms, automated feedback tools, and content generators—the AI Act introduces targeted obligations (e.g., transparency on training data via public summaries, copyright compliance, risk assessment and mitigation for models with systemic risk), complemented by Commission guidance and a Code of Practice to smooth compliance pathways while the regime matures; these GPAI rules matter for schools precisely because they externalize upstream model accountability, making it more feasible for deployers (schools, universities, ministries) to rely on contractually enforceable representations about model safety, security, and data provenance ([European Commission, 2025](#)).

The timeline also conditions sectoral planning: the Act entered into force on 1 August 2024, with prohibitions (and AI-literacy duties) applicable from 2 February 2025, GPAI obligations applicable from 2 August 2025, full applicability slated for 2 August 2026, and extended deadlines (e.g., certain high-risk systems embedded in regulated products) out to 2 August 2027—sequencing that allows ministries and institutions to schedule gap assessments, budget cycles, and teacher professional development in step with legal milestones ([European Commission, 2025](#)), ([European Parliament, 2025](#)). Read normatively, the AI Act reframes adaptive systems from a narrow narrative of hyper-personalization and efficiency into a governed pedagogy of adaptation: personalisation remains legitimate when underwritten by explainability, contestability, and documented human oversight; dataset quality is not a technical nicety but a legal duty; and learning analytics pipelines become subject to post-market monitoring in which serious incidents (e.g., systemic misgrading, biased placement recommendations) can trigger reporting and corrective action ([European Commission, 2025](#)), ([European Parliament, 2025](#)).

For assessment, the implications are even more direct: automated scoring, proctoring, and predictive risk flagging now sit within a compliance perimeter where model-risk analysis, bias testing, robustness to adversarial inputs, and audit-ready documentation are not optional; in practice, this encourages assessment designs that pair algorithmic scoring with calibrated human adjudication (“human-in-the-loop”),

set clear escalation protocols for edge cases, and publish intelligible grading rationales to preserve due process for students ([European Commission, 2025](#)), ([European Parliament, 2025](#)).

Because risk classification is the linchpin on which obligations hang, education policymakers can profitably integrate the OECD Framework for the Classification of AI Systems into their institutional risk triage: mapping an AI system along the framework's dimensions—People & Planet, Economic Context, Data & Input, AI Model, Task & Output—helps expose where a given adaptive tutor or proctoring tool implicates fundamental rights (e.g., equality in access), what data flows create privacy and representativeness risks, which model characteristics stress explainability, and how task-output couplings might enable or foreclose pedagogically sound override mechanisms; the OECD framework's lifecycle orientation (planning/design; data collection; model building/validation; deployment and monitoring) dovetails with the Act's demands for pre-market conformity assessment and post-market vigilance, offering a shared vocabulary by which school systems, vendors, and regulators can align on evidence of safety and accountability ([OECD, 2022](#)).

In domains where generative AI becomes a co-author of learning artefacts (lesson plans, formative feedback, worked examples), UNESCO's global Guidance for Generative AI in Education and Research supplies a concrete programmatic complement to the Act's legal minimalism: beyond compliance, UNESCO calls for age-appropriate guardrails, teacher-led validation workflows, and explicit curricular integration of AI literacy to sustain a human-centered, rights-preserving adoption; taken together, the UNESCO guidance and the AI Act suggest that the legitimate pedagogical uses of GPAI are those that preserve human agency, embed privacy-by-design, and make the limits of model knowledge visible to learners ([Miao and Holmes, 2023](#)).

Notably, the AI Act's emphasis on robustness, cybersecurity, and accuracy is not decorative: the ENISA Threat Landscape underscores that ransomware, DDoS, phishing, and related attacks remain persistent and adaptive, while AI-enabled threats (e.g., automated phishing, synthetic-voice fraud, data exfiltration targeting model pipelines) complicate institutional risk profiles; in education settings—where networks carry sensitive learner data and assessment systems may present tempting extortion targets—the Act's requirements for logging, incident reporting, and resilience should translate into concrete procurement clauses (e.g., secure model-update channels, adversarial-robustness testing, rate-limiting and anomaly detection) and operational controls (e.g., separation of duties, recovery plans, red-team exercises) commensurate with the sector's exposure ([European Commission, 2025](#)), ([ENISA, 2024](#)).

On governance, the Act's distributed architecture—European AI Office, national market-surveillance authorities, AI Board and advisory bodies—creates escalation and coordination pathways likely to affect education in three ways: first, by generating codes of practice and interpretive guidance that lower transaction costs for small ministries and school networks; second, by standardizing registries and documentation templates that enable comparability of high-risk deployments (e.g., exam-scoring engines); and third, by anchoring enforcement in administrative routines that encourage continuous improvement (e.g., corrective action plans, proportional penalties) rather than one-off punitive gestures ([European Commission, 2025](#)), ([European Parliament, 2025](#)). Strategically, aligning adaptive learning with the Act means institutionalizing a pedagogy-first compliance culture:

- needs analysis and theory-of-change before tool selection;
- explainability criteria tied to specific learner decisions (placement, feedback, progression);
- calibration studies that check algorithmic recommendations against expert judgments across student subgroups;
- a duty to furnish students with recourse, including the right to a human review of impactful automated decisions ([European Commission, 2025](#)), ([OECD, 2022](#)).

