May 9, 2026

The AI Double-Edge: Balancing Learning Analytics and Student Privacy

Privacy+ Security Forum

Speakers

Privacy+ Security Forum

Inna Barmash
Chief Legal Officer & Corporate Secretary
Amplify

Rhonda Powell
Executive Vice President, Chief Legal Officer, and Corporate Secretary
Strada

Jim Siegl
Senior Technologist, Youth & Education Privacy
Future of Privacy Forum

Robyn Mohr
Partner and Deputy Chair Privacy, Security, and Data Innovations
Loeb & Loeb

Chanda Marlowe
Associate Privacy, Security, and Data Innovations
Loeb & Loeb

Topics for Today

Privacy+ Security Forum

I. Introduction

II. Current Landscape of AI and EdTech

III. Overview of Applicable Laws, Regulations, and Self-Regulatory Regimes

IV. Best Practices for EdTech Companies

V. Conclusion

Defining Artificial Intelligence

Privacy+ Security Forum

What is AI?

Basics:

- Computerized ability to perform tasks commonly associated with human intelligence – reasoning, discovering patterns, etc.

Machine Learning:

- Algorithms improve through experience

Generative AI:

- AI techniques that train models on existing data and then generate new content

Neural Networks and Deep Learning:

- AI techniques that teach computers to process data in a way inspired by the human brain

Overview of Applicable Laws, Regulations, and Self-Regulatory Regimes

Privacy+ Security Forum

Family Educational Rights and Privacy Act

Privacy+ Security Forum

What is FERPA?

- Federal student privacy law enacted in 1974
- Applies to schools that receive federal funding from the U.S. Department of Education.
- Gives parents (and eligible students) the right to **access** and **amend** their child's education records.
- Prohibits a school from disclosing **personally identifiable information** from a student's **education record** to a third party without written consent from the **parent or eligible student** (unless an exception applies).
- Requires **reasonable security methods**.

FERPA Enforcement

- U.S. Department of Education enforces FERPA
- Penalties
  - Schools can lose federal funding
  - School can be banned from using a vendor for 5 years

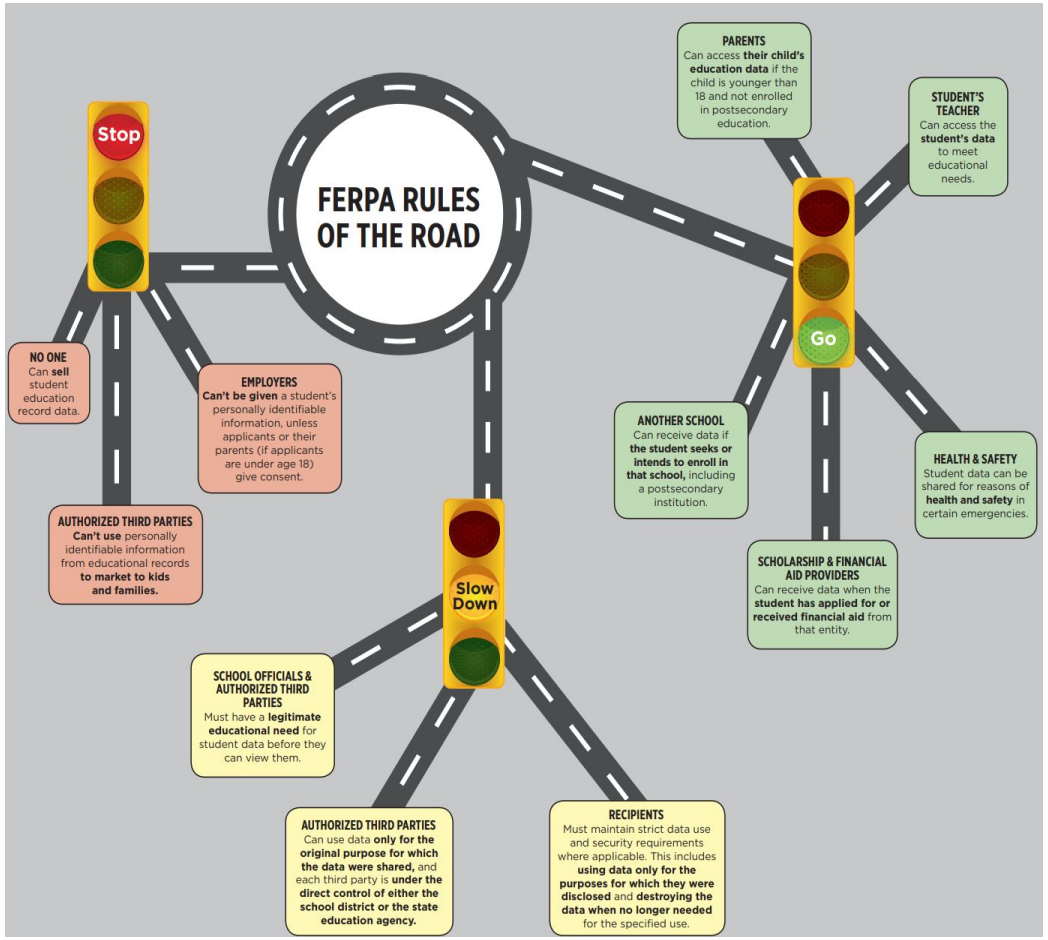

Source: Data Quality Campaign

What Type of Data Does FERPA protect?

Privacy+ Security Forum

- Directory information
- School official exception
- Audit and evaluation and exception
- Studies exception
- Health and safety exception
- Law enforcement and subpoena
- Lawsuit

Directory Information

- Information contained in an educational record that would not generally be harmful if disclosed
- **Examples:**
  - Name
  - Address
  - Photograph
  - Phone number
  - Date and place of birth
  - Grade level
  - Dates of attendance
  - Participation in officially recognized sports and activities
  - Degrees, honors and awards received
- Annual notice must be given to parents
- Parents must have a way to **opt out** of sharing directory information.

Children's Online Privacy Protection Act

Privacy+ Security Forum

Congress in 1998 enacted COPPA better protect children and their personal information while online.

Applies to "any operator of a website or online service directed to children, or any operator that has actual knowledge that it is collecting personal information **from** a child."

Who is a "child"?

Minor Under the Age of 13.

What is "personal information" in the context of COPPA?

Broadly defined to include persistent identifiers that can recognize users over time and across different websites or online services, photo/video/voice, and precise geolocation

Requirements, subject to certain exceptions, include:

"Notice" of operator data collection, use, and disclosure practices, and

"verifiable parental consent" before the operators collects, uses, or discloses the personal information.

Penalties?

Up to $46,517 per violation

Data destruction

20 year reporting requirements.

State Attorneys General may also enforce the Rule

Comprehensive State Privacy Laws with Requirements for EdTech

Privacy+ Security Forum

California

- Must obtain opt-in consent from minors age 13 to 15 and from parents of minors under the age of 13 before selling or sharing their personal information.
- ADMT Regulations Pending

New Jersey

- Must obtain opt-in consent from parents of minors under the age of 13 before selling or sharing their personal information.
- Must obtain consent for minors aged 13 to 16 to sale or use their personal information for targeted advertising or profiling.

