

POLICY BRIEF



YOUTH &
EDUCATION

Vetting Generative AI Tools for Use in Schools

David Sallay, Future of Privacy Forum, April 2024



AUTHORED BY

David Sallay

Director for Youth & Education Privacy

ACKNOWLEDGEMENTS

The author would like to thank Jim Siegl, Alexa Mooney, Daniel Hales, Anne Flanagan, and the many experts and stakeholders whom were consulted for their contributions to the report



The Future of Privacy Forum (FPF) is a non-profit organization that serves as a catalyst for privacy leadership and scholarship, advancing principled data practices in support of emerging technologies. Learn more about our work by visiting fpf.org.

Table of Contents

Executive Summary	3
Introduction	5
What should a school include in a legal compliance review?	6
What is unique about AI tools in edtech?	8
Understand the most common use cases	10
Does the use case require student PII?	11
Is the technology able to meet requirements for transparency and explainability?	14
Will the company use student PII for product improvement?	16
Ensuring that student PII does not appear in the AI edtech tool's output	17
Will the tool be used for substantive decision making?	18
Takeaways	19
Incorporating Generative AI Into Your School's App Vetting Checklist	21

Executive Summary

Edtech tools using artificial intelligence have been in schools for years, but due to the recent widespread release of AI tools that can generate text, images, audio, and video, the topic has risen to the top of public discourse. Schools are developing policies to cover adopting new AI tools, and many organizations have released frameworks of what those policies should cover. These frameworks typically have been light on detail when it comes to privacy, stating that schools need to follow privacy laws, but not explaining how those laws should be implemented as teachers, students, and others use rapidly-evolving AI technologies in the classroom. To address this gap, the Future of Privacy Forum's **Vetting Generative AI Tools for Use in Schools** explains the steps schools should consider incorporating into their more general edtech vetting and AI policies. It is crucial to keep in mind that three main classroom communities are using AI tools - students, teachers, and institutions - and AI by these stakeholders can implicate different privacy equities.

Laws, policies, and agreements that should be included in a legal compliance review

The landscape of privacy laws applying to schools includes federal and state privacy laws, which provide parents and eligible students with rights related to sensitive data collections, transparency, correcting inaccurate information, and place obligations on schools to require specific contract provisions with vendors and maintain prescribed data retention and deletion practices. Adding to this are laws that cover vendor responsibilities, including newer ones relating specifically to AI and high risk decision making. Along with applicable laws, many school districts have local policies in place governing app vetting that should be considered when assessing the appropriateness of an AI tool for school use. Additionally, many companies offering AI tools are already contracted under a written agreement with schools for other services. It is imperative for schools to review these agreements and determine if the use of an AI tool aligns with the terms of the original agreement or if additional terms should be negotiated for specific AI tool use.

Unique attributes of AI in edtech

Existing student privacy laws have many requirements that a school needs to consider when contracting with any edtech vendor, including those using AI. Because of this, it is important for schools to know what is unique about AI tools, which can then be added to existing review processes. The main differences schools be aware of are

- **Use case dependent.** Since generative AI edtech tools can take more user input to perform any number of tasks as output compared to traditional edtech tools, schools will need to consider the specific use cases for the tool.
- **Data collection.** Student privacy laws typically will cover use cases where the tool requires student personally identifiable information (PII) as input or where the output from the tool will become part of the student's record. There are many use cases that do not require student PII, which the school can use without implicating most student privacy laws. Even still, there are many use cases where a school may not be able to control all the information the tool collects, so schools should consider whether the data collection risk can be mitigated or avoided altogether.
- **Transparency and explainability.** For tools that use student PII, the school will need to consider how it will meet requirements for transparency and explainability to teachers, parents, and students. State privacy laws frequently require schools to publicly share information on what student data they share and its recipients. Many edtech companies are creating AI transparency pages to better explain the data their tools use and how they make decisions.
- **Product improvement.** Many Generative AI tools rely on large amounts of data to continuously train the underlying model that generates responses. Other tools train a model initially but do not use student data to further train the tool. An important question schools need to ask is whether the vendor will use student PII to train the model, and if so, if any additional products the vendor creates with the model are educational or commercial, and if that additional use is permitted under state law.
- **Unauthorized disclosure of student PII.** If student PII is used to train the model, then there exists the chance that snippets of the PII will appear in future output from the tool. The school will need to understand the steps the company takes to prevent these sorts of unauthorized disclosures.
- **High risk decision making.** Some proposed use cases that involve substantive decision making may be governed by long standing rules or new AI laws. Other uses may have such a high risk of harm to students that schools should be cautious in pursuing them. Potential options schools may consider are only permitting these cases with parental consent, requiring that a human be in the loop, or prohibiting the use case.

