


A report from

cdt | Research



Opening the Book

**A Rubric to Support Effective
Transparency for EdTech
Products That Incorporate AI**

**Hannah Quay-de la Vallee
Morgan Badurak
Dhanaraj Thakur**

November 2025



The Center for Democracy & Technology (CDT) is the leading nonpartisan, nonprofit organization fighting to advance civil rights and civil liberties in the digital age. We shape technology policy, governance, and design with a focus on equity and democratic values. Established in 1996, CDT has been a trusted advocate for digital rights since the earliest days of the internet. The organization is headquartered in Washington, D.C., and has a Europe Office in Brussels, Belgium.

HANNAH QUAY-DE LA VALLEE

Senior Technologist at the Center for Democracy & Technology

MORGAN BADURAK

Ph.D. Student at Louisiana State University

DHANARAJ THAKUR

Research Director at the Center for Democracy & Technology



Opening the Book

A Rubric to Support Effective Transparency for EdTech Products That Incorporate AI

Hannah Quay-de la Vallee, Morgan Badurak,
and Dhanaraj Thakur

Suggested Citation: Quay-de la Vallee, H., Badurak, M. and Thakur, D. (2025) *Opening the Book: A Rubric to Support Effective Transparency for EdTech Products That Incorporate AI*. Center for Democracy & Technology. <https://cdt.org/insights/opening-the-book-a-rubric-to-support-effective-transparency-for-edtech-products-that-incorporate-ai/>

References in this report include original links or links archived and shortened by the Perma.cc service. The Perma.cc links also contain information on the date of retrieval and archive.



This report is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

Contents

Executive Summary	5
Introduction	8
A Rubric for Comprehensive Transparency of EdTech and AI	10
Use and Context Limitations	10
Underlying AI Models	12
Training Data and Methodologies	14
Domain Adaptation	15
Testing and Evaluation	17
Data Governance	19
Governance Structures	21
Information Accessibility	22
Transparency in EdTech Products That Use AI: Key Findings	24
Conclusion	28
Appendix A: Comprehensive Transparency Rubric of EdTech and AI	29
Appendix B: Method and Procedure	32
References	34

Executive Summary

The use of artificial intelligence (AI) in educational technology (edtech) in the K–12 space has grown significantly ([Grand View Research, 2025](#)), and schools are being encouraged to adopt AI by a range of actors from states ([Dwyer, 2025](#)) to the federal government ([White House, 2025](#)) to industry ([Gulezian, 2025](#); [Microsoft, 2025](#)). Used well, AI may provide significant educational benefits. But the use of AI in a high-stakes context such as K–12 education also raises concerns about how well these systems will work, whether they will be worth the investment, whether they will introduce bias or other harms, and more ([Quay-de la Vallee, 2025](#)). One commonly recommended, broadly endorsed approach to support the responsible use of edtech products that incorporate AI is transparency from the developers and designers of these systems.¹

Transparency about edtech products that incorporate AI can help school administrators (anyone at a school or local education agency who is making decisions about whether and how AI is used) select systems that will provide them with the functionality they need; ensure that those systems are suited for their contexts; identify potential risks and ways to mitigate them; and push the field at large toward improved development practices, outcomes, equity, and more. Recognizing these potential benefits, different edtech industry associations have begun pushing their members to commit to transparency about the use of AI in their products. However, little information is available about what constitutes meaningful transparency in this context and to what extent edtech companies are meeting the needs of schools, educators, students, and parents.

To better assess the current state of transparency in edtech products in the K–12 context that incorporate AI, the Center for Democracy & Technology (CDT) has developed a rubric of eight key elements of transparency that contribute to a full picture of an edtech product and can be used to inform decisions such as whether to procure the product and how to use it:

- **Use and Context Limitations:** AI systems are often designed for a specific context or task, and using them outside these confines can lead to unexpected and potentially harmful behavior. Being able to assess these parameters is critical for educators to determine which tool to select and the appropriate confines and safeguards needed.

¹ While transparency plays a critical role in supporting informed decision-making, it should not be the end of the process for addressing potential risks in edtech. As the Center for Democracy & Technology has extensively documented, certain products that incorporate AI raise significant risks related to privacy, accuracy, and student discipline that should be seriously evaluated before and during their deployment. See, e.g., [Quay-de la Vallee, 2025](#).

- **Underlying AI Models:** Many edtech products that incorporate AI are built on top of pre-existing models. The strengths and weaknesses of an underlying AI model may be present in the end product, contributing to its efficacy and risks. Knowing which model was used and how it affects the final tool is an important part of assessing how the tool will work in a given deployment.
- **Training Data and Methodologies:** How a model is trained, including the data used, affects how an AI system performs as the model will reflect any underlying assumptions, biases, gaps, and data quality issues. Decision-makers will need to assess whether the training data and methods make the system likely to perform well for their specific needs.
- **Domain Adaptation:** For products incorporating a pre-existing model, fine-tuning the system for an education context can help ensure that the product is effective (e.g., results and outputs reflect best practices in teaching and learning, not just a synthesis of any available information on a given topic) and mitigate risks that are of particular concern for educational use. As a result, decision-makers need to know what adaptations were made to assess if those changes are sufficient for their needs.
- **Testing and Evaluation:** Testing and evaluating products that incorporate AI is important to ensure that they are performing as expected, producing the intended results, and not causing unintended consequences. Providing decision-makers with access to a description of the tests and evaluations performed and to the results of these analyses supports informed choices and effective product adoption.
- **Data Governance:** AI systems consume and produce considerable data that, if not protected, can place students and schools at risk of privacy and other harms, so understanding how data is kept secure and managed appropriately is crucial to not only fulfilling legal obligations but also keeping students safe.
- **Governance Structures:** Governing AI systems is a complex process, incorporating steps such as determining data use policies, ensuring that staff are trained, communicating with stakeholders, and more. Companies should have frameworks in place for assessing their products from an ethical standpoint and ensuring that their products are used in an ethical manner. Companies should be clear about how they make these decisions, including who is involved and what decision-making structures they have in place, to ensure that their systems and data are appropriately managed.
- **Information Accessibility:** The previous elements are critical for decision-makers trying to select appropriate AI systems, and the information needs to be made available to them in an accessible and comprehensive way, such as by ensuring that the information is available in one prominent place; is written in plain, accessible language; and is well organized.

Using this rubric, CDT reviewed publicly available information for more than 100 edtech companies and found that:

- On average, edtech companies offer little transparency about their products, with companies receiving an average transparency score of 4 out of a possible 16.
- When edtech companies attempt to be transparent, the information is often focused on select categories, namely Use and Context Limitations and Information Accessibility.
- Companies with middling transparency make information about AI easy to access.

A lack of robust transparency from edtech vendors hinders the ability of educators and educational institutions to find effective, safe, and equitable AI-based edtech products. The rubric CDT sets out in this report provides guidance to vendors on what information they should provide about their products that incorporate AI, as well as what school administrators should demand and expect before they purchase and use such products.