Best Practices for EdTech Companies

Privacy+ Security Forum

Best Practices / AI Principles for EdTech

Privacy+ Security Forum

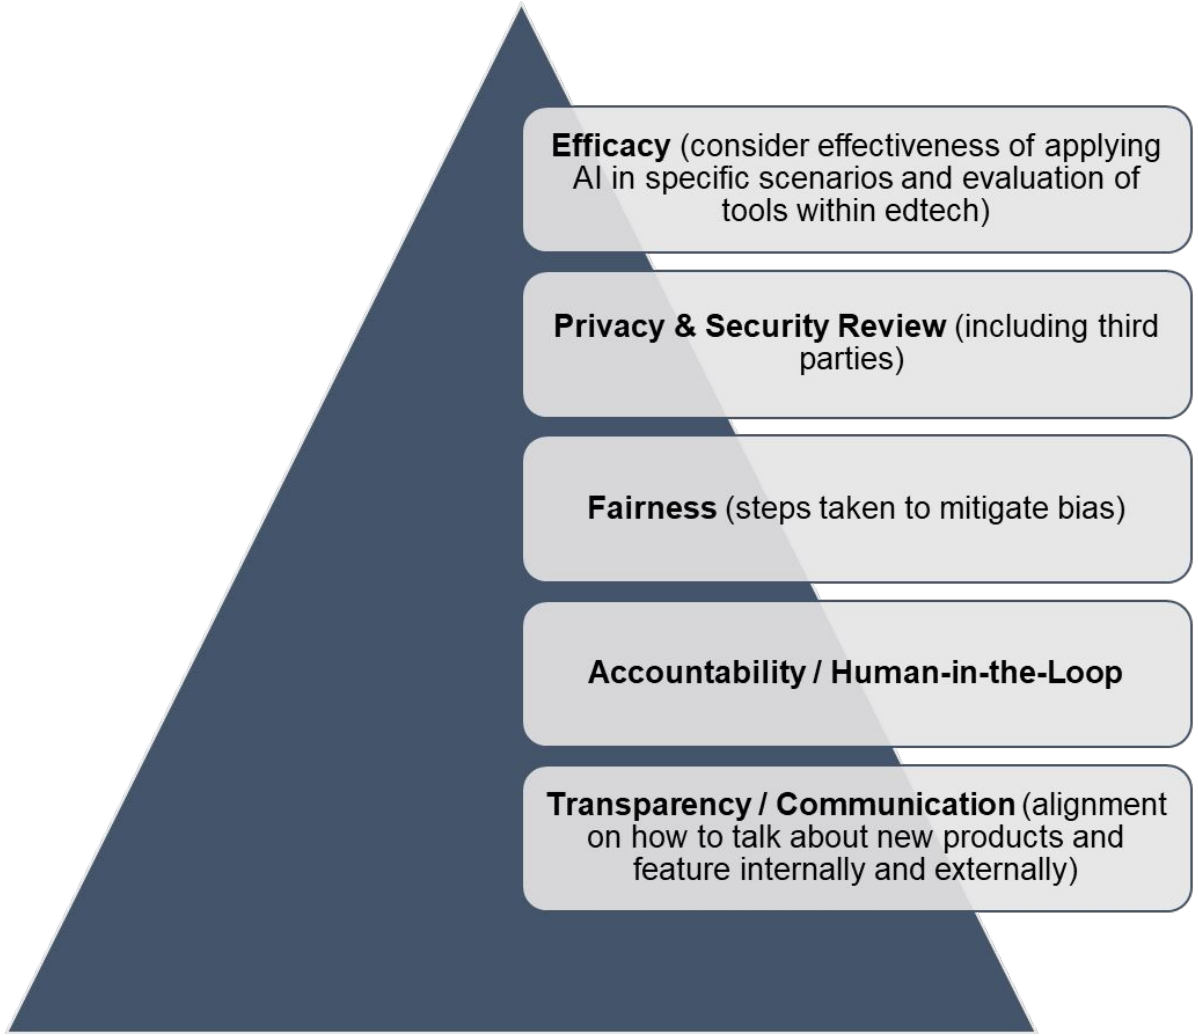

AI Transparency Example: McGraw Hill AI Reader

Privacy+ Security Forum

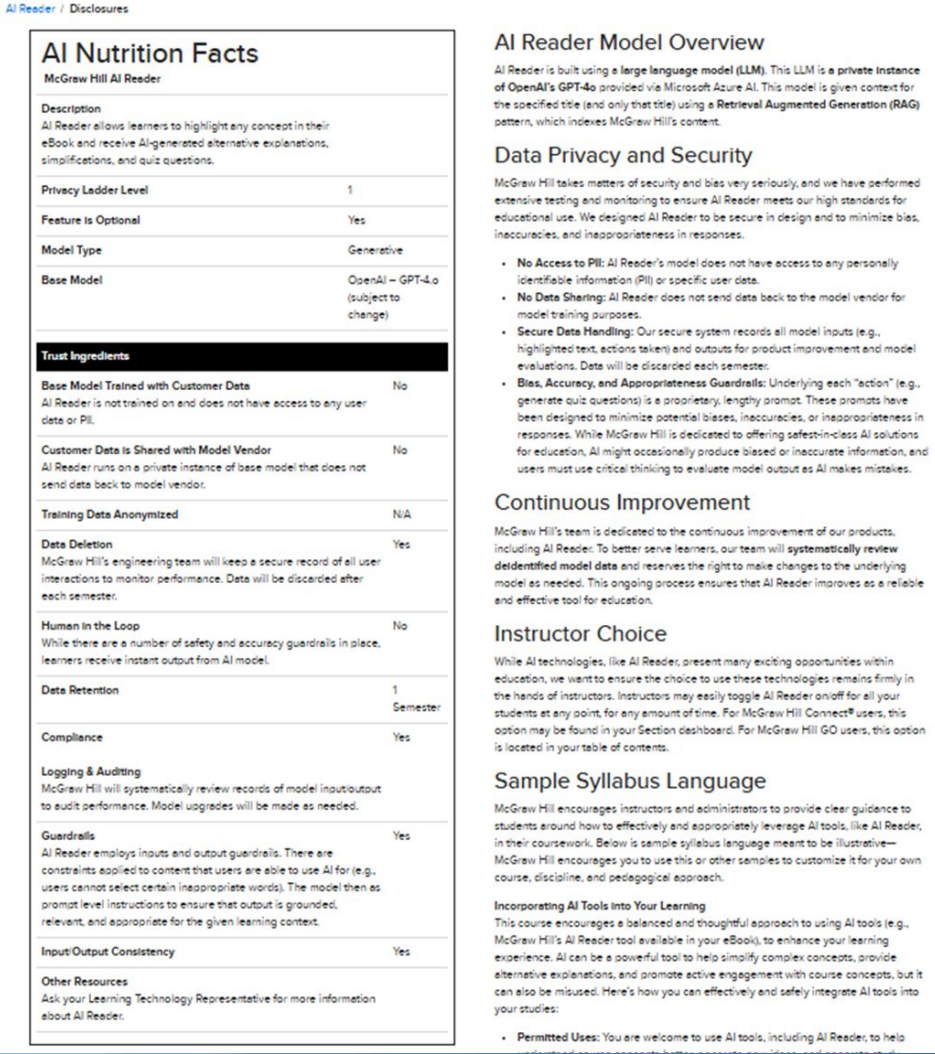

AI Reader / Disclosures

AI Nutrition Facts

McGraw Hill AI Reader

Description AI Reader allows learners to highlight any concept in their eBook and receive AI-generated alternative explanations, simplifications, and quiz questions.	
Privacy Ladder Level	1
Feature is Optional	Yes
Model Type	Generative
Base Model	OpenAI – GPT-4.o (subject to change)
Trust Ingredients	
Base Model Trained with Customer Data AI Reader is not trained on and does not have access to any user data or PII.	No
Customer Data is Shared with Model Vendor AI Reader runs on a private instance of base model that does not send data back to model vendor.	No
Training Data Anonymized	N/A
Data Deletion McGraw Hill's engineering team will keep a secure record of all user interactions to monitor performance. Data will be discarded after each semester.	Yes
Human in the Loop While there are a number of safety and accuracy guardrails in place, learners receive instant output from AI model.	No
Data Retention	1 Semester
Compliance	Yes
Logging & Auditing McGraw Hill will systematically review records of model input/output to audit performance. Model upgrades will be made as needed.	
Guardrails AI Reader employs inputs and output guardrails. There are constraints applied to content that users are able to use AI for (e.g., users cannot select certain inappropriate words). The model then as prompt level instructions to ensure that output is grounded, relevant, and appropriate for the given learning context.	Yes
Input/Output Consistency	Yes
Other Resources Ask your Learning Technology Representative for more information about AI Reader.	

