Google Cloud

From Principles to Practice: Responsible AI in Academia

Multimodal generative AI (GenAI) offers higher-ed institutions the opportunity to turn the mountains of big data across students, academics, staff data into always-on intelligence and unified campus operations. With generative AI, colleges and universities can deliver personalized study materials to students, help their students navigate bureaucracies, increase student retention, and accelerate research breakthroughs.

Many colleges are focused on opportunities like generative AI tutoring, adaptable curricula, and microcredentialing. Some universities are already using AI chatbots to connect first-generation and ESL college students to resources like FAFSA, the registrar, student housing, and more.

With AI-powered tools, faculty and administrators can spot data trends that help predict student outcomes, analyze early warning signs of falling behind, and identify effective interventions. But, if not managed responsibly, there are potential risks associated with AI, including introducing bias, jeopardizing data privacy, and eroding human connection.

That's why it's important for higher-ed institutions to take a step back and conduct rigorous tests to consider how to design, develop, and deploy campus-wide AI in a responsible, secure, and intelligent way.

An effective AI policy must include a strategy of robust data governance, stringent security measures, ethical AI design, and a commitment to transparency and accountability.

A Central Place for Unifying Structured and Unstructured Data

Understanding the Function of Institutional Data

To put AI to its most efficient and effective potential, higher-ed technology leaders face a tall order. They must identify high-quality data sources — both structured and unstructured — then connect them. Once the data is accessible, they need to create an interface that makes it easy for nontechnical users to ask questions and get back useful answers.

Unfortunately, two in three organizations reported at least half of their data is dark, meaning the data isn't readily usable. Fewer than half of data leaders are fully confident in their organization's data, according to Google's 2024 Data and AI Trends Report.

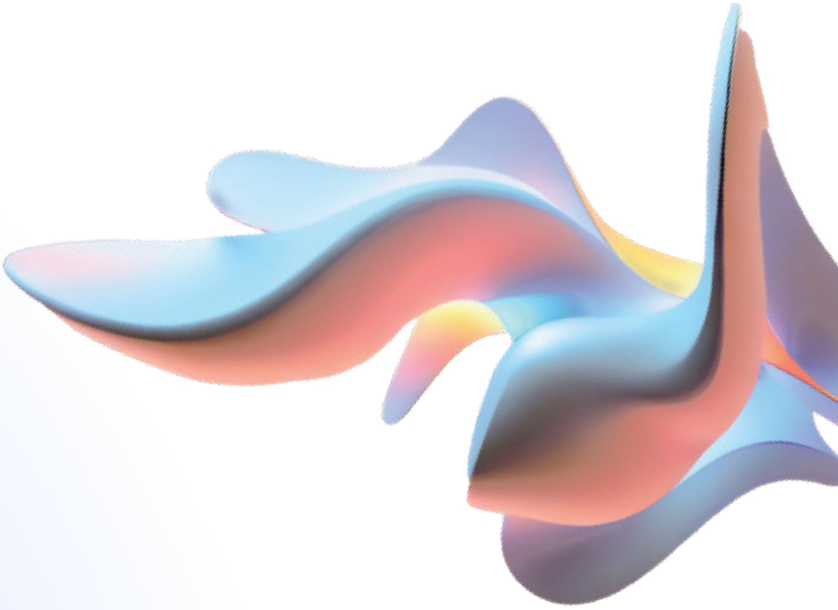

Google Cloud

Unifying Structured and Unstructured Data

The formats that have the most interesting data are not going to be text and it's not going to be structured. It's going to be videos, audio files. The ability to access those unstructured data sources the same as you're using your structured data sources, that is the kind of thing that really builds value really quickly."

Chris Hein
Director of Public Sector Engineering
Google Public Sector

"A lot of the IT shops that run on traditional campuses...exist to support infrastructure, not to build products," Hein said.

Artificial intelligence can help busy IT teams make this data visible and useful by quickly combing through the information, categorizing and classifying it, and identifying patterns, trends, and insights.

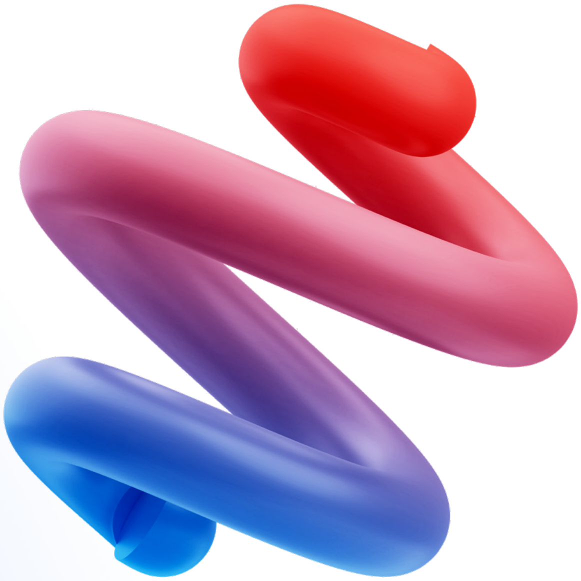

Resiliency Throughout Your Institution’s Data

Choosing the Right Partners and Platforms

Like any new technology, artificial intelligence introduces new fronts in the ongoing battle for cybersecurity. AI’s expanding capabilities in higher ed necessitates that colleges and universities address its potential vulnerabilities thoughtfully and proactively.

