

# From Chalkboards to Algorithms

EdTech Readiness for the  
Colorado AI Act

The Privateers:

Haley Smith

Wisdom Obinna

Bridget Reineking

Khadija Mian

Jessica Jamous

Tech Impact Lab  
Georgetown University  
May 2025

# Table of Contents

- Executive Summary ..... 4
- I. Introduction ..... 4
  - I.A. Introducing the Colorado Artificial Intelligence Act.....4
  - I.B. Core AI Technologies of EdTech .....7
- II. Methodology..... 10
  - II.A. Selection of Industry.....10
  - II.B. Selection of Representative Companies .....11
  - II.C. Analysis of Compliance .....12
  - II.D. Limitations of our Methods .....12
- III. Representative Companies and their Compliance- Part I..... 13
  - III.A. Representative Companies, Categories I-III.....13
    - III.A.i Category I: Student Surveillance Safety.....13
    - III.A.ii Category II: Feedback Generation .....19
    - III.A.iii Category III: Predictive Analysis .....22
  - III.B. Compliance Insights, Part I: Review of Public-Facing Materials .....27
    - Transparency and Disclosure Readiness.....27
    - Algorithmic Fairness and Reasonable Care Concerns.....27
- IV. Representative Companies and their Compliance, Part II..... 28
  - IV.A. Methodology - Technical Analysis .....28
    - IV.A.i. Research Design and Data Collection .....28
    - IV.A.ii. Analytical Framework.....29
  - IV.B. Representative Companies, Categories I-III.....30
    - IV.B.i. Category IV: Special Education.....30
  - IV.C. Compliance Insights, Part II: Review of Technical and Public-Facing Materials .....38
    - Transparency and Disclosure Readiness.....38
    - Algorithmic Fairness and Reasonable Care Concerns.....38
- V. Recommendations ..... 39
- VI. Future Research Ideas..... 40
- VII. Conclusion..... 40
- ACKNOWLEDGMENTS ..... 41
- APPENDICES ..... 42
  - Appendix I - Frequently Asked Questions (FAQs) .....42
  - Appendix II - Test Prompts .....42
  - Appendix III - Code used for analysis of IEP outputs .....42

# Executive

## Summary

Artificial intelligence (AI) is becoming increasingly integrated into educational technology (EdTech), offering the promise of reduced administrative burdens for educators and more individualized instruction for students. Facing limited resources, many school districts and teachers are adopting these tools in hopes of enhancing educational outcomes without overextending themselves. However, this use of AI in education has raised serious concerns of algorithmic discrimination. In response, Colorado passed the AI Act, a law designed to hold developers and deployers accountable for identifying and minimizing such harms.

Our review of several EdTech Companies operating in Colorado revealed that many EdTech companies demonstrate readiness to comply with the AI Act. Several deployers provide transparent, public-facing information regarding their AI offerings, responsible uses, and apparent limitations. However, our technical assessment revealed that certain EdTech AI tools may produce discriminatory outcomes that disproportionately impact students based on protected characteristics.

We encourage regulators and EdTech companies to engage in productive dialogue to clarify compliance expectations and address regulatory constraints ahead of the AI Act's February 2026 effective date. Moreover, we urge educators and administrators to hold off against widespread adoption and implementation of sensitive AI technologies in educational settings until they receive thorough training on these tools' limitations and potential biases against protected classes. These technologies offer tremendous opportunities for improved student outcomes but must be adopted for classroom use in a manner designed to balance their benefits against the dangers of automation and overreliance.

# I. Introduction

Technology is transforming K-12 education across Colorado, introducing both promising opportunities and significant challenges. Where teachers once relied primarily on chalkboards and paper, adaptive learning platforms and artificial intelligence systems now increasingly shape students' educational experiences. As the next school year approaches, Colorado schools are adopting promising technological solutions to enhance learning, ease administrative burdens, and address safety concerns, while simultaneously preparing for compliance with the state's new AI legislation.

Despite these benefits, the integration of artificial intelligence into educational contexts may introduce or exacerbate certain risks. Responding in part to this concern, Colorado enacted the Colorado AI Act in May 2024, which imposes affirmative obligations on developers of AI technologies to exercise reasonable care in preventing algorithmic discrimination within educational settings (among other specified domains). The legislation is scheduled to take effect in February 2026. This report analyzes the preparedness of a representative sample of Colorado's EdTech companies as they approach the effective date, based on publicly available information.

What follows is a compliance assessment of observable aspects of one industry in advance of the AI Act's effective date, limited by our constrained visibility into proprietary systems and internal company practices. In this, we also describe current EdTech offerings and their potential impacts, and highlight considerations for administrators, educators, regulators, and parents alike as districts continue to adopt AI technologies into classrooms.

This report breaks down the AI Act, as well as common AI technologies to set the scene for the EdTech products described in the following sections. We also set forth the general methods used to select the representative companies, specific methods for the document-based and technical analyses. For each company, we provide a general overview of the specific AI tool, potential areas for algorithmic discrimination, as well as the perceived compliance of each developer with the AI Act. For two of the companies, we were able to directly access and use the system, and provided a technical analysis of the outputs. Based on our findings, we provided suggestions to regulators, educators, and parents, as well as future areas of research to further delve into this topic.

## I.A. Introducing the Colorado Artificial Intelligence Act

The Colorado AI Act takes effect in February 2026, establishing obligations for companies that develop or deploy AI systems in educational contexts.

### Central Purpose and Key Terms

The AI Act aims to protect Coloradans from **algorithmic discrimination**, defined by the law to mean:

any condition in which the use of an **artificial intelligence system** results in an unlawful differential treatment or impact that disfavors an individual or group of individuals on the basis of their actual or perceived **age, color, disability, ethnicity, genetic information, limited**

**proficiency in the English language, national origin, race, religion, reproductive health, sex, veteran status,** or other classification protected under [Colorado] or federal law.<sup>1</sup>

An **artificial intelligence system** (“AI system”) means “any machine-based system that, for any explicit or implicit objective, infers from the inputs the system receives how to generate outputs, including **content, decisions, predictions, or recommendations,** that can influence physical or virtual environments.”<sup>2</sup>

In turn, a “**high-risk artificial intelligence system**” (“HRAIS”) is any AI system “that, when deployed, makes or is a substantial factor in making, a **consequential decision.**”<sup>3</sup> In the educational context, a **consequential decision** means a “decision that has a material legal or similarly significant effect on the provision or denial to any consumer, or the cost or terms of . . . [e]ducation enrollment or an education opportunity

The AI Act excludes from the definition of HRAIS certain technologies, including anti-virus software, calculators, data storage, cybersecurity, and chatbot-like technology “subject to an accepted use policy that prohibits generating content that is discriminatory or harmful,” among others,<sup>4</sup> “unless the technologies, when deployed, make, or are a substantial factor in making, a consequential decision.”<sup>5</sup>

A **substantial factor** means a “factor that (1) assists in making a consequential decision; (2) is capable of altering the outcome of a consequential decision; and (3) is generated by an [AI system].”<sup>6</sup> In addition, the term “substantial factor” includes any “use of an [AI system] to generate any “content, decision, prediction, or recommendation concerning a consumer that is used as a basis to make a consequential decision concerning the consumer.”<sup>7</sup>

As we detail further in this report, the broad aperture of the term “substantial factor” is meaningful given the many EdTech offerings that provide recommendations related to student opportunities. Equally significant is the ambiguity regarding how much an HRAIS must “assist” in a consequential decision or to what extent it must be “capable” of altering outcomes to constitute a substantial factor.

The AI Act applies to both **developers** (a “person doing business in [Colorado] that develops or intentionally and substantially modifies an [AI system]”) and **deployers** (a person “doing business in [Colorado] that deploys a [HRAIS]”).<sup>8</sup> This report focuses on the preparedness requirements for **developers**, as detailed in subsequent sections.

## Key Developer Obligations

The AI Act imposes a duty of “reasonable care” on developers of HRAIS to “protect consumers from any **known or reasonably foreseeable risks of algorithmic discrimination**” resulting from their technology’s use.<sup>9</sup> This duty extends to such risks “arising from the intended and contracted

---

<sup>1</sup> See Colo. Rev. Stat. § 6-1-1701(1)(a) (emphasis in bold added).

<sup>2</sup> See Colo. Rev. Stat. § 6-1-1701(1)(2) (emphasis in bold added).

<sup>3</sup> *Id.* § 6-1-1701(9)(a) (emphasis in bold added).

<sup>4</sup> *Id.* § 6-1-1701(9)(b).

<sup>5</sup> *Id.*

<sup>6</sup> *Id.* § 6-1-1701(11)(a).

<sup>7</sup> *Id.* § 6-1-1701(11)(b).

<sup>8</sup> See *id.* §§ 6-1-1701(6), (7); see also *id.* § 6-1-1701(10) (defining “intentional and substantial modification”).

<sup>9</sup> *Id.* § 6-1-1702(1) (emphasis in bold added).

uses”<sup>10</sup> of their AI product, and the law requires developers to make available certain disclosures and documentation to accomplish the mandate.

**Public disclosures:** Each developer must “make available, in a manner that is clear and readily available on [its] website or in a public use case inventory, a statement summarizing (1) the types of [HRAIS] that the developer has developed or intentionally and substantially modified and currently makes available to a deployer or other developer; and (2) how the developer manages known or reasonably foreseeable risks of algorithmic discrimination that may arise from the development or intentional and substantial modification” of its HRAIS.<sup>11</sup>

Notably, the AI Act does not further define “clear and readily available,” thus granting developers discretion in determining how these public disclosures should be presented.

**Non-public disclosures.** Under the AI Act, each developer must also “make available” to deployers and other developers of its HRAIS the following information:

- A “general statement” describing “reasonably foreseeable uses and known harmful or inappropriate uses” of the HRAIS;<sup>12</sup>
- Documentation disclosing:
  - “High-level summaries of the type of data used to train the [HRAIS]”
  - “Known or reasonably foreseeable limitations of the [HRAIS],” including “known or reasonably foreseeable risks of algorithmic discrimination arising from intended uses of the [HRAIS]”
  - The “purpose” and “intended benefits and uses” of the HRIAS, and
  - “All other information necessary” for a deployer to meet its obligations under the Act.<sup>13</sup>
- Documentation describing the material listed below, as well “any additional documentation reasonably necessary to assist a deployer in understanding the outputs and monitor the performance of the [HRAIS] for risks of algorithmic discrimination:”<sup>14</sup>
  - How the HRAIS was “evaluated for performance and mitigation of algorithmic discrimination” before being made available to the deployer
  - The “data governance measures used to cover the training datasets and measures used to examine the suitability of data sources, possible biases and appropriate mitigation”
  - The “intended outputs” of the HRAIS

---

<sup>10</sup> *Id.*

<sup>11</sup> *Id.* § 6-1-1702(4).

<sup>12</sup> *Id.* § 6-1-1702(2)(a).

<sup>13</sup> *Id.* § 6-1-1702(2)(b).

<sup>14</sup> *Id.* § 6-1-1702(2)(d).

- Measures taken by the developer to “mitigate known or reasonably foreseeable risks of algorithmic discrimination that may arise from the reasonably foreseeable deployment of the [HRAIS]” and
- “How the [HRAIS] should be used, not be used, and monitored by an individual” when “used to make, or is a substantial factor in making, a consequential decision.”<sup>15</sup>
- Finally, “documentation or information” “necessary for a deployer . . . to complete an impact assessment.”<sup>16</sup>

**Reporting to the Attorney General (AG):** If a developer learns through testing that its HRAIS “has been deployed and has caused or is reasonably likely to have caused algorithmic discrimination” or if the developer receives a “credible report” from a deployer that its HRAIS “has been deployed and caused algorithmic discrimination,” the developer must disclose to the AG and “all known deployers or other developers of the HRAIS” any “known or reasonably foreseeable risks of algorithmic discrimination arising from the intended uses of the [HRAIS] without unreasonable delay but no later than 90 days thereafter.”<sup>17</sup> The AG may also request the information.<sup>18</sup>

## I.B. Core AI Technologies of EdTech

This short explainer discusses common AI technologies used by EdTech companies in tools designed to safeguard students and personalize learning journeys. This information is helpful for appreciating the functionality and effects of technologies described further in sections III and IV of this report.

### Large Language Models (LLMs) in Education

LLMs are AI systems trained on vast amounts of text data to recognize patterns in language. When deployed, they process and generate human-like language. In educational settings, these systems can be used by educators for many purposes, including to:

- Generate lesson plans, student assignments, and special education materials such as individualized education plans (IEPs)<sup>19</sup>
- Draft communications for distribution to parents
- Create ideas and material for extracurricular activities
- Support content creation for differentiated instruction

---

<sup>15</sup> *Id.* § 6-1-1702(2)(c).

<sup>16</sup> *Id.* § 6-1-1702(2)(3).

<sup>17</sup> *Id.* § 6-1-1702(5)

<sup>18</sup> *Id.* § 6-1-1702(7).

<sup>19</sup> As described *infra* at section V, an IEP is a document developed for public school students who require special education services, outlining personalized learning goals, necessary accommodations, and specific support services tailored to address the student’s educational needs. The IEP serves as both an educational roadmap and accountability tool, created through collaboration between parents, teachers, and specialists, with regular reviews to track progress and make adjustments as the student develops. *See generally* Colorado Department of Education, IEP Procedural Guidance: Exception Student Services Unit Technical Assistance (July 2017), *available at* [https://www.cde.state.co.us/cdesped/iep\\_proceduralguidance](https://www.cde.state.co.us/cdesped/iep_proceduralguidance).

Students, too, can use LLMs in the classroom and at home, including to request explanations for a difficult topic, generate customized practice problems in science, coding, reading comprehension or math, or receive targeted study guidance based on practice problems or exam results.

One possible drawback is LLM outputs may contain inaccuracies or reflect biases present in the patterns of an LLM's training data. Users of LLMs must also be mindful of privacy considerations and be careful about how and when to use the tools in connection with material reflecting sensitive student data.

