

Time for a Pause: Without Effective Public Oversight, AI in Schools Will Do More Harm Than Good.



Ben Williamson
University of Edinburgh

Alex Molnar and Faith Boninger
University of Colorado Boulder

March 2024

National Education Policy Center

School of Education
University of Colorado Boulder
nepc.colorado.edu

Acknowledgements

NEPC Staff

Faith Boninger
Publications Manager

Patricia Hinchey
Academic Editor

Elaine Duggan
Production Design

Alex Molnar
NEPC Director

Kevin Welner
NEPC Director

Suggested Citation: Williamson, B., Molnar, A., & Boninger, F. (2024). *Time for a pause: Without effective public oversight, AI in schools will do more harm than good*. Boulder, CO: National Education Policy Center. Retrieved [date] from <http://nepc.colorado.edu/publication/ai>

Funding: This policy brief was made possible in part by funding from the Great Lakes Center for Educational Research and Practice.




Peer Review: *Time for a Pause: Without Effective Public Oversight, AI in Schools Will Do More Harm Than Good* was double-blind peer-reviewed.

This publication is provided free of cost to NEPC's readers, who may make non-commercial use of it as long as NEPC and its author(s) are credited as the source. For inquiries about commercial use, please contact NEPC at nepc@colorado.edu.



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

The National Education Policy Center (NEPC), a university research center housed at the University of Colorado Boulder School of Education, sponsors research, produces policy briefs, and publishes expert third-party reviews of think tank reports. NEPC publications are written in accessible language and are intended for a broad audience that includes academic experts, policymakers, the media, and the general public. Our mission is to provide high-quality information in support of democratic deliberation about education policy. We are guided by the belief that the democratic governance of public education is strengthened when policies are based on sound evidence and support a multiracial society that is inclusive, kind, and just. Visit us at: <http://nepc.colorado.edu>



Time for a Pause: Without Effective Public Oversight, AI in Schools Will Do More Harm Than Good.

Ben Williamson
University of Edinburgh

Alex Molnar and Faith Boninger
University of Colorado Boulder

March 2024

I. Executive Summary

Ignoring their own well-publicized calls to regulate AI development and to pause implementation of its applications, major technology companies such as Google, Microsoft, and Meta are racing to fend off regulation and integrate artificial intelligence (AI) into their platforms. Concerns underlying calls for such a pause run the gamut; here, we focus on democracy and privacy. The weight of the available evidence suggests that the current wholesale adoption of unregulated AI applications in schools poses a grave danger to democratic civil society and to individual freedom and liberty.

Public education is a public and private good essential to democratic civic life. The public must, therefore, be able to provide meaningful direction over schools through transparent democratic governance structures. Yet important discussions about AI's potentially negative impacts on education are being overwhelmed by relentless rhetoric promoting its alleged ability to positively transform teaching and learning. The result is that AI, with little public oversight, is on the verge of becoming a routine presence in schools.

These risks and harms follow years of warnings and precedents. The widespread use of pre-AI digital technologies in teaching and administration has already helped obscure educational decision-making and allowed student data to be exploited for non-school purposes. Without effective public oversight, the introduction of opaque and unproven AI systems and applications will likely intensify these problems and create many more.

The integration of AI into our computers, phones, cars, and homes is well on its way. Inevitably, this integration will impact schools, too—but we have choices about the nature and

extent of these changes. The core elements of learning and the types of interactions that take place in schools should be thoughtfully considered, with the resulting policies and practices following from deliberation decisions.

Advocates for AI claim that it will transform teaching and learning for the better. This will not happen, however, if integrating AI into the pedagogy of schools degrades the relationship between teachers and students. It will not happen if the AI imposes a rigid mechanistic model of instruction, corrupts curriculum with misinformation, and biases consequential decisions about student performance. It will not happen if integrating AI into schools' administrative processes locks schools and districts into an expensive "stack" of corporate tech systems for many of their everyday operations, with the result that funds available for other uses—including the teachers who can develop deep connections with nation's students—are increasingly shifted to corporate vendors. It will not happen if the AI exacerbates violations of student privacy, increases surveillance, and further reduces the transparency and accountability of educational decision-making.

All of these harms are likely if lawmakers and others do not step in with carefully considered regulations.

As existing school-focused products are updated to include newly created AI-based products, the immediate danger facing schools is not a future apocalypse as predicted by technology industry leaders. Rather, the danger is that—in the absence of responsible development, proper evaluation, or regulatory oversight—untested, opaque AI models and applications will become enmeshed in routine school processes.

The evidence set forth in this policy brief alarms us. The adoption of largely unregulated AI systems and applications would, we conclude, force students and teachers to become involuntary test subjects in a giant experiment in automated instruction and administration that is sure to be rife with unintended consequences and potentially negative effects. Allowing AI to become inextricably enmeshed in school processes and procedures invites both short- and long-term harms, because once it is enmeshed in those processes the only way to disentangle from it would be to completely dismantle those systems.

To forestall the far-reaching dangers posed by hasty implementation of AI in pedagogical and administrative systems, we recommend that school leaders pause adoption of AI applications until policymakers have had adequate time to fully educate themselves about AI and to formulate legislation and policy ensuring effective public oversight and control of its school applications. Any development of AI for schools should be conducted under "responsible AI" frameworks, in partnership with schools. In addition, we recommend that:

Federal and state policymakers:

- Stop promoting AI as a way to transform and modernize schools' pedagogical and administrative practices.
- Prohibit schools from adopting AI-based educational applications until appropriate regulatory structures are established.
- Adopt regulations that prohibit schools from using any technology, including "black

box” AI models, whose workings are not transparent to state regulators, unless they have provided those regulators a well-developed rationale and justification for why the particular technology is the only way to achieve a clearly defined and valid school purpose, and how it offers an improvement over existing education practices.