Introduction

Educational technology (edtech) in the K–12 space has increasingly incorporated artificial intelligence (AI) in recent years, and the trend is likely to continue ([Grand View Research, 2025](#)). This shift has been encouraged by actions from both the federal and state governments ([Dwyer, 2025](#); [White House, 2025](#)) and from tech companies ([Gulezian, 2025](#); [Microsoft, 2025](#)). These supporters argue that AI can improve the quality and availability of K–12 education, but the use of AI also carries risks. In a high-stakes context such as education, failures in AI systems can carry significant consequences including inequitable treatment of students, wasted resources, loss of community trust, and more ([Quay-de la Vallee, 2025](#)).

Transparency from the developers and designers of these systems can help school administrators better understand the AI-based products that are available, including the potential benefits and risks of those products. With this understanding, they can select the products that are best suited to their needs and effectively manage potential risks ([Office of Educational Technology, 2023](#); [Zhang & Aslan, 2021](#)). In some cases, transparency can also help parents and students better understand how edtech products that incorporate AI are used in their schools and how these products affect their education. Additionally, researchers argue that transparency can improve the edtech field as a whole, supporting improved development practices, outcomes, equity, and more ([Vincent-Lancrin & van der Vlies, 2020](#)).

To be sure, transparency is only one component of an overall AI management framework ([Laird et al., 2025a](#); [Laird et al., 2025b](#); [Quay-de la Vallee, 2025](#); [Quay-de la Vallee & Dwyer, 2023](#)). It alone cannot mitigate risks ranging from privacy to inaccuracy to inequity. Nevertheless, recognizing the potential benefits, many edtech companies have committed to being transparent about the use of AI in their products ([EDSAFE AI, 2025b](#); [Software & Information Industry Association, 2023](#)).

This report examines what transparency in AI-based edtech products looks like in practice. Based on existing research and emerging best practices, it sets forth a multi-faceted eight-part rubric of what constitutes meaningful transparency ([Anik & Bunt, 2021](#); [Bommasani et al., 2023](#); [Chmielinski et al., 2024](#); [King & Meinhardt, 2024](#); [Kucirkova et al., 2023](#)). The rubric does not aim to evaluate the *content* of the information available but focuses solely on the completeness and availability of the given categories of information. So, for example, it allows for the evaluation of whether a company provides information

about its data governance policies but does not assess whether those policies themselves are robust or effective.

The rubric was used to assess publicly available information from more than 100 edtech products incorporating AI that were targeted for K–12 uses. This examination revealed that, despite numerous calls, pushes, and pledges focused on transparency in AI-based edtech, the state of play in the field leaves much to be desired, with most companies providing very limited information about their products.

This lack of transparency generally leaves school administrators without the information they need to evaluate if AI is appropriate for their purposes, to select tools that are appropriate for their needs and context, and to adequately prepare for the impacts of incorporating a given tool. It also limits the ability of school administrators to help parents and students understand the nature and role of these technologies in schools. Additionally, limited transparency restricts the ability of experts to assess the field as a whole, spot areas of concern, and push for improvements or best practices that could result in better outcomes across the board ([Corrigan, 2025](#)).

Although the overall status of transparency of edtech tools that incorporate AI currently has significant gaps, vendors can use the rubric created for this assessment as a guide to the information they should be providing to fulfill their commitment to implement meaningful transparency. School administrators can use it as a framework for what information they should demand before purchasing and using edtech tools that incorporate AI.



A Rubric for Comprehensive Transparency of EdTech and AI

Effective transparency for AI-based edtech tools requires robust information across a number of categories. The following rubric, based on existing research and emerging approaches to transparency, describes how information across eight dimensions of AI can aid decisions such as whether to purchase an AI product and how to best use it.

Use and Context Limitations

Types and uses of AI vary widely, as do the contexts in which it is deployed. Vendors and developers should be able to clearly articulate what educational context they expect their tool to be deployed in, what it should be used for, and tasks for which it is effective ([Chmielinski et al., 2024](#)). Providing this information proactively helps prevent shoehorning a product's capabilities into the school's described context and needs. Vendors should also provide information about the limitations of their system: situations where it should *not* be used and tasks it *cannot* perform. *Figure 1* provides an example of the types of information that might be disclosed under this category.

Using AI to Support Student Creativity and Communication

Wixie takes advantage of AI and advanced predictive algorithms to support and enhance a student's multimodal communications.

Natural Language Processing



Voice dictation

When students use voice dictation, machine language pulls from a range of data to produce actual English words (and Spanish words if your OS is set to Spanish). If a device is user-specific, dictation becomes even more accurate as the platform learns how the user speaks.



Animated, talking characters

Wixie's animated, talking images support AI-powered lip-syncing to generate visemes from text or voice recordings.

Machine Learning



Interpreting search requests

When young learners search for images, Wixie uses advanced algorithms to determine the best response based both on actual typing as well as what the user's input is likely to mean. For example, searching "jrfe" will return giraffe images.



Removing backgrounds

Removing backgrounds from images is done using AI machine learning to identify human figures in images and turning the rest of the image transparent so that it can be placed over other images.

Machine Translation



Translating text

Wixie uses neural machine translation, to translate text, instructions, and entire projects from one language to another, providing users the opportunity to read information in their native language.

Wixie is a tool for students to demonstrate their learning by combining writing with images, voice narration, video and more. Tech4Learning works hard to ensure that Wixie's AI implementations do not do a student's work for them, but serve to make multimodal communication accessible to more users, easier to achieve, and enhance their ability to effectively share their ideas and learning.

There will likely be many more AI features in Wixie, but from our design and development perspective, they have to match the purpose of Wixie, which is to help young learners build and create as they learn to read, write, and build math foundations.

◀ **Figure 1.** Example of a page documenting the specific tasks that AI performs (for example, removing backgrounds) and how those tasks fit into the vendor's product (identifying and isolating human figures so they can be placed over other images)

Source: (wixie.com — captured Sept. 22, 2025)

Underlying AI Models

Edtech products are often built on top of pre-existing AI models. For example, off-the-shelf large language models such as OpenAI's ChatGPT, Anthropic's Claude, and Google's Gemini may be embedded into other tools because they are time and data intensive to develop and maintain. Therefore, using these large language models instead of creating one that is specific to the education sector frequently makes more sense for efficiency and performance reasons for an edtech vendor. However, as discussed in more detail in the box on *Transparency for General Purpose Models Used in the Classroom* on page 23, school administrators need to understand the functionality and limitations of underlying AI models so they can consider how that system may affect the tool's operation in their context (Winecoff & Bogen, 2024).

The vendor should share what systems its product is built on top of and provide relevant information about these base systems to enable school administrators to consider how the choice of an underlying AI model may affect the product's efficacy and any risks it may introduce. For instance, an underlying model trained on a wide corpus of information may be more accurate in some contexts than a model with a narrower training dataset, but it may also be more likely to provide inappropriate information in its outputs. Decision-makers can use this information to assess the different risks a given underlying model may create and how manageable or concerning the risks are in their particular deployment context. *Figure 2* provides an example of describing the underlying AI models in an edtech tool, including proactively addressing common concerns about underlying models such as the privacy risks that stem from sharing any student information with another party. If relevant, the vendor should explain why it chose the base model it did, such as if the base model performs well against best practices in learning science or allows for more privacy-preserving behavior.