Introduction

Algorithms, analytics, and Artificial Intelligence have been used by K–12 and higher education for some time,¹ but recently rose to the top of the public discourse due to the proliferation of new powerful AI tools powered by generative pretrained transformers (GPTs), commonly referred to as generative AI. Generative AI tools rely on large language models that crunch data from multiple sources to produce AI-generated content at unprecedented speed, scale and accuracy. Initially, concerns about this new generation of AI tools related to questions about academic honesty and plagiarism,² though by the start of the 2023 school year, K–12 organizations seemed more willing to adopt³ them while calling for the development of policies⁴ and frameworks^{5,6} to safely use them. Issues to be addressed included protecting copyright, addressing the inaccurate “hallucinations” the tools frequently produce, and also complying with existing law, including privacy laws like the Family Educational Rights and Privacy Act (FERPA) and state laws like California’s Student Online Personal Information Protection Act (SOPIPA). As of 2024, schools are struggling to address policies that determine the use of AI tools by multiple audiences, including edtech vendors and service providers, students, teachers, and school administrative staff. The dawn of generative AI tooling and the reality of the proliferation of their use has made the conversation about AI use policies seem more urgent than ever

Policies and frameworks proposed by US states and professional organizations supporting schools typically look at the use of AI tools holistically (not just the use of generative AI tooling), answering big-picture questions related to pedagogy, procurement, equity, and many others.⁷ Privacy is often included in these frameworks, though when it comes to legal compliance, they

¹ <https://www.elgaronline.com/edcollbook/book/9781800375413/9781800375413.xml>

² <https://ed.stanford.edu/news/what-do-ai-chatbots-really-mean-students-and-cheating>

³ <https://www.edweek.org/technology/180-degree-turn-nyc-schools-goes-from-banning-chatgpt-to-exploring-ais-potential/2023/10>

⁴ <https://fpf.org/blog/fpf-weighs-in-on-the-responsible-use-and-adoption-of-artificial-intelligence-technologies-in-new-york-city-classrooms/>

⁵ <https://www.edweek.org/technology/schools-want-guidance-on-ai-use-in-classrooms-states-are-not-providing-it-report-says/2023/09>

⁶ <https://www.edsafeai.org/safe>

⁷ For examples, see CoSN’s checklist at <https://www.cgcs.org//cms/lib/DC00001581/Centricity/Domain/417/K-12%20Generative%20AI%20Readiness%20Checklist%20October%2011%202023%20V1.1.pdf>, guidance from the California Department of Education at https://drive.google.com/file/d/1k8kjbLRoIKOB7pu5s4wh-4_CufUNJEA/view, guidance from the Oregon Department of Education at [https://www.oregon.gov/ode/educator-resources/teachingcontent/Documents/ODE_Generative_Artificial_Intelligence_\(AI\)_in_K-12_Classrooms_2023.pdf](https://www.oregon.gov/ode/educator-resources/teachingcontent/Documents/ODE_Generative_Artificial_Intelligence_(AI)_in_K-12_Classrooms_2023.pdf), and the Software & Information Industry Association at <https://edtechprinciples.com/principles-for-ai-in-education/>

often say little more than that schools should ensure that their adoption complies with state and federal privacy laws, like FERPA, and ensure a process is in place to vet the tools. However, if a school does not know how laws like FERPA and state student privacy laws that protect student personally identifiable information (PII) would apply to the complexities of machine learning systems, this very general direction does not provide adequate guidance.

To address this emerging gap, FPF's **Vetting Generative AI Tools for Use in Schools** publication explains the steps a school should typically take to vet an AI tool for compliance with student privacy laws. AI refers to a wider spectrum of tools and technologies,⁸ some of which have been in schools for years and belonging to the broader field of AI in education.⁹ This publication will look specifically at text-based generative AI tools—the sort of technology that is top of mind for many school districts right now—because it best illustrates the process for vetting any AI tool, highlighting the questions the reviewer would ask and the ones they likely could not answer without further explanation from the vendor. Our approach, with minor tech-specific modifications, would also be useful when analyzing emerging AI technologies that generate images, audio, video, or other content.

What should a school include in a legal compliance review?

Most of the privacy laws a reviewer should consider apply directly to the school. For example, FERPA “applies to all schools that receive funds under an applicable program of the U.S. Department of Education.”¹⁰ As FERPA does not apply to companies, a given edtech product cannot be “FERPA-compliant,”¹¹ so vetting considers how a tool can be used by a school in a FERPA-compliant manner.¹² There are several laws and other legal instruments that a school should consider:

⁸ For example, predictive AI processes data to forecast probable outcomes based on historical data. This has been used in student monitoring, particularly in early warning systems to predict which students may need interventions. Generative AI, on the other hand, utilizes large data sets to create new, original content, such as text or images. In some cases, it may not be apparent that a tool even uses AI. For more details on the variety of AI tools, see <https://fpf.org/blog/newly-updated-report-the-spectrum-of-artificial-intelligence-companion-to-the-fpf-ai-info-graphic/>

