

AI in Education

The Key Challenges Institutions Must Solve Now

Introduction

Artificial Intelligence (AI) holds real promise for education, from enabling unprecedented personalization to improving operational efficiency. However, realizing those benefits is anything but straightforward.

Educational institutions, from K–12 schools and universities to certification bodies and corporate learning providers, face tough questions when introducing AI: Can we protect student privacy? Will the tools actually improve outcomes? Are we ready for the governance it requires?

From unclear regulations to system fragmentation, the road to responsible AI is filled with pitfalls that organizations often discover only after costly missteps.

This whitepaper aims to illuminate those pitfalls before you're in too deep. Through sector-specific insights, we explore the most pressing challenges education organizations encounter when adopting AI, not to sell you a specific solution, but to help you ask the right questions and prepare your organization for what lies ahead.



1 The Rising Expectations for Personalized Learning

We live in a world where personalization is the norm, from Netflix on-demand recommendations to Spotify tailored playlists. It's no surprise that students and professionals alike now expect the same from their learning environments. Learners want to be seen, understood, and supported as individuals, with unique goals, backgrounds, and engagement levels, and they expect it in real time. But delivering individualized learning experiences at scale remains a daunting task for most educational institutions.

Personalization goes beyond adaptive tests or recommended videos. It means a full understanding of each learner's context, their goals, challenges, background, and current level of engagement, and, of course, responding dynamically. AI has the potential to make this level of insight possible at scale, but only when it is trained on diverse, relevant data and embedded into platforms that can actually talk to each other in real time.

Many EdTech platforms start with good intentions: provide smart recommendations, adapt to student behavior, and identify when someone needs help. But without a full learner profile, including historical performance, socio-emotional signals, and instructor feedback, AI tools fall short. They end up reinforcing existing patterns rather than helping students break through them.

There's also a risk of personalization becoming a proxy for surveillance. When AI monitors every click, students may feel evaluated instead of supported. Without strong, clear governance and thoughtful communication, personalized AI can easily cross the line into intrusive territory, eroding trust, the very thing it aims to build.

2 Compliance and Regulatory Pressures

Data privacy is a foundational element of trust in any educational setting. And that trust is being tested. As EdTech platforms grow more sophisticated, collecting more data and implementing AI more deeply into learning environments, the regulatory landscape is evolving fast and becoming more complex.

The European Union's General Data Protection Regulation (GDPR) has already set a high bar for privacy, consent, and data portability. In the United States, the Family Educational Rights and Privacy Act (FERPA) governs how student records are accessed and shared. The Children's Online Privacy Protection Act (COPPA) adds further layers for K–12. And many institutions add their own layer of internal ethics policies and data governance frameworks on top of national laws.

Regulations aren't just about keeping data secure. They also ensure that individuals understand how their data is being used, and have meaningful control over it. While many EdTech platforms build features with privacy in mind, few are designed to be auditable, explainable, and adaptable enough to handle the shifting regulatory landscape.

One growing concern is how AI systems make decisions, especially when those outcomes affect academic standing, access to resources, or recommendations for future learning paths. GDPR and similar laws give users the right to an explanation for automated decisions. But many machine learning models, particularly black-box systems, can't provide that. Institutions must either redesign models for explainability or risk noncompliance.

The upcoming EU AI Act raises the bar even further: educational AI systems will be treated as high-risk, placing new demands on documentation, monitoring, and human oversight. These requirements won't be limited to EU-built products; they'll apply to any system operating in EU markets or serving EU citizens. For global EdTech platforms, this means rethinking infrastructure, governance, and vendor contracts.

Agentic AI, a system that can independently perform tasks, make decisions, and initiate actions on behalf of institutions, adds an even more significant compliance challenge. Unlike traditional AI, agentic systems must navigate regulations autonomously, continually adjusting their actions based on evolving data protection and privacy frameworks. This autonomous behavior heightens the stakes: the compliance mechanisms must be embedded directly into the AI's decision-making logic.

Ensuring compliance of agentic AI requires a new level of rigor: continuous oversight, embedded governance, transparent documentation, and regular audits. It also demands close collaboration among EdTech providers, regulatory bodies, educators, and technology experts to ensure that agentic AI not only meets current standards but is also resilient enough to adapt to future regulatory shifts.

The pressure isn't only external. Boards want clear policies. Parents and students expect control, visibility into how systems work, and clarity that their information won't be misused. Instructors want tools that support—not replace—their judgment. And legal teams want clear, defensible policies that don't require heroics to enforce.