As many decision-makers considering whether and how AI should be used in the classroom are not technical experts in AI, information about the base system should be easy to understand, including information that otherwise is written for a technical audience such as the model cards (a tool providing information about a model such as its intended deployment context and performance against relevant benchmarks) (Hugging Face, 2025; Mitchell et al., 2019). If the information is not accessible to a layperson, the vendor should provide summaries of the most pertinent information for its customers.

Key principles for usage of AI in Kahoot!

Kahoot! use public AI-models from third party vendors and does not build its own large language models (LLMs). **No data from the Kahoot! platform is used to train or refine any models.**

Kahoots!'s main partner and LLM provider for its AI features is Microsoft Azure based in Europe. Kahoot! has integrated with Microsoft through our backend to avoid unknown information from the browser or client being sent to Microsoft.

Data inserted into AI features in Kahoot! is sent to Microsoft for processing and to generate AI output. The information Kahoot! sends to the APIs of Microsoft is what we call the prompt. Kahoot! does not knowingly provide any user information, personal data or previous prompts to Microsoft in providing its AI features and advises users against providing the same in prompts or upload files.

The processing of the information sent by Kahoot! to Microsoft is performed according to Microsoft's standard terms and conditions, which may be found [here](#).

Microsoft describes how it processes, used and stores the data it receives here:
<https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy>

In order to ensure high quality content on our platform, Kahoot! has elected not to opt out of Azure OpenAI Abuse Monitoring. More info on how Microsoft performs abuse monitoring can be found here:
<https://learn.microsoft.com/en-us/azure/ai-services/openai/concepts/abuse-monitoring>

Without warranting the content of the information, Kahoot! assumes that this information is correct and accurate also regarding customer data transferred by Kahoot! to Microsoft for processing:

Important

Your prompts (inputs) and completions (outputs), your embeddings, and your training data:

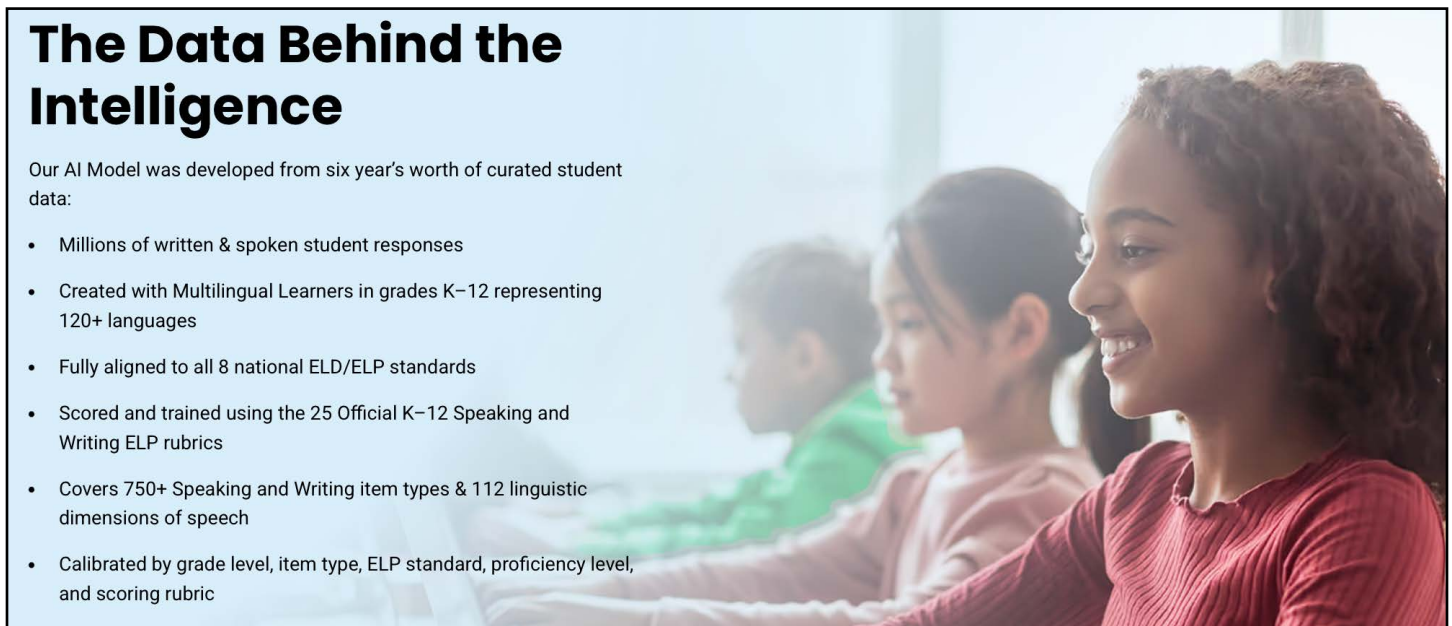
- are NOT available to other customers.
- are NOT available to OpenAI.
- are NOT used to improve OpenAI models.
- are NOT used to improve Microsoft or 3rd party product services.
- are NOT used for automatically improving Azure OpenAI models for your use in your resource (The models are stateless, unless you explicitly fine-tune models with your training data).
- Your fine-tuned Azure OpenAI models are available exclusively for your use.

▲ **Figure 2.** Example of information about the models used and a description of data processing

Source: (kahoot.com — captured Sept. 22, 2025)

Training Data and Methodologies

Vendors who build their own AI models should provide information about how their system was developed or trained (such as whether it used a supervised learning approach in which humans provided feedback to the model as it was trained or an unsupervised approach in which the model was not given feedback from humans during the training process). Vendors should also disclose the training data they relied on to develop their model or relevant characteristics of that data such as the size of the dataset, demographic breakdowns of data subjects, etc. (Anik & Bunt, 2021). This information can be valuable to compare systems or for third parties providing guidance to customers as it speaks to when the system will perform best and the types of challenges it might face. *Figure 3*, for example, highlights relevant aspects of the training data such as that it came from students specifically.



The Data Behind the Intelligence

Our AI Model was developed from six year's worth of curated student data:

- Millions of written & spoken student responses
- Created with Multilingual Learners in grades K–12 representing 120+ languages
- Fully aligned to all 8 national ELD/ELP standards
- Scored and trained using the 25 Official K–12 Speaking and Writing ELP rubrics
- Covers 750+ Speaking and Writing item types & 112 linguistic dimensions of speech
- Calibrated by grade level, item type, ELP standard, proficiency level, and scoring rubric

▲ **Figure 3.** Example of disclosures related to training data and methodologies that are located on a dedicated AI webpage describing how the model was trained

Source: (summitk12.com — captured Sept. 22, 2025)

Domain Adaptation

Whether an edtech tool develops its own AI model or incorporates an underlying AI model, the education sector has unique opportunities and challenges that would benefit from additional tailoring and adaptation. Developers should specify what steps they have taken to adapt the base system to address both functionality and risks related to the education context, including the training data and methodologies they used as described in the previous section ([Chmielinski et al., 2024](#)). This information allows school administrators to assess whether the tool accounts for the nuances of the education context, including both specific educational needs such as adherence to rubrics or standards and compliance with education-specific legal frameworks such as the Family Educational Rights and Privacy Act (FERPA) or the Individuals with Disabilities Education Act. This adaptation may be through adjustment of the model itself (fine-tuning) or through cabining around the model (such as by placing restrictions on responses or askable questions for a generative AI tool). *Figure 4* provides an example of adapting a model by instructing it to behave in certain ways, such as producing short sentences in its responses to meet the needs of younger users.