⁹ <https://www.elgaronline.com/edcollbook/book/9781800375413/9781800375413.xml>

¹⁰ <https://www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html>

¹¹ Former director of the US Department of Education’s Student Privacy Policy and Assistance Division, Michael Haws, said, “There is no such thing as a ‘FERPA seal of approval.’” <https://marketbrief.edweek.org/marketplace-k-12/5-tips-protecting-student-data-living-ferpa/>

¹²

<https://www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html#:~:text=The%20Family%20Educational%20Rights%20and,the%20U.S.%20Department%20of%20Education>

- **Federal privacy laws**
 - **Family Educational Rights and Privacy Act (FERPA)**, which includes requirements for schools to allow parents and eligible students to inspect and review education records, have a fair process for contesting the accuracy of records, and only share records with consent or by ensuring certain protections are in place.
 - **Protection of Pupil Rights Amendment (PPRA)**, which requires that schools provide parents with the ability to review instructional materials and places restrictions on the collection of certain sensitive information.
 - **Children’s Online Privacy Protection Act (COPPA)**, which is administered by the Federal Trade Commission (FTC) and applies to for-profit companies that collect personal information from children younger than 13. In 2023, the FTC proposed updated regulations to COPPA,¹³ which would codify long standing guidance on how it applies in the school context.
- **State student privacy laws.** There are more than 128 state student privacy laws,¹⁴ which cover additional requirements relating to schools sharing student PII with edtech vendors. These vary state-to-state, but frequently place requirements on the schools (e.g., provide more transparency to parents about sharing of student PII, require specific contract provisions with vendors) and on companies (e.g., prohibitions on creating profiles on students for noneducational purposes, selling student data, using student PII for targeted advertising). State laws also address record retention and deletion by requiring schools to classify records they maintain to adhere to different retention schedules and by placing requirements on vendors to delete records containing student PII at the school’s request or following the termination of the contract.
- **Other state consumer privacy laws.** These laws typically include provisions that exclude data protected by FERPA; however, some include provisions related to automated decision making, which may cover some use cases proposed by schools.
- **Local policies.** Since many schools and districts will already have policies in place that cover app vetting, the reviewer will want to ensure that they follow those existing policies.
- **Terms of service and existing vendor agreements.** The reviewer should also consider the tool’s terms of use and ensure that the proposed use is in alignment. For instance, ChatGPT’s terms of service prohibits children under 13 from using the service and only permits teens to use it via parental consent.¹⁵ Furthermore, many of the AI tools that a reviewer considers will be offered by companies the school already has written

¹³<https://www.federalregister.gov/documents/2024/01/11/2023-28569/childrens-online-privacy-protection-rule>

¹⁴ <https://studentprivacycompass.org/state-laws/>

¹⁵ <https://openai.com/policies/terms-of-use>

agreements with. In this case, the reviewer will want to consider whether the new AI tool and use case are already covered by the agreement or if additional terms should be negotiated.

What is unique about AI tools in edtech?

Federal and state student privacy laws and local policies already place specific requirements on schools sharing student PII with edtech companies. The process by which a school reviews any new edtech vendor for their compliance with these laws should be part of a larger app vetting process.¹⁶ Any legal compliance review will require schools to understand how the data flows from the student to the vendor throughout the data lifecycle, starting with providing any required notice to parents and potentially getting consent before moving onto data collection, data usage and storage, any additional sharing of the data, and finally ending with the eventual deletion of the data.

In most ways, the same principles for vetting edtech tools that collect, use, protect and share student data apply to the emerging technologies of generative AI (e.g., determine if the data exchange follows the requirements of FERPA's school official exception, have the vendor sign a data protection agreement), but there are some specific unique considerations. Schools should understand what is unique about generative AI tools in the K–12 education privacy context and add that to an existing app vetting process where one exists, rather than inventing a brand new standalone policy for AI. If the school does not have an app vetting process in place yet, they should develop a comprehensive process that considers how to review all apps, including those with AI, instead of proposing two separate processes.¹⁷

Due to the specific characteristics of the technology, generative AI tools are more likely to create unique legal compliance issues in the following areas:

- **Use case dependent.** Generative AI tools are powerful in the sense that the user can use plain language to request the tool to perform any number of tasks compared to traditional edtech tools, which generally have more limited outputs (e.g., a math program that only tells the student whether they answered the question correctly or not). Because of this, it is essential that the reviewer consider the specific use cases for the tool as well as whether the tool will prevent unauthorized use cases. Some use cases will relate to when the input entered into the tool contains student PII. The output of a generative AI tool may also be protected, though it will depend on whether the output contains student PII, which

¹⁶ https://studentprivacycompass.org/wp-content/uploads/2020/06/Adopting-EdTech_-Privacy-Vetting.pdf

¹⁷ For guidance on creating a comprehensive app vetting process, see this from CoSN: https://www.cosn.org/wp-content/uploads/2023/01/CoSN-Student-Data-Privacy-Toolkit-Part-2-0323_v5.pdf

it may not (e.g., if an educator entered student PII including test scores and requested output that summarized how the class performed without including any identifiers). The reviewer will want to consider both the inputs and outputs of specific use cases to determine which protections the law requires.

