Beyond the Sales Pitch

A Practical Guide to Questioning EdTech Vendors

USC Rossier
School of Education

USC Center for
Generative AI and Society

Introduction

This summer school administrators need to decide if they are going to deploy more capital on AI EdTech products that promise everything from teachers operating on auto-pilot to ensuring that each student receives exactly what they need, minute-by-minute.

Principles, superintendents, and other school leaders will be subjected to sleek demos from presenters who know all of the right buzzwords. Salespeople will showcase useful technologies and hint at the idea that, should administrators pass on their product, students will be left behind. The digital divide will grow—under their watch.

Intense pressure to do the right thing, or not "fall behind," will lead some school leaders to sign agreements that will not serve their students well. While most mistakes won't rise to the level of supporting a failed startup, some administrators will nonetheless let anxiety rule the day. They will spend money that could be deployed elsewhere and do more good on tried and true methods.

What can prevent this? Asking the right questions. That is the impetus behind this report. It contains key questions developed by leaders at the USC Rossier School of Education. We hope that every school leader considers these questions before they sign on the dotted line and commit resources to AI powered EdTech.

Contributors

Dean Pedro Noguera

Dean Pedro Noguera is a distinguished scholar of education whose work focuses on the ways in which socioeconomic and racial factors influence schools and student outcomes. As an institutional leader, his perspective on educational technology centers on equity, systemic impact, and whether new tools genuinely elevate teaching quality and support marginalized student populations. His inquiries push vendors to prove that their systems foster meaningful skill acquisition and provide differentiated support for all learners.

Dr. Stephen J. Aguilar

Dr. Stephen Aguilar is an Associate Professor of Education specializing in educational psychology and educational technologies, with a research focus on how students adopt and utilize generative AI tools. His work advocates for "human-in-the-loop" designs that prevent technology from hijacking the learning process or encouraging passive, executive help-seeking. His technical and methodological screening questions demand rigorous transparency from vendors regarding pilot data, foundational model choices, external validation, and direct institutional access to raw data.

Dr. Gale Sinatra

Dr. Gale Sinatra is a leading expert in educational psychology, particularly known for her research on climate change education, science denial, and how individuals process public misunderstandings of science. Her focus on generative AI centers on cognitive resilience, digital literacy, and student safety. Her screening questions target the critical mechanics of vendor tools: how models are trained, how student data is protected, and how platforms actively insulate learners from misinformation and inappropriate content.

Dr. Kenneth Yates

Dr. Kenneth Yates is a specialist in instructional design, cognitive task analysis, and workforce development, focusing on the cognitive architectures that underpin effective learning. His perspective on educational technology is rooted in cognitive efficiency—ensuring that tools actively scaffold complex mental processes rather than short-circuiting them. His targeted questions force AI developers to explicitly define which cognitive processes their technology supports and, crucially, which mental faculties risk being diminished by its use.

Questions when viewing demos of AI tools

Introduction

When evaluating generative AI products for educational settings, live demonstrations often showcase seamless interfaces and ideal use cases. However, school leaders and educators must look beyond marketing features to assess the pedagogical, technical, and ethical foundations of these tools. This section outlines critical questions proposed by leading educational experts. These prompts are designed to guide administrators during vendor demos, ensuring that any adopted technology actively enhances teaching quality, respects student differences, and maintains rigorous scientific and data validation.

Dean Pedro Noguera

- Does it enhance the quality of teaching?
- Does it further differentiation support for students?
- Does it enhance the skills and knowledge students acquire?

Dr. Stephen Aguilar

- How much pilot testing was done before the product went live?
- What populations of students is this designed for?
- How does your tool/system account for students' individual differences?
- Was educator input used when this tool/system was designed? If so, how was it used?
- Can I see any white paper or internal report that discuss this tools efficacy?
- Has there been any external validation of this tool/system?
- Which foundation model does your system rely on (e.g., ChatGPT, Gemini, etc.)? Why did you choose that model over others?
- Can we access our own raw data used by the sytem? How? If not, why not?

Dr. Gale Sinatra

- How are my students' data used?
- How are the models trained?

- How can you keep students safe from misinformation?
- How can you keep students safe from inappropriate content?
- How much does this cost?

Dr. Kenneth Yates

- What cognitive processes in learning does your technology support?
- What cognitive processes in learning are diminished by using your technology?

Additional Questions

- How do you address bias, accessibility, and unequal impact across student groups?
- What decisions is the tool not allowed to make without human review?

Key Takeaways

- **Pedagogical Impact Over Novelty**: AI tools must do more than automate tasks; they must demonstrably enhance the quality of teaching and support the differentiation of student learning without diminishing core cognitive processes.
- **Rigorous Validation and Transparency**: Districts should demand evidence of pilot testing, external validation, and clarity regarding the underlying foundational models (e.g., ChatGPT, Gemini) being utilized.
- **Student Safety and Equity**: Guarding against bias, misinformation, and inappropriate content must be built into the tool's design, requiring clear boundaries on what decisions the AI can make without human oversight.

Contracts, Terms and Conditions

Introduction

Procuring educational technology involves navigating complex legal frameworks, especially when generative AI is involved. Beyond standard software-as-a-service (SaaS) agreements, AI vendors introduce unique challenges regarding data harvesting, model training, and algorithmic liability. This section compiles essential screening questions regarding contracts and terms and conditions. It focuses on safeguarding student and teacher privacy, ensuring compliance with federal regulations like FERPA and HIPAA, and protecting educational institutions from legal and financial risks.

Dr. Stephen Aguilar

- Will student or teacher data be stored? If so, how will it be used?
- Will content generated by students or teachers be owned by your company?
- Does your company, or the companies you use to handle data, have a history of data breaches?
- How will you protect FERPA or HIPPA data if it is accidentally entered into your system?
- Will students' or teachers' use of your tool/system be used to train additional iterations of it?
- Can we opt out of allowing our communities data being used to train future versions of this software? (How/Why not?)
- How does your company deal with the possibility of inappropriate content is shown to students?
- Can student or teacher data collected through their use of this product later be sold? If so, can we opt out?

Dr. Gale Sinatra

- Can we use other AI platforms?
- What support do we get for teachers?
- Can we cancel the contract if it does work out?

Dr. Kenneth Yates

- Who owns the data collected from use of the technology?
- What data security exists on the platform?
- Can we opt out at any time without penalty?

Additional Questions

- What resources will the EdTech company provide for incidents resulting from their AI tool?
- Who bears liability if the AI goes off the guard rails?
- How can districts leverage outcomes based contracts to protect students and future budgets?
- How will the community be impacted by data centers, if at all, and if the local community will not be impacted, which community will?
- How will schools be able to track the impact of their AI usage on both students and the environment?
- Are the TOS compliant with existing state regulations? (Check, these are changing rapidly as state legislatures create guidelines.)

Key Takeaways

- **Data Ownership and Privacy Sovereignty**: Where possible, school districts should maintain strict ownership of student and teacher data, with clear provisions preventing the vendor from selling data or using it to train future commercial AI iterations.
- **Compliance and Security Infrastructure**: Vendors must demonstrate robust data security, a clean history regarding breaches, and explicit protocols for handling accidental entries of protected FERPA or HIPAA data.
- **Contractual Flexibility and Liability**: Agreements should include penalty-free opt-out clauses, comprehensive teacher support, and clear legal liability frameworks specifying who is responsible if the AI system fails or produces harmful output.
- **Environmental and Community Footprints**: A forward-looking procurement process should take into account the broader impacts of AI adoption, including tracking its environmental footprint and assessing its strain on existing tech infastructure.