Safety

Alegra Learning's latest AI-enabled activity—Kady Chain—leverages the API from OpenAI, the team that pioneered the LLM (large language model) known as ChatGPT. Joy School English users are the beneficiaries of extensive guardrails pre-built into the chatbot, including:

- **Content Filtering**

ChatGPT has been trained on a wide range of text, but it includes filters and content moderation systems that help detect and avoid generating harmful, offensive, or inappropriate content. If a query or response is flagged as potentially problematic, the system is designed to avoid it or provide a neutral, non-offensive answer.

- **Ethical Guidelines**

The training of ChatGPT includes ethical guidelines that discourage the generation of content that is violent, sexually explicit, hateful, or otherwise harmful. The model is designed to prioritize respect, inclusivity, and appropriateness in all interactions.

- **User Reports and Feedback**

If an inappropriate response does slip through, users can report it, and such feedback is used to improve future versions of the model. This helps in continuously refining the system's safety measures.

- **Prompt Reframing**

The model is designed to reframe or redirect conversations that may veer into sensitive or inappropriate topics. Instead of directly addressing certain requests, it may steer the conversation in a more constructive direction.

- **Inherent Safeguards**

During the training process, data containing offensive or inappropriate content is minimized to reduce the likelihood of the model producing such responses. Additionally, ongoing updates and improvements to the model aim to enhance its ability to handle sensitive topics responsibly.

In addition to the above, Alegra Learning filters the content even more in Kady Chain, by providing guidelines to the AI bot which it then adheres to in its conversations with students. At present, we instruct it to speak as it would to a 5 year old, using simple language, short sentences, and short responses. This demographic-directed concision makes it much more likely that Kady's responses aren't able to stray into topical areas of impropriety or objectionable content. And by studying and observing how it is being used, we plan to continually adjust those restrictions and guidelines by adding to them and widening the filtering net to be even more effective.


▲ **Figure 4.** Example of domain adaptation including a listing of the guardrails of the underlying AI system and the steps taken to make the model more suitable to younger users, such as by directing it to produce short sentences and use simple language

Source: (joyschoolenglish.com — captured Sept. 25, 2025)

Testing and Evaluation

To enable school administrators to evaluate the efficacy of the tool, the vendor should provide information about how it has tested its product (or had the product tested by a third party) pre-deployment for both efficacy and risk mitigation or disclose if it has not rigorously tested its products. This information allows decision-makers to see what risks and challenges have been accounted for and addressed and to determine whether any concerns they have about their specific context were considered ([Bommasani et al., 2023](#)). Vendors should include the findings of that testing, including adjustments they have made to their products in response to test results; whether their testing is performed entirely in house or through external firms; and how they monitor for performance and other concerns post-deployment (including how they collect data to enable such monitoring while preserving user privacy and how frequently post-deployment monitoring is conducted).

Companies that make specific claims about outcomes driven by the use of their tools should provide evidence of those claims so that school administrators can have confidence in the claims, as seen in *Figure 5*, which provides a dedicated page for the findings of evaluations of the StudyFetch tool ([Kucirkova et al., 2023](#)). This information should include the testing methodologies used and the context in which the testing was done. Additionally, the vendor should provide information about when the testing was done and how it re-evaluates when it updates its system.



▲ **Figure 5.** Example of evaluation and testing information that is located on a dedicated research webpage and provides data-driven reports and results from evaluations and case studies

Source: (studyfetch.com — captured Sept. 25, 2025)

Data Governance

AI systems both rely on and generate large amounts of data that vendors must govern and manage. In the education sector, this data can include sensitive information about students that also requires robust legal and policy protections. Therefore, transparency around data related to the operation of an edtech product that uses AI is important to enable schools to protect students online in the same way they protect them in the classroom (King & Meinhardt, 2024). Important information about data governance and management that vendors should disclose includes:

- What data the system collects or uses, how that data is used, how it is stored and secured, and other governance elements such as retention schedules, access control policies, etc. (See *Figure 6*, which lays out the type of information that is used by the AI tools and notes security practices such as encrypting information at rest.)
- Whether the data fed into the system is stored, shared or sold, or reused by the vendor and, if so, what for and in what form (anonymized, deidentified, aggregated, etc.). (See *Figure 7*, which explicitly states that the vendor does not use personal data about children to train the vendor's AI systems.)
- If and how the school is able to control what data the vendor collects. For example, does the vendor use student data or responses to further train its model? (See *Figure 7*, which explicitly states that the vendor does not allow partners to repurpose its users' data for training models.)

This information helps school administrators understand how data from and about students is used and the level of risk presented. It is also critical for school administrators' ability to assess their compliance with legal frameworks such as FERPA.

 **ClassDojo**

What Information is Used to Provide the AI Tools

In serving these new AI Tools, ClassDojo leverages the following information:

- Inputs from teachers, such as text for posts, messages or other contextual text input.
- Prompts or instructions from ClassDojo, such as “Write a thank you note”.
- Outputs which include AI-generated text suggestions and assessment questions, which are reviewed by teachers before being shared with the intended audience.

All information provided to the LLM Service Provider is securely transmitted to the LLM Service Provider for processing, stored transiently in the United States, and encrypted both in transit and at rest. Our data retention policies ensure Inputs, Prompts, and Output are deleted after processing by the LLM Service Provider, and that Outputs retained by ClassDojo are safeguarded and stored securely in alignment with ClassDojo’s existing policies.

The following chart provides more detail on the information used to provide the various AI Tools (broken down into our AI Classroom Tools and AI Productivity Tools) as well as the processes or “Prompts” performed on the “Inputs” to obtain the results or “Output”. We will update this chart as we update our features to provide schools with the information necessary for them to act as the agent of a parent and provide any necessary [consent](#) under the Children’s Online Privacy Protection Act and its implementing regulations under the Children’s Online Privacy Protection Rule (collectively “**COPPA**”) or to otherwise approve this use of the AI Tools under applicable laws and regulations for any student [personal information](#) that is used (for example, under state student privacy laws and the Family Educational Rights and Privacy Act (“**FERPA**”).

AI Classroom Tools

May use Student PII

ClassDojo AI Classroom Tools and Processing Done on Inputs	User Information Used to Provide the AI Tools (Inputs)	Output Generated*	User Types Input is Collected from and Output Disclosed to	How this Information is provided to the LLM	Retention Period / Storage / Security
--	--	-------------------	--	---	---------------------------------------

◀ **Figure 6.** Example of AI data governance transparency on a dedicated page that lays out the general types of information used for the product’s AI components and then dives into more specifics in a table

Source: ([classdojo.com](#) — captured Oct. 6, 2025)

Table of Contents

How ClassDojo Uses AI Tools for Parents and Children

Safety, Accountability, Fairness & Efficacy in AI.


As part of our commitments, ClassDojo does not allow any AI third-party [service provider](#) (e.g. Open AI or Anthropic) it may utilize to provide the Service, to use any information ([personal information](#) or [de-identified](#)) to train, improve, or develop their [AI systems](#). ClassDojo also does not use any children’s [personal information](#) contained in any Inputs, Prompts, or Output to train any of ClassDojo’s [AI systems](#).