- **Data collection and processing including potentially proprietary information.** Similar to the example above, since generative AI tools generally use an open-ended textbox interface that will prompt the user to provide certain information (i.e., input), the school may not be able to control all the information the tool collects. This risk will be exacerbated if the tool is put in front of students, who may be more inclined to provide information they should not, including student PII and proprietary information from third party sources.
- **Transparency and explainability.** Generative AI tools rely on massive amounts of data to train their models, and that training compounds over time. Even the tool's designers may not be able to explain why it made the decisions it did as a result of a discrete input from a student. Schools will need to consider how they will meet federal and state requirements to provide parents with transparency and reasonable explanations of what data the tools receive and how they use, maintain, and ultimately dispose of it. Since the vendor will know their product best, they should be prepared to share this information with schools or offer built in safeguards for users as is the commercial trend with business-to-business (B2B) generative AI systems.
- **Product improvement.** Large language model (LLM) tools continue to learn by adding new input into the training data set. Existing student privacy law typically allows for student PII or de-identified data to be used for product improvement for the services provided to the schools, but would not permit it without parental consent if the student PII is used to develop noneducational products. Because of this, the reviewer will need to know if the vendor intends to use student PII to further train their model and if so, what kind of products will use the training data set. This is more likely to be an issue with a vendor providing a general purpose business-to-consumer (B2C) tool than one that is specifically designed for K–12 schools and follows a B2B style model.
- **Unauthorized disclosure of student PII.** Student privacy laws prohibit edtech companies from additional disclosures of student PII except under limited circumstances as directed by the school. Many edtech vendors offering generative AI are contracting with a third party to provide the LLM, and so the school will need to ensure that the proper contract terms are in place for this subcontracting, the same as with any other edtech tool. Where this becomes especially relevant is generative AI output since if student PII is included in the tool's training set, then there exists the possibility that it will be included in identifiable format in responses to unauthorized users. In this case, the reviewer will want to ensure that the tool takes steps to either prevent student PII from entering the training dataset or from having it be reproduced in identifiable form.

- **High-risk decision making.** Many proposed use cases for generative AI are to save users time by quickly processing documents. Existing federal and state privacy laws do not typically restrict what schools use data for or how they process it, provided that the use case is within the educational context. Newer state laws that target AI in general, however, recognize the risks in entirely automated decision making and have started adding restrictions to their use.

Understand the most common use cases

Since generative AI tools can be used in a variety of use cases—ranging from saving time, personalizing learning, providing support (e.g., tutoring), and assisting with creativity—it is essential that the reviewer first understand the context that the school plans to deploy the tool.

According to Holmes and Tuomi,¹⁸ AI edtech tools generally fall into three distinct categories: student-focused, teacher-focused, and institution-focused. This analysis will consider examples of each, namely:

Focus of AI edtech tool	Generative AI Examples
Student-focused	<ul style="list-style-type: none"> ● Intelligent tutoring system that personalizes learning based on student inputs ● Mental health chatbot that will discuss the student’s personal problems ● Tool to assist student in writing a personal essay ● Creativity tool (e.g., photo editing program) that uses AI to quickly generate new images ● Tool that helps a student generate computer code and does not require a login
Teacher-focused	<ul style="list-style-type: none"> ● Tool used to assist teacher in grading papers, including identifying plagiarism

¹⁸ Holmes, W., & Tuomi, I. (2002). State of the art and practice in AI in education. *European Journal of Education, Research, Development and Policy* 57(4). <https://doi.org/10.1111/ejed.12533>

Focus of AI edtech tool	Generative AI Examples
	<ul style="list-style-type: none"> ● Tool used to analyze spreadsheet of student’s scores ● Wizard used to quickly design the teacher’s landing page in the Learning Management System ● Tool used to generate lesson plan ideas ● Tool used to generate letters of recommendation
<p style="text-align: center;">Institution-focused</p>	<ul style="list-style-type: none"> ● Tool used to review admissions applications ● Tool used to generate text for announcements to send to parents ● Tool used to analyze student attendance and behavior

Each of these proposed uses will carry different risks, and the purpose of vetting the AI tool is to consider these risks and ultimately determine which course of action to take (e.g., approve, approve only with parental consent, not approve). This publication will largely focus on the risk of not being compliant with federal and state law but will also note some of the risks not covered by current law including perception risks, which are ones where the proposed use may be technically legal, but would likely lead to concerns or complaints from parents, students, or teachers.

Does the use case require student PII?