## Algorithmic Student Monitoring Systems

AI monitoring systems fundamentally rely on algorithms - structured sets of rules and statistical models - to analyze digital content and identify patterns. In the educational context, teachers, administrators, or school security personnel can leverage these systems to analyze student online activity and communications to identify potential safety concerns. For example, these systems can:

- Scan student writing, social media interactions, digital activities, and online searches for indicators of concerning content, including expressions of self-harm or distress, or intentions to harm others
- Alert designated staff when the algorithm detects concerning patterns or as students attempt to access dangerous or unauthorized material (such as prohibited websites)
- Track trends to inform school wellness initiatives

These monitoring systems can provide helpful support to administrators as they seek to ensure student wellness and school safety. But the tools can sometimes misinterpret figurative language, cultural expressions, or creative writing, or flag legitimate research as concerning content. The tools also tend to trigger important considerations around privacy and surveillance practices.

## Facial Recognition Technology (FRT)

Facial Recognition Technology uses computer vision algorithms to identify or verify a person's identity by analyzing facial features. In schools, administrators, teachers, and security personnel may use FRT applications in a few contexts, such as:

- Building access control and monitoring of school grounds
- Identifying students in potential distress or conflict situations
- Streamlining administrative processes like library or lunchroom checkouts and attendance

Common tasks for FRT tools include "one-to-one" matching, where the tool is used to match a photo of someone to another photo of the same person in a database, and "one-to-many" searching, which involves determining whether a photo of someone has any match in a database. In the educational context, FRT systems typically work by capturing images from security cameras, comparing facial geometry to a database of known individuals (which might include yearbook photos, or photos taken of persons entering a school), and alerting staff when necessary. These tools also trigger important privacy

and surveillance considerations, as use of FRT in a school may create an atmosphere of distrust or potentially chill speech of students.<sup>20</sup>

## Predictive Algorithms for Student Success

Predictive algorithms use historical data to identify patterns that may help forecast future outcomes such as graduation likelihood, academic performance, or intervention needs. Administrators and educators can use these systems to:

- Analyze a corpus of data points (attendance, grades, behavior records, demographic information) to identify students who may need additional support
- Inform the allocation of student support resources for greatest impact
- Facilitate earlier interventions when students show signs of academic struggles

These algorithms typically work by generating "risk scores" or categorizations that help educators prioritize attention and resources, but the origin of a risk score can be something of a black box. The algorithms are traditionally trained (before use) on a large data set, which may or may not generally reflect the student body of the school in which it is deployed. The factors that most strongly influence predictions are usually unknown to a user. There is therefore some danger that the predictions might be overly deterministic or unhelpful, leading to the differential treatment of students based on algorithmic forecasts that do not match the reality of individual circumstances.

## AI-Powered Feedback Generation Technologies

Feedback generation technologies use algorithms, often in concert with LLMs, to provide personalized responses to student work. Both educators and students themselves can take advantage of this growing category of educational tools to:

- Assess student responses in real-time and provide feedback, including with scaffolded learning techniques that provide progressive hints rather than immediate answers
- Identify skill gaps from student work
- Track improvement patterns over time to adjust difficulty

These systems typically work by comparing student responses against expected patterns and identifying deviations that suggest misunderstandings. The platforms offer several benefits for student learning experiences but may struggle with unconventional (but valid!) responses that don't match expected patterns – which might have the unintended consequence of quashing creativity. When implemented thoughtfully into an instructional approach, however, these technologies offer potential for enhanced teacher capacity and improved timeframes for intervention and feedback.

---

<sup>20</sup> See Mwyer, M. (2023). *What was once science fiction is now reality: Orwellian uses of safety tech in K-12 schools*. Center for Democracy & Technology. <https://cdt.org/insights/what-was-once-science-fiction-is-now-reality-orwellian-uses-of-safety-tech-in-k-12-schools/>; Nevradakis, M. (2023). Normalizing surveillance from a young age: More schools using facial recognition, AI technologies to monitor kids. *Children's Health Defense*. <https://cdt.org/insights/what-was-once-science-fiction-is-now-reality-orwellian-uses-of-safety-tech-in-k-12-schools/>

## II. Methodology

### II.A. Selection of Industry

This report examines the preparedness of the EdTech sector within Colorado in anticipation of the effective date of the AI Act. We choose to evaluate this industry for a few reasons.

#### **First, the use of AI in EdTech is rapidly growing.**

In 2023, the Department of Education’s Office of Education Technology observed the rapid introduction of AI into classrooms—a trend that has only accelerated. A 2023 survey of 1,200 teachers found that about one in five used AI tools daily.<sup>21</sup> The same report showed that by the end of the 2023-24 school year, approximately 60 percent of districts surveyed reported plans to offer AI training to their educators—with 37 percent of those districts offering this training for the first time.<sup>22</sup> This significant expansion in professional development will likely contribute to a rise in the number of teachers who become regular AI users.<sup>23</sup>

Another factor likely to contribute to the increasing use of AI in the classroom is the shortage of K-12 teachers in the United States. A survey of 30 states and the District of Columbia (DC) reported 41,920 unfilled teaching positions.<sup>24</sup> In the 2023-34 school year, Colorado has approximately 635 vacancies out of a total of about 53,500 teaching positions statewide.<sup>25</sup> Staffing shortages often lead to larger class sizes and reduced individual attention for students. Notably, another study found that teachers at schools in higher-poverty neighborhoods and at schools with larger populations of minority students were more likely to report feeling like their school was understaffed compared to schools with fewer disadvantaged students.<sup>26</sup>

#### **Second, Colorado readily identifies school technologies.**

Many Colorado public schools provide a high degree of transparency regarding the technologies used in classrooms, which made it relatively easy to identify AI vendors operating within the state—and thus potentially subject to regulation under the AI Act. Numerous school districts offer clear resources for students, parents, and teachers outlining the AI systems in use or recommended for classroom use.<sup>27</sup> While not every school district offers such a resource, this appears to be the case at least for many larger school districts located in metropolitan areas across the US. We were able to leverage this high level of visibility into potential developer activity in Colorado for our analysis.

---

<sup>21</sup> See Rand Corporation, *Using Artificial Intelligence Tools in K-12 Classrooms* (Apr. 17, 2024), available at [https://www.rand.org/pubs/research\\_reports/RR956-21.html](https://www.rand.org/pubs/research_reports/RR956-21.html), at 2

<sup>22</sup> See Rand, *supra* at 10

<sup>23</sup> See Wething, H., Working Economics Blog, *Today’s Teacher Shortage is Just the Tip of the Iceberg, Part I* (Oct. 9, 2024), available at <https://www.epi.org/blog/teacher-shortage-part1/>; see also See The White House, Exec. Order: Advancing Artificial Intelligence Education for American Youth (April 23, 2025) (“we must also invest in our educators and equip them with the tools and knowledge to ... utilize AI in their classrooms to improve educational outcomes. Professional development programs focused on AI education will empower educators to confidently guide students through this complex and evolving field”).

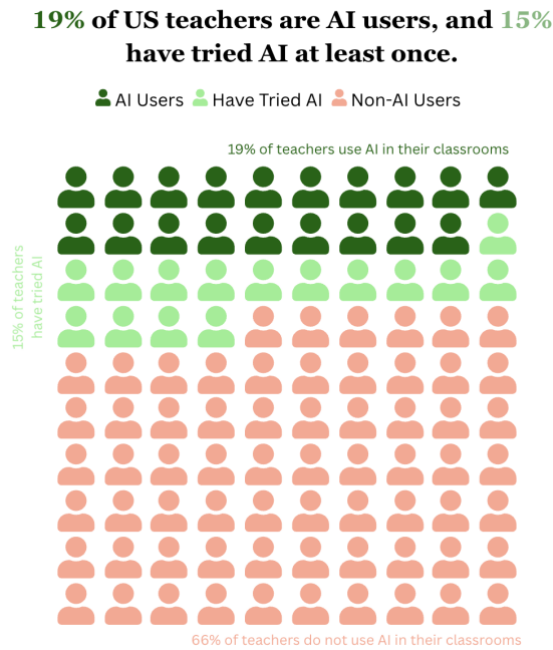
<sup>24</sup> See Tiffany S. Tan et al, *State Teacher Shortages 2024 Update*, Learning Policy Institute (August 2024), available at [https://learningpolicyinstitute.org/media/4412/download?inline&file=State\\_Vacancy\\_2024\\_RESOURCE.pdf](https://learningpolicyinstitute.org/media/4412/download?inline&file=State_Vacancy_2024_RESOURCE.pdf) at 1

<sup>25</sup> See *id.* at 4

<sup>26</sup> See Wething, *supra* note 29.

<sup>27</sup> See AI Educator Choice Board. (n.d.). AI educator choice board. Google Drive. <https://drive.google.com/file/d/14FeHDWUKLOM0DJkiOUTVXs65HvKvLB3/view>; DPS EdTech & Library Services. *DPS ETLs resources*. <https://tinyurl.com/DPSEdTechAll-In-Ones>

**FIGURE 1.**<sup>28</sup>



### **Third, the high-stakes nature of EdTech AI.**

A meaningful portion of AI used in educational technology is likely to constitute HRAIS under the AI Act. Given the broad interpretation set forth above, we view the industry as offering a multitude of AI systems that could readily impact student educational terms and opportunities.

Factors such as teacher shortages, the absence of federal AI regulation, and school administrators' shifting attitudes about AI are all likely to continue to shape the use of AI and the degree of its effects on students across Colorado. Many of the AI tools examined in this study promise to automate and accelerate time-consuming tasks or assist directly with instruction; others will monitor activities and help determine educational planning and placement. Education-focused AI could help close equity gaps, but also has the potential to exacerbate them.

## **II.B. Selection of Representative Companies**

We initially identified about twenty EdTech companies with AI offerings. We identified them through three main approaches. First, we leveraged information made available by Colorado school districts about vendors in the 2023-2025 school years. We searched broadly for “Colorado school vendors” using the internet, appreciating that public schools in particular might name “approved vendors” in public-facing lists. As we began to compile many similar types of companies using this technique, we sought to diversify the types of companies we could evaluate. Accordingly, we next searched for “EdTech in Colorado schools” and “AI in EdTech,” from which results we pulled a few additional EdTech companies that we confirmed operated in Colorado. Finally, we identified potential facial recognition technology companies by searching online specifically for a FRT developer operating in Colorado schools. We were

<sup>28</sup> See Rand Corporation, *Using Artificial Intelligence Tools in K-12 Classrooms* (Apr. 17, 2024), available at [https://www.rand.org/pubs/research\\_reports/RR956-21.html](https://www.rand.org/pubs/research_reports/RR956-21.html), at 2

interested in including such a technology within this report because Colorado legislators have aggressively pursued legislation to protect the biometric data of Colorado school students.

We selected eight companies from within our initial group to deeply analyze, choosing vendors that provided sufficient public information about their technologies to indicate they may develop HRAIS used by Colorado schools. Finally, to achieve a diversity in the AI technology types we evaluated, we selected two companies each from four categories: Student Safety, Predictive Analysis, Feedback Generation, and Special Education. We believe these categories reflect the primary ways AI is currently used in schools to directly engage with students and influence their experiences and outcomes.

## **II.C. Analysis of Compliance**

Our initial analysis began with a thorough review of each deployer’s website and other public-facing materials, including webinars, corporate blogs, and public-facing policies. We searched for information that might satisfy the disclosure requirements of the AI Act, including information in public-facing vendor agreements, policy documents, and submissions to regulators and legislators.

In many cases, relevant details were scattered across different webpages, included on government websites, or embedded within blog posts, YouTube video transcripts, and corporate policies, making them difficult to locate. Some company websites included a search function, which was useful for identifying key terms (in addition to company-specific terms, we searched for “bias” and “AI responsibility” as well as for policy repositories).

When available, we reviewed each company's privacy policies, AI responsibility policies, and technology Terms and Conditions for pertinent information. Finally, we gathered insights from product demos, webinars demonstrating use of the technologies, and from direct access to technology platforms.

The information gathered included details on each of the developer’s companies as well as on the specific technologies each developer provides. We used this public-facing information to assess whether the developer’s technology might be considered a HRAIS and under what circumstances. We also assessed the existence of policies and practices that might be relevant for compliance under the AI Act. Where possible, we identified potential areas where a developer’s technologies may implicate risks for algorithmic discrimination. As we describe below in sections III and IV, in some cases, developers provided descriptions or disclosures that tended to improve their compliance posture under the AI Act, but our work nonetheless identified vulnerabilities in terms of potential for algorithmic discrimination.

## **II.D. Limitations of our Methods**

The primary limitation of our research was our lack of consumer access to many of the AI systems. Several tools require paid subscriptions or restricted full account creation to verified school faculty or staff, making direct access difficult. As a result, the information presented below—unless otherwise noted—is based solely on publicly available materials provided by the deployers, paired with demos in many cases. This limits our ability to fully assess the selected companies’ readiness to comply with the AI Act’s disclosure requirements to other developers and deployers, as such information is not required to be publicly accessible.

A second key limitation stems from the fact that the AI Act does not take effect until February 2026, therefore developers are not yet obligated to comply with its provisions. Additionally, a task force in

Colorado is currently reviewing the statute with an eye toward potential reforms.<sup>29</sup> While there appears to be consensus on some minor definitional updates, several more substantive issues remain unresolved—but at bottom, whether any changes are made, and what they are likely to be, is still unclear. As a result, developers may be hesitant to plan for certain compliance obligations set forth under the current version of the AI Act.

### III. Representative Companies and their Compliance- Part I

The following three categories have only been subject to document-based review, as we were not able to access the technologies directly. The fourth category will be discussed later in Part II, both in terms of document-based review and a technical analysis, as we were able to use the AI systems.