▲ **Figure 7.** Example of AI data governance transparency stating that the company does not allow third-party service providers, including underlying AI models, to use its information to train or improve their AI systems

Source: ([classdojo.com](#) — captured Oct. 21, 2025)

Governance Structures

The vendor should provide information on how its assesses its products from an ethical standpoint (such as the advisory board shown in *Figure 8*), as well as any processes it has in place to ensure ethical use ([Cibralic et al., 2025](#)). This information could include any feedback channels it has in place to solicit feedback from customers or others, any ethics or advisory boards it relies on and their makeup, any consultation with subject matter experts, and internal policies to allow employees to raise ethical considerations.




▲ **Figure 8.** Example of a governance structure that announced a dedicated AI Advisory Board, which is composed of K–12 educators and meets monthly to provide ongoing guidance and feedback on the development and ethical deployment of Otus AI tools

Source: ([otus.com](#) — captured Oct. 6, 2025)

Information Accessibility

Effective transparency requires that the information described in the previous sections be easily accessible and findable by procurers, users, and potential users of the system. A visitor to a company's website should be able to find it fairly easily without extensive searching, and information relevant to the AI components of the tool should be centralized, as in the example from Merlyn Mind shown in *Figure 9*. This example provides links to relevant AI information and is easily reachable from a link in the website footer. Additionally, the information should be comprehensible to non-experts so that school administrators and other audiences such as parents can understand it and assess the system based on their needs (Chmielinski et al., 2024).



▲ **Figure 9.** Example of information accessibility in an AI hub that links to much of the relevant information about how it uses AI and is accessible via a link in the standard footer on the website, making it easy to reach from anywhere on the site

Source: (merlyn.org — captured Sept. 25, 2025)

Transparency for General Purpose Models Used in the Classroom

Many edtech products that incorporate AI are built on top of general purpose AI models provided by other technology companies such as Anthropic, Google, Microsoft, and OpenAI, but schools and educators may also use these systems directly. These companies provide general purpose AI models that primarily fall into one of three categories:

- **Free off the shelf:** These models, such as ChatGPT or Gemini, are available to all users free of charge. The free versions may give users lower priority, requiring them to wait to use the system, and they may have different data policies than their paid counterparts.
- **Targeted toward the education sector:** These products are specifically targeted version of models, such as [ChatGPT Edu](#), [Claude for Education](#), and [Gemini for Education](#), and some are focused on higher education rather than a K–12 context. These models, like paid versions, may offer stricter control on data privacy as well as other features that make them suited to education, such as “homework modes” that ask leading questions instead of providing direct answers.
- **Paid premium:** These paid products (often premium versions of freely available models such as ChatGPT) may have different data policies than their free counterparts (such as the ability to opt out of use for training or longer data retention to produce more personalized results).

The transparency rubric created for this report is targeted at assessing products specifically designed for K–12 use. When these general purpose models are used directly in the education context, school administrators will still need to assess the products and will have to grapple with additional questions about whether they will be effective in an education context ([Langreo, 2024](#)). For example, different versions of general purpose models or different types of accounts may have different data use policies (e.g., as of August 2025, ChatGPT Free and Plus use data from conversations for training by default, while ChatGPT Edu does not [[OpenAI, 2025a](#); [OpenAI, 2025b](#)]), so transparency around the privacy and security practices of general purpose models is essential to legal compliance.

Developers of general purpose models should take steps to improve transparency to help school administrators understand if their tool will be appropriate, such as providing in one place clear and easily accessible information about the tool’s use in education, explicitly clarifying data protection practices and controls that go along with specific versions of their tool, and clarifying whether and how their tool was adapted and tested for educational settings. These steps can help stakeholders make informed choices and understand the technology being used in education contexts.

Transparency in EdTech Products That Use AI: Key Findings

To understand what transparency looks like in practice, this study applied the rubric described in the previous section to the public websites of 108 different edtech companies that offer products that incorporate AI. Each company was scored along the eight categories of the rubric for transparency between June and July 2025 (see *Appendix B: Method and Procedure* on page 32 for more detail). The assessment focused only on the extent to which information in each dimension of the rubric was available, and no further analysis was done on the accuracy of that information.

The assessment showed that, on average, edtech companies offer very little transparency about their products that incorporate AI. When transparency is provided, the information most commonly falls in the Use and Context Limitations and Information Accessibility categories. Information about Domain Adaptation and Testing and Evaluation is the least common.

To score the companies, each company was assigned one of three ratings per transparency category:

- **GREEN:** the highest level of transparency with an equivalent numerical score of 2;
- **YELLOW:** the intermediate level with some amount of transparency and an equivalent score of 1; or
- **RED:** the lowest level with no information or transparency and an equivalent score of 0.

The resulting total score for each company ranged from 0 to 16 (with a maximum of 2 for each of the eight categories). Descriptions used to guide which rating a company received are described in more detail in *Appendix A: Comprehensive Transparency Rubric of EdTech and AI* on page 29.

On Average, Edtech Companies Offer Little Transparency about Their Products

Across all companies, overall transparency was very low. On average, companies received a total score of 4 (out of 16). As seen in *Figure 10*, the distribution of these scores was such that more than half of the companies (58) had a total score of 4 or less. This result means that most companies received a score of 0 for multiple categories. Indeed, for 71 out of the 108 companies (more than 65%), the modal, or most common, score received in each of the eight transparency categories was 0.