FERPA protects education records, which are records that directly relate to a student and are maintained by the school or by a party acting under the direct control of the school on the school’s behalf.¹⁹ If the reviewer determines that the use case involves disclosing existing or creating new FERPA protected education records, the school will need to ensure that they are following the protections required by FERPA and able to provide parents and eligible students with the rights granted to them by FERPA.

¹⁹ See the definition of “Education Record” in 34 CFR § 99.3

Typically, records maintained by an edtech tool that the school has students use meet the definition of a record under FERPA and in state student privacy laws if these records include student PII in the input or output. To determine if the record directly relates to the student, the reviewer would want to know whether the record contained student PII as defined in FERPA or state law. Typically, this includes direct identifiers, such as a student’s name, or student ID, as well as indirect identifiers like a birth date. Use cases that would involve direct identifiers include an educator copying a spreadsheet of student scores and requesting that the AI tool perform an analysis or where the educator enters an individual student’s PII and requests that the tool quickly produce an individualized education program (IEP). At this point, several of the example use cases are clearly more likely to include student PII in either the input or output.

Focus of AI edtech tool	Likely to include student PII in the input/output	Likely does not include student PII in the input/output
Student focused	<ul style="list-style-type: none"> ● Intelligent tutoring system that personalizes learning based on student inputs ● Mental health chatbot that will discuss the student’s personal problems ● Student uses the tool to assist in writing a personal essay 	<ul style="list-style-type: none"> ● Creativity tool (e.g., photo editing program) that uses AI to quickly generate new images ● Tool that helps a student generate computer code and does not require a log in
Teacher focused	<ul style="list-style-type: none"> ● Tool used to assist teacher in grading papers ● Tool used to analyze spreadsheet of student’s scores ● Tool used to generate letters of recommendation 	<ul style="list-style-type: none"> ● Tool used to generate lesson plan ideas ● Wizard used to quickly design the teacher’s landing page in the Learning Management System

Focus of AI edtech tool	Likely to include student PII in the input/output	Likely does not include student PII in the input/output
Institution focused	<ul style="list-style-type: none"> • Tool used to review admissions applications • Tool used to analyze student attendance and behavior 	<ul style="list-style-type: none"> • Tool used to generate text for announcements to send to parents • Tool used to assist in hiring educators and other staff

There may be some gray area use cases because they only use de-identified records, which are not covered by federal and state student privacy laws. In fact, well intentioned informal guidance at the school may be to tell educators to simply not enter any student PII into the tool, but de-identification is rarely that simple.²⁰ Without clear guidelines and training, it is not clear that all educators know what constitutes student PII. For example, if instead of a student’s name, the educator uses a substitute, such as a student ID, those are also PII in FERPA and in many state laws. If the educator is having students create their own accounts to use the tool, the account creation often requests an email address or other credentials, which are also by definition PII.²¹ While not a specific step in the vetting process, this demonstrates that establishing and communicating an app vetting policy needs to be complimented by regular privacy training for staff and educators who will use the tools.²²

Assuming the educator figures out a way to remove all direct and indirect identifiers, there is still one last criteria in FERPA, which is the reasonable person standard. This includes any other information that alone or in combination could allow a reasonable person in the school community, without any personal knowledge of the situation, to link the record back to the student.²³ Since users tend to be conversational with these tools, if the plan is to have students use the tool directly, the reviewer should understand the exact use cases the educator had in

²⁰ <https://studentprivacycompass.org/resource/student-data-and-de-identification/>

²¹ See the definition of “Personally Identifiable Information” in 34 CFR § 99.3

²²

<https://studentprivacycompass.org/resource/the-best-way-to-protect-students-personal-data-teacher-training/>

²³ See (f) under the definition of “Personally Identifiable Information” in 34 CFR § 99.3

mind to determine whether the student’s input meets this final criteria. If, for example, the student is using the tool as a writing partner to brainstorm ideas for a personal essay for class—even without direct identifiers like a name or student ID—that is more likely to involve student PII that would meet the reasonable person standard than a creativity tool where the student can request that specific images be generated or the use case of having the student generate computer code. In the case of the former, the reviewer would be right to say that a reasonable member of the school community could identify the student, and that the data entered should follow the required protections of state and federal law.

Depending on the specific use case, the reviewer may also need to consider if another federal K-12 student privacy law, the Protection of Pupil Rights Amendment (PPRA), applies.²⁴ PPRA gives parents the right to opt in or opt out of certain sensitive data collections that fall under specific protected categories, such as political or religious views, sexual behaviors, or mental health issues. Whether the data collection requires an opt in or opt out depends on whether the usage is funded by the US Department of Education or not, though certain state laws may require parental consent regardless of funding source.²⁵ Regardless of whether the usage technically falls under the scope of these laws or not, due to the unpredictable nature of generative AI tools, the reviewer should address the perception risk and ask the vendor how the product will handle these and similar sensitive topics should a student spontaneously start asking the question that lead to responses that cover these and similar topics.²⁶

Is the technology able to meet requirements for transparency and explainability?