#### III.A. Representative Companies, Categories I-III

##### III.A.i Category I: Student ~~Surveillance~~ Safety

###### **Verkada**

*Representative Colorado Deployers: Steamboat Springs School District (Steamboat Springs, CO); East Grand School District, Granby, CO*<sup>30</sup>

###### HOW DOES VERKADA’S TECH WORK?

Verkada offers an AI-powered cloud-based security system that leverages language and vision models to allow users to search for relevant camera footage of people or events by uploading images or using everyday language.<sup>31</sup> Users can monitor camera systems in real-time<sup>32</sup> or search historical footage with phrases like “find Caucasian man with red hat and backpack on second floor on April 10” to conduct investigations.<sup>33</sup> In addition to using freeform text with AI-powered search, organizations can upload a photo of a person, vehicle, or object to the system to obtain relevant search results. This functionality pulls multiple aspects of an image, then delivers the likeliest set of matches from the school’s camera fleet (a “one-to-many” search).<sup>34</sup> Users can also prompt the system to flag when a specific person, such as an expelled student or a non-custodial parent, enters the premises (a “one-to-one” type search).<sup>35</sup>

---

<sup>29</sup>Legislative Council Staff. (2025). *Report and recommendations: Artificial Intelligence Impact Task Force. Colorado General Assembly.* [https://leg.colorado.gov/sites/default/files/images/report\\_and\\_recommendations\\_0.pdf](https://leg.colorado.gov/sites/default/files/images/report_and_recommendations_0.pdf)

<sup>30</sup> See Verkada, *Steamboat Springs School District*, <https://www.verkada.com/customers/steamboat-springs-school-district/> (last visited May 1, 2025); Verkada, *East Grand School District*, <https://www.verkada.com/customers/east-grand-school-district/> (last visited May 1, 2025). Unlike with respect to other identified representative companies, we identified Verkada by intentionally searching online for an FRT developer operating in Colorado. We were interested in including such a technology within this report because Colorado legislators have aggressively pursued legislation to protect the biometric data of Colorado school students. As of this writing, a bill that would regulate the use of facial recognition in K-12 education awaits Gov. Polis’ signature. See Anthony Kimery, *Colorado moves to regulate FRT in schools with groundbreaking legislation* (Apr. 10, 2025), <https://www.biometricupdate.com/202504/colorado-moves-to-regulate-frt-in-schools-with-groundbreaking-legislation>

<sup>31</sup> See Verkada, *Introducing Verkada’s AI-Powered Search: Setting a New Standard for Conducting Investigations* (May 8, 2024), available at <https://www.verkada.com/blog/ai-powered-search/>; Verkada, *Introducing AI-powered search in Verkada Command (Webinar)*, available at <https://www.youtube.com/watch?v=2jRZ3mjXryU>.

<sup>32</sup> Verkada, *One Platform to Simplify Campus Security*, available at <https://docs.verkada.com/docs/higher-education-solutions.pdf>.

<sup>33</sup> See Verkada, *Introducing History Player Search: A New Feature for Finding Subjects of Interest Across Wide Areas* (Feb. 13, 2025), available at <https://www.verkada.com/blog/introducing-people-search-history-player/>.

<sup>34</sup> See Verkada, *Search on Verkada Command*, available at <https://docs.verkada.com/docs/verkada-search-overview.pdf>.

<sup>35</sup> See Verkada, *AI-Powered Search at Verkada* at 6, available at <https://docs.verkada.com/docs/ai-powered-search-whitepaper.pdf>.

**FIGURE 2.**



**Person of Interest notifications**

Proactively set alerts for when a person is detected onsite who has a face that matches a preselected individual.



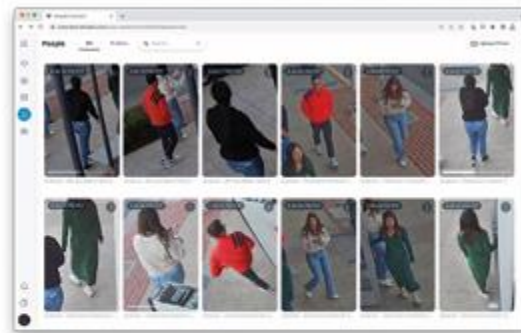
**Face search**

Search across all your organization's cameras for people that have been detected as a match to an existing or uploaded face.



**Person attributes**

Quickly search through your cameras to find people that match selected attributes, including clothing color, gender appearance and more.



**Person history**

View all people that have been captured by your Verkada cameras to speed up investigations.

## THE TECHNOLOGY UNDER THE AI ACT

**Verkada's technology may qualify as a HRAIS** because when deployed, it may constitute a substantial factor in making certain consequential decisions.<sup>36</sup> As described above, administrators can use the technology to identify individuals through reverse image searches or natural language searches. These searches, which make predictions about the identity of students, are poised to inform or assist in making important decisions relating to discipline and educational opportunities when leveraged in administrative investigations relating to crimes or violations of school disciplinary codes (e.g., investigations into vandalism, bullying, vaping, certain pranks, or trespassing activities).<sup>37</sup>

<sup>36</sup> See Colo. Rev. Stat. § 6-1-1701(3), (9)(a).

<sup>37</sup> See *id.* § 6-1-1701(11).

## POTENTIAL FOR ALGORITHMIC DISCRIMINATION

Administrators have reason to be careful about deploying the tool. A 2019 National Institute of Standards and Technology (NIST) study, which examined several leading FRT systems,<sup>38</sup> concluded that such systems have demographic differentials with varying accuracy, depending on a person's gender, age, and race.<sup>39</sup> US-developed algorithms were consistently worse at accurately identifying African-Americans, Asians, and native groups (including Native American, American Indian, Alaskan Indian and Pacific Islanders), relative to Caucasians.<sup>40</sup> With respect to one-to-many matching, false positive rates were highest for Black faces.<sup>41</sup> The researchers also found false positives to be between two and five times higher in women than men.<sup>42</sup> Other studies report results that comport with NIST's findings regarding demographic differential between races.<sup>43</sup>

Given empirical evidence of the existence of demographic differentials in the facial recognition algorithms studied by NIST and others, it is not unreasonable to question whether Verkada's FRT technology also misidentifies students of color more often than their white classmates. Consequences for those students can reasonably range from unnecessarily missing class to due to participation in unwarranted investigations to being suspended or expelled based on misidentification.

## VERKADA'S COMPLIANCE POSTURE

Verkada, for its part, appears to recognize the historical challenges to accuracy within FRT. Verkada has an "AI-Powered Search Moderation Policy," and has made other public-facing statements noting the potential risks of using AI and internal steps taken to avoid them.<sup>44</sup> Specifically, Verkada noted they intentionally chose not to train their models to identify sensitive characteristics.<sup>45</sup> Verkada reports that it deploys "proprietary protections" within its models, including a list of prohibited terms.<sup>46</sup> But it is not clear how these "prohibitions" work in practice. Our team was able to view a demo of the AI-powered search product, and during that demo, a Verkada representative used search terms that referenced skin color with responsive image results. It is possible that search would not have been permitted on a consumer's product. Regardless, Verkada reports to provide administrators with audit logs that can be used to review queries from their school staff and help ensure that users have not attempted inappropriate queries.<sup>47</sup>

Ultimately, Verkada appears to have already taken many of the steps required by the Colorado AI Act to identify known or reasonably foreseeable risks of algorithmic discrimination stemming from the intended use of its technology and provide deployers with the information required under the AI Act.

---

<sup>38</sup> We did not have the opportunity to test Verkada's algorithm in connection with this report; the NIST study did not include Verkada either.

<sup>39</sup> See NIST FRT Report, *supra* note 21.

<sup>40</sup> *Id.*

<sup>41</sup> *Id.*

<sup>42</sup> *Id.*

<sup>43</sup> *Id.* at 11; see also Cook et al., *supra* note 21; GAO, *supra* note 21.

<sup>44</sup> Verkada, *AI-Powered Search Moderation Policy and FAQ* at 1, available at <https://docs.verkada.com/docs/ai-powered-search-moderation-policy.pdf>; see also Verkada, *AI-Powered Search at Verkada* at 6, available at <https://docs.verkada.com/docs/ai-powered-search-whitepaper.pdf>.

<sup>45</sup> Verkada, *User Guide for People Analytics* at 2, available at <https://docs.verkada.com/docs/people-analytics-guide.pdf>.

<sup>46</sup> Verkada, *AI-Powered Search at Verkada* at 9-10, available at <https://docs.verkada.com/docs/ai-powered-search-whitepaper.pdf>.

<sup>47</sup> Verkada *Announcing Audit Logs: Ensuring Compliance and Chain of Custody* (Feb. 20, 2019), available at <https://www.verkada.com/blog/audit-logs-compliance-chain-custody/>.

## GoGuardian (Beacon)

*Representative Colorado Deployers: Aurora Public Schools (Aurora, CO),<sup>48</sup> Boulder Valley Public Schools (Boulder, CO)<sup>49</sup>*

### HOW DO GOGUARDIAN AND BEACON WORK?

GoGuardian provides a suite of tools that enable educators and administrators to monitor student activity in over 10,000 schools when students are logged into their student accounts, generally when using school-issued devices.<sup>50</sup> Beacon is GoGuardian's specialized suicide and self-harm prevention tool for K-12 schools that flags indicia of self-harm, bullying, and contemplated harm to others across all online browser activities, including web searches, email, chats/messaging, videos, documents, and social media.<sup>51</sup>

When Beacon flags “concerning activity,” the tool will take a screenshot of the student’s page and reflect the browsing history “from 15 minutes before and after alerts are generated”, and “highlights specific terms” triggering the alert.<sup>52</sup> Beacon’s highest tier offering is paired with 24/7 human monitoring services.<sup>53</sup> However, schools may customize Beacon to adjust notifications and limit alerts to appropriate administrators or staff.<sup>54</sup> Beacon’s parent resources indicate that students are likely monitored when browsing both in school and at home on a school device but are not likely monitored on personal devices.<sup>55</sup>

---

<sup>48</sup> Aurora Public Schools. (2025). *GoGuardian*.

[https://www.aurorak12.org/divisions/teaching\\_learning/equity\\_in\\_learning/jennifer\\_sheldon/learning\\_resources\\_and\\_educational\\_technology/technology\\_tool\\_support/teachers/essential/go\\_guardian](https://www.aurorak12.org/divisions/teaching_learning/equity_in_learning/jennifer_sheldon/learning_resources_and_educational_technology/technology_tool_support/teachers/essential/go_guardian)

<sup>49</sup>Boulder Valley School District. (2019). More apps added to Single Sign On! <https://www.bvdsd.org/about/news/news-article-spanish/~board/teacher-intranet/post/more-apps-added-to-single-sign-on>

<sup>50</sup>Maass, D., et al. (2023). *GoGuardian: A red flag machine by design*. Electronic Frontier Foundation. <https://redflagmachine.com/RFM-whitepaper.pdf>

<sup>51</sup>GoGuardian. (2021). *GoGuardian Admin product update: Beacon Starter*. <https://www.goguardian.com/blog/product-update-beacon-starter>; GoGuardian, *New Alert Categories Added to GoGuardian Beacon* (2021), available at <https://www.goguardian.com/blog/new-categories-in-goguardian-beacon>.; GoGuardian, *Transcript: Keeping Students Safe Online with GoGuardian Beacon Webinar 2024*, available at [https://goguardian.wistia.com/medias/hpredwrpbk?\\_\\_hstc=181186424.ab43a3f51edb8dc24ed7bd5837e8396b.1742848597177.1743006460720.1744997599622.9&\\_\\_hssc=181186424.3.1744997599622&\\_\\_hsfp=2859479032](https://goguardian.wistia.com/medias/hpredwrpbk?__hstc=181186424.ab43a3f51edb8dc24ed7bd5837e8396b.1742848597177.1743006460720.1744997599622.9&__hssc=181186424.3.1744997599622&__hsfp=2859479032).

<sup>52</sup> GoGuardian, *Understanding Why an Alert was Generated* (2023), <https://support.goguardian.com/s/article/Understanding-Why-an-Alert-was-Generated-1630105917322>.

<sup>53</sup> *Id.*

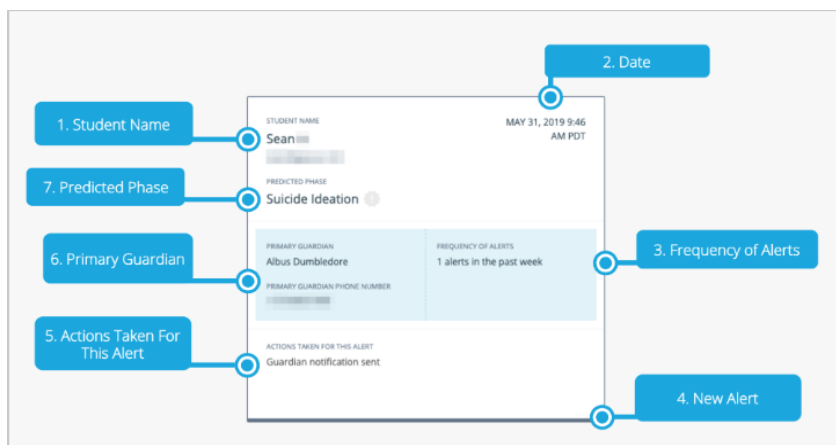
<sup>54</sup>GoGuardian. *Beacon*. available at <https://www.goguardian.com/beacon>.

<sup>55</sup> GoGuardian. (2023). *Beacon parent and guardian communication resources*. <https://support.goguardian.com/s/article/Beacon-Parent-and-Guardian-Communication-Resources-1629762296066>

**FIGURE 3.**<sup>56</sup>

The figure shows a user interface for alert investigation. At the top, there are tabs for 'Rationale', 'Browsing History', 'Previous Actions Taken', and 'Other Alerts'. Below this is a 'SCREENSHOTS' section with a carousel of images. The first image is highlighted and labeled '7:02 PM'. The second image is labeled '7:02 PM' and 'SOURCE OF THE ALERT'. Below the screenshots is an 'ADDITIONAL CONTEXT' section with the text: 'This is the original page that generated the alert: [https://docs.google.com/document/d/12Hj99RMoUh\\_2K6ZxfdSY1OmkvV-HOzyNK6yD/](https://docs.google.com/document/d/12Hj99RMoUh_2K6ZxfdSY1OmkvV-HOzyNK6yD/)'. To the right, a 'WORDS ON THE PAGE' section shows a snippet of text from a document: 'untitled document saved to drive share fileeditviewinsertformattoolsadd-onshelpi. roboto editing to enable screen reader support, press ⌘+option+z to learn about 6 5 4 3 2 1 1 2 3 4 5 6 7 8 i tried it once before, but i didn't get too far i felt a lot of i really wanted was someone to give a little fuck but i waited there forever and not before, and i think i might have messed up i struggled with the veins, and i guess i alive 'cause i didn't really wanna die but nothing very special ever happens in my l a freak, i am afraid that all the blood escaping me won't end the pain and i'll be h, died to be the white ghost of the man that i was meant to be i tried it like before, thought about my friends and the way i didn't give enough and i should've told my son but this life is overwhelming, and i'm ready for the next one take the blade av that all the blood escaping me won't end the pain and i'll be haunting all the lives white ghost. of the man that i was meant to be. yeah i tried it once again, and i thi'. Below the text is a 'TEXT HIGHLIGHT' toggle switch.