AI Reader Model Overview

AI Reader is built using a **large language model (LLM)**. This LLM is **a private instance of OpenAI's GPT-4o** provided via Microsoft Azure AI. This model is given context for the specified title (and only that title) using a **Retrieval Augmented Generation (RAG)** pattern, which indexes McGraw Hill's content.

Data Privacy and Security

McGraw Hill takes matters of security and bias very seriously, and we have performed extensive testing and monitoring to ensure AI Reader meets our high standards for educational use. We designed AI Reader to be secure in design and to minimize bias, inaccuracies, and inappropriateness in responses.

- **No Access to PII:** AI Reader's model does not have access to any personally identifiable information (PII) or specific user data.
- **No Data Sharing:** AI Reader does not send data back to the model vendor for model training purposes.
- **Secure Data Handling:** Our secure system records all model inputs (e.g., highlighted text, actions taken) and outputs for product improvement and model evaluations. Data will be discarded each semester.
- **Bias, Accuracy, and Appropriateness Guardrails:** Underlying each "action" (e.g., generate quiz questions) is a proprietary, lengthy prompt. These prompts have been designed to minimize potential biases, inaccuracies, or inappropriateness in responses. While McGraw Hill is dedicated to offering safest-in-class AI solutions for education, AI might occasionally produce biased or inaccurate information, and users must use critical thinking to evaluate model output as AI makes mistakes.

Continuous Improvement

McGraw Hill's team is dedicated to the continuous improvement of our products, including AI Reader. To better serve learners, our team will **systematically review deidentified model data** and reserves the right to make changes to the underlying model as needed. This ongoing process ensures that AI Reader improves as a reliable and effective tool for education.

Instructor Choice

While AI technologies, like AI Reader, present many exciting opportunities within education, we want to ensure the choice to use these technologies remains firmly in the hands of instructors. Instructors may easily toggle AI Reader on/off for all your students at any point, for any amount of time. For McGraw Hill Connect® users, this option may be found in your Section dashboard. For McGraw Hill GO users, this option is located in your table of contents.

Sample Syllabus Language

McGraw Hill encourages instructors and administrators to provide clear guidance to students around how to effectively and appropriately leverage AI tools, like AI Reader, in their coursework. Below is sample syllabus language meant to be illustrative—McGraw Hill encourages you to use this or other samples to customize it for your own course, discipline, and pedagogical approach.

Incorporating AI Tools into Your Learning

This course encourages a balanced and thoughtful approach to using AI tools (e.g., McGraw Hill's AI Reader tool available in your eBook), to enhance your learning experience. AI can be a powerful tool to help simplify complex concepts, provide alternative explanations, and promote active engagement with course concepts, but it can also be misused. Here's how you can effectively and safely integrate AI tools into your studies:

- **Permitted Uses:** You are welcome to use AI tools, including AI Reader, to help

Link to Resource:
https://www.mheducation.com/highered/digital-products/ai/disclosures

Thank you!

Privacy+ Security Forum

Inna Barmash
Chief Legal Officer & Corporate Secretary
Amplify

Rhonda Powell
Executive Vice President, Chief Legal Officer, and Corporate Secretary
Strada

Jim Siegl
Senior Technologist, Youth & Education Privacy
Future of Privacy Forum

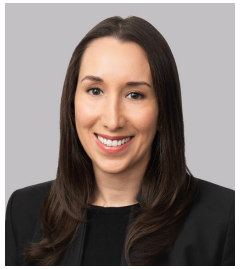

Robyn Mohr
Partner and Deputy Chair Privacy, Security, and Data Innovations
Loeb & Loeb

Chanda Marlowe
Associate Privacy, Security, and Data Innovations
Loeb & Loeb

Additional Information

Privacy+ Security Forum

What is AI?

Basics:

- Computerized ability to perform tasks commonly associated with human intelligence – reasoning, discovering patterns, etc.

Machine Learning:

- Algorithms improve through experience

Generative AI:

- AI techniques that train models on existing data and then generate new content

Neural Networks and Deep Learning:

- AI techniques that teach computers to process data in a way inspired by the human brain

Risks

- Data privacy concerns
- Bias and fairness issues
- Cybersecurity threats
- Lack of accountability
- Lack of explainability and transparency
- Misinformation

Opportunities

- Supporting students through use of AI for personalized learning
- Supporting teachers by automating time-consuming tasks
- Supporting administrators with faster and more nuanced data insights about students.

Overview of Applicable Laws, Regulations, and Self-Regulatory Regimes

Privacy+ Security Forum

What Laws Govern?