Having determined that records created by these tools include student PII either in the input or output, the reviewer next needs to determine how the school will meet requirements for transparency and explainability. Under FERPA, this include providing access to inspect and review education records²⁷ and to provide a fair process to request that inaccurate records be corrected.²⁸ State laws may take this further, requiring that the school proactively share information, often on the school or district’s website, about the recipient of student PII and the specific PII that is shared.

²⁴ <https://studentprivacycompass.org/faqs-ppra/>

²⁵ See Utah Code Annotated 53E-9-203, as an example.

²⁶ The importance of this is highlighted by a study from the Center for Democracy and Technology, which found that 29% of students have used ChatGPT for dealing with anxiety or mental health issues, see p. 33 in <https://cdt.org/wp-content/uploads/2023/09/091923-CDT-Off-Task-web.pdf>

²⁷ 34 CFR Subpart B

²⁸ 34 CFR Subpart C

FERPA requires schools to respond to reasonable requests for explanations and interpretations of the records, which in use cases where the tool creates new records as part of the output may be challenging since how AI technology works can be a black box, and in many cases even the technology's own developers cannot always explain how it makes decisions.²⁹ Beyond legal compliance, studies have shown that this black box creates a perception risk, and students and teachers are more likely to accept and trust AI technologies when given explanations of how they work.³⁰ Therefore, the reviewer will want to reach out to the vendor for more clarification on how they can work together to process these requests.³¹ Many edtech vendors are currently considering the use of AI transparency pages to proactively provide schools with this information.³²

FERPA also requires schools to have a process in place wherein parents can request that the school amend inaccurate records. Since generative AI tools can fabricate information, called hallucinations, this should cause the reviewer to more deeply scrutinize use cases where the tool will create new education records as part of the output, especially where the planned use is for substantive decisions (e.g., grading a student's paper, designing a student's IEP). Generative AI tools are marketed as being time savers, but the school would need to weigh the time saved in quickly generating these decisions against the time spent looking for and correcting hallucinations or in processing requests to amend inaccurate records. Furthermore, using generative AI tools to make substantive decisions may be in violation of the company's terms of service.³³

Will the company use student PII for product improvement?

The general rule under FERPA is that schools may only share records with an edtech company with written parental consent; however, there are several common-sense exceptions to this rule, which the reviewer should consider first when approving before going the parental consent

²⁹ <https://umdearborn.edu/news/ais-mysterious-black-box-problem-explained>

³⁰ Kizilcec, R. To Advance AI Use in Education, Focus on Understanding Educators. *Int J Artif Intell Educ* 34, 12–19 (2024). <https://doi.org/10.1007/s40593-023-00351-4>

³¹ Common approaches to explain AI in education are described in Khosravi, H., Shum, S. B., Chen G., Conati, C., Tsai Y-S., Kay, J., Knight, S., Martinez-Maldonado, R., Sadiq, S., & Gašević, D. (2022). Explainable Artificial Intelligence in education. *Computers and Education: Artificial Intelligence*, 3(100074). doi: doi.org/10.1016/j.caeai.2022.100074

³² For examples, see https://support.zoom.com/hc/en/article?id=zm_kb&sysparm_article=KB0057861 or <https://tutor.classdojo.com/#/ai-transparency-note>

³³ From ChatGPT's terms of use: "You must not use any Output relating to a person for any purpose that could have a legal or material impact on that person, such as making credit, educational, employment, housing, insurance, legal, medical, or other important decisions about them."

route. The most common exception when sharing with edtech is “school official,” which permits schools to share with services that perform a service the school would otherwise use its employees for and that are under the direct control of the school. Many state laws add to this requirement by specifying that the direct control should be achieved by a contract with the school that guarantees specific privacy protections.³⁴

One of these requirements that is relevant to generative AI relates to product improvement, specifically whether the service (or its third parties) will use student PII to improve the underlying Large Language Model (LLM). Many edtech products incorporate services like ChatGPT using an Application Programming Interface (API) and have terms³⁵ that state that their product does not use user input to train the ChatGPT AI model.

If the generative AI tool will use data inputs to further train its larger model, the reviewer will need to understand whether the data will be de-identified and what kind of additional tools the company will create from the LLM. Existing FERPA guidance holds that a company “may use data (even in individually identifiable form) to improve its delivery of these applications. The provider may also use any non-PII data, such as metadata with all direct and indirect identifiers removed, to create new products and services.”³⁶ State laws frequently include similar language, such as when California’s Student Online Personal Information Protection Act clarifies that none of the restrictions on vendors “shall be construed to prohibit the operator’s use of information for maintaining, developing, supporting, improving, or diagnosing the operator’s site, service, or application,”³⁷ though in some states the requirements may be more restrictive. For example, New York’s Ed Law 2-D specifies that vendors may not use data “for any other purposes than those explicitly authorized in its contract,”³⁸ and Florida’s student privacy law only permits product improvement with de-identified data.³⁹ Therefore, the reviewer would need to know which types of products the training set was going to help improve, whether they all are educational products, and whether those uses are explicitly authorized in the contract. If the company uses student PII to train noneducational products, then it likely would not pass any of these requirements.