**FIGURE 4.**<sup>57</sup>



## THE TECHNOLOGY UNDER THE AI ACT

Beacon’s **technology may qualify as a HRAIS** under the AI Act when it assists in making a consequential decision or generates content or predictions that form the basis of consequential decisions relating to whether a student should be disciplined or investigated in connection with his or her technology use, or when a student is recommended for mental health or similar interventions as a result of online behaviors. Such disruptions and interventions may foreseeably lead to missing class for emergency counseling, involuntary psychiatric evaluations, in-school suspensions, or otherwise impact educational opportunities and terms of enrollment.

<sup>56</sup> GoGuardian, *Understanding Why an Alert was Generated* (2023), <https://support.goguardian.com/s/article/Understanding-Why-an-Alert-was-Generated-1630105917322>.

<sup>57</sup>Maass et al., *supra* note 64, at 19.

## POTENTIAL FOR ALGORITHMIC DISCRIMINATION

GoGuardian publicly claims Beacon has helped prevent "over 18,623 students from physical harm" since 2020; however, this figure reflects notifications generated, not verified incidents requiring intervention.<sup>58</sup> This distinction is crucial when considering the impact of potentially unwarranted notifications. GoGuardian tools have been publicly repudiated for algorithmic inaccuracies resulting in excessive flagging of student activities.<sup>59</sup> The Electronic Frontier Foundation (EFF) has conducted research indicating that GoGuardian's algorithms, including that which powers the Beacon tool, are prone to over-inclusive activity alerts and could flag content that is actually invisible to students.<sup>60</sup>

Such algorithmic errors can lead to unnecessary investigations or interventions, potentially creating disproportionate impacts on students in protected classes. According to a 2022 Center for Democracy & Technology (CDT) report on student activity monitoring (which did not explicitly name or discuss GoGuardian), students from lower socioeconomic backgrounds—who are more frequently students of color<sup>61</sup>—experience monitoring at higher rates than their peers due to greater reliance on school-issued devices for educational access.<sup>62</sup> Students with disabilities similarly depend more heavily on school-issued assistive and adaptive technology and are also likely to face high rates of monitoring relative to their peers.<sup>63</sup>

## COMPLIANCE POSTURE

GoGuardian has made public statements about the "responsible AI" principles it follows, describing that it has evaluated its tools for biases and reliability and has established "Responsible AI Principles" to define and scope its commitment to AI solution development.<sup>64</sup> In a public response to Congressional inquiries, GoGuardian provided descriptions of how it has trained its AI-powered tools and the steps it has taken to mitigate bias in its training data and models.<sup>65</sup> In addition, GoGuardian provides customers with best practices for setting monitoring policies and templates for notifying students, parents or guardians, and the community (such as law enforcement partners) of specific alerts.<sup>66</sup>

---

<sup>58</sup>GoGuardian, *Beacon*, <https://www.goguardian.com/>.

<sup>59</sup> See Maass et al., *supra* note 64, *n* at 5 (describing GoGuardian's algorithm as an "undiscerning yet aggressive webpage scanning formula [that] . . . routinely misidentifies perfectly run-of-the-mill websites as potentially harmful).

<sup>60</sup> *Id.* at 20 ("Beacon was . . . prone to error, reporting several students for actively planning suicide who clearly weren't"); see generally *id.* at 5-6, 18-20.

<sup>61</sup> Ayer, L. et al., *Artificial Intelligence-Based Student Activity Monitoring for Suicide Risk* (2023) at 32, available at [https://www.rand.org/content/dam/rand/pubs/research\\_reports/RRA2900/RRA2910-1/RAND\\_RRA2910-1.pdf?utm](https://www.rand.org/content/dam/rand/pubs/research_reports/RRA2900/RRA2910-1/RAND_RRA2910-1.pdf?utm).

<sup>62</sup> Laird, E. et al., *Hidden Harms: The Misleading Promise of Monitoring Students Online* (2022) at 23, available at <https://cdt.org/wp-content/uploads/2022/08/Hidden-Harms-The-Misleading-Promise-of-Monitoring-Students-Online-Research-Report-Final-Accessible.pdf>.

<sup>63</sup> See Cornell, D. & Maeng, J., Nat'l Crim. Just. Reference Serv., *Student Threat Assessment as a Safe and Supportive Prevention Strategy: Final Technical Report* 24 (2020) <https://www.ojp.gov/pdffiles1/nij/grants/255102.pdf> ("students receiving special education services were 3.9 times more likely to be referred for threat assessment than those not receiving special education services"); Center for Democracy & Technology, *Ableism and Disability Discrimination in New Surveillance Technologies* 15-16 (2002) ("[Surveillance technology t]hreat assessment poses a high risk of disproportionate impact on disabled students . . . who may be more likely to be referred to threat assessment teams . . . research shows that students with disabilities are more likely to be referred to threat assessment than nondisabled students, and that schools"). This second CDT report also did not specifically name or discuss GoGuardian.

<sup>64</sup> GoGuardian, *GoGuardian Reveals Responsible AI Principles* (Sept. 26, 2023), available at <https://www.goguardian.com/blog/goguardian-reveals-responsible-ai-principles>.; GoGuardian, *AI in the Classroom: How GoGuardian Balances Innovation with Student Protection* (Jan. 24, 2025), available at <https://www.goguardian.com/blog/ai-in-the-classroom-how-goguardian-balances-innovation-with-student-protection>.

<sup>65</sup> Ltr. From GoGuardian to Senator Elizabeth Warren, et al (Oct. 26, 2021), [https://www.documentcloud.org/documents/23695654-goguardian-response-\\_re\\_edtech-letter/](https://www.documentcloud.org/documents/23695654-goguardian-response-_re_edtech-letter/).

<sup>66</sup> GoGuardian, *What to Consider: How To Implement GoGuardian Beacon for System-Wide Success* (Mar. 21, 2024), available at <https://www.goguardian.com/blog/how-to-implement-goguardian-beacon-for-system-wide-success>.

## III.A.ii Category II: Feedback Generation

### Amira Learning

*Representative Colorado Deployer: Poudre School District*<sup>67</sup>

#### WHAT IS AMIRA?

Amira Learning is a tool designed to help students improve their reading and speaking skills. Through three steps of Assessing, Instruction, and Tutoring, Amira provides individualized instruction to each student, while also keeping teachers keyed into student’s progress. Amira also provides dyslexia screening, scores students against state and nationwide standards, and helps teachers develop lesson plans.

Amira provides little disclosure as to what specific technology is used in its AI system. However, the general operation of the technology indicates that it leverages algorithms to automate speech recognition, natural language processing, and intelligent tutoring. Amira’s tool “listens” to students read aloud and engages them in open-ended conversations to analyze their speech patterns, vocabulary levels, and listening comprehension<sup>68</sup>. For complete accuracy in assessment, Amira claims to be able to understand any accent or dialect.<sup>69</sup> The Instruct phase uses AI to automate repetitive tasks for teachers, by creating and assigning sequences of lessons to each student or student group to help them grow their reading abilities.<sup>70</sup> Through these steps, Amira is also providing updated, real time insights for the teachers in the form of a dashboard (see Figure 5 below) that can also be used to communicate with the families of students.

#### THE TECHNOLOGY UNDER THE AI ACT

Amira Learning **may qualify as a HRAIS** due to the impact its assessments may have on the educational opportunities of Colorado students. The platform's reading assessments may influence instructional decisions that shape a child's educational trajectory. When these assessments are inaccurate—either overestimating or underestimating a student's reading level or dyslexia risk—they can lead to inappropriate instructional placements.

#### POTENTIAL FOR ALGORITHMIC DISCRIMINATION

While Amira Learning appears to present a relatively low risk of algorithmic discrimination in regards to gender, there are some concerns that warrant consideration. Schools serving predominantly minority populations—where regional accents or dialectical variations are more common—could potentially register artificially lower reading assessment scores by use of the tool. Further, students that have learning differences that affect their ability to read, and difficulties with speech and comprehension, are likely to be negatively impacted by the resulting assessment, with corresponding effects on educational opportunities. However, it is important to emphasize that this analysis represents a

---

<sup>67</sup> See Poudre School District, School Services Contract and On-Demand Providers (listing Amira as a School Service Contract Provider) *available at* <https://www.psdschools.org/academics/academic-resources/student-data-transparency-security/school-services-contract-demand>

<sup>68</sup> Amira Learning. (2025). *Amira ISIP Assess*. Retrieved May 1, 2025, from <https://amiralearning.com/amira-isip-assess>

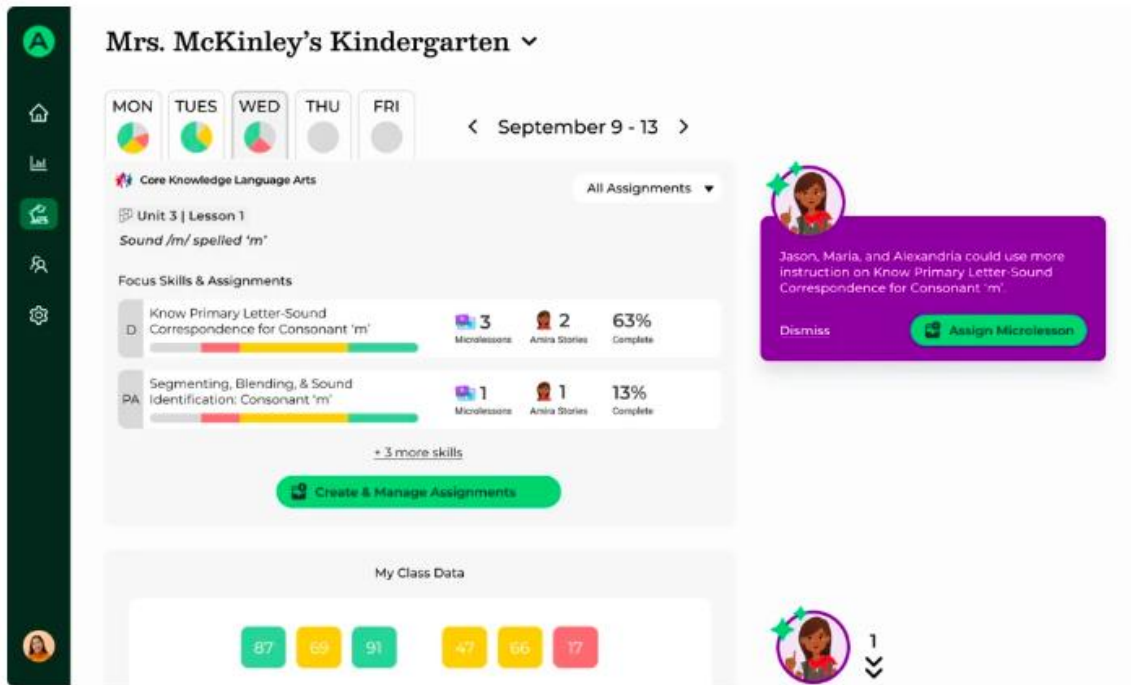
<sup>69</sup> Amira Learning. (2025). *Amira Reading*. Retrieved May 1, 2025, from <https://amiralearning.com/amira-reading>

<sup>70</sup> Amira Learning. (2025). *Amira Instruct*. Retrieved May 1, 2025 from <https://amiralearning.com/amira-instruct>

theoretical concern based on the technology's design and implementation context, not documented evidence of discriminatory outcomes in actual practice.

The platform's teacher dashboard interface raises additional concerns through its use of visual "red flags" (see Figure 5 above, under "My Class Data") to identify struggling readers or students with potential dyslexia. This stigmatizing visual indicator can create lasting negative perceptions, even for students who make significant progress but remain below grade-level benchmarks. Such labeling mechanisms risk creating self-fulfilling prophecies that affect teacher expectations and student self-perception throughout the academic year.

**FIGURE 5.**<sup>71</sup>



## COMPLIANCE POSTURE

Amira Learning has room for improvement with respect to its public facing disclosures. Its public material offers nothing in the way of a clear and succinct summary of the types of HRAISs that Amira has developed. As noted above, Amira is clearly marketed as an AI tool, but it fails to offer public descriptions of how Amira operates. Similarly, there is no mention of any possible risks of discrimination, or what was done to manage these risks as the HRAIS was developed.

<sup>71</sup> Amira Learning. (2025). *Amira ISIP Assess*. Retrieved May 1, 2025, from <https://amiralearning.com/amira-isip-assess>

## Top Hat

*Representative Colorado Deployers:* University of Denver,<sup>72</sup> Colorado State University<sup>73</sup>

### WHAT ARE TOP HAT AND ACE?

Top Hat is a comprehensive education technology platform designed to bolster student engagement and support educators through its collection of offerings. In addition to courseware generation and assessment feedback, it offers automated attendance tracking by geolocating students' devices to confirm attendance. The platform uses this data to provide instant insights to help “at-risk students” and “discipline-specific tools” to educators.<sup>74</sup> According to its website, Top Hat serves over 750 institutions and more than 3 million students across North America.

Introduced in 2023, Ace is Top Hat's AI assistant designed to help instructors with course preparation and enhance lectures with student engagement options like AI-generated assessments. Ace also acts as a study resource for students by generating questions and providing feedback on assessments.

Ace is powered by LLM models developed by OpenAI, the same models that power ChatGPT.<sup>75</sup> Ace operates differently from ChatGPT because course materials on the platform provide Ace with contextual knowledge. Top Hat does warn that, like other GPTs, Ace may generate inaccurate responses, and answers may also contain stereotypes and societal biases.<sup>76</sup> According to a Top Hat representative, AI performs an indirect role in assessing at-risk students and discipline-specific tools by helping analyze data to identify trends and inform recommendations, but it does not make decisions independently.<sup>77</sup>

### THE TECHNOLOGY UNDER THE AI ACT

Top Hat's Ace technology **may qualify as a HRAIS** under the AI Act in limited circumstances when it generates recommendations that inform student assessments for educational placement or academic status, enrollment, or other educational opportunities. When using Top Hat, AI-powered features were clearly labeled and in some cases included disclaimers, such as “AI responses may occasionally be inaccurate” (see Figure 6 below). We were not able to determine how Top Hat's discipline-specific tools and any other AI-generated insights are displayed to educators or administrators.