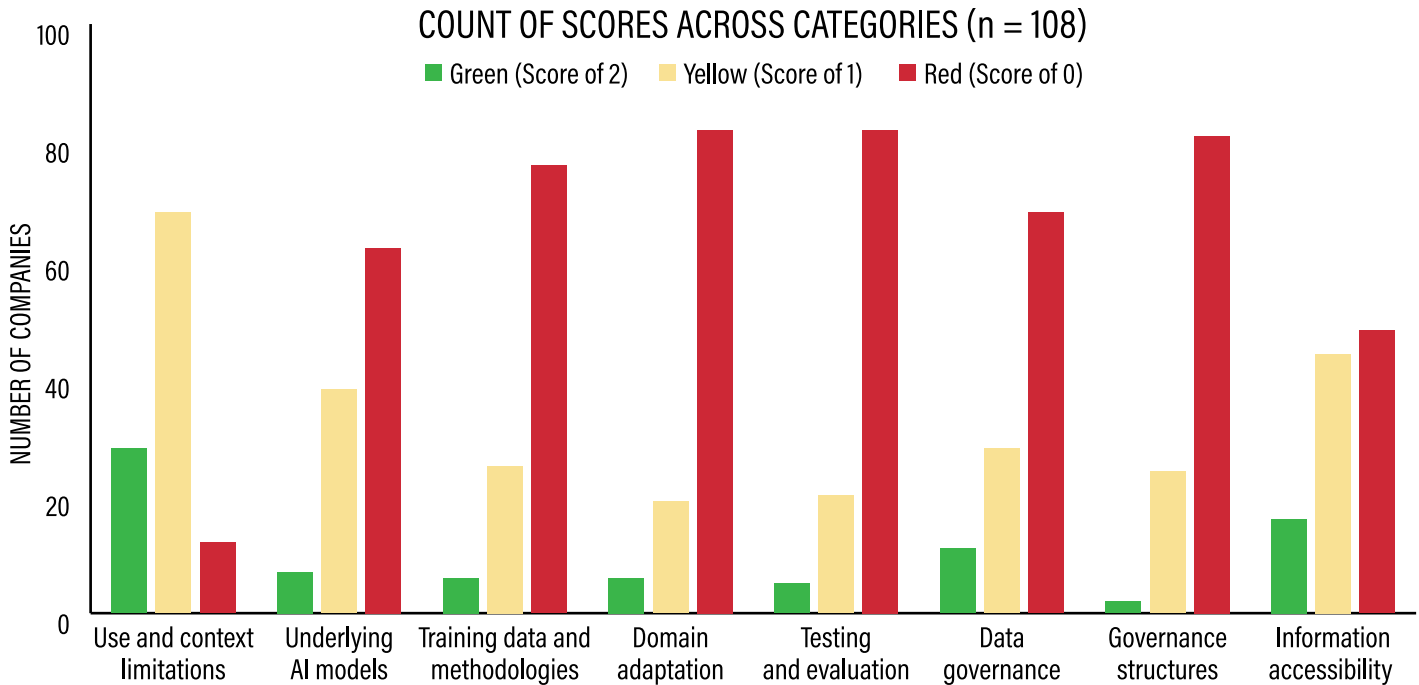


▲ Figure 10. Distribution of total transparency scores for companies Source: (CDT)

When Edtech Companies Attempt to Be Transparent, the Information Is Often Focused on Select Categories, Namely Use and Context Limitations and Information Accessibility

While overall transparency was low, both strong and weak transparency practices were clustered around a subset of the overall categories examined. The categories with the highest average scores across all 108 companies were Use and Context Limitations (1.20) and Information Accessibility (0.70). The categories with the lowest average scores were Testing and Evaluation (0.28) and Governance Structures (0.26).

Another way to understand these results is to look at the number of ratings (green, which equals a score of 2; yellow, which equals a score of 1; and red, which equals a score of 0) companies received for each category. *Figure 11* shows that 28 companies received a rating of green in the Use and Context Limitations category, and 16 companies did so under Information Accessibility, more than any other categories. Similarly, these two categories received the highest number of yellow ratings and the lowest number of red ratings (where no information was available). Use and Context Limitations was the highest scoring category across companies, possibly because for marketing purposes, they were more likely to explain the use case for their product. However, even in that category, the most common rating was yellow, meaning school administrators are still left with limited information about when and for what a given product should be used.




▲ Figure 11. Count of company ratings by transparency category Source: (CDT)

Companies with Middling Transparency Make Information about AI Easy to Access

In addition to overall scores, companies could be grouped based on the similarities of their scores across categories. This analysis resulted in three separate groups of companies:

- **Strong transparency** (20 companies): This group had the highest average scores across all categories and was also the smallest of the three, which underlines the finding that, overall, transparency among these companies is very limited.
- **Middling transparency** (40 companies): This group did relatively well in a few select categories, but their overall average was not among the best.
- **Poor transparency** (47 companies): This group was the largest and had the lowest average scores across all categories.



▲ Figure 12. Average scores for each transparency category, by company group Source: (CDT)

As shown in *Figure 12*, the middling transparency group was generally closer to the poor transparency group across most transparency categories, except for Information Accessibility. In the Information Accessibility category, the middling group scored much closer to the strong transparency companies (meaning that what information they do provide is easy to find and presented in a comprehensible way). This finding suggests that the middling transparency companies are interested in making information about their tools easy to find and understand, as scores in this dimension were determined by the extent to which companies made information across other dimensions easily accessible and understandable. These companies could be attempting to be transparent about their AI products but may lack a roadmap for doing so effectively. If this is the case, a rubric such as the one used in this report could help these companies improve their overall transparency, enabling school administrators to more easily select products from companies acting in good faith to build effective AI edtech.

■

Conclusion

The meager transparency currently offered by the bulk of the companies evaluated makes it difficult to determine the extent to which the developers of these tools are taking steps such as intentionally designing their systems, rigorously assessing those systems for effectiveness, and carefully managing and securing student data — all issues schools need to evaluate to assess whether a system will meet their needs and perform effectively for their students. The absence of this data forces educators to forgo the use of AI or gamble with adopting a tool they have not been able to fully evaluate, hampering the responsible, effective, and equitable use of AI in schools. Edtech vendors that incorporate AI into their products should use the rubric as a guide to improve their transparency in meaningful ways that enable school administrators to make informed decisions that center what is best for students.



Appendix A: Comprehensive Transparency Rubric of EdTech and AI

The following rubric was used to assess the website of each tool.

	RED	YELLOW	GREEN
Use and Context Limitations	The vendor does not provide information, or the stated purpose is unhelpfully broad. For example, “personalized learning” is a purpose but is not sufficiently specific to be useful to a procurer.	The vendor provides information that is somewhat narrower but still leaves a wide range of possible functionalities. For example, “student-engagement-focused personalized learning” provides bread crumbs, but there are still quite a few possible specific uses.	The vendor provides information about specific tasks the system performs. For example, the vendor might say, “Our tool personalizes examples in word problems to appeal to students’ particular areas of interest.”
Underlying AI Models	The vendor does not provide information about whether the tool was built on an existing system or only states that an existing system was used with no further information about that system. For example, “our system is built on top of leading large language models” does not provide useful information.	The vendor provides general information about the base system, such as which system was used, and links to information about that system. Or the vendor provides the information directly (such as what sort of company developed the base system and what it was designed to do).	The vendor provides detailed information about the base system, including which system was used, and provides an explanation about why it chose that base system. For example, the vendor might say, “Our tool uses Company A’s Engine123 because we found it produced the most accurate results.”
Training Data and Methodologies	The vendor does not provide information about how its system was designed and trained, and it does not provide training data information. For example “our model uses state-of-the-art AI technologies” is not meaningful information about the system.	The vendor provides high-level information that may allow procurers or other stakeholders to draw inferences, but it does not provide specifics. For instance, the vendor might say, “Our system uses a large language model to interact with students about their academic work.” For data, the vendor provides some limited information about the data used to train the system, such as the size of the dataset, but no demographic information.	The vendor provides specific information about training methodologies and data, such as, “Our model was trained using supervised learning over a corpus of student essays to provide first-pass grading, with interactive rounds during which humans assessed the system’s output and provided feedback and updated labels. This information was then fed back into the system to improve outcomes. The essays were sourced from a range of middle and high schools, with a variety of student races and economic backgrounds as well as school sizes.”