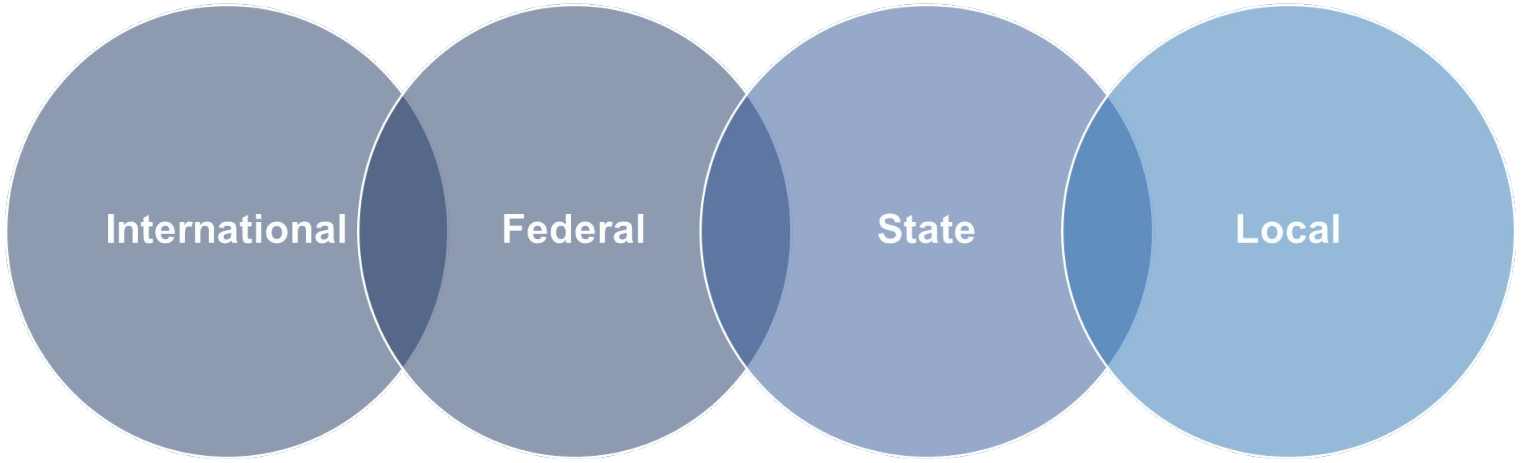

Old laws still apply → New laws are coming

President Trump Signs Executive Order to Advance AI Education in K-12 Schools

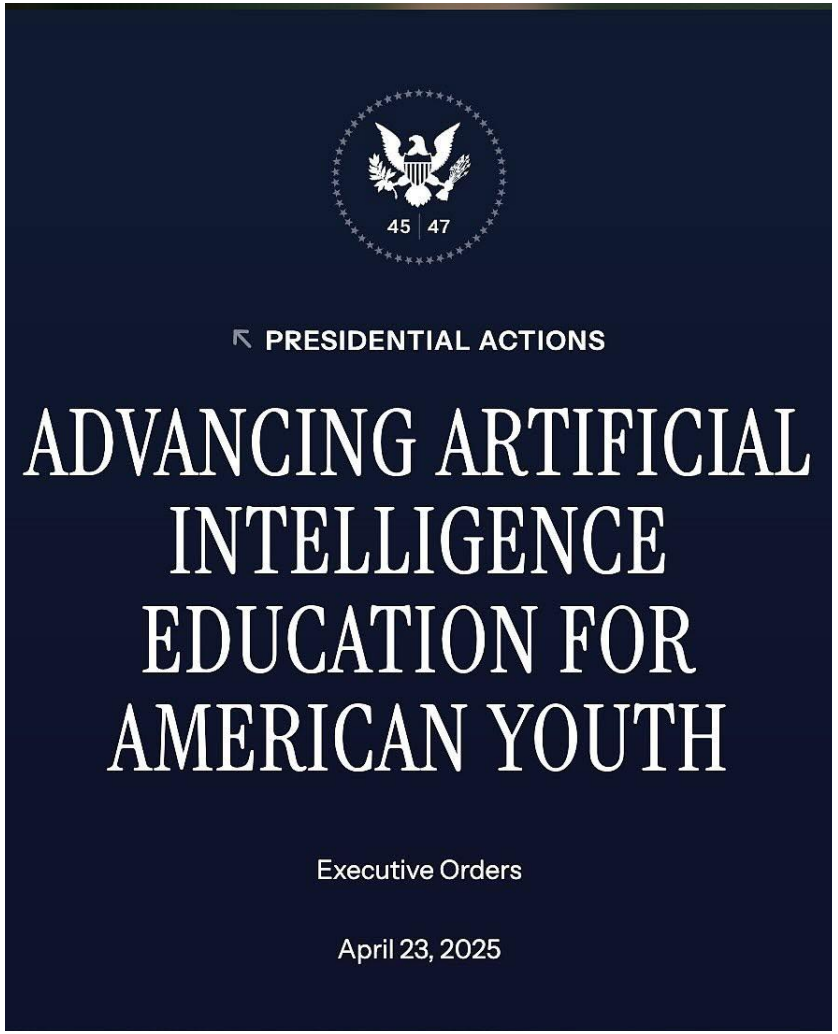

The Order instructs federal agencies to:

- teach students how to use AI
- train teachers to incorporate it into their tasks, and
- partner with the private sector to develop relevant programs in schools.

Family Educational Rights and Privacy Act

Privacy+ Security Forum

What is FERPA?

- Federal student privacy law enacted in 1974
- Applies to schools that receive federal funding from the U.S. Department of Education.
- Gives parents (and eligible students) the right to **access** and **amend** their child's education records.
- Prohibits a school from disclosing **personally identifiable information** from a student's **education record** to a third party without written consent from the **parent or eligible student** (unless an exception applies).
- Requires **reasonable security methods**.

FERPA Enforcement

- U.S. Department of Education enforces FERPA
- Penalties
  - Schools can lose federal funding
  - School can be banned from using a vendor for 5 years

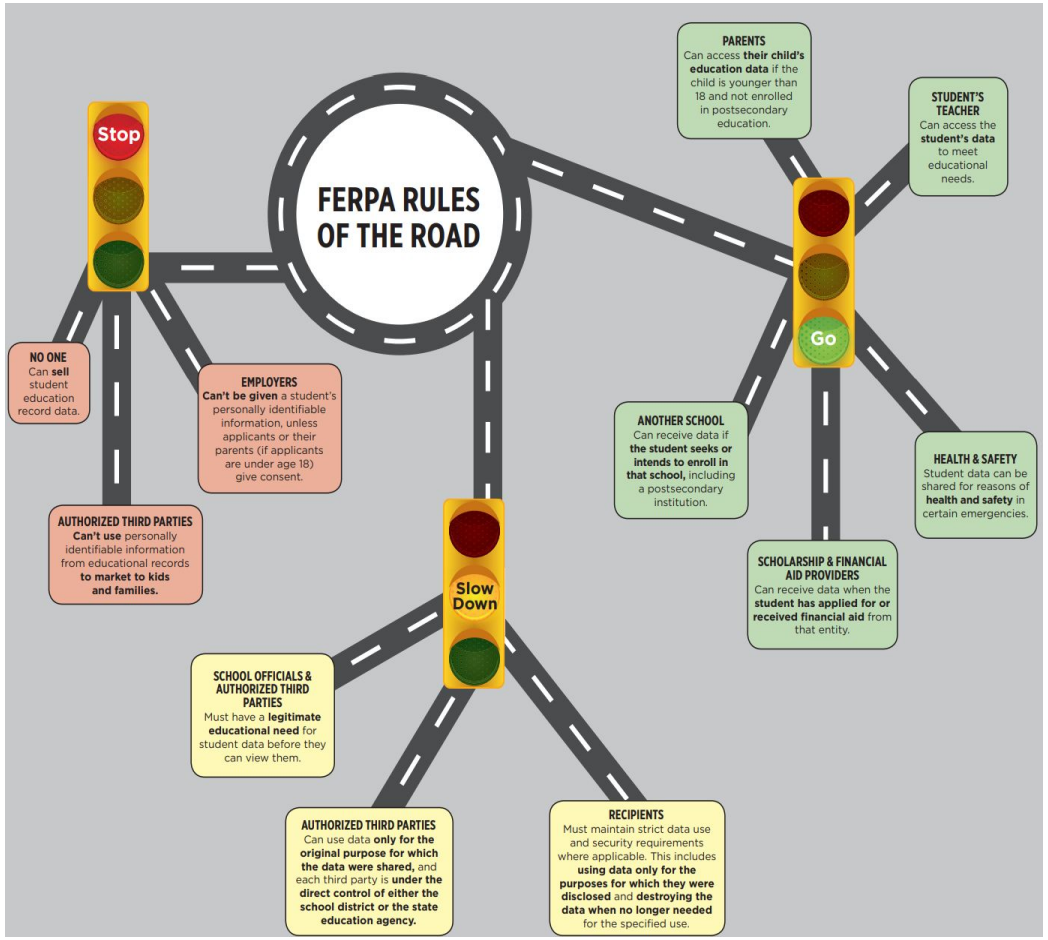

Source: Data Quality Campaign

Top Exceptions to FERPA's Parental Consent Requirements

Privacy+ Security Forum

Education Records:

- Records that are **directly related** to a student and are **maintained by** an educational agency or institution or by a party acting of the educational agency or institution.

Personally Identifiable Information (PII):

- Name and address.
- Parents' or other family members' names and addresses.
- Personal identifier, including a Social Security number or student number.
- Biometric records, including fingerprints, retina, and iris patterns, voiceprints, DNA sequence, facial characteristics, or handwriting.
- Indirect identifiers, including a birth date and location or mother's maiden name.
- **Other information that**, alone or in combination, is **linked or linkable** to a specific student that would allow a **reasonable person in the school community,** who does not have personal knowledge of the relevant circumstance**s, to identify the student with reasonable certainty**
- Information requested by a person who the educational agency or institution reasonably believes knows the identity of the student to whom the education record relates.

What's Not Covered?

Deidentified information –requires the removal of all personally identifiable information and for the educational institution to make "a reasonable determination that a student's identity is not personally identifiable, whether through single or multiple releases, and taking into account other reasonably available

What Type of Data Does FERPA protect?

Privacy+ Security Forum

- Directory information
- School official exception
- Audit and evaluation and exception
- Studies exception
- Health and safety exception
- Law enforcement and subpoena
- Lawsuit

Directory Information

- Information contained in an educational record that would not generally be harmful if disclosed
- **Examples:**
  - Name
  - Address
  - Photograph
  - Phone number
  - Date and place of birth
  - Grade level
  - Dates of attendance
  - Participation in officially recognized sports and activities
  - Degrees, honors and awards received
- Annual notice must be given to parents
- Parents must have a way to **opt out** of sharing directory information.