For assessment, the combination of transparency duties and high-risk obligations invites renewed attention to validity and fairness: if a proctoring classifier generates false positives differentially by lighting, skin tone, or disability status, then both the Act's dataset-quality requirement and equity commitments in education policy are implicated, compelling iterative retraining, alternative accommodations, or—where irreparable—retirement of the system; conversely, where automated scoring aids formative feedback without deciding high-stakes outcomes, limited-risk transparency may suffice, provided institutions clearly signal to learners when AI is involved and how its suggestions are moderated by instructors ([European Commission, 2025](#)), ([European Parliament, 2025](#)), ([Miao and Holmes, 2023](#)).

Finally, the chapter contends that responsible governance post-Act is a project of institutional capability as much as legal conformity: ministries and universities should inventory AI-mediated decisions, classify them using OECD criteria, decide—policy-by-policy—where high-risk thresholds are crossed, and build UNESCO-aligned AI literacy into teacher education and student orientation; procurement should prefer vendors who evidence adherence to GPAI guidance, publish training-data summaries, and support auditability; and CISOs should plan for ENISA-profiled threats with controls that map to the Act's robustness and post-market monitoring obligations; done well, the result is not compliance theatre but a defensible, learner-first operationalization of the Act's animating principle: trustworthy AI in education that enhances learning while protecting rights (European Commission, 2025), (ENISA, 2024).

While this paper outlines foundational principles for responsible AIED, further empirical research is needed to evaluate real-world impacts across diverse educational systems. Future work should investigate longitudinal learning outcomes, the effectiveness of human-in-the-loop models, and equity effects of algorithmic interventions. Policymakers should prioritize capacity-building and iterative evaluation as core components of AI adoption in education.

Conclusion

Education is changing rapidly, and AI is one of its most powerful catalysts—but it is far from a neutral instrument, and deploying it responsibly demands sustained attention to ethical, pedagogical, and legal standards. This paper has assembled a framework that weaves together theory, real-world applications, and regulatory alignment, foregrounding the roles of human agency, sound data governance, and transparency at every turn. With the EU AI Act now redefining the ground rules for how adaptive systems and automated assessments can be designed and deployed, educators and policymakers face both an obligation and a rare opportunity to shape a pedagogy that embraces innovation without abandoning its responsibilities. Going forward, what the field needs most is empirical work—studies that evaluate these systems across different institutional contexts, with a persistent focus on whether they actually promote equity and whether the institutions using them are ready for the challenge.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Conflict of interests

The authors declare no conflict of interest.

Author Contributions

Conceptualization, G.D., Z. S and D.V; methodology, Č.V.; software, R.I., D.S. and M.N.; formal analysis, R.I. and Č.V.; writing—original draft preparation, G.D, Z.S. and R.I.; writing—review and editing, D.S., Č.V. and D.V. All authors have read and agreed to the published version of the manuscript.

References

- ENISA. (2024, October). Threat landscape. European Union Agency for Cybersecurity. <https://www.enisa.europa.eu/topics/cyber-threats/threat-landscape>
- European Commission. (2025, August 1). AI Act | Shaping Europe's digital future. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- European Commission, Directorate-General for Education, Youth, Sport and Culture. (2022). Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators. Publications Office of the European Union. <https://data.europa.eu/doi/10.2766/153756>
- European Parliament. (2025, February 19). EU AI Act: First regulation on artificial intelligence. <https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>

- European Union. (2024, June 13). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828. *Official Journal of the European Union*, L (12 July 2024), 1–144. <http://data.europa.eu/eli/reg/2024/1689/oj>
- Friedman, L., Black, N. B., Walker, E., & Roschelle, J. (2021, November 8). Safe AI in education needs you. *Communications of the ACM Blog*. <https://cacm.acm.org/blogs/blog-cacm/256657-safe-ai-in-education-needs-you/fulltext>
- Gartner. (2025). Augmented intelligence. *Gartner Glossary*. <https://www.gartner.com/en/information-technology/glossary/augmented-intelligence>
- Ivković, R. (2025). Transformation and economic aspects of software engineering through the implementation of the EU AI Act. *Pravo – Teorija i Praksa*, 42(1), 186–200. <https://doi.org/10.5937/ptp25011861>
- Maslej, N., et al. (2023). AI Index 2023 annual report. Stanford University, Institute for Human-Centered AI (HAI). <https://aiindex.stanford.edu/report/>
- Miao, F., & Holmes, W. (2023). Guidance for generative AI in education and research. UNESCO. <https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>
- OECD. (2022). OECD framework for the classification of AI systems. OECD Publishing. https://www.oecd.org/content/dam/oecd/en/publications/reports/2022/02/oecd-framework-for-the-classification-of-ai-systems_336a8b57/cb6d9eca-en.pdf
- The White House, Office of Science and Technology Policy. (2022, October). Blueprint for an AI Bill of Rights: Making automated systems work for the American people. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
- U.S. Department of Education, Office of Educational Technology. (2022). Advancing digital equity for all: Community-based recommendations. <https://files.eric.ed.gov/fulltext/ED624655.pdf>
- U.S. Department of Education, Office of Educational Technology. (2023). Artificial intelligence and the future of teaching and learning: Insights and recommendations. <https://tech.ed.gov>
- Wang, S., Wang, F., Zhu, Z., Wang, J., Tran, T., & Du, Z. (2024). Artificial intelligence in education: A systematic literature review. *Expert Systems with Applications*, 252, 124167. <https://doi.org/10.1016/j.eswa.2024.124167>
- IEEE-USA Board of Directors. (2017, February 10). Artificial intelligence research, development and regulation [Position statement]. IEEE-USA Global Public Policy. <http://globalpolicy.ieee.org/wp-content/uploads/2017/10/IEEE17003.pdf>