	RED	YELLOW	GREEN
Domain Adaptation	The vendor does not provide information about how the system was updated for an education context, or it simply states that adaptation was done but provides no further information. For example, “our tool was tuned with education in mind” does not give a procurer enough information to assess the system.	The vendor provides partial information about its adaptation approach, such as information about the specific goals of its domain adaptation, but it does not provide specifics about how domain adaptation was achieved.	The vendor provides specific information about how the system was updated for an education context, including what use cases it has accounted for, what risks it has accounted for, and whether those risks have been addressed through updates to the model itself or scaffolding around the model. For example, “our system takes responses from the base model and strips any inappropriate information from the response based on keywords and sentiment analysis” helps a procurer assess whether the adaptation will meet their needs.
Testing and Evaluation	The vendor does not provide information about its testing approach, or it provides reassurances without specifics, such as, “We have robustly tested our system to ensure high performance.” The vendor does not make specific claims.	The vendor provides some information about how it tested its systems, such as results but no methodology. The vendor makes specific, falsifiable claims about the performance of its technology but does not provide evidence for those claims. The vendor notes that it tests its system but does not provide information about when it retests.	The vendor provides specific information about how it has tested its technology, including what it was testing for, methodology, and results. The vendor makes specific falsifiable claims and provides evidence to back them. Additionally, the vendor shares process information about when and how it updates and retests, including the most recent evaluation date.
Data Governance	The vendor has a privacy policy that does not specifically address its AI data use. It does not discuss what data its system generates or how it manages and protects that data.	The vendor’s privacy policy specifically covers AI data use. It discusses data generated by its system and makes unsupported claims about how it maintains that data (e.g., “we ensure that our clients’ data is always secured and protected”).	The vendor provides specific information about what data its system generates and how it secures the data (including specific technologies such as encryption at rest, least-privilege access control, etc.). It has a separate privacy policy for its AI tool and/or a plain-language page addressing the AI’s use of data.

	RED	YELLOW	GREEN
Governance Structures	The vendor does not provide information about any internal or external governance structures it relies on to ensure that its products are ethically designed and deployed. The vendor has a stated interest in ethics but provides no information about how it acts on those interests (e.g., “We design products with students’ best interests in mind,” without further information).	The vendor provides partial information about any governance or ethics structures it maintains. For instance, it provides information about the existence of such structures but not the makeup (e.g., a governance board without member information), or it makes specific ethical commitments but does not provide information about how it ensures that it adheres to them.	The vendor makes specific commitments and has information about the processes it has in place to ensure adherence to those commitments. Where possible, the vendor enables buyers to communicate with any ethics personnel the company has. Ideally, the vendor provides resources to procurers to ensure ethical deployment.
Information Accessibility	Relevant data is buried on pages that are deep in the website, such as old blog posts, and data is scattered across pages.	The data is only a few clicks away from the main page or is on a clearly named page that would turn up in a targeted web search.	The information is pervasively available on the webpage via a header or footer link or some other omnipresent link. The information page is descriptively named such that it is intuitive to find and/or search for. The information is consolidated.

Appendix B: Method and Procedure

Sample Selection

To create the sample of edtech companies, an initial list was developed using several public sources focused on edtech and AI, some of which discuss company commitments to transparency. These sources included companies that were members of the EDSAFE AI Industry Council ([EDSAFE AI, 2025a](#)), the Software & Information Industry Association's EdTech Principles Committee members ([Education Technology Industry's Principles for the Future of AI in Education, 2025](#)), the Future of Privacy Forum's Student Privacy Pledge signatories ([Student Privacy Pledge, 2025](#)), and companies that exhibited at the 2025 ASU+GSV Summit (based on exhibitor listings and photos) ([ASU+GSV Summit, 2025](#)). Additionally, *Time's* list of the Top 100 EdTech Companies of 2024 ([Time, 2025](#)) was consulted to ensure that the most prevalent companies were captured in the analysis.

This process initially yielded 154 companies. The list was then filtered by spot-checking every website to confirm that companies specifically targeted K–12 education and incorporated AI technologies. Companies primarily focused on higher education, enterprise, or career solutions were excluded, as were products that did not appear to use AI, used AI only in experimental or pilot phases, had not yet released their product, were not designed for classroom use, or lacked sufficient product information in English. This process resulted in a final sample of 108 companies for in-depth review using the transparency rubric.

Validity

As part of an effort to ensure consistency and clarity in the coding process, an interrater reliability check was conducted during the early stages of the review on a randomly selected sample of 10 companies from the initial list of 154. Two coders independently reviewed and coded them using the preliminary rubric. After completing assessments, the coders met to compare results, resolve discrepancies, and identify areas in which the rubric needed clarification. Based on this discussion, the coding schema were refined to include more specific definitions and clearer instructions. With the updated rubric in place, the remaining companies were divided between the two coders for independent review.

Analysis

For the analysis, each company's transparency and public-facing information on AI was assessed through a systematic review of its websites, guided by the transparency rubric. The process began with browsing each company's webpages directly and using any available on-site search functionality with targeted keywords such as "AI," "Artificial Intelligence," "Generative AI," "Large Language Model (LLM)," "Machine Learning (ML)," and "Natural Language Processing (NLP)." The available privacy and terms of use policies were searched for these same keywords, with a primary focus on the terms "AI" and "Artificial Intelligence." If no results were returned, the policies were read in full to ensure that content was not overlooked. The key navigational areas of each site were reviewed, particularly the header and footer menus, for links to dedicated

pages discussing AI use, transparency, or data practices. Additionally, “About Us” pages or mission statement sections were reviewed to identify any references to AI-related strategies or principles.

Beyond on-site searching, Google’s advanced site search function was employed with two targeted keyword groups to locate relevant mentions of AI technologies throughout the website that may not have been found through simple site navigation. The first group included general AI-related keywords, including “Artificial Intelligence,” “AI,” “GenAI,” “Generative AI,” “Large Language Model (LLM),” “Training,” “Transparency,” “Machine Learning (ML),” “Natural Language Processing (NLP),” and “Text Classifier.” The second group focused on specific companies and model names (selected from [Bommasani et al., 2024](#)): Adept Fuyu, AI21 Labs (Jurassic), Aleph Alpha (Luminous Supreme), Titan Text Express (Amazon), Anthropic (Claude), BigCode/HuggingFace/ServiceNow, Gemini (Google), IBM Granite, Meta Llama, Microsoft (Phi-2, CoPilot), Mistral, OpenAI (GPT, Dall-e), Stability AI (Stable Video Diffusion), and Writer Palmyra.

Following this review procedure, information was gathered for each company across the eight categories outlined in the transparency rubric. Information related to each category was recorded in a spreadsheet and evaluated based on the presence or absence of key characteristics defined in the rubric, described in detail in *Appendix A*. Each category was then assigned a color designation (green, yellow, or red) reflecting the level of transparency standards demonstrated in that area, as noted previously. Scores were assigned as follows:

- **GREEN:** the highest level of transparency with an equivalent numerical score of 2;
- **YELLOW:** the intermediate level with some amount of transparency and an equivalent score of 1; or
- **RED:** the lowest level with no information or transparency and an equivalent score of 0.

Given that the same range was used for each category, the scores are similarly scaled, and there was no need to standardize the data. This report includes descriptive analysis of these scores in *Transparency in EdTech Products That Use AI: Key Findings*. For the analysis on groups of companies (e.g., poor transparency, middling transparency, and strong transparency), CDT used the k-means clustering algorithm — an unsupervised algorithm that attempts to find data points that are as similar as possible in terms of variance and size. More specifically, this approach was used to group the companies based on the similarities of their scores.