State Student Privacy Laws

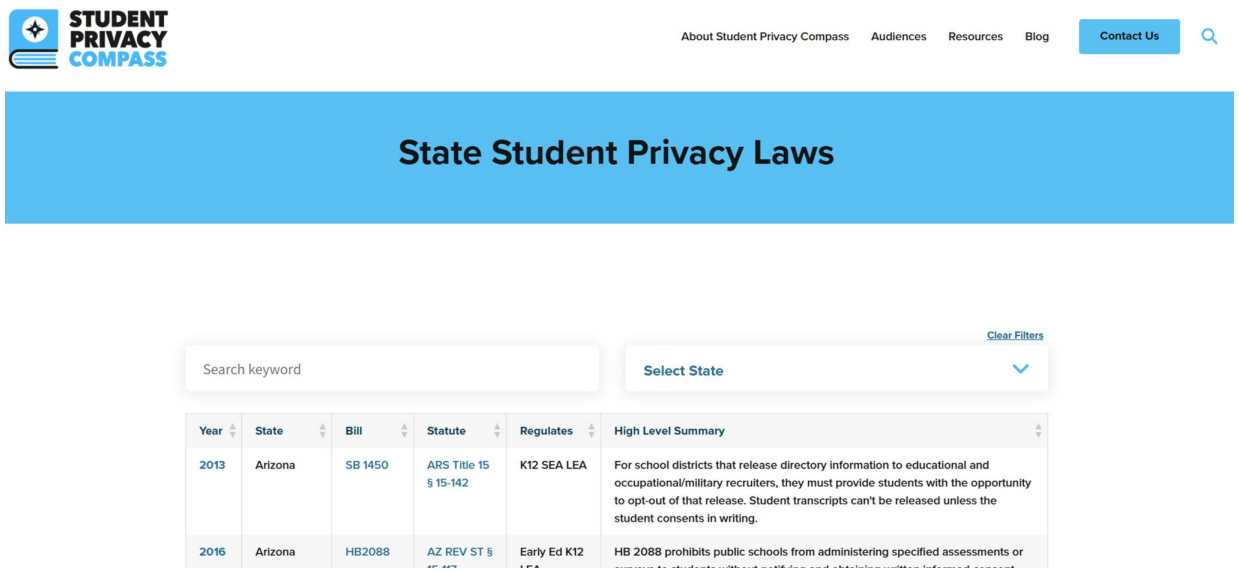

Resource: [State Student Privacy Laws - Student Privacy Compass]

Most states have enacted student privacy laws

- Generally, all impose restrictions on the use and disclosure of student data.

Common Frameworks

- Laws that govern schools
  - E.g., **Illinois School Student Records Act** is similar to FERPA.

Laws that govern vendors

- California's **Student Online Personal Information Protection Act ("SOPIPA")** was the first law to comprehensively address student privacy.
- Unlike FERPA, it imposes direct liability on operators.
- It prohibits Targeting advertising; Creating student profiles except for K-12 purposes; and Selling student information (except for M&A).

- Laws that govern both
  - **New York State Education Law §2-d** ("Ed Law 2D") requires both schools and vendors to comply with the NIST Cybersecurity Framework in an effort to secure student data.

Children's Online Privacy Protection Act

Privacy+ Security Forum

Congress in 1998 enacted COPPA better protect children and their personal information while online.

Applies to "any operator of a website or online service directed to children, or any operator that has actual knowledge that it is collecting personal information **from** a child."

Who is a "child"?

Minor Under the Age of 13.

What is "personal information" in the context of COPPA?

Broadly defined to include persistent identifiers that can recognize users over time and across different websites or online services, photo/video/voice, and precise geolocation

Requirements, subject to certain exceptions, include:

"Notice" of operator data collection, use, and disclosure practices, and

"verifiable parental consent" before the operators collects, uses, or discloses the personal information.

Penalties?

Up to $46,517 per violation

Data destruction

20 year reporting requirements.

State Attorneys General may also enforce the Rule

How Does COPPA Apply in the School Setting?

COPPA allows, but does not require, schools to act as agents for parents in providing consent for the online collection of students' personal information.

The school's ability to consent on behalf of parents is **limited to the educational context**, and no other commercial purposes.

The operator must continue to **meet all other COPPA requirements.**

Best practice: obtain consent from the school, rather than the teacher.

COPPA Enforcement – AI Focus

Privacy+ Security Forum

Why Compliance is Key

Kurbo (formerly Weight Watchers) was required to delete personal information illegally collected from children under 13, **destroy any algorithms derived from the data**, and pay a **$1.5 Million** because of COPPA violations. The FTC's complaint alleged that Kurbo's signup process encouraged users to falsely claim they were over the age of 13.
March 2022

Amazon (Alexa) was required to pay **$25 Million** for violating COPPA by retaining voice recordings of children indefinitely by default and failing to honor deletion requests. Notably, the FTC called out the company for using kids' data to feed its algorithms to aid in developing artificial intelligence.
May 2023

NGL Labs was required to pay **$5 Million**, has been banned from offering its anonymous messaging app to kids under 18, and is **prohibited from making misrepresentations about capabilities of its AI technology**. NGL Labs allegedly violated COPPA and Section 5 of the FTC Act by actively marketed their service to kids despite being aware of the harms from similar services without obtaining verifiable consent and made false claims about their AI content moderation program.
July 2024

FTC Publishes Final COPPA Rule

Privacy+ Security Forum

2000: COPPA Rule in effect

2013: FTC Revised COPPA Rule in Effect

2019: FTC Announced COPPA Rule Review

2023: FTC Proposes Updates to COPPA Rule

2025: FTC Publishes to COPPA Rule

FEDERAL REGISTER
The Daily Journal of the United States Government

Rule

Children's Online Privacy Protection Rule

A Rule by the Federal Trade Commission on 04/22/2025

PUBLISHED DOCUMENT: 2025-05904 (90 FR 16918)

PDF
Document Details
Document Dates
Table of Contents
Public Comments
Regulations.gov Data
Sharing
Print
Document Statistics
Other Formats
Public Inspection

DOCUMENT HEADINGS

Federal Trade Commission
16 CFR Part 312
RIN 3084-AB20

AGENCY:

Federal Trade Commission.

ACTION:

Final rule amendments.

SUMMARY:

The Federal Trade Commission amends the Children's Online Privacy Protection Rule (the "Rule"), consistent with the requirements of the Children's Online Privacy Protection Act. The amendments to the Rule, which are based on the FTC's review of public comments and its enforcement experience, include one new definition and modifications to several others, as well as updates to key provisions to respond to changes in technology and online practices. The amendments are intended to strengthen protection of personal information collected from children, and, where appropriate, to clarify and streamline the Rule since it was last amended in January 2013.

DATES:

Effective date: The amended Rule is effective June 23, 2025.

Compliance date: Except with respect to § 312.11(d)(1), (d)(4), and (g), regulated entities have until April 22, 2026 to comply.

Key Amendments to the COPPA Rule

Privacy+ Security Forum

- Broader Definitions (Personal Information)
- Expanded Factors for Determining When a Website is Directed to Children
- Mixed Audience Clarification
- Stronger Parental Notice and Consent Requirements
- New Verifiable Consent Methods
- Enhanced Security and Data Retention Requirements
- Safe Harbor Programs

Age Appropriate Design Codes

Privacy+ Security Forum

California, Colorado, Connecticut, Florida, and Maryland have enacted AADCs

- Thresholds: Florida, Colorado and Connecticut do not have minimum business thresholds (California and Maryland do).
- Age: Definition of child varies from <13 to <18.
- Harm: Prohibitions on Prohibitions against:
  - using data in a manner that harms the child or is not in the best interests of the child.
  - profiling a child (subject to limited exceptions)
  - the collection, sharing or selling of a child's personal information (subject to limited exceptions in some states)
  - the collection of precise geolocation information (subject to limited exceptions)
  - using dark patterns
- DPIAs: Required in most states.
- Litigation: Many have been challenged and/or enjoined on First Amendment grounds.