### POTENTIAL FOR ALGORITHMIC DISCRIMINATION

Ace may provide inaccurate feedback or reinforce any gender- or race-based stereotypes reflected in its LLM training data. If instructors become over-reliant on AI-generated course content, assessments,

---

<sup>72</sup> DU Ed-Tech Knowledge Base. (2025). *Getting Started: Top Hat for Students*. Retrieved May 3, 2025, from <https://otl.du.edu/knowledgebase/getting-started-students/>

<sup>73</sup>Top Hat. (2017). *Top Hat Partners with Unizin Consortium to Bring Affordable and Accessible Digital OER Content to 25 Universities*. Retrieved May 3, 2025, from <https://tophat.com/press-releases/top-hat-partners-with-unizin-consortium-to-bring-affordable-and-accessible-digital-oer-content-to-25-universities/>

<sup>74</sup> Top Hat. (2025). *Features*. Retrieved May 1, 2025, from <https://tophat.com/features/>

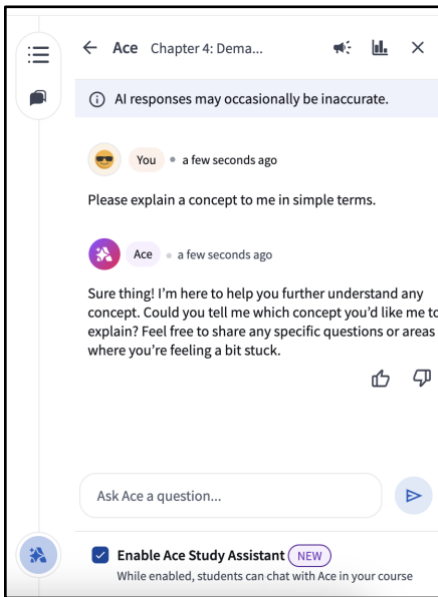
<sup>75</sup>Top Hat. (2025). *Top Hat Ace*. Retrieved May 2, 2025, from <https://tophat.com/features/ace-ai/>

<sup>76</sup> Top Hat. (2025). *Features*. Retrieved May 1, 2025, from <https://tophat.com/features/>

<sup>77</sup> Top Hat Sales Representative, email message dated April 16, 2025 (on file with author).

and feedback to students, they may leave students with insufficient guidance and allow potential flaws in these automated tools to have an outsized effect on education.

**FIGURE 6.**



## COMPLIANCE POSTURE

Top Hat publicly discloses the nature of its AI offering and how AI is used in its tools, including through in-product indications of AI use. Top Hat has published “AI Guiding Principles” on its website, stating its commitment to “ethical and responsible AI-powered innovation.”<sup>78</sup> It acknowledges the risks of improper use of AI to its customers, including false or inaccurate AI-generated information, bias, and intellectual property infringement.

### III.A.iii Category III: Predictive Analysis

#### Civitas Learning

*Representative Colorado Deployers: University of Colorado Denver*<sup>79</sup>

#### WHAT IS CIVITAS?

Civitas Learning is an edtech platform that offers tools for data analytics, workflows, and predictive modeling to inform planning and decision-making at higher education institutions. The company’s stated purpose is to help colleges and universities improve student outcomes, retention, and graduation rates using data-driven insights.

Civitas’ platform integrates with institutional systems to provide AI-driven predictive modeling to forecast student outcomes, including persistence and completion rates. The insights derived from its

<sup>78</sup>Top Hat. (2025). *AI Guiding Principles*. Retrieved May 1, 2025, from <https://tophat.com/ai-guiding-principles/>

<sup>79</sup> University of Colorado Denver. (2025). Success Technologies. Retrieved May 3, 2025, from <https://www.ucdenver.edu/offices/sess/student-success/Initiatives/Technologies>

predictions are designed to identify at-risk students by flagging academic performance or progress to enable timely interventions. The platform also predicts retention and program completion based on various predictors, including GPA, background, demographics, financial aid status, and standardized test scores, along with an option to add custom categories.<sup>80</sup> A table of categories and typical associated predictors is shown in Figure 7.

Civitas publicly discloses the data types, data sources, and primary algorithm used to train its prediction model. Its data modeling and machine learning platform uses data from each institution’s historical data to reflect the unique student experience on different campuses. Individual student predictions are therefore somewhat established based on historical data from institutions. When scored, predictions are grouped into five categories: Very Low, Low, Moderate, High, and Very High, and results can be filtered to target priority groups for intervention.

**FIGURE 7.**<sup>81</sup>

These are the categories and typical predictors for each category:

CATEGORY	Examples of Powerful Predictors
Academic Performance (GPA)	GPA (Cumulative), Change in GPA
Academic Progress	Credits Attempted (Cumulative), Terms Completed (Cumulative)
Area of Study	Degree Program, Major Changes (Cumulative)
Background/Demographics	Age, Gender, Is Veteran
Engagement (LMS)	Average Count of Discrete Days of Any LMS Activity
Enrollment	Average Number of Days Enrolled Before Start (Current Term)
Financial Aid	Loan Amount (Current Term), Financial Aid (Cumulative)
Sentiment	Average Course Net Promoter Score (Cumulative)
Standardized Tests	Max SAT Writing Score
Standing	Academic Standing (Current Term)
Custom	Examples vary

<sup>80</sup> Civitas Learning. (2025). *Civitas Learning: AI-Driven Student Success Platform*. Retrieved May 1, 2025, from <https://docs.civitaslearning.com/app/page/13j9hqU1mG6NCOHhbKBAlfreHCKlz5wnAIQGnCPZZP38?p=1DCISg1eziBcTfBIcksEoY6Fe8qNgrrZn>

<sup>81</sup> Civitas Learning. *Powerful Predictors*. Retrieved May 3, 2025, from <https://docs.civitaslearning.com/app/page/13j9hqU1mG6NCOHhbKBAlfreHCKlz5wnAIQGnCPZZP38?p=1DCISg1eziBcTfBIcksEoY6Fe8qNgrrZn>

## THE TECHNOLOGY UNDER THE AI ACT

If Civitas Learning's AI outputs are used in a manner that directly makes or informs decisions like course enrollment, academic probation, or graduation eligibility, the system **may qualify as a HRAIS** under the AI Act. The classification will likely depend in part on how a given institution leverages the AI tools's insights within its decision-making processes. The insights generated from its predictions are intended to encourage interventions with students and could result in biased treatment or consequential decisions related to students' educational opportunities or outcomes.

## POTENTIAL FOR ALGORITHMIC DISCRIMINATION

Civitas directly addressed algorithmic discrimination in a 2023 white paper titled "Addressing Questions About Bias in Predictive Modeling."<sup>82</sup> The report names mismatched data, poor models, and malicious intent as potential ways predictive models could become biased, then addresses five ways it avoids these pitfalls. Civitas maintains that "demographic variables are not generally among the most predictive and actionable variables in our models."<sup>83</sup> Therefore, eliminating demographics would not likely substantially affect potential bias in predictions. Rather, Civitas Learning emphasizes that its models prioritize real-time behavioral data, like classroom engagement, enrollment patterns, academic performance, and degree alignment. By focusing on factors that students can influence, the models are more likely to provide equitable and actionable information.

How Civitas accounts for bias in institutional historical data while training its predictive models is unclear based on their public-facing materials. Data used in model training (at any stage) could perpetuate stereotypes or bias, because it may reflect past inequities or systemic patterns of exclusion.<sup>84</sup> Finally, it is not clear that this tool could or should be used to inform decisions relating to students with learning differences or disabilities. A disabled student may have, for example, behavioral data and standardized test scores that do not reflect his or her true abilities, and which cannot be properly accounted for when measured against data from other students on his or her campus.

## COMPLIANCE POSTURE

Civitas Learning provides a robust and relatively comprehensive public-facing repository of documentation about its analytics tools, predictive model, and best practices for using its platform. It acknowledges the risks involved in the development and deployment of its predictive models and offers explanations into its mitigating actions, as well as emphasizing the importance of critical human analysis of all results.

---

<sup>82</sup>Civitas Learning. (2025). *Addressing Questions About Bias in Predictive Modeling*. Retrieved May 1, 2025, from <https://docs.civitaslearning.com/app/page/1XMznwyrKfAmqS6dAr6LEHGzQFYvHTy-Xpx04rA8SvF0?p=1DCISg1eziBcTfBlcksEoY6Fe8qNgrZn>

<sup>83</sup> Id.

<sup>84</sup> Gaskins, N. (2023). Interrogating Algorithmic Bias: From Speculative Fiction to Liberatory Design. *TechTrends*, 67(3), 417–425. <https://doi.org/10.1007/s11528-022-00783-0>

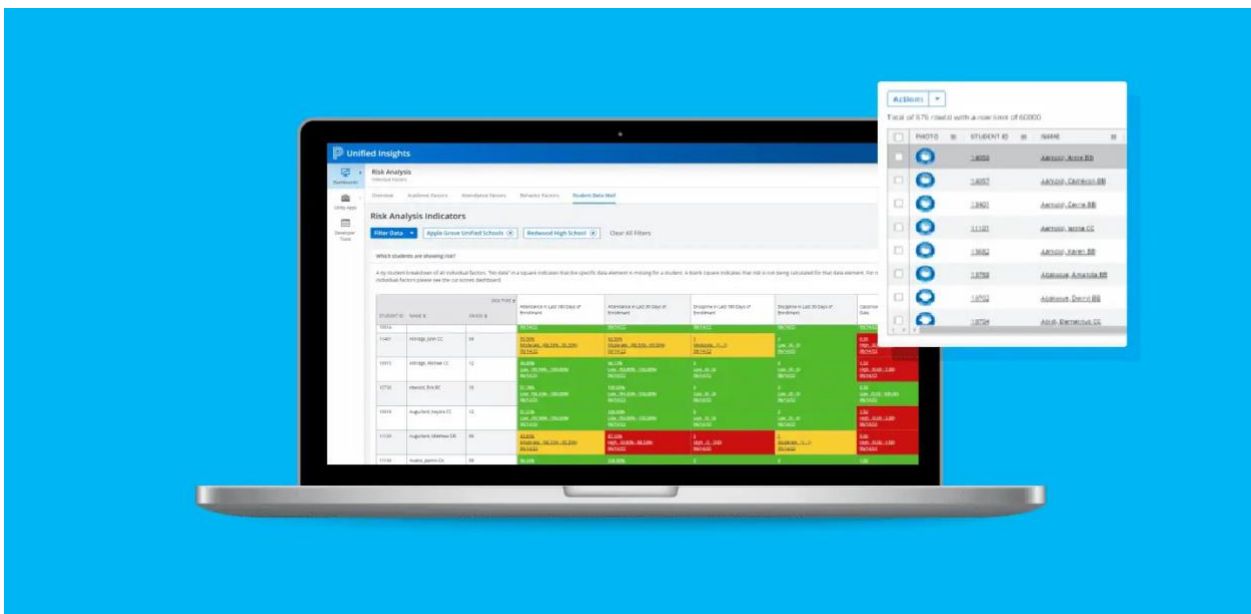
# PowerSchool

*Representative Colorado Deployers: Colorado Military Academy,<sup>85</sup> Colorado Springs School District 11,<sup>86</sup> Eagle County School District<sup>87</sup>*

## WHAT ARE POWERSCHOOL AND RISKANALYSIS?

PowerSchool offers a range of AI offerings aimed at data management and enhancing student learning experiences. At the time of this writing, the only PowerSchool tool that appears likely to implicate the AI Act is PowerSchool RiskAnalysis, which we will focus on. PowerSchool RiskAnalysis uses machine learning algorithms to predict a student's likelihood of graduating on time. Several algorithms use factors such as GPA, attendance rates, credits taken, scores on standardized state assessments, and behavior.<sup>88</sup> The algorithms put students in three categories: high risk, moderate standing, and good standing (see Figure 8 below). School districts can decide what factors to include in the analysis, as well as set their own thresholds for risk levels. RiskAnalysis uses longitudinal data sets, allowing teachers and administrators to segment students and compare their risk at different points in time.<sup>89</sup> The RiskAnalysis platform also allows educators to see what factors have contributed to the risk of students not graduating on time.

**FIGURE 8.<sup>90</sup>**



<sup>85</sup> Colorado Military Academy, *Parent-Cadet Handbook* (2024) available at <https://coloradomilitaryacademy.org/wp-content/uploads/2023/10/23-24-Parent-Handbook-UPDATES-Approved-10.10.23.pdf>

<sup>86</sup> Colorado Springs School District 11, *PowerSchool* (2024) available at <https://www.d11.org/employees/powerschool>

<sup>87</sup> Eagle County School District, *PowerSchool* available at <https://www.eagleschools.net/resources/powerschool>

<sup>88</sup> PowerSchool. (2025). How to boost your students' on-time graduation rates. Retrieved May 1, 2025, from <https://www.powerschool.com/blog/how-to-boost-your-students-on-time-graduation-rates/>

<sup>89</sup> PowerSchool. (2025). *Unified Insights – Risk Analysis*. Retrieved May 1, 2025, from <https://www.powerschool.com/insights/risk-analysis/>

<sup>90</sup> PowerSchool, *PowerSchool Risk Analysis: Watch a Demo!*, available at <https://www.powerschool.com/product-demo/risk-analysis/>

## THE TECHNOLOGY UNDER THE AI ACT

If PowerSchool’s predictive offerings are used to inform decisions like class eligibility, academic probation, or graduation eligibility, those systems **may qualify as HRAIS** under the AI Act. Students who are flagged as high or moderate risk are more likely to be subject to different interventions, which can often disrupt the educational environment or lead to different educational terms or opportunities.