References

- Ariful Islam Anik & Andrea Bunt, *Data-Centric Explanations: Explaining Training Data of Machine Learning Systems to Promote Transparency*, Human Factors in Computing Systems (May 2021) <https://dl.acm.org/doi/abs/10.1145/3411764.3445736>.
- ASU+GSV, *ASU+GSV Summit* (accessed Oct 2025) <https://perma.cc/RA2C-24Q6>.
- Rishi Bommasani et al., *Improving Transparency in AI Language Models: A Holistic Evaluation*, Stanford University HAI (Feb 2023) <https://perma.cc/5UEG-AKHE>.
- Rishi Bommasani et al., *The Foundation Model Transparency Index v1.1*, Stanford University Center for Research on Foundation Models (May 2024) <https://perma.cc/BV4B-DJBZ>.
- Kasia Chmielinski et al., *The CLeAR Documentation Framework for AI Transparency: Recommendations for Practitioners & Context for Policymakers*, Harvard Shorenstein Center on Media, Politics and Public Policy (May 2024) <https://perma.cc/L64E-26SP>.
- Beba Cibralic et al., *Principled Practice: A Playbook for Operationalizing Responsible AI*, Center for Democracy & Technology (Sept 2025) <https://cdt.org/insights/principled-practice-a-playbook-for-operationalizing-responsible-ai/>.
- Jack Corrigan, *AI Monopolies Are Coming. Now's the Time to Stop Them.*, Tech Policy Press (May 2025) <https://perma.cc/4XNC-RA3H>.
- Maddy Dwyer, *Looking Back at AI Guidance Across State Education Agencies and Looking Forward*, Center for Democracy & Technology (Apr 2025) <https://cdt.org/insights/looking-back-at-ai-guidance-across-state-education-agencies-and-looking-forward/>.
- EDSAFE AI, *EDSAFE AI Industry Council* (accessed Oct 2025a) <https://perma.cc/4LUD-3BE6>.
- EDSAFE AI, *S.A.F.E. Benchmarks* (accessed Oct 2025b) <https://perma.cc/LM2A-CVEV>.
- Education Technology Industry's Principles for the Future of AI in Education, *AI Education Principles* (accessed Oct 2025) <https://perma.cc/3DHN-2HR4>.
- Grand View Research, *AI in Education Market (2025–2030)* (accessed Oct 2025) <https://perma.cc/5CTE-EMPV>.
- Kirk Gulezian, *K–12: Building Custom GPTs to Automate Tasks*, OpenAI Academy (Jun 2025) <https://perma.cc/N42C-EZUT>.
- Hugging Face, *Model Cards* (accessed Sept 2025) <https://perma.cc/W9UZ-KK6W>.
- Jennifer King & Caroline Meinhardt, *Rethinking Privacy in the AI Era: Policy Provocations for a Data-Centric World*, Stanford University HAI (Feb 2024) <https://perma.cc/F7CP-D6X7>.

- Natalia Kucirkova et al., *Applying the Science of Learning to EdTech Evidence Evaluations Using the EdTech Evidence Evaluation Routine (EVER)*, npj Science of Learning (2023) <https://www.nature.com/articles/s41539-023-00186-7>.
- Elizabeth Laird et al., *Out of Step: Students, Teachers in Stride with EdTech Threats While Parents Are Left Behind*, Center for Democracy & Technology (Jan 2025a) <https://cdt.org/insights/out-of-step-students-teachers-in-stride-with-edtech-threats-while-parents-are-left-behind/>.
- Elizabeth Laird et al., *Hand in Hand: Schools' Embrace of AI Connected to Increased Risks to Students*, Center for Democracy & Technology (Oct 2025b) <https://cdt.org/insights/hand-in-hand-schools-embrace-of-ai-connected-to-increased-risks-to-students/>.
- Elizabeth Laird & Maddy Dwyer, *Off Task: EdTech Threats to Student Privacy and Equity in the Age of AI*, Center for Democracy & Technology (Sept 2023) <https://cdt.org/insights/report-off-task-edtech-threats-to-student-privacy-and-equity-in-the-age-of-ai/>.
- Lauraine Langreo, *Don't Buy the AI Hype, Learning Expert Warns*, EdWeek (Aug 2024) <https://perma.cc/B8N3-9SW4>.
- Microsoft, *Elevating Communities | Microsoft Elevate* (accessed Aug 2025) <https://perma.cc/6NC7-ETGC>.
- Margaret Mitchell et al., *Model Cards for Model Reporting, Fairness, Accountability, and Transparency* (Jan 2019) <https://dl.acm.org/doi/10.1145/3287560.3287596>.
- Office of Educational Technology, *Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations*, U.S. Department of Education (May 2023) <https://perma.cc/6G3T-TQ8X>.
- OpenAI, *How Your Data Is Used to Improve Model Performance* (accessed Sept 2025a) <https://web.archive.org/web/20250925190352/https://help.openai.com/en/articles/5722486-how-your-data-is-used-to-improve-model-performance>.
- OpenAI, *ChatGPT Education* (accessed Oct 2025b) <https://web.archive.org/web/20251011181708/https://chatgpt.com/business/education/>.
- Hannah Quay-de la Vallee, *A for Effort, Needs Improvement on Execution: Lessons Learned from Failures of AI EdTech Tools*, Center for Democracy & Technology (Aug 2025) <https://cdt.org/insights/a-for-effort-needs-improvement-on-execution-lessons-learned-from-failures-of-ai-edtech-tools/>.
- Hannah Quay-de la Vallee & Maddy Dwyer, *The Shortcomings of Generative AI Detection: How Schools Should Approach Declining Teacher Trust in Students*, Center for Democracy & Technology (Dec 2023) <https://cdt.org/insights/the-shortcomings-of-generative-ai-detection-how-schools-should-approach-declining-teacher-trust-in-students/>.
- Software & Information Industry Association, *SIIA Releases Principles for the Future of AI in Education* (Oct 2023) <https://perma.cc/2UEA-SSGA>.
- Student Privacy Pledge, *Former Signatories* (accessed Jul 2025) <https://web.archive.org/web/20250716182533/http://studentprivacypledge.org/signatories/>.

Time, *World's Top EdTech Companies of 2024* (accessed Oct 2025) <https://perma.cc/ST8B-695E>.

Stéphan Vincent-Lancrin & Reyer van der Vlies, *Trustworthy Artificial Intelligence (AI) in Education: Promises and Challenges*, OECD Education Working Papers (2020) <https://doi.org/10.1787/a6c90fa9-en>.

White House, *Executive Order — Advancing Artificial Intelligence Education for American Youth* (Apr 2025) <https://perma.cc/9ZN9-UVQW>.

Amy Winecoff & Miranda Bogen, *Improving Governance Outcomes Through AI Documentation: Bridging Theory and Practice*, Center for Democracy & Technology (Sept 2024) <https://cdt.org/insights/report-improving-governance-outcomes-through-ai-documentation-bridging-theory-and-practice/>.

Ke Zhang & Ayse Begum Aslan, *AI Technologies for Education: Recent Research & Future Directions*, Computers and Education: Artificial Intelligence (2021) <https://www.sciencedirect.com/science/article/pii/S2666920X21000199>.

 cdt.org

 cdt.org/contact

 **Center for Democracy & Technology**
1401 K Street NW, Suite 200
Washington, D.C. 20005

 202-637-9800

 @CenDemTech