Other states have proposed modified AADCs

Illinois	South Carolina	Vermont
Nebraska	Oklahoma	Rhode Island
South Carolina	Vermont	Washington

New York Child Data Protection Act

Privacy+ Security Forum

- The New York Child Data Protection Act (CDPA) prohibits operators of online services from processing the personal data of minors ages 13 – 17 without informed consent from the minor, or unless doing so is strictly necessary for the service. Effective June 20, 2025
- Operators must:
  - Respect age flags—i.e., treat a user as a covered user if the user's device communicates or signals that the user is a minor
  - Enter into a written, binding agreement with third-party operators and processors.
  - Delete personal data within 30 days if the operator learns that such data was improperly collected from a covered user and notify any third-party recipients of that data
- New York's office of attorney general has rulemaking and enforcement authority
- Penalties include disgorgement of unlawfully obtained data and other ill-gotten profits or gains, and up to $5,000 in civil penalties per violation.

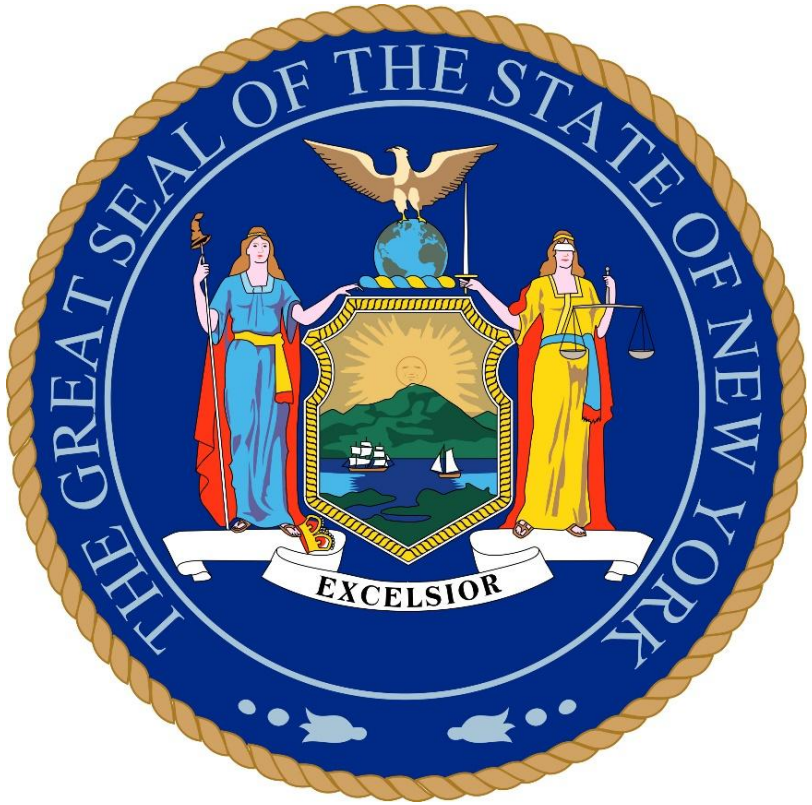

Comprehensive State Privacy Laws with Requirements for EdTech

Privacy+ Security Forum

California

- Must obtain opt-in consent from minors age 13 to 15 and from parents of minors under the age of 13 before selling or sharing their personal information.
- ADMT Regulations Pending

New Jersey

- Must obtain opt-in consent from parents of minors under the age of 13 before selling or sharing their personal information.
- Must obtain consent for minors aged 13 to 16 to sale or use their personal information for targeted advertising or profiling.
-

State AI Laws

Privacy+ Security Forum

Utah AI Act
- Effective May 1, 2024

Connecticut AI Act
- Effective July 1, 2025

Colorado AI Act
- Effective February 1, 2026

California ADMT Regulations
- Effective date: TBD

Note: AI laws like the Colorado AI Act have a challenging undefined language of Educational impact opportunity