## POTENTIAL FOR ALGORITHMIC DISCRIMINATION

Predictive algorithms trained on sets of data will reflect biases in that data. It is possible that the PowerSchool algorithms will reflect biases directly or through proxies. For example, students who might otherwise have responsibilities - sibling care, a job, elder care - could be missing more classes than students without these responsibilities. One respected study has shown that Asian, Black, and Latinx students take care of family members more frequently than White and non-Latinx students, and further that students from Black, Latinx, and Native American households have lower annual incomes.<sup>91</sup> Students with these additional responsibilities therein could be disproportionately affected by this risk flagging system due to these factors outside of their control.

Finally, this tool may lead to problematic results if used to inform decisions relating to students with learning differences or disabilities. A disabled student may have, for example, GPA, attendance rates, standardized state assessments, or behavioral data that places him or her in a high or moderate risk category only because of deviations from data patterns of non-disabled students.

## COMPLIANCE POSTURE

PowerSchool has a significant amount of literature regarding AI and its safe use. They offer public-facing information on the risks of AI as well as steps that PowerSchool is taking to make sure their systems are operating in a way that mitigates risk. For Educators, resources include a self-test for district readiness to implement AI,<sup>92</sup> a webinar series called “Are You AI Ready?”,<sup>93</sup> as well as a guide for educators to use when assessing whether an AI tool is proper for them.<sup>94</sup> In regards to risk mitigation, PowerSchool has published their “6 Guiding Principles” which address bias in AI, the ethical use of AI, and the importance of keeping humans in the loop.<sup>95</sup>

---

<sup>91</sup> Armstrong-Carter, E., Osborn, S., Smith, O., Siskowski, C., & Olson, E. A. (2023). Middle and high school students who take care of siblings, parents, and grandparents: Associations with school engagement, belonging, and well-being. *AERA Open*, 9, 23328584221140337.

<sup>92</sup> PowerSchool. (2025). *AI Readiness Self-Assessment*. Retrieved May 1, 2025, from <https://www.powerschool.com/ai-readiness-self-assessment/>

<sup>93</sup> PowerSchool. (2025). Are You AI Ready? On-Demand Webinar Series. Retrieved May 1, 2025, from <https://www.powerschool.com/powerschool-ai/are-you-ai-ready/on-demand-webinar-series/>

<sup>94</sup> PowerSchool. (2025). *Evaluating AI applications: Questions to ask* [Infographic]. Retrieved May 1, 2025, from <https://www.powerschool.com/infographic/evaluating-ai-applications-questions-to-ask/>

<sup>95</sup> PowerSchool. (2025). *PowerSchool AI*. Retrieved May 1, 2025, from <https://www.powerschool.com/powerschool-ai/#guiding-principles>

## III.B. Compliance Insights, Part I: Review of Public-Facing Materials

### Transparency and Disclosure Readiness

Across the Colorado EdTech sector, several developers appear well-positioned to meet the disclosure and documentation requirements under the AI Act. While current public-facing materials are fragmented and sometimes incomplete, the level of transparency already present is promising. Developers like Verkada, Civitas Learning, and PowerSchool have made substantial efforts to articulate their AI tools' functions, intended uses, and bias mitigation approaches. Some developers (e.g., Top Hat, PowerSchool, GoGuardian) have also published “AI Principles” or policy statements that align with the public disclosure requirements of the AI Act. This information tends to be decentralized across blogs, white papers, and embedded in demonstration materials, making the companies' AI compliance posture difficult to assess comprehensively as the AI Act's 2026 effective date approaches. However, the fact that these companies operate in an environment already shaped by data privacy, education policy, and federal guidelines appears to support a baseline level of regulatory awareness and responsiveness.

### Algorithmic Fairness and Reasonable Care Concerns

The representative companies have inconsistently addressed the AI Act-specified duty to exercise reasonable care to prevent algorithmic discrimination, largely due to the varying functions of their respective AI-powered technologies. Several EdTech tools demonstrate concerning patterns that may produce adverse effects for students. GoGuardian's Beacon has faced criticism for generating false positives triggered by contextually irrelevant text strings or culturally sensitive searches. Civitas Learning relies on historical datasets that potentially encode pre-existing institutional biases, particularly when such data reflects inequitable treatment of marginalized student populations. Tools like Top Hat's Ace and the IEP generators that will be discussed in section V incorporate LLMs into classroom tools and operate in a manner that illustrates how biases from model training can seep into educational contexts if uncritically applied. These findings raise concerns about educational institutions potentially over-relying on automated systems that generate biased recommendations that shape student outcomes without appropriate human oversight and intervention protocols.

## IV. Representative Companies and their Compliance, Part II

We used a technical approach to complement our public-facing material review of the two special education deployers selected as representative companies. Both deployers offer AI tools to address the needs of students with disabilities - particularly by developing Individualized Education Plans (IEPs).

IEPs are legally binding documents created for students who qualify for special education services under the Individuals with Disabilities Education Act (IDEA).<sup>96</sup> In Colorado, IEPs are developed through a collaborative process involving parents, teachers, special education providers, and sometimes the student, where the team identifies the child's present levels of performance, establishes measurable annual goals, and determines necessary accommodations and services. Once established, the IEP guides the delivery of specialized instruction and related services that help students access the general education curriculum. Colorado's Department of Education provides oversight and guidance on IEP implementation, while requiring schools to review each IEP at least annually and conduct comprehensive reevaluations every three years to ensure continued appropriateness of services.<sup>97</sup>

### IV.A. Methodology - Technical Analysis

#### IV.A.i. Research Design and Data Collection

Our evaluation followed a mixed-methods approach. First, we generated 200 IEPs based on 20 different prompts, each requesting an IEP for a sample student with a distinct disability. Disability types were drawn from the categories recognized under the U.S. Department of Education's regulations on the IDEA.<sup>98</sup> To test the tailoring of accommodations, we selected 10 primarily neurological disabilities and 10 primarily physical disabilities. For each category, we generated multiple IEPs by systematically varying demographic inputs while keeping all other parameters consistent.

The demographic variables included:

- Gender: female, male
- Race/ethnicity: Asian, Black, White
- Grade level: elementary (K-5), middle (6-8), and high school (9-12)
- Religion: Christian, Muslim, Jewish, and Atheist.

We generated each IEP using identical disability descriptions and behavioral patterns, with only demographic variables altered between iterations. This controlled approach allowed us to isolate the effects of demographic factors on the AI tool's outputs.

---

<sup>96</sup> 20 U.S. Code § 1400 et seq., see U.S. Department of Education. (2025). *Individuals with Disabilities Education Act (IDEA)*. Retrieved May 1, 2025, from <https://www.ed.gov/laws-and-policy/individuals-disabilities/idea>

<sup>97</sup> Colorado Department of Education, *IEP Procedural Guidance* (2017) available at [https://www.cde.state.co.us/cdesped/iep\\_proceduralguidance](https://www.cde.state.co.us/cdesped/iep_proceduralguidance)

<sup>98</sup> U.S. Department of Education. (2018). *Sec. 300.8 Child with a disability*. Retrieved May 1, 2025, from <https://sites.ed.gov/idea/regs/b/a/300.8>

**FIGURE 9.**

Grade level: \*  
9th grade

Disability Category: \*  
Tourette Syndrome

Description of student behaviors, needs, and strengths: \*  
15-year-old female student with Tourette syndrome requiring comprehensive support in an educational setting. The student shows strong cognitive abilities but faces challenges with tic management and related social difficulties. They need support with:  
Tic management strategies  
Classroom accommodations for focus and distractions  
Peer education and sensitivity training  
Coping strategies for stress and anxiety

*female*

Grade level: \*  
9th grade

Disability Category: \*  
Tourette Syndrome

Description of student behaviors, needs, and strengths: \*  
15-year-old male student with Tourette syndrome requiring comprehensive support in an educational setting. The student shows strong cognitive abilities but faces challenges with tic management and related social difficulties. They need support with:  
Tic management strategies  
Classroom accommodations for focus and distractions  
Peer education and sensitivity training  
Coping strategies for stress and anxiety

*male*

## IV.A.ii. Analytical Framework

Our evaluation framework consisted of two analytical phases:

### Phase 1: Adversarial Prompt Testing

Firstly, we implemented adversarial prompt testing, which involves deliberately crafting inputs to probe the AI tool’s response mechanisms and ethical boundaries.<sup>99</sup> This included:

- Subjecting the tool to a series of tests to assess its robustness and identify vulnerabilities.<sup>100</sup>
- Identifying how the system responded to personal student information.
- Documentation of responses to various edge cases.

### Phase 2: Comparative Bias Analysis

We evaluated the system for potential biases by comparing IEP language, goal structures, and recommended accommodations across demographic profiles. Each IEP underwent coding for:

- Language patterns and framing (deficit-based versus strengths-based approaches)
- Academic expectations
- Behavioral characterizations
- Recommended interventions

<sup>99</sup> Xu, X., Kong, K., Liu, N., Cui, L., Wang, D., Zhang, J., & Kankanhalli, M. (2023). An LLM can fool itself: A prompt-based adversarial attack. *arXiv*. Retrieved from <https://arxiv.org/abs/2310.13345>; Gu, Y. (2024). A peek behind the curtain: Using step-around prompt engineering to identify bias and misinformation in GenAI models. *arXiv*. Retrieved from <https://arxiv.org/pdf/2503.15205>

<sup>100</sup> Grosse, K., Manoharan, P., Papernot, N., Backes, M., & McDaniel, P. (2017). On the (statistical) detection of adversarial examples. *arXiv*. Retrieved from <https://arxiv.org/abs/1702.06280>

## Phase 3: Content Analysis

We implemented qualitative and quantitative analytical methods, which included:

- Thematic coding of content patterns by two independent reviewers (achieving 87% inter-rater reliability)
- Sentiment analysis to identify recurring patterns in language, tone, and recommendations
- Measurement of key linguistic metrics, including keyword analysis and agency attribution analysis (whether challenges were framed as internal or external to the student)

## Phase 4: Practical Usability Check of the of User Interface (UI/UX Audit)

We reviewed the AI tool's interface and user experience design against established usability principles:

- Usability and User Support
- Error Prevention and Handling
- Flexibility and Input Limitations

This methodological approach provided a comprehensive framework for evaluating both the technical performance and ethical implications of AI-generated special education documentation.

## IV.B. Representative Companies, Categories I-III

### IV.B.i. Category IV: Special Education

#### MagicSchool

*Representative Colorado Deployers: Aurora Public Schools, Denver's South High School<sup>101</sup>*

#### WHAT IS MAGIC SCHOOL AI?

MagicSchool is one of the leading educational AI platforms. It serves five million teachers from nearly every US school district and 160 countries globally. According to MagicSchool, educators using its generative AI tools save up to 10 hours each week through automating administrative tasks.<sup>102</sup>

Magic School implements these generative capabilities through two main approaches: chat-based interactions and template-based text generation. To collect student data, it uses drop-down menus and text fields. It then funnels these personal details about a child's disability, behaviors, and demographics into proprietary black-box algorithms. These algorithms automatically assign different AI models to various parts of each task. Behind a user-friendly interface, MagicSchool routes sensitive student data through multiple AI models like ChatGPT, Claude, and Gemini.

---

<sup>101</sup> MagicSchool AI. *Aurora Public Schools case study*. <https://www.magicschool.ai/case-studies/aurora-public-schools>, Brady, N. (2024, August 26). *Denver Public Schools focusing on safety as schools incorporate more artificial intelligence*. Denver7. <https://www.denver7.com/lifestyle/education/denver-public-schools-focusing-on-safety-as-schools-incorporate-more-artificial-intelligence>

<sup>102</sup> MagicSchool. (2024). *Use cases for special education teachers*. <https://www.magicschool.ai/use-cases/special-education-teachers>; MagicSchool (2025). *master tools list: List of all MagicSchool tools, tool name and description*. Google Docs. [https://docs.google.com/spreadsheets/d/1ZTK7X2D8R\\_uqYW8SF9sejnkpBCQzMz\\_qVXS9G\\_Rhag0/edit?gid=0#gid=0](https://docs.google.com/spreadsheets/d/1ZTK7X2D8R_uqYW8SF9sejnkpBCQzMz_qVXS9G_Rhag0/edit?gid=0#gid=0)

Our analysis focuses specifically on MagicSchool's IEP generator tool. We signed up on MagicSchool AI's platform, which granted us access to evaluate the tool directly. This hands-on testing approach enabled us to observe how the IEP tool generates drafts based on real input parameters, including grade level, disability classification, and student-specific strengths and challenges. The generated drafts typically include present levels of academic performance, measurable goals, and suggested accommodations and modifications. This practical examination provided valuable insights into how the tool functions in real educational contexts.

## THE TECHNOLOGY UNDER THE AI ACT

The MagicSchool AI IEP tool **may qualify as a HRAIS** because it directly influences consequential educational decisions. IEPs determine essential services and accommodations for students with disabilities, and this platform collects and processes sensitive data, including detailed behavioral, diagnostic, and demographic information about children with disabilities.

## POTENTIAL FOR ALGORITHMIC DISCRIMINATION

### *Standardized Metrics*

Our analysis of measurable goals and objectives reveals a concerning pattern of standardization that undermines the fundamental "individualized" intent of IEPs. A key to understanding the recommendations of an IEP is to understand the use of a "success metric." A success metric is the specific criterion used to measure whether a student has mastered a goal. It is typically expressed as a percentage of correct responses, frequency of behavior, or level of independence required. For example, "75% success" means the student must perform the specified skill correctly 75% of the time over a designated period or number of attempts. Other common forms include frequency (how often), duration (how long), or level of support/prompting needed. These metrics are usually tailored to each student's unique needs, baseline performance, and the specific skill being targeted. However, we found that an overwhelming majority of AI-generated IEPs default to the same arbitrary threshold:

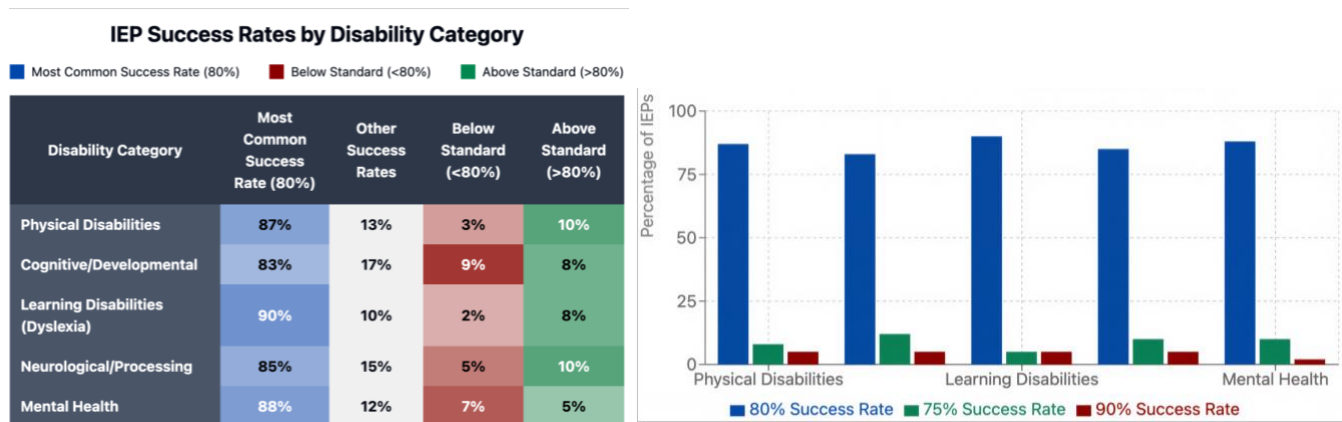
- 87% of IEPs for physical disabilities use an 80% success metric
- 90% of IEPs for learning disabilities use an 80% success metric
- 88% of IEPs for mental health conditions use an 80% success metric.

Notably, these standardized 80% metrics appear in the AI-generated IEPs even though the prompts do not explicitly suggest or instruct them. This finding raises an important question: What determines when the system deviates from this default? The implications are significant: 13% of students with physical disabilities, 10% with learning disabilities, and 12% with mental health conditions are assigned goals either below or above 80% without any documented contextual justification or individualized rationale.

Human-written IEPs also frequently use the 80% benchmark. This pattern likely stems from a combination of institutional training, regulatory pressures, and the desire for measurable, defensible

goals, rather than thoughtful personalization based on individual student needs.<sup>103</sup> Education scholars and legal experts have long challenged defaulting to 80% success metrics in IEPs, with research showing that truly individualized education requires metrics tailored to each student's specific abilities and the nature of the skill being taught rather than arbitrary standardized thresholds.<sup>104</sup> For instance, when 90% of learning disability IEPs use this rigid metric, students face a stark dilemma: some encounter unchallenging goals well below their capabilities, while others struggle with unrealistic targets they cannot reasonably achieve. Ultimately, this one-size-fits-all approach, especially when driven by AI systems, reflects algorithmic convenience over educational customization, which may compromise the educational and emotional well-being of students the system is meant to support.

**FIGURE 10.**



### ***Disability-Specific Language Patterns***

In understanding how different disability types are framed, our analysis adopted established frameworks from special education research that distinguish between strength-based and deficit-based language.<sup>105</sup> The strengths-to-deficits ratios (S:D) represent the relative balance between positive, capability-focused, and problem-focused language, measured on a 10-point scale. Mental health conditions such as anxiety disorder and Tourette syndrome sometimes demonstrated the strongest deficit-focused language (5:9 ratio), where statements like "experiences heightened stress" and "avoidance behaviors impact attendance" significantly outweighed strength-focused language such as "demonstrates understanding when supported." By contrast, physical, learning, and autism spectrum

<sup>103</sup> Goran, L., Harkins Monaco, E. A., Yell, M. L., Shriner, J., & Bateman, D. (2020). Pursuing academic and functional advancement: Goals, services, and measuring progress. *TEACHING Exceptional Children*, 52(5), 333-343

<sup>104</sup> Diehm, E. (2017). Writing measurable and academically relevant IEP goals with 80% accuracy over three consecutive trials. *Perspectives of the ASHA Special Interest Groups*, 2(16), 34-44.

<sup>105</sup> Elder, B. C., Rood, C. E., & Damiani, M. L. (2018). Writing strength-based IEPs for students with disabilities in inclusive classrooms. *International Journal of Whole Schooling*, 14(1), 116-155.; Garwood, J. D., & Ampuja, A. A. (2019). Inclusion of students with learning, emotional, and behavioral disabilities through strength-based approaches. *Intervention in School and Clinic*, 55(1), 46-51.

disabilities showed a moderate deficit emphasis (6:8 ratio), while complex neurological conditions uniquely maintained a balanced approach (7:7 ratio).

This recorded pattern likely reflects broader societal stigma surrounding mental health. Prior studies have shown that conditions like anxiety and Tourette syndrome are often viewed through a more pathologizing lens than other disability types.<sup>106</sup> Such differential framing may contribute to stigmatization and reinforce a hierarchy of disability, where some conditions are presented more negatively than others.<sup>107</sup> This finding aligns with our broader analysis of how the model treats different disability categories systematically rather than individually.

Some generated IEPs show differences in male and female student outputs. For students with identical disabilities, the AI sometimes describes male students using terms related to performance, achievement, and independence. At the same time, it portrays female students as needing support tools and emotional regulation.

For instance, let's compare two responses about autism from the same prompt:

*"[Student Name] is a 13-year-old female student in 7th grade diagnosed with Autism Spectrum Disorder. She demonstrates strong interests in specific subjects, which can be leveraged to enhance her learning experience. However, she faces challenges with social communication....."(Document 16, **Female Autism IEP**)*

*"[Student Name] is a 13-year-old male student in 7th grade who demonstrates strengths in specific interest areas, particularly in subjects related to technology and science. However, they experience difficulties with social communication..."(Document 16, **Male Autism IEP**)*

Consider the same for dyslexia, both responses from the same prompt.

*"[Student Name] will utilize graphic organizers to plan and write paragraphs, achieving at least 4 out of 5 on a writing rubric." (Document 19, **Female Dyslexia IEP**)*

*"[Student Name] will read grade-level texts aloud with 90% accuracy and appropriate pacing during assessments." (Document 19, **Male Dyslexia IEP**)*

Despite concerns in other areas, the application demonstrates a level of accuracy in creating tailored accommodations specific to each disability. For instance, dyslexia IEPs appropriately focus on reading support, writing support, and self-advocacy, autism spectrum disorder IEPs correctly emphasize social

---

<sup>106</sup> Shogren, K. A., Wehmeyer, M. L., Schalock, R. L., & Thompson, J. R. (2016). Reframing educational supports for students with intellectual disability through strengths-based approaches. In *Handbook of research-based practices for educating students with intellectual disability* (pp. 25-38). Routledge.

<sup>107</sup> Volkow, N. D., Gordon, J. A., & Koob, G. F. (2021). *Choosing appropriate language to reduce the stigma around mental illness and substance use disorders*. *Neuropsychopharmacology*, 46(13), 2230-2232.; Malli, M. A., & Forrester-Jones, R. (2022). *Stigma and adults with Tourette's syndrome: "Never laugh at other people's disabilities, unless they have Tourette's—Because how can you not?"*. *Journal of Developmental and Physical Disabilities*, 34(5), 871-897.

communication, sensory processing, structured routines, and physical disability IEPs accurately incorporate mobility, adaptive equipment, and accessibility needs.

## COMPLIANCE POSTURE

When MagicSchool's IEP generator provides disability-specific accommodations, it sometimes frames recommendations for students of certain demographic groups and disability types with indicia of underlying bias, a fact which should trigger scrutiny by deployers. The standardization of success metrics (87-90% of IEPs defaulting to 80% thresholds) and gender-based language patterns documented in our analysis also suggests that the tool could be used in a manner that may result in algorithmic discrimination. For their part, MagicSchool has implemented some mitigation efforts, including a set of platform best practices on its webpage that advises teachers to "check for bias and accuracy" and to view AI outputs as starting points rather than final solutions. However, this warning appears only the first time an educator opens the tool. Similarly, MagicSchool's Terms of Service recommend following the 80/20 rule, with the AI handling most of the work (80 percent) and the teacher responsible for refining the output (20 percent).<sup>108</sup>

### Playground IEP

*Representative Colorado Deployer: Pinnacle Charter School*<sup>109</sup>

#### WHAT IS PLAYGROUND IEP?

Playground IEP is a newer entry into the AI EdTech market - it is distinctly specialized as a support tool for the education of children with special needs or disabilities. This AI deployer markets itself as a tool to streamline, simplify, and "democratize" required processes for special educators, involved gen-ed educators, and administration. The subject of our analysis is Playground IEP's 'AI IEP CoPilot', which generates individualized education plans (IEPs), behavior intervention plans (BIPs), and Multi-Tiered System of Support (MTSS) interventions.<sup>110</sup>

As indicated above, the creation of a child's IEP generally requires IEP teams of multiple support personnel and several days of consultations and evaluations. IEP CoPilot conversely promises to generate an IEP within "seconds" using advanced language models. Teachers input information about students or select pre-set goals and attributes, to receive AI-generated standard lesson plans or interventions. Playground IEP notes on its website that these outputs are purely supplementary tools, but this disclaimer is not present within generated responses. In fact, during the testing process of this tool, the original "AI response" label preceding the output was replaced by "original response").

---

<sup>108</sup> MagicSchool *Terms of Service* (Aug. 1 2024) available at <https://www.magicschool.ai/privacy-security/terms-of-service>

<sup>109</sup> Pinnacle Charter School hosted a Collaborative for Exceptional Education Learning Lab event sponsored by Playground IEP on March 12, 2024. The collaborative supports special populations in Colorado schools. See [https://www.linkedin.com/posts/playground-iep\\_eld-spced-coteaching-learning-lab-pinnacle-activity-7161827725001891841-au-F](https://www.linkedin.com/posts/playground-iep_eld-spced-coteaching-learning-lab-pinnacle-activity-7161827725001891841-au-F) and <https://exceptionalcollaborative.org/special-populations>.

<sup>110</sup> *EP 156: Using AI to write IEP goals and manage service plans (with Sean Klamm) – Dr. Karen Dudek-Brannan.* (2024, April 10). <https://drkarendudekbrannan.com/ep-156-using-ai-to-write-iep-goals-and-manage-service-plans-with-sean-klamm/>

Though the public-facing Playground IEP website does not specify which language models it uses for generative AI output, the associated GitHub page for IEP CoPilot notes that it leverages Chat GPT<sup>111</sup>.

## THE TECHNOLOGY UNDER THE AI ACT

For similar reasons to those discussed above in the context of analyzing MagicSchool's AI IEP tool, Playground IEP's technology **may qualify as a HRAIS** when it contributes to consequential educational decisions through its automated IEP generation.

## POTENTIAL FOR ALGORITHMIC DISCRIMINATION

### *Evaluation of Playground IEP's AI IEP CoPilot tool*

Consistent with the methodology described above, our Playground IEP tool evaluation first focused on adversarial prompt testing: logging and comparing how the system responded to or incorporated sensitive personal information (such as that relating to race and religion). We documented the system's standard responses to each prompt variation, and documented its responses when asked specifically to incorporate any potential gender-, race-, or religion-based needs into IEP goals. We also conducted tests of responses to identical prompts, empty prompts, and prompt injections.

This was supplemented by a heuristic (experience-based) evaluation of the IEP generator user interface, to identify whether this AI tool met key relevant usability principles that would ordinarily assist in limiting disproportionate impacts on educational outcomes: maximal user support, feedback, clarity, and error prevention / handling.<sup>112</sup>

As a general overview of the tool, it should be noted that:

- 100% of observed outputs generated by IEP CoPilot followed the same general structure: an annual goal, three quarterly benchmarks, and three activities per benchmark.
- 90% of Playground IEP's generated outputs could be categorized based on format into one of only three templates.
- 0% of outputs were identical, even when the same prompt was input into the system multiple times.<sup>113</sup>

We flagged the lack of uniformity specifically as a significant potential risk. Despite the standardized templates, the extent of suggested allowances, support, and resources varied significantly across outputs. For example, when prompted to generate IEP goals for a 10-year old Jewish student with ADHD, CoPilot produced one IEP that recommended incorporating Jewish holidays in classroom discussions, including culturally-relevant class material, and providing makeup work for days missed

---

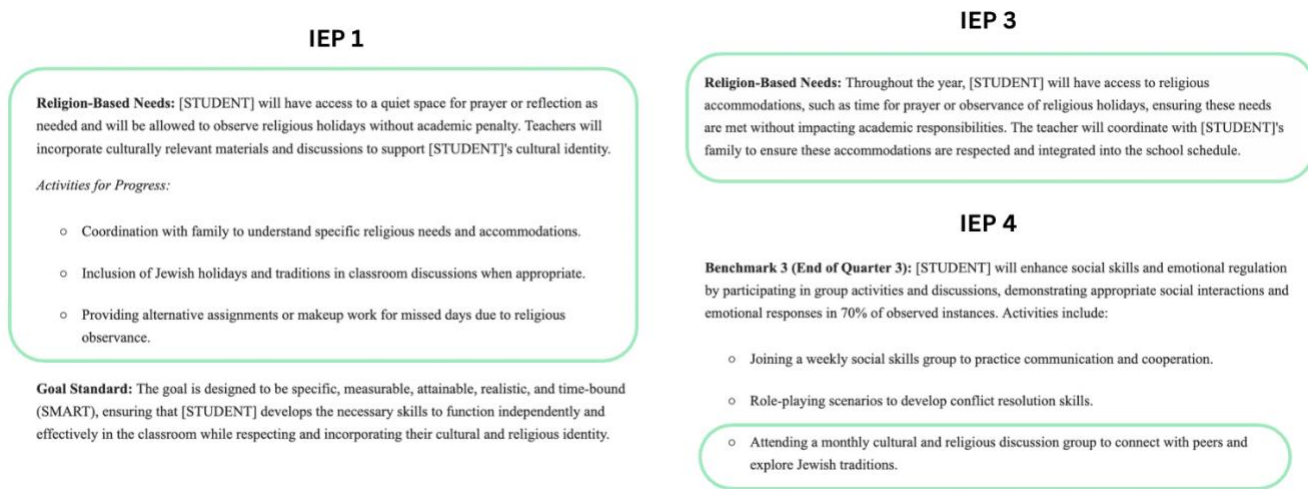
<sup>111</sup>TEDAI-Hackathon. (2025). *IEPCoPilot.AI: AI Co-Pilot for Special Needs Education Planning* [Computer software]. GitHub. <https://github.com/tedai-hackathon/IEPCoPilot.AI>

<sup>112</sup>Stanford University. *Usability principles*. Improvement, Analytics, and Innovation Services. <https://improvement.stanford.edu/resources/usability-principles>

<sup>113</sup>Lack of uniformity is a common result of using LLMs such as ChatGPT, which Playground IEP leverages, known as non-determinism. LLMs are known to produce various outputs even when given the same input - though AI developers can reduce randomness of responses by turning down the LLM's 'temperature' parameter.

due to religious observance - “without academic penalty”. Given the same prompt again, CoPilot produced an IEP that suggested some religious accommodations should be provided, but must be met “without impacting academic responsibilities”. And when prompted again, CoPilot generated an IEP with no religious accommodations at all - instead putting the onus on the student by suggesting they attend a religious discussion group. These consistent variations in accommodations may facilitate unintended preferential treatment across student populations, and would be highly likely to consequently impact a child’s educational outcome.

**FIGURE 11.**



Though variation remained prevalent across prompts, there were no specific trends in keywords, accommodation levels, or language choices based on students’ racial or religious demographic groups.

Notably, CoPilot outputs did still include students’ race or religion when included in the prompt - but outputs had no mention of “male” or “female”. While outputs did include the associated gender pronouns, this avoidance of explicit gendered identification was present even when CoPilot was specifically prompted to consider gender-based needs. This suggests the possibility that the developers of Playground IEP utilized some AI guardrails – such as filters or fine-tuning – to suppress gender in outputs. This prevention of customizing IEPs by gender was possibly intended to avoid discrimination, but may ultimately be a limitation, given that neurological disabilities tend to present differently in students of different genders.

There was one exception to this lack of gender consideration: when the goal category was related to social science. There was no difference between output for students without gender specification and students identified as male. However, if the student was female, CoPilot’s output appeared to focus less on disability accommodations, and centered IEP goals around lessons on the contributions of women. A few tests were also run with the student’s gender specified to be non-binary instead - this similarly resulted in CoPilot output with reduced note of accommodations, and a recommendation to focus on the contributions of non-binary and LGBTQ+ individuals instead. Further work may need to be done to avoid content-based disparity between the IEPs targeted at students of different genders.

## FIGURE 12.

Topic: Female Student with ADHD

Goal Standard: Analyze the roles and contributions of women in historical and contemporary societies, focusing on their impact on social, political, and economic developments.

### IEP Goal:

By the end of the academic year, [STUDENT] will demonstrate an understanding of the roles and contributions of women in history by researching and presenting on three influential women from different historical periods, achieving a score of at least 80% on a teacher-created rubric that assesses research quality, presentation skills, and content understanding.

Topic: Understanding Diverse Perspectives and Contributions in History

Goal Standard: By the end of the academic year, [STUDENT] will demonstrate an understanding of diverse perspectives and contributions in history, focusing on underrepresented groups, including non-binary and LGBTQ+ individuals, by accurately identifying and describing the significance of at least three historical figures or events in written or oral presentations, with 80% accuracy in content and context, as measured by teacher-created rubrics and assessments.

*female*

*non-binary*

Another key concern that arose was IEP CoPilot's response when student information was missing. We ran ten prompts through the system that consisted of only a spacebar click (effectively an empty input box). Not only did the CoPilot still generate an output, it did not prompt the user to fill in additional information - the AI tool instead filled in the missing information itself:

**“To create a comprehensive IEP goal for [STUDENT], I'll need to fill in the missing information. Let's assume the grade level is 5th grade, the topic is “Self-Regulation,” and the goal standard is aligned with CASEL's SEL competencies. Here's a SMART goal with benchmarks”**

## FIGURE 13.

Original Response

[Edit](#) [Copy](#)

To create a comprehensive IEP goal for [STUDENT], I'll need to fill in the missing information. Let's assume the grade level is 5th grade, the topic is “Self-Regulation,” and the goal standard is aligned with CASEL's SEL competencies. Here's a SMART goal with benchmarks:

**Grade Level:** 5th Grade

**Student First Name:** [STUDENT]

**Goal Subject:** Social Emotional

**Topic:** Self-Regulation

**Goal Standard:** Aligned with CASEL's SEL competencies, specifically focusing on self-management skills.

### IEP Goal:

By the end of the academic year, [STUDENT] will demonstrate improved self-regulation skills by independently utilizing at least three coping strategies to manage emotions and behaviors in 4 out of 5 observed situations, as measured by teacher observations and self-reflection journals.

This mirrors issues regarding input restrictions found while conducting heuristic evaluation of Playground IEP's interface. Strict prompt character limits and the need to select pre-identified goal category options limit the potential for educators and administration to use this tool to its fullest potential - and to ensure that sufficient information about their students is available to best inform CoPilot's output. And though each output is followed by an option to “refine” the AI generated

response, selecting this only results in another standardized AI-prompt box with character limits. There is no method of manually refining or correcting one section of the output.

Most relevant to the AI Act, Playground IEP would likely not meet reporting requirements. The tool itself does not offer disclaimers regarding the limitations of AI (unlike Magic School), and does not outline deployer requirements to use the tool effectively or the need to supplement CoPilot's output.

## COMPLIANCE POSTURE

Playground IEP lacks many of the materials offered by other companies featured in this report. It does not provide AI responsibility principles, offer best practices for deployers or disclaimers on limitations within the tool itself, or disclosures about bias or moderation practices. Its Terms of Use state that a user is responsible for “review[ing] all content generated by Playground IEP and modify[ing] it as necessary,” but this requirement is not clearly linked to algorithmic discrimination, nor is it visible while using the tool.<sup>114</sup>

## IV.C. Compliance Insights, Part II: Review of Technical and Public-Facing Materials

### Transparency and Disclosure Readiness

The evaluation of our two final representative companies in Category IV similarly reveals varied levels of apparent preparedness for AI Act disclosure requirements. MagicSchool demonstrates some indicia of transparency by referring to its AI system as an AI tool and by offering platform advice to teachers to “check for bias” during use; it also discourages teachers from viewing the tool's AI outputs as final solutions. However, these warnings are limited in function and potential effectiveness: they appear only once, before the tool's first use. The tool's Terms of Service also recommend that teachers employ an “80/20” approach to drafting but offer no explanation for this rationale, including as to whether any particular model training or bias-testing informs that breakdown. In turn, Playground IEP exhibits concerning transparency gaps, offering neither AI responsibility principles nor disclaimers about the limitations of its tool. Its Terms of Use place responsibility on users to review and modify generated content, and do not explicitly refer to bias concerns or algorithmic discrimination. The companies appear as if they could benefit from a more thoughtful approach to potential upcoming compliance obligations in anticipation of the AI Act's effective date.

### Algorithmic Fairness and Reasonable Care Concerns

Our technical evaluation, paired with our document review, identified certain algorithmic fairness issues that require attention in anticipation of the AI Act's reasonable care obligation. MagicSchool's IEP generator shows concerning standardization patterns that undermine the individualized intent of special education plans. This standardization occurs without contextual justification or individualized rationale, potentially compromising educational customization. The tool also displays differences in framing recommendations based on demographic factors, with varying strength-to-deficit language ratios across disability types and gender-based language patterns that describe male students in terms of performance and independence while portraying female students as needing support tools.

---

<sup>114</sup> Playground IEP *Terms of Service* (May 23, 2024) available at <https://www.playgroundiep.com/terms-of-use>

Playground IEP's AI IEP tool similarly raises algorithmic fairness concerns through its inconsistent outputs. Despite using standardized templates, generated accommodations vary when given identical inputs. While the tool appears to implement guardrails against gender-based discrimination by avoiding explicit gender identification, this approach inadvertently creates apparent limitations, particularly for neurological disabilities that present differently across genders. Moreover, when prompted with student information related to social sciences, the tool produces gender-disparate outputs that reduce focus on disability accommodations for female and non-binary students, instead emphasizing gender-related educational content. Playground IEP also reacts to empty prompts by autonomously generating assumptions about missing student information rather than prompting educators to provide complete data, potentially reinforcing algorithmic bias without human oversight.

These findings highlight challenges these and similar companies may face in meeting the AI Act's reasonable care standards as they work to identify and address potential algorithmic discrimination.

## V. Recommendations

### Educational Resources for Industry in Anticipation of February 2026

#### Effective Date

Clear educational materials and industry engagement opportunities would benefit EdTech developers seeking to understand their obligations under the AI Act. Many developers may not fully appreciate which of their AI offerings qualify as HRAIS under the Act's provisions. Informational resources similar to the model FAQ document provided in **Appendix I** or the sample test prompts provided in **Appendix II** could help companies identify regulated AI systems and implement appropriate compliance measures before the law's February 2026 effective date.

#### Ongoing Stakeholder Engagement

The EdTech industry and regulatory stakeholders would benefit from regular communication forums over the next year through conferences, workshops, and feedback sessions. These structured interactions could help vendors understand implementation requirements in anticipation of the effective date.

Given that interpretations of the law are still evolving, and the law itself may change, these conversations could also serve to inform regulatory approaches and ensure the AI Act's language and application meet its intention: to ensure necessary protections against algorithmic discrimination. Collaborative approaches have proven effective in other regulatory contexts where technology often outpaces traditional regulatory frameworks (*e.g.*, food and drug devices, export controls, and foreign direct investment review regulation).

#### Identify Best Practices as Compliance Models

Identifying and publicizing "best in class" compliance approaches from leading companies across the four different EdTech sectors identified in this report could provide valuable guidance. These types of case studies can serve as practical benchmarks for the broader industry, offering concrete models for documentation and transparency measures.

By highlighting effective disclosures, or strategies for bias detection and impact assessments, the industry can develop clear standards while acknowledging sector-specific challenges. This approach encourages a “race to the top” in compliance quality rather than merely meeting minimum requirements, ultimately better serving Colorado students and educational institutions.

## VI. Future Research Ideas

### Assess “Strategic” Patterns in Compliance Disclosures of HRAIS

Our analysis of current disclosure practices suggests that once the AI Act comes into full effect, companies may employ strategic opacity in their compliance disclosures—similar to patterns observed following the implementation of cookie consent regulations.<sup>115</sup> The AI Act’s lack of specificity regarding disclosure format and prominence creates potential for companies to technically comply while minimizing visibility of critical information. Currently, as mentioned above, most of the disclosures from the developers are spread out inconsistently across their website and user interfaces rather than found in one consistent location, and a user must know what they are searching for to find it. PowerSchool is one of the most clear and centralized, with all relevant information kept under a “PowerSchool AI” tab. Civitas, while it does disclose a great deal of information, does so across blog posts and different documents, with the documents page hidden in the footer of the website. Amira, on the other hand, makes no mention of risks, but does suggest that its AI system performs better than a human counterpart. While deployers may not want a bright warning sticker slapped onto their home pages, parents, students, and educators do have a right to clearly be able to identify which AI systems are high risk.

### Conduct Field Research on Classroom Implementation Patterns

The classification of educational AI tools as HRAIS under the current regulatory framework is substantially contingent upon their practical implementation in educational settings. These technologies cross the threshold into high-risk territory when they functionally replace or substantially inform human judgment rather than augment it. A study or survey of educators utilizing these tools in the classroom could inform an assessment of frequency of utilization and decision-making patterns when AI recommendations conflict with educator judgment.

## VII. Conclusion

This report provides an assessment of the Colorado EdTech industry’s preparedness to meet certain requirements set forth by the Colorado AI Act in anticipation of that law’s February 2026 effective date. We submit that several developers within the industry are positioned to satisfy the disclosure and transparency requirements of the Act, while others require additional preparation and may benefit from clarification regarding their regulatory obligations.

Our also work uncovered daylight between the current functionality of certain AI tools in the space and the apparent goals of the AI Act: in particular, while IEP generators offer time saving opportunities, and their companies appear positioned to make available the disclosures and

---

<sup>115</sup> Li, Danyang, *The FTC and the CPRA’s Regulation of Dark Patterns in Cookie Consent Notices* (2022) *The University of Chicago Business Law Review*: Vol. 1: No. 1, Article 19. Available at: <https://chicagounbound.uchicago.edu/ucblr/vol1/iss1/19>

documentation required under the AI Act, the technologies themselves generate IEPs suggestive of underlying biases. Moreover, these tools appear to offer standardized outputs that may strip students of their rights to receive authentically “individualized” education plans to meet the needs of their disabilities.

These insights aim to support both regulators and regulated entities as those stakeholders prepare for the AI Act’s effective date. Moreover, we hope it may be a tool leveraged by administrators and other educational professionals as they consider and adopt new AI technologies in classrooms across the state of Colorado and the nation. This research contributes to critical discussions about the prudent use of technologies in and outside of the classroom and the harms of algorithmic discrimination more broadly. Preserving the ability for students to secure educational opportunities without discriminatory harms is a critical goal of educators and regulators alike.

## **ACKNOWLEDGMENTS**

We leveraged existing research and insights of additional researchers and technical specialists within Georgetown University to inform our research and analysis. Our heartfelt appreciation extends to everyone who helped make this report possible, including Stevie DeGross and Zoë Kirchoff at the Colorado AG Office, who provided feedback and direction for our research project; and Meg Leta Jones, Paul Ohm, and Jon Brescia of Georgetown University, who provided legal and technical expertise on the Colorado AI Act.

# APPENDICES

## Appendix I - Frequently Asked Questions (FAQs)

See sample FAQs [here](#).

## Appendix II - Test Prompts

See sample test prompts that could be used to test for algorithmic discrimination in the EdTech tools discussed in Categories I-III [here](#).

## Appendix III - Code used for analysis of IEP outputs

**[disability framing]** [https://github.com/wisdombinna/IEP-Analysis/tree/main/disability\\_framing](https://github.com/wisdombinna/IEP-Analysis/tree/main/disability_framing)

**[gender analysis]** [https://github.com/wisdombinna/IEP-Analysis/tree/main/gender\\_analysis](https://github.com/wisdombinna/IEP-Analysis/tree/main/gender_analysis)

**[standard metrics]** <https://github.com/wisdombinna/IEP-Analysis/tree/main/metrics>

**[sample list of IEP prompts, responses]** [IEP Work - MagicAI](#)