³⁴ The US Department of Education has this guide, which goes over best practices in contracts with vendors:
https://studentprivacy.ed.gov/sites/default/files/resource_document/file/TOS_Guidance_Mar2016.pdf

³⁵ <https://openai.com/policies/business-terms>

³⁶

<https://tech.ed.gov/wp-content/uploads/2014/09/Student-Privacy-and-Online-Educational-Services-February-2014.pdf>, specifically example 4

³⁷ https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201320140SB1177

³⁸ <https://codes.findlaw.com/ny/education-law/edn-sect-2-d.html>

³⁹ See 1006.1494(5)(a), <https://www.flsenate.gov/Session/Bill/2023/662/BillText/er/HTML>

The distinction between the use of student PII to improve the product the school has authorized rather than to create new products is also seen in the recent proposed COPPA rule, which limits the use of student PII for improvement to “the specific educational service that the school has authorized.”⁴⁰ Furthermore, the FTC has stated that the school may provide consent in place of the parent only if the usage is entirely within the educational context; therefore, if children under 13 will use the generative AI tool, a key question the reviewer should ask is whether the data exchange falls entirely within the educational context or if any of it is in a commercial context. If, for instance, the vendor will use data collected from the child to further train the model and then build noneducational products based on that model, then the school likely cannot provide verifiable consent in place of the parent, and parental consent would be the only approved route for the school to take.

Ensuring that student PII does not appear in the AI edtech tool’s output

When receiving data under FERPA’s school official exception, vendors are not permitted to reshare student PII except under the direction of the school and when following another FERPA exception. State laws frequently include similar provisions specifically prohibiting resharing or selling student PII.

Since generative AI tools repurpose user input into the training model and use that to create new output, the reviewer would want the company to be explicit with how they would prevent student records from being shared further. This is especially important if student PII will enter the tool’s LLM. For example, in at least one case, Google Bard was accidentally leaking snippets of user chats into public Google search results.⁴¹ Though Google has since resolved this specific issue, it highlights the need to ask the question when approving the use of a new AI tool.

Will the tool be used for substantive decision making?

The newest area of the law that schools should consider governs AI technologies in general, beyond just generative AI tools. In many instances, use cases that involve automated substantive decision making require some amount of heightened scrutiny, such as requiring risk assessments,⁴² parental consent, or being outright banned. For example, with the proposed use

⁴⁰<https://www.ftc.gov/legal-library/browse/federal-register-notices/16-cfr-part-312-childrens-online-privacy-protection-rule-nprm>

⁴¹<https://www.fastcompany.com/90958811/google-was-accidentally-leaking-its-bard-ai-chats-into-public-search-results>

⁴² <https://studentprivacycompass.org/congresswoman-lori-trahans-new-student-privacy-discussion-draft/>

case of using a generative AI tool to facilitate admissions decisions, the reviewer should also be aware of more general state consumer privacy laws that cover legal or similarly significant effects of any automated decision making. Both Colorado⁴³ and Virginia's⁴⁴ laws explicitly list education enrollment in their list of use cases that would require an opt out. Depending on the tool, these use cases may also be in violation of the terms of service.⁴⁵ Schools using AI tools to make substantive decisions should ensure that there are internal processes in place to manage internal and external risk and ensure that the use of such tools are fit for purpose.

Use cases that are not fully addressed by these types of laws may be in the near future as more states are considering AI specific laws to address any gaps not covered by a patchwork of privacy laws. Even absent specific laws that address AI, there may be an elevated risk of harm to a student. For example, even if a mental health app met all of the privacy requirements in FERPA,⁴⁶ the unpredictable nature of the tool's answers could lead to giving potentially dangerous advice. In one case, a generative AI chatbot was quick to give advice that was essentially advocating for an eating disorder.⁴⁷ Another area of concern is the level of bias these tools may reintroduce. For example, a recent Stanford University study found that when used for writing letters of recommendation, generative AI tools have written letters that discriminate against women meaning the model is biased in its design.^{48,49} In cases where these use cases are permitted by the law, the reviewer should still consider whether parental consent or ensuring that having a human be in the loop are appropriate to mitigate the risk of harm to students.

⁴³ https://leg.colorado.gov/sites/default/files/documents/2021A/bills/2021a_190_rer.pdf

⁴⁴ <https://lis.virginia.gov/cgi-bin/legp604.exe?211+ful+SB1392ES1+pdf>

⁴⁵ See ChatGPT's terms of use, which states, "You must not use any Output relating to a person for any purpose that could have a legal or material impact on that person, such as making credit, educational, employment, housing, insurance, legal, medical, or other important decisions about them."