Best Practices for EdTech Companies

Privacy+ Security Forum

Best Practices / AI Principles for EdTech

Privacy+ Security Forum

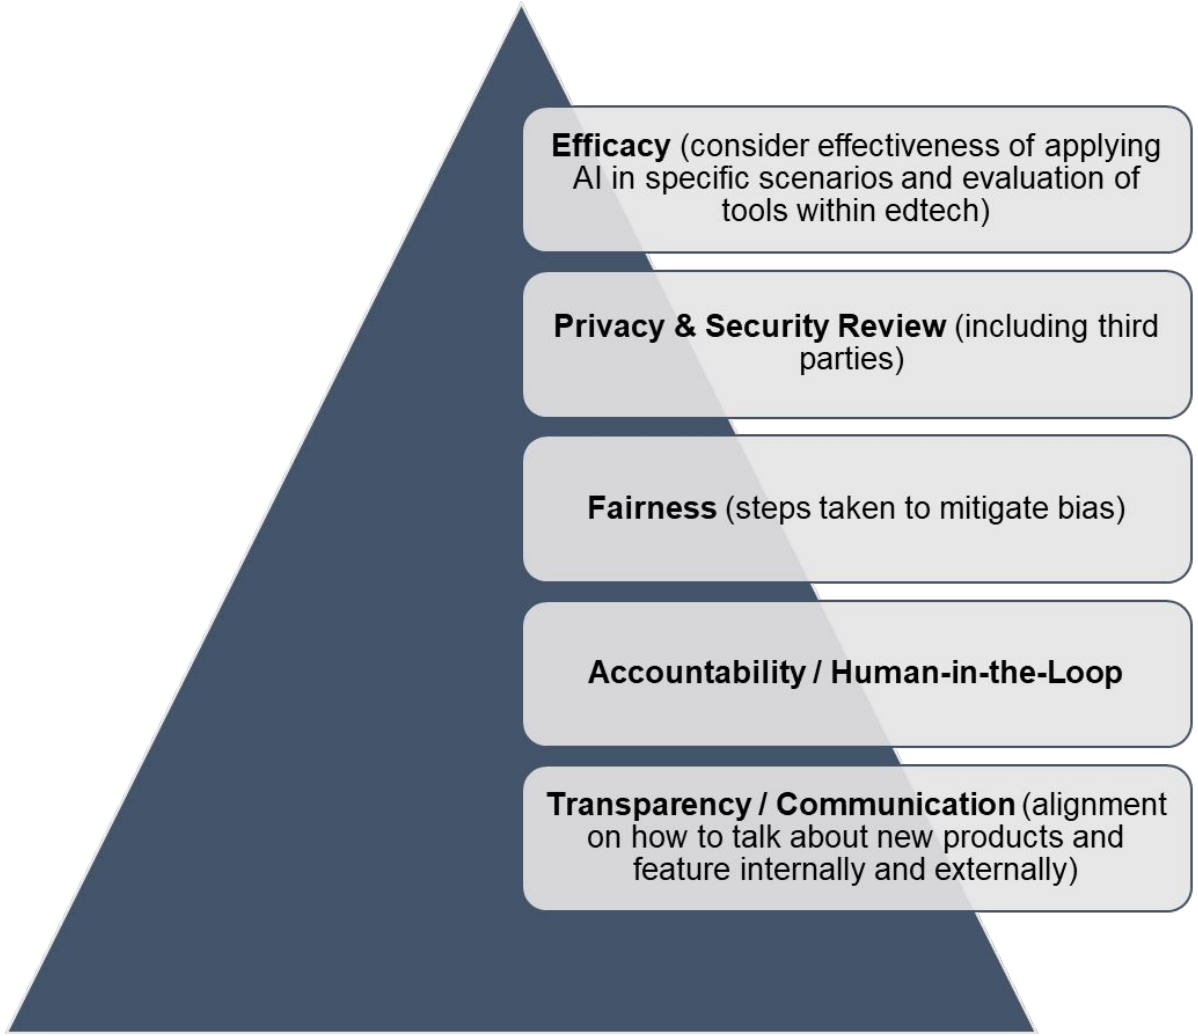

FPF Resource Provides Considerations for AI Use Cases

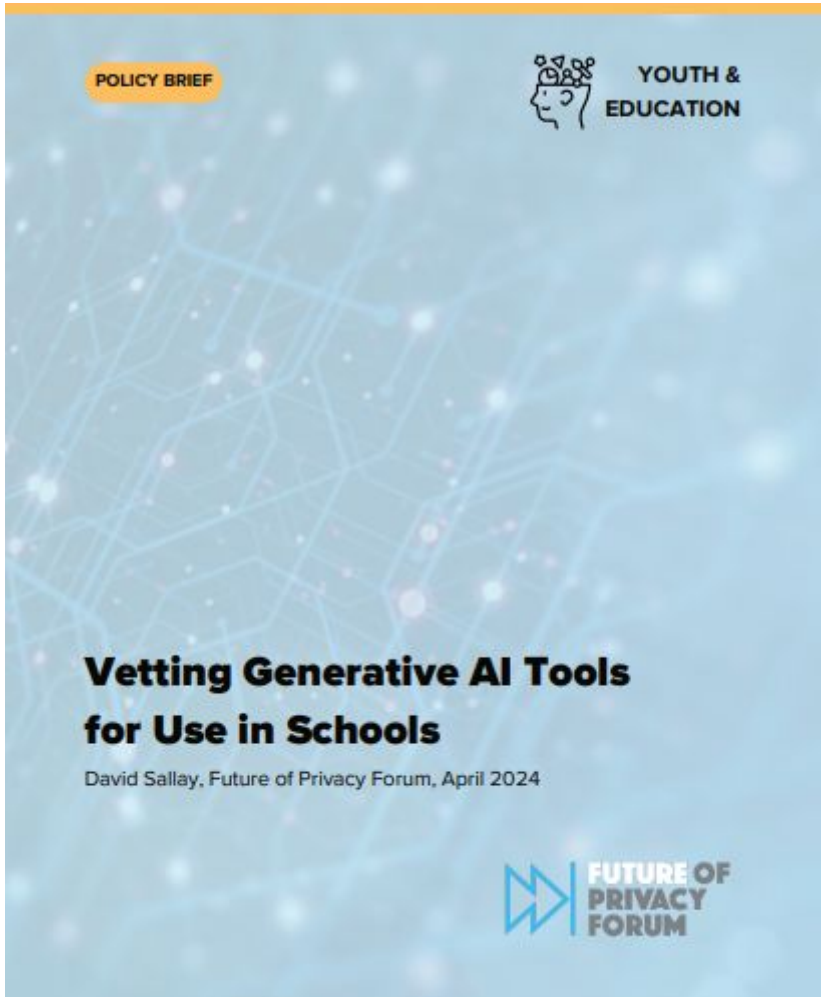

- Does the use case require student PII?
- Is the technology able to meet to meet requirements for transparency and explainability?
- Will the edtech company use student PII for product improvement?
- Will the tool be used for substantive decision making?

Link to Resource:
https://fpf.org/wp-content/uploads/2024/10/Ed_AI_legal_compliance.pdf_Final_OCT24.pdf

Resource: EDSAFE AI SAFE Framework

Privacy+ Security Forum

what is the EDSAFE AI **SAFE** Framework?

EDSAFE AI ALLIANCE

Founded in 2020, the EDSAFE AI Alliance is a global initiative coordinated by InnovateEDU and powered by a coalition of organizations representing stakeholders across the education sector to provide global leadership for developing a safer, more secure, more equitable, and more trusted AI education ecosystem through a focus on research, policy, and practice.

By joining forces and complementing rather than competing with stakeholders in the space, we can address one of the most pressing educational policy challenges of our time: how to build and develop an ecosystem that reflects the best practices for AI use in education.

With a shared mission to leverage AI to facilitate better, more equitable student outcomes, save time for teachers and other educators, and increase efficiencies for those invested in teaching and learning ecosystems, ED SAFE AI is an uncommon alliance dedicated to furthering safe, accountable, fair and transparent, and efficacious AI use within the K-12 education space through the SAFE framework.

We believe that AI should not be used to replace educators in classrooms or in decision-making processes: instead, AI can be used to reduce educator paperwork and help to synthesize and make suggestions based on learner data. In this fashion, AI can be leveraged to allow teachers more time for direct instruction with their learners, using technology as a "co-pilot" that supports teachers in doing the human, emotional, affective work that is at the heart of the learning process.

Led by a Steering Committee of organizations and their leaders who are already working at the intersection of AI and education, EDSAFE AI aims to bring diverse voices to help educate policymakers and shape an equitable and safe AI education ecosystem.

STEERING COMMITTEE

AASA, The School Superintendents Association
AI in Education
aiEDU
American Association of Colleges for Teacher Education (AACTE)
American Federation of Teachers (AFT)
CAST
Code.org
Consortium for School Networking (COSN)
Council of Great City Schools
Data Science 4 Everyone
Digital Promise
Dxtera Institute
EdTrust
Education Counsel
Federation of American Scientists (FAS)
ISTE/ASCD
National Center for Learning Disabilities (NCLD)
National Digital Inclusion Alliance (NDIA)
National Education Association (NEA)
National Parents Union
Opportunity Labs
Project Evident
SETDA
Software & Information Industry Association (SIIA)
The New Teacher Project (TNTP)
Transcend

THE FRAMEWORK

The work of the EDSAFE AI Alliance centers on the SAFE Benchmarks Framework. The framework creates a policy process and roadmap for the essential issues in creating a SAFE AI ecosystem. The framework was built starting in 2021 and brings together more than 24 global AI safety, trust and market frameworks. Frameworks and benchmarks are essential to innovation as a means of targeted guidance, focusing disparate efforts towards shared language, objectives, and outcomes and ensuring the development of appropriate guidelines and guardrails for use. By working together through the Framework, EDSAFE aims to accomplish two things: achieve equitable outcomes for students and improve working conditions for teachers.

EDSAFE AI **SAFE** Framework

S	**SAFETY** Security, Privacy, Do Not Harm
A	**ACCOUNTABILITY** Defining Stakeholder Responsibilities
F	**FAIRNESS** Equity, Ethics, and Mitigating Bias
E	**EFFICACY** Improved Learning Outcomes

SAFETY

Safety is a primary consideration for ensuring edtech users can be active in a digital environment that prioritizes protecting their data and privacy while managing cybersecurity risks. At the same time, solution providers have a shared commitment to responsibly building innovative education solutions. Solution providers must be able to continue developing and deploying tools and the evolution of their product roadmaps while responsibly building innovative learning solutions.

ACCOUNTABILITY

Accountability is a cornerstone for establishing benchmarks that are collaboratively defined by a diverse group of constituents, encompassing subject matter experts, edtech solution providers, educators from Pre-K through higher education, and learners of all ages. This approach ensures the creation of standards that not only foster accountability but also bolster it through the integration of pertinent existing policies and regulatory undergirding. Concurrently, all parties involved are dedicated to the transparent development and implementation of these standards, ensuring they are in line with the evolving educational landscape and effectively address the needs and expectations of all stakeholders.

FAIR and TRANSPARENT

Achieving ethical, unbiased, and equitable learning opportunities necessitates a conscious and deliberate effort from both solution providers and users to scrutinize the quality of data being utilized, including the acquisition of datasets, the application of their products, and the monitoring for any inadvertent biases. AI products and experiences must be accessible for all individuals and there needs to be strategies for creating standards, guidelines, and/or quality indicators to gauge this accessibility. This endeavor extends to ensuring fairness and transparency, particularly in the procurement of materials produced by AI outputs, underlining the importance of vigilance and intentionality in every aspect of educational technology engagement.

EFFICACY

Deliberations are essential regarding the effectiveness of applying AI in specific scenarios, coupled with the commitment by edtech solution providers to integrate comprehensive and transparent evaluation tools within educational technologies. Such integration aims to precisely measure advancements and provide educators and learners with insightful feedback on usage and progress. Efficacy is understood to be deliberately tied to an equity in student experiences as well as outcomes. This approach underscores the significance of carefully considering the utility of AI and the importance of inclusivity and clarity in the mechanisms used to gauge and communicate educational outcomes.