- Reduce the pressure on schools to adopt AI in their administrative systems by reducing the data reporting requirements placed on schools and teachers.

State policymakers:

- Establish an independent government entity charged with ensuring the quality of digital educational products used in schools. Charge this entity with reviewing and approving the pedagogy and programming of any digital educational product a school proposes to use, both prior to implementation and periodically thereafter. Require that the programming of any digital educational product—explicitly including products that incorporate AI—used in schools be transparent and amenable to review.
- Create classroom contexts that allow teachers to spend more time with their students, such as enacting legislation to limit class size, so that teachers are not pressured to find ways to keep students quietly occupied with digital products.

District policymakers:

- Refrain from adopting AI-based educational applications until:
 - Strict transparency and accountability requirements are put in place as part of an overall technology accountability plan.
 - The public has been provided with compelling evidence or thoughtful and clear explanations as to how those applications are an improvement over other education practices that do not require digital technology.



Time for a Pause: Without Effective Public Oversight, AI in Schools Will Do More Harm Than Good.

Ben Williamson
University of Edinburgh

Alex Molnar and Faith Boninger
University of Colorado Boulder

March 2024

II. Introduction

Public education is a public and private good essential to democratic civic life. The public must, therefore, be able to provide meaningful direction over schools through transparent democratic governance structures. And yet artificial intelligence (AI)¹, with little public oversight, is on the verge of becoming a routine presence in schools. The weight of the available evidence suggests that the current wholesale adoption of unregulated AI applications in schools poses a grave danger to democratic civil society and to individual freedom and liberty.

Background

In the spring of 2023, digital technology titans warned that artificial intelligence posed “profound risks to society and humanity,”² and called for regulation of its development and a pause in implementation of related applications.³ But concurrently, major players such as Google, Microsoft, Meta, and Amazon raced both to integrate AI into their platforms⁴ and to fend off regulation.⁵ Tech industry marketing went into high gear⁶ and soon the popular press overflowed with industry hype and speculation about AI’s potential and pitfalls.⁷ The effect of all the attention-grabbing predictions, self-interested corporate behavior, flamboyant marketing claims, and uncritical reporting has been to obscure the immediate dangers posed by AI’s rapid implementation.⁸ Given this backdrop, it is not surprising that rhetoric promoting AI’s alleged ability to positively transform teaching and learning has dominated discussions about its impacts on education.⁹

School administrators and teachers already use an array of digital educational technologies in teaching and management.¹⁰ Their use has increasingly obscured educational decision-making, made a mockery of student privacy rights, and allowed student data to be exploited for non-school purposes.¹¹ In the absence of effective public oversight, the introduction of AI systems and applications in education will likely intensify these problems and create many more.^{12,13}

As existing school-focused platforms and applications are updated to include AI, the immediate danger facing educators is not a future apocalypse. Instead, the danger is that AI models and applications will become enmeshed in school processes and procedures in ways that allow private entities to increasingly control the structure and content of public education, to reinforce surveillance practices, and to amplify existing biases and inequalities.¹⁴ For decades, academic researchers have worked on AI models for use in schools.¹⁵ Today, however, it is commercial enterprises that are aggressively pushing AI (and its attendant risks) into classrooms.¹⁶

The campaign to promote AI in education follows the logic of a half century of commercial, political, and ideological efforts to privatize and commercialize education.¹⁷ Given this logic it is not surprising that, despite the known dangers, corporations, private researchers, and governments are aggressively promoting the use of AI¹⁸ before a statutory and regulatory framework has been put in place to ensure that AI programs are transparent and subject to effective public scrutiny and control.¹⁹ This puts schools under tremendous pressure to accept AI as an inevitable upgrade to existing processes.²⁰

Computer scientists and software developers focus primarily on technical engineering questions²¹ and corporate leaders and investors prioritize profit²² over the common good. Nevertheless, educators are being asked to trust that these people, who have no educational expertise and who stand to financially benefit when AI is used in schools, are best suited to imagine and lead educational transformation.²³

III. Review of the Literature

The term “artificial intelligence” (AI) was first introduced by computer scientists in the 1950s, though many of its underlying mathematical processes and mechanistic procedures can be traced back to early models of computing and manufacturing in the 1800s.²⁴ Standard definitions of AI usually refer to computers performing tasks that only humans could normally do.²⁵ Today, however, AI has become a slippery term without a widely agreed-upon meaning.²⁶ This enables marketers to apply the label “AI” to almost any digital process or product they are selling.²⁷

Development of Artificial Intelligence—From Rule-Based Systems to Generative AI

AI has been developed over many decades by people and corporations with distinctive agendas and strategies.²⁸ Their efforts are poised to shape AI implementation in schools, especially since AI development and marketing efforts are rapidly expanding.²⁹

Three principal phases of AI's evolution began in the 20th century, when scientists and technicians focused on building systems that followed a sequence of pre-set rules, or algorithms, derived from the knowledge of experts.³⁰ Eventually, over the past two decades, a *machine learning* approach emerged. This approach involves using sophisticated mathematical processes (known as *learning algorithms*) to analyze massive amounts of data (known as *big data*) to identify trends or commonalities within it.³¹ On the assumption that the vast amount of data analyzed confers validity on the trends identified, such programs generate predictions about behavior in complex events and phenomena—including human behavior.³² The success of such analyses, however, has been mixed.³³ For example, one study asked six teams of data scientists to predict children's life outcomes with huge quantities of big data and cutting-edge machine learning tools. None of them came close to a reasonable level of accuracy, casting serious doubt on the use of AI in social policy areas like education.³⁴

Since late 2022, a further evolution of machine learning, *generative AI*,³⁵ has become an object of intense public, media, and political interest.³⁶