⁴⁶ It is also worth noting that many of these tools when tested have rated poorly on their privacy controls in general. See <https://www.consumerreports.org/health/health-privacy/mental-health-apps-and-user-privacy-a7415198244/>

⁴⁷ <https://www.talkingmentalhealth.com/post/tessa-chatbot-cautionary-tale-ai-mental-health#:~:text=A%20dangerous%20lapse%20in%20advice,them%20distressed%20and%20ill%2Dadvised.>

⁴⁸ Wan, Y., Pu, G., Sun, J., Garimella, A., Chang, K., & Peng, N. (2023). "Kelly is a Warm Person, Joseph is a Role Model": Gender biases in LLM-generated reference letters. *EMNLP 2023 Findings*. <https://doi.org/10.48550/arXiv.2310.09219>

⁴⁹ Bias in the model can be engineered out, but all models are likely to contain some bias since no data set is ever truly random and evenly distributed. This is why redteaming is so essential for improving models. User feedback is an important contributor.

Federal and state student privacy laws typically only cover data about students, but some uses will cover educators and other school staff, especially in the area of hiring. For schools considering this use case, FPF has published a separate best practices guide.⁵⁰

Takeaways

The main takeaway for schools is if they do not currently have a process for vetting edtech tools, this AI moment around generative AI demonstrates the need for such policies. Schools that already have a process should consider what makes AI unique and add that to the existing process instead of creating a new one. To ensure that schools have adequate capacity to review edtech, state legislatures should ensure that these initiatives are funded and staffed.⁵¹ As this review makes clear there are many use cases for generative AI tools that are unlikely to use student PII as an input or create protected records as output, and so student privacy laws are not implicated. Many other use cases will include student PII and will need to follow legal requirements, but if they consider what makes this technology unique as part of their app vetting, then this should be achievable. Finally, there are some use cases where the risks to the student's privacy are so high that the school should only allow them in limited cases with parental consent, if not outright prohibiting them.

For vendors looking to have their products used in schools, the main takeaway should be that there are existing requirements and protections that schools are going to expect to see, and vendors are going to need to demonstrate that they meet those requirements, or if the existing product currently does not pass muster, vendors will need to create a K–12 education-specific product that does. In order to establish direct control, schools are likely going to request that vendors sign a data protection agreement that covers the relevant legal requirements. If the vendors are not willing to do this, schools will not allow the products in their schools. So that schools can be transparent with parents and students regarding their use of AI, vendors should also consider ways to proactively be transparent to schools about their use of AI and take advantage of edtech professional development conferences to train school staff on the appropriate use of their tools.

⁵⁰<https://fpf.org/blog/future-of-privacy-forum-and-leading-companies-release-best-practices-for-ai-in-employment-relationships/>

⁵¹ For more on how what happens when states fund and staff student privacy, see FPF's case studies on Utah at <https://studentprivacycompass.org/resource/utah-case-study/> and New York at <https://studentprivacycompass.org/resource/ny-case-study/>

Incorporating Generative AI Into Your School's App Vetting Checklist

- Determine your local requirements for vetting all edtech:** Vetting a generative AI tool will have much in common with vetting any edtech tool. Ideally, your school/district already has policies in place for this review. If not, the school should develop a more comprehensive policy for vetting edtech. A comprehensive policy based on existing federal and state laws may include the following:
 - Keeping student data safe from commercial purposes (e.g., not selling data, prohibiting targeted advertising)
 - Requiring data breach notification
 - Requirements for data minimization, de-identification, and aggregation
 - Specific requirements for vendor contracts
 - How to address what is unique about AI technologies
- Describe the proposed use case(s):** Since generative AI tools can be used for a variety of applications, the reviewer will need to understand the specific use cases and whether they are covered by existing student privacy laws, consumer privacy law, or AI laws.
 - Will the use case require student PII as input?
 - Will the tool's output also be considered student PII?
 - Does the use case involve substantive decision making? If so, determine which precautions are adequate to address risk:
 - human-in-the-loop
 - consent/opt out
 - prohibiting the use case
- Prepare to address transparency and explainability:** If the use case involves student PII, the school needs to be prepared to answer the following:
 - Are you able to explain how the tool will be used and how the data will flow to teachers, parents, and students?
 - Does the vendor provide an AI transparency page that explains how the tool works?
- Determine if student PII from the tool will train the large language model (LLM):**
 - If so, determine if additional tools built on the model will be educational products.
 - Check if your state law will permit additional product development with student PII or if data must be de-identified.
 - Ask the vendor how they ensure that student PII will not be redisclosed in output



1350 EYE STREET NW | SUITE 350 | WASHINGTON, DC 20005

info@fpf.org | FPF.ORG