Link to Resource:
https://www.edsafeai.org/_files/ugd/5be6a9_0dffff673cd042578c25cc098b2929fc.pdf

AI Transparency Example: McGraw Hill AI Reader

Privacy+ Security Forum

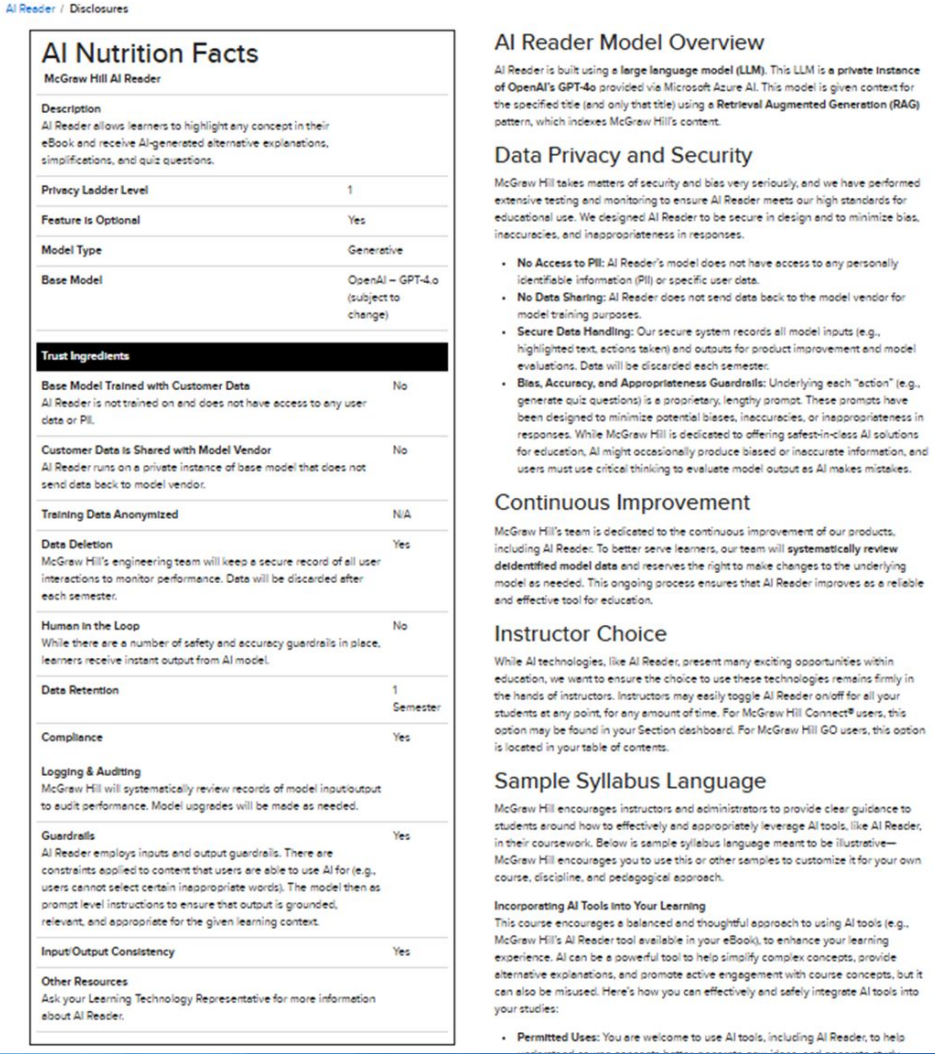

AI Reader / Disclosures

AI Nutrition Facts

McGraw Hill AI Reader

Description AI Reader allows learners to highlight any concept in their eBook and receive AI-generated alternative explanations, simplifications, and quiz questions.	
Privacy Ladder Level	1
Feature is Optional	Yes
Model Type	Generative
Base Model	OpenAI – GPT-4.o (subject to change)
Trust Ingredients	
Base Model Trained with Customer Data AI Reader is not trained on and does not have access to any user data or PII.	No
Customer Data is Shared with Model Vendor AI Reader runs on a private instance of base model that does not send data back to model vendor.	No
Training Data Anonymized	N/A
Data Deletion McGraw Hill's engineering team will keep a secure record of all user interactions to monitor performance. Data will be discarded after each semester.	Yes
Human in the Loop While there are a number of safety and accuracy guardrails in place, learners receive instant output from AI model.	No
Data Retention	1 Semester
Compliance	Yes
Logging & Auditing McGraw Hill will systematically review records of model input/output to audit performance. Model upgrades will be made as needed.	
Guardrails AI Reader employs inputs and output guardrails. There are constraints applied to content that users are able to use AI for (e.g., users cannot select certain inappropriate words). The model then as prompt level instructions to ensure that output is grounded, relevant, and appropriate for the given learning context.	Yes
Input/Output Consistency	Yes
Other Resources Ask your Learning Technology Representative for more information about AI Reader.	

AI Reader Model Overview

AI Reader is built using a **large language model (LLM)**. This LLM is **a private instance of OpenAI's GPT-4o** provided via Microsoft Azure AI. This model is given context for the specified title (and only that title) using a **Retrieval Augmented Generation (RAG)** pattern, which indexes McGraw Hill's content.

Data Privacy and Security

McGraw Hill takes matters of security and bias very seriously, and we have performed extensive testing and monitoring to ensure AI Reader meets our high standards for educational use. We designed AI Reader to be secure in design and to minimize bias, inaccuracies, and inappropriateness in responses.

- **No Access to PII:** AI Reader's model does not have access to any personally identifiable information (PII) or specific user data.
- **No Data Sharing:** AI Reader does not send data back to the model vendor for model training purposes.
- **Secure Data Handling:** Our secure system records all model inputs (e.g., highlighted text, actions taken) and outputs for product improvement and model evaluations. Data will be discarded each semester.
- **Bias, Accuracy, and Appropriateness Guardrails:** Underlying each "action" (e.g., generate quiz questions) is a proprietary, lengthy prompt. These prompts have been designed to minimize potential biases, inaccuracies, or inappropriateness in responses. While McGraw Hill is dedicated to offering safest-in-class AI solutions for education, AI might occasionally produce biased or inaccurate information, and users must use critical thinking to evaluate model output as AI makes mistakes.

Continuous Improvement

McGraw Hill's team is dedicated to the continuous improvement of our products, including AI Reader. To better serve learners, our team will **systematically review deidentified model data** and reserves the right to make changes to the underlying model as needed. This ongoing process ensures that AI Reader improves as a reliable and effective tool for education.

Instructor Choice

While AI technologies, like AI Reader, present many exciting opportunities within education, we want to ensure the choice to use these technologies remains firmly in the hands of instructors. Instructors may easily toggle AI Reader on/off for all your students at any point, for any amount of time. For McGraw Hill Connect® users, this option may be found in your Section dashboard. For McGraw Hill GO users, this option is located in your table of contents.

Sample Syllabus Language

McGraw Hill encourages instructors and administrators to provide clear guidance to students around how to effectively and appropriately leverage AI tools, like AI Reader, in their coursework. Below is sample syllabus language meant to be illustrative—McGraw Hill encourages you to use this or other samples to customize it for your own course, discipline, and pedagogical approach.

Incorporating AI Tools into Your Learning

This course encourages a balanced and thoughtful approach to using AI tools (e.g., McGraw Hill's AI Reader tool available in your eBook), to enhance your learning experience. AI can be a powerful tool to help simplify complex concepts, provide alternative explanations, and promote active engagement with course concepts, but it can also be misused. Here's how you can effectively and safely integrate AI tools into your studies:

- **Permitted Uses:** You are welcome to use AI tools, including AI Reader, to help

Link to Resource:
https://www.mheducation.com/highered/digital-products/ai/disclosures