EDUCAUSE’s 2024 AI Landscape Study found 95 percent of those surveyed report AI impacting their institutions’ teaching & learning policies, while a much lower 72 percent are seeing AI’s impact on their cybersecurity and data privacy policies. This is a blind spot institutions must address to future-proof their technology infrastructure.

To safeguard the vast amount of sensitive data within these systems, it’s crucial to select a robust, secure infrastructure and to align with partners whose security values and practices match your organization’s standards. These shared fate principles between higher-ed institutions and their chosen partners can help ensure success without compromising security.

Resiliency Throughout Your Institution's Data

> **Numerous higher education institutions might not fully account for the extensive ecosystem of partners and platforms they work with. Their attention is often internal, centered on their own staff and resources, which can lead to overlooking the myriad of external stakeholders with access to their data."**

Charles Elliott
Field CTO for Research and Education
Google Public Sector

Elliott recommends pairing an aggressive program of penetration testing, training staff to recognize things like phishing attacks, and implementing a Zero Trust model within your vendor ecosystem

Zero Trust is a security model that operates on the concept of "never trust, always verify." Where traditional data security is like locking the main gate of a building, Zero Trust locks every door inside the building and installs security badge readers at every entry point.

Zero Trust Concepts

1. Assume all network traffic is a threat, at all times.
2. Enforce least-privileged access.
3. Always monitor.

Verified, Contextual Information Is the Foundation to a Trusted AI Model

Reducing Hallucinations to Ensure Accurate Knowledge for Students

It’s not enough for AI to be accurate. To be an effective tool in an educational setting, AI must provide contextual answers. It’s a concept called “grounding”. The technique entails feeding information like the course textbook, class lectures, and other information into the AI model.

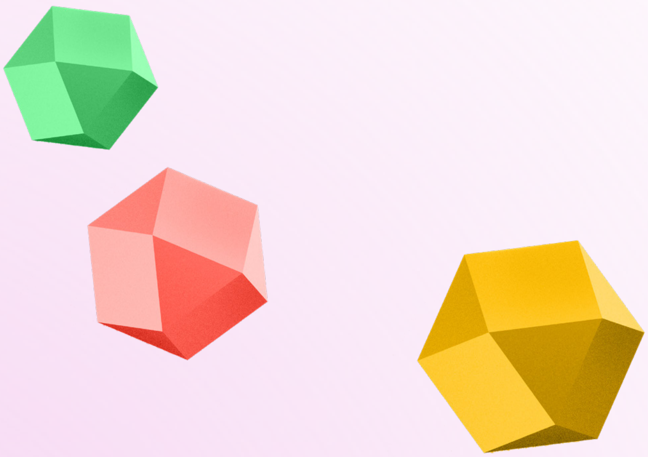

ASU Arizona State University

Bethany Weigel, chief information officer at Arizona State University, points to long division as just one example.

There are more than a dozen ways to do long division, Weigel explained. Say a student who has been taught the partial quotients method of learning long division. If they ask an AI tutor program for help on a division problem, and the model returns an answer using the grid method, that answer is accurate, but not appropriate.

“We saw math proficiency rates go down [during the pandemic],” Weigel said. “We don’t want to see that go down again because we’ve got AI incorrectly helping students toward the wrong answer or the wrong underlying principle.”

Verified, Contextual Information Is the Foundation to a Trusted AI Model

The contextual grounding ensures the AI model's responses are tailored to a specific educational context, aligned with both the instructor's teaching methods and the course's learning objectives. That way, when you ask a model a question, it refers to that information before it provides an answer. Advanced models further build trust in their responses by citing where they pulled the information from and giving users the opportunity to confirm that the information the model provided was accurate.

In generative AI, grounding involves connecting model outputs to verifiable sources of information. It's important for:

- Reducing model hallucinations (instances where the model generates content that's not factual).
- Anchoring model response to specific information.
- Enhancing the trustworthiness and applicability of the generated content.

Data Hygiene and Security Across the AI Lifecycle

Setting Up Your Institution for Responsible AI Implementation

Higher-education institutions can create a foundation for responsible and effective AI implementation by incorporating practices that can ensure safety, security, transparency, and equity.

This work includes starting with high-quality data sources and putting safeguards in place to make sure the responses AI models give are appropriate and contextual.

IT teams should test AI tools in controlled, isolated environments and monitor those tools carefully even after they are deployed. Once live, users should be given control over their own data and tools should provide notice and receive consent for how that data is used.

After users are trained and comfortable with the tools, IT teams should invite and offer plenty of opportunities for user feedback.

For more insights and practical strategies for implementing responsible AI practices, watch our webinar "Better Data Governance With Responsible AI."

Accelerate your institution's AI journey with our 10-step plan. Generate personalized content across text, video, images, and code, gain valuable insights from student data, and achieve measurable ROI — all with the enterprise-grade security and scalability of Google Cloud.

Discover 5 generative AI Use Cases in Education

Download our 10-step, 30-day Public Sector Generative AI Guide to kickstart your first generative AI use case.

Watch these on-demand sessions from Google Cloud Next '24 to jumpstart your Responsible AI journey:

A guide for enterprises: How to implement generative AI applications

AI: From proof of concept to impact

You can only secure what you can see: How observability empowers security

Secure the future: How government agencies lead cybersecurity innovation

Accelerate the value of generative AI with three secret ingredients

What's next for security professionals in under 4 minutes