The shift is clear: deploying AI quickly is no longer good enough. Institutions must deploy it responsibly. That means privacy by design, explainability by default, and compliance as a continuous process, not a one-time audit.

3 Transparency, Explainability, and Trust

As AI systems become more embedded in educational workflows, recommending content, predicting students' risk, or prioritizing interventions, transparency is no longer optional. Stakeholders want more than results; they want to understand the logic behind them.

Data protection laws increasingly demand explainability. Yet, many AI models remain opaque. Students, instructors, and administrators are often handed decisions or predictions with no insight into how they were generated. This black-box approach can erode trust quickly.

For AI to be credible, institutions must prioritize interpretability. This means not only designing systems that provide understandable reasoning but also training users to interpret outputs appropriately. For example, when a system flags a student as "at risk," faculty need to know whether that's based on attendance patterns, quiz performance, behavioral cues, or a mix of all three, and how to respond constructively.

Transparent systems foster confidence and engagement. They invite dialogue and allow for critique, course correction, and continuous improvement. More importantly, they signal that institutions respect users enough to explain the systems shaping their academic journeys. That's what builds trust and makes innovation sustainable.

4 Operational Challenges and Instructor Burnout

AI tools are often introduced with the promise of reducing workload, automating repetitive tasks, surfacing key insights, and streamlining processes. However, when poorly implemented, they can have the opposite effect.

Instructors already juggle large class sizes, high administrative loads, and the expectation to be simultaneously tech-savvy, empathetic, and data-driven. If AI systems add more dashboards, more notifications, and more complexity without removing friction elsewhere, they become yet another burden.

For AI to deliver value, it must fit within existing workflows or improve them. It should be intuitive, easy to use and understand, low-maintenance, and genuinely helpful, without demanding technical expertise or turning educators into testers. Most importantly, it should free up instructors' time to focus on what matters: mentoring, content development, and pedagogy.

Success starts with inclusion. Instructors must be part of the AI systems design, rollout, and refinement. Professional development, feedback loops, and opt-in models (not mandates) are key to building sustainable, supportive AI experiences.

5 Fragmentation and Integration Challenges

Educational technology environments are often made up of disjointed systems: Learning Management Systems (LMS), Student Information Systems (SIS), Customer Relationship Management (CRM) tools, Learning Experience Platforms (LXP), and countless third-party applications. Each serves a purpose, but together, they rarely function as a cohesive whole.

AI depends on interconnected, clean, and timely data. Yet, in most organizations, that data is fragmented and siloed. This fragmentation makes it difficult to build a full picture of the learner journey or deliver timely, context-aware interventions. AI systems trained on partial data produce partial insights. Worse, when different systems operate on conflicting assumptions or definitions, like what counts as an “engaged student,” trust in outputs breaks down.

Institutions must invest in integration, standardizing data models, ensuring real-time interoperability, and creating unified APIs or data lakes. But the technical work alone won’t fix the problem; addressing the cultural and operational silos that prevent teams from collaborating on shared data governance is just as crucial. Without a strategic foundation, even the smartest AI will fall short.

6 Scaling AI Safely and Effectively

Piloting AI in a classroom or department is one thing; scaling it across an institution is another. Many issues that seem manageable at a small scale—manual monitoring or individual oversight—become untenable as the scope expands.


Scalability isn't just about infrastructure. It's about governance, consistency, and foresight. As AI is deployed across more use cases, institutions need clear policies for system behavior, escalation paths for unexpected outcomes, and metrics for success. Without these, adoption will either stagnate or spiral.

Institutions also face the challenge of maintaining ethical and pedagogical integrity at scale. There's also a critical human dimension. One-size-fits-all AI models are unlikely to serve diverse learner populations. Responsible expansion requires adaptive mechanisms, feedback loops, and human review layers that must grow along with the deployment footprint.


Responsible scaling is iterative, not linear, and success depends as much on organizational maturity as technological readiness.

About DataArt


DataArt is a leading global software engineering firm that delivers breakthrough data, analytics, and AI platforms for the world's most demanding organisations. We're your partners for progress!



experts across
20+ countries



340+ clients
surveyed



satisfied
clients



community-
focused

Book a free 30-min strategy session to discuss your strategy.



edtech@dataart.com

We'll walk you through a high-level maturity assessment, identify quick wins, and share how top Education organizations are using data to stay ahead.