Problems with the DOE's AI Guidance

The DOE deadline is May 8, 2026 for feedback via their survey at
on.nyc.gov/AiFeedbackNYCPS
Feel free to borrow any of the below points or provide your own.

The Guidance itself is posted here; our one-page critique is here, with suggested survey responses. Also posted online is an annotated pdf of the AI Guidance with our comments.

1. Lack of public input

The Privacy Working Group was appointed in February 2025 by Chancellor Melissa Ramos to provide input on the Chancellor's Regulations A-820 on student privacy which had not been updated since 2009. Subsequently, Chancellor Ramos appointed an AI Working Group (AIWG) in June 2025 that grew out of the Privacy Working Group. Its members were told that they were supposed to help the DOE develop their guidance on AI, with a special focus on student privacy, and were promised several times that they would have input on any DOE guidance before it was released. That never happened. Instead, the AIWG met only a handful of times and their work was severely hampered because DOE officials denied even their most basic requests for information.

One of the first requests made by AIWG members was to get a copy of the AI training that DOE claimed to have given to thousands of NYC teachers and other school staff. The members also asked for a list of AI products currently used in schools, along with their privacy policies. Both requests were denied. Instead, DOE said they could not provide the names of AI products currently used in schools because they had non-disclosure agreements with their vendors.

Because of parent opposition, the Panel for Educational Policy has refused to approve contracts for AI products at least four times in recent months. Yet DOE has allowed several of these products to be used in schools anyway. More than 2,000 NYC parents, teachers and concerned community members have signed a petition calling for a two year moratorium on the use of AI in the classroom.[1] Five Community Education Councils have

[1] https://actionnetwork.org/petitions/please-enact-a-two-year-moratorium-on-the-use-of-ai-in-nyc-public-schools?clear_id=true

passed resolutions, also calling for a moratorium.[2] This growing opposition has been largely dismissed by DOE officials, who have characterized their concerns as unjustified.

Despite DOE claims in the AI Guidance that they would "Activate stakeholder engagements" by holding "events and webinars," thus far officials have not scheduled any such sessions open to the public. The survey created for stakeholder input on the AI guidance is extremely limited and asks respondents to check off their three top concerns from a pre-selected list. [3]

To our knowledge, the only group that was able to provide actual input on these guidelines before their release was an AI Advisory Council, whose members were appointed by Chancellor David Banks before he left office and were composed primarily of DOE officials and representatives from the ed tech industry. The AI Advisory process and the development of the guidance was also subsidized by Google, according to an email from DOE to a CCHS member.

2. Lack of transparency concerning the use of AI products

The AI guidance provides no clarity or transparency about which AI products can be used with students, and which have gone through the DOE privacy vetting process known as ERMA, or Enterprise Request Management Application. Perhaps as a result, teachers continue to assign and students to use off-the-shelf AI products like Chat GPT that data-mine personal student information to improve their products –a commercial use specifically prohibited by the New York student privacy law, Ed Law 2D. Its developer Open AI admits that Chat GPT has "the potential to be used …to identify private individuals when augmented with outside data" and that it should never be used by any child under 13, and teenagers between 13 and 18 only with parent consent.[4]

Yet according to the DOE guidance, the pivotal question of age restrictions in the use of these tools is instead left up to teachers, with the following advice: *"Know and apply tool-specific age restrictions. Teach students and families where to find age requirement information*." But where they are supposed to find these age requirements is not explained.

While the guidance claims that "NYCPS requires that all AI tools be explainable to every stakeholder, including what the tool does, why it produces a given output, and how human judgment can intervene," no such explanation has been made available, though many parents have asked their principals for this sort of information. In fact, even AI executives admit that they do not fully understand how Generative AI produces its outputs. Moreover, there is zero transparency at the school level about which products used in schools actually employ AI, or which outputs are being generated through AI processes.

The DOE's refusal to fully disclose the names of AI products currently used in schools also violates Local Law 35 passed by the City Council in 2022, which requires every city agency to annually report on the names of algorithmic products it employs. In the most recent such report released in March 2026, the DOE included only seven algorithmic tools, while omitted most of the AI products we have heard from teachers, students and parents have been used in their schools. (For an incomplete list of these products, see the Appendix.)

---

[2] https://classsizematters.org/resolution-for-an-ai-moratorium-for-nyc-schools/

[3] Feedback form here: on.nyc.gov/AiFeedbackNYCPS

[4] https://cdn.openai.com/papers/gpt-4-system-card.pdf

Many AI companies, including both Open AI and Google, have acknowledged they can insert watermarking and other sorts of indicators to identify AI-generated texts and images, and every school and district should be requiring that this sort of indicator be incorporated in any output of teachers or students that was created with AI.[5]

3. The AI guidance fails to rigorously protect student privacy

The DOE vetting process ERMA is ineffective and primarily composed of a series of checklists which vendors are asked to mark in order to be approved. This process has not worked to sufficiently protect student privacy, as shown by breaches of Illuminate and PowerSchool that exposed the personal data of more than a million current and former NYC students, as well as the ongoing illegal use of student data for commercial purposes, as indicated by recent court settlements and consent decrees regarding Naviance and the College Board. Yet even though AI represents an even higher risk to student privacy and safety than most ed tech programs, the DOE rejected the recommendations from the AI working group and others to strengthen its processes, which included requiring independent privacy impact assessments, data security audits, and tests for algorithmic biases.

New York's student privacy law, Ed Law 2D, prohibits the use of personal student data for commercial purposes, and the regulations clearly state that using the data to improve the vendor's product is a commercial purpose and thus illegal. Yet most AI products do exactly this: data-mine to improve the product, and some of the vendors of the products used on NYC schools admit this occurs, and yet the DOE has allowed their use anyway. For example, DOE has approved the use of at least ten HMH products that utilize AI, although the HMH privacy agreement posted on the DOE page links to their public Privacy Policy which openly admits that data will be used for "*product improvement*" in violation of Ed Law 2D.[6]

HMH "partners" with Amira, another product used in many NYC schools, despite the fact that the Panel for Educational Policy Amira collects student voices, considered especially sensitive biometric data. About biometric data, NYSED says: ***"schools must consider the privacy implications thereof; the impact on civil rights, if any; the effectiveness of the biometric tool; and parental input."***[7]

MagicSchool AI programs are used in many NYC schools and offers a product that teachers can use to create IEPs. Their publicly facing Privacy Agreement says the product may use AI

[5] https://ai.google.dev/responsible/docs/safeguards/synthid and https://blog.google/innovation-and-ai/products/ai-image-verification-gemini-app/ and https://www.wsj.com/tech/ai/openai-tool-chatgpt-cheating-writing-135b755a

[6] https://www.schools.nyc.gov/about-us/policies/data-privacy-and-security-policies/supplemental-information-for-parents-about-doe-agreements-with-outside-entities/vendors-a-h links to https://www.hmhco.com/policy/prek-12-products-privacy-policy

[7] https://www.nysed.gov/news/2023/state-education-department-issues-determination-biometric-identifying-technology-schools

sub-contractors, including OpenAI, Anthropic, Perplexity.ai and Google Gemini, among others.[8]

Another member of the AI Working Group proposed that all vendor software operate on DOE servers, including any AI tools. If NYC DOE used either a private cloud or on-premises facilities to host the software used by schools, personal student data could not be data mined or escape into vendors' hands. As a large purchaser of services, the DOE could set the terms of service to protect the privacy of students, families and staff. This proposal was rejected as well.

4. The AI "traffic light" instructions are inadequate, sometimes inaccurate, often confusing, and even contradictory as to which AI uses are authorized and which are not. In too many cases, the guidance merely leaves the most difficult judgements up to teachers, without giving them direction or support.

The guidance separates the possible uses of AI into three categories, based on traffic lights: red light for "What AI will never be allowed to do in our schools," yellow light for "Proceed with caution- where professional judgement is essential" and green light for "Proceed with Confidence - Approved, Encouraged, and Supported."

Most of the most problematic issues are cited in the "yellow light" section. For example, it says that the thorny question of how students should use AI will depend on teacher judgement, and that "*students may use AI for research, exploration and creative projects"* without defining those terms*, adding, "Educator guidance, critical evaluation of outputs, and age-appropriate context are required."*

Moreover, while it is widely acknowledged that more than half of students nationwide use AI to do their homework, such as solving math problems and writing essays,[9] there is no real discussion in the guidance as to whether this is acceptable. Instead, the guidance punts the issue of potential plagiarism to some future time, writing that in a section entitled "Homework and academic integrity**"** that "NYCPS is developing guidance on assessment design and academic honesty in an AI-enabled environment, grounded in existing academic policies and research on effective practice."

The "red light" section also is full of confusing and contradictory guidance. While it says that AI cannot be used for high stakes decision-making in the cases of "eligibility, promotion, graduation, and program access," it also says that "Every student has full access to advanced coursework. Any algorithmic pathway can be overridden by educators, leaders, or students.."

Being able to override algorithmic pathways and access to advanced coursework assumes that these decisions can indeed be made by AI, contradicting the previous statement. Moreover, there are no assurances in the document as to how students will be alerted to which decisions have been made by AI and which by human beings, so it is difficult to

---

[8] https://cdn.prod.website-files.com/645187265d5e5e386be40629/68eec5681dc42158304b004c_Magic%20School%20Data%20Privacy%20Addendum%20v1.4%2010142025.pdf

[9] https://www.nytimes.com/2026/02/24/technology/schoolwork-chatbot-cheating-pew.html

understand how they would know which could be overridden. Finally, the omission of parents as also eligible to challenge these decisions is troubling.

Another statement in the red light section is misleading, if not outright erroneous: "Behavioral monitoring and student surveillance are prohibited." Yet more than 150 NYC public schools were recently reported to be using SmartPass, a digital program that employs AI to track the actions of students when they go to the bathroom, and to analyze whether they may be meeting up with the other students at the same time.[10] In addition, many NYC schools have used "Go Guardian" since 2017, a program that surveils students through their computers, or as the agreement on the DOE website describes it, "provides teachers with a way to view student online activity during their class sessions."[11]

In the greenlight section, many questionable assumptions are made. For example, it says that educators can use AI to "explore lesson ideas, approaches and unit planning, aligned with intellectual property guidance" without defining the latter term.

The document links to something called the IAUSP [for Internet Acceptable Use Policy] on the DOE website, which says "Intellectual property rights related to AI systems, including algorithms, models, and data, must be clearly defined and any proprietary AI technologies must be disclosed and licensed appropriately." [12] Thus, this section appears to be warning teachers not to plagiarize the AI tools they may use, rather than consider whether their own intellectual property rights will be respected if they create lesson plans with the use of AI --- a topic of concern to many educators that goes unexplored.

Also greenlighted by DOE is the use of AI for "scheduling, formatting, and summarizing nonsensitive information" without defining nonsensitive information. Is a student's name, along with their special education status "nonsensitive" if used to schedule them into Inclusion or self-contained classes for example? Most parents would probably disagree.

5. There is no attempt in the DOE guidance to address many of the most serious concerns about AI use.

Up front, the DOE admits that it has not evaluated AI for "algorithmic bias, equity impact, or instructional effectiveness." Many researchers have pointed out how AI can undermine students' cognitive development, their acquisition of fundamental skills, weaken their critical thinking and creativity, and harm their mental health and the global environment.[13]

---

[10] Detect & Automatically Block Students From Meeting Up" https://gothamist.com/news/new-digital-hall-passes-track-bathroom-breaks-gather-data-in-nyc-schools

[11] https://nycpublicschoolparents.blogspot.com/2021/12/doe-using-goguardian-and-other.html https://www.schools.nyc.gov/about-us/policies/data-privacy-and-security-policies/supplemental-information-for-parents-about-doe-agreements-with-outside-entities/vendors-a-h ; https://www.nyclu.org/commentary/software-could-be-spying-nyc-students

[12] IAUSP Data Privacy and Security Compliance Process (Intellectual Property) https://infohub.nyced.org/in-our-schools/policies/data-privacy-and-security-compliance-process

[13] A good list of citations is available on the Fairplay letter, urging a five-year moratorium on the use of AI in the classroom, signed onto by more than 200 organizations and experts, https://fairplayforkids.org/wp-content/uploads/2026/04/Coalition-of-Organizations-and-Experts-Calls-for-Pause-on-Generative-AI-in-PreK-12-schools-1.pdf Another comprehensive bibliography is available at the recent Brookings report, cautioning against AI in education at

Nor does the guidance propose how to address the fact that AI products regularly feature high levels of errors and/or sycophancy, both of which can lead students into dangerous alleyways of error and self-harm.

While the DOE admits that the use of AI may represent "risks to society" including "reinforcing inequity or reducing human judgment in civic institutions, and exacerbating the climate crisis, use of limited natural resources, strain on the electric grid, and mental health challenges," they do not even begin to evaluate whether these risks outweigh the benefits, with the implication that mentioning them is somehow sufficient.

Among other problems, AI products have been shown to contain high rates of factual errors, known as "hallucinations," as large as 77%, as DOE acknowledges. And yet the guidance fails to provide suggestions on how misleading or errors can be avoided or minimized.

While 66% of AI products have been shown to potentially contribute to bias, hate speech and discrimination,[14] the DOE merely states that "Identifying and addressing bias in AI tools is an active area of development and focus," and that at some point in the future, DOE will "strengthen AI tool evaluation practices, addressing algorithmic bias." Yet on another page, they appear to leave this responsibility up to teachers, as "Educators and school leaders must critically evaluate all AI-generated output for accuracy, appropriateness, and potential bias…" without explaining how that can be achieved.

Given that the research is substantial and growing that AI use leads to "cognitive off-loading" and impairs the development academic skills, DOE merely counters with the statement that "NYCPS is developing research-informed guidance on instructional design that ensures AI supports, rather than substitutes for, student thinking."

None of this is reassuring; and all of this together reinforces the conviction of parents, educators, advocates and many experts that the Mayor and the Chancellor should hold off using AI in our schools until these evaluations are complete, and effective policies are established to prevent serious harm to our children's education and our planet.

https://www.brookings.edu/wp-content/uploads/2026/01/A-New-Direction-for-Students-in-an-AI-World-FULL-REPORT.pdf

[14] https://korabench.substack.com/p/we-tested-32-ai-models-for-child

6. Appendix: An Incomplete list of AI products that are being used in NYC schools

Many of these were added to TeachHub in the last few weeks. Also some teachers have told us that Google Gemini has recently been uploaded into student Chromebooks without parent knowledge or consent. if you know of AI-enabled products used in NYC schools to add, please let us know at info@studentprivacy.org; your name will be kept confidential.

21stCenturyEd
Amira learning (grades K-3)
Algebra Teaching Assistant (Microsoft) -- LLM
Book Creator
Brisk Teaching
Codio
Coteach
CS4ALL Curriculum
Discovery Education
Edsoma
EPS Reading Assistant
Glint
Google Arts
Google Elementary School Drive
Google Gemini
Google Notebook LM
HMH Reading Counts!® (Grades K–12)
HMH Math Inventory® (Grades K–12)
HMH Reading Inventory® (Grades K–12)
HMH Phonics Inventory® (Grades 3–12)
HMH iRead® (Grades K–2)
HMH Waggle® (Grades 2–8)
HMH Writable® (Grades 3–12)
HMH System 44® (Grades 3–12)
HMH READ 180® (Grades 4–12)
HMY MATH 180® (Grades 5–12)
HMH English 3D (includes four different products for different grade levels)
Illustrative Math
Kiddom
Lexia
MagicSchool AI
National Geographic (powered by Spark AI)
Open Gen AI and Teaching Assistant Tool
Snorkl
ST Math
ThinkCerca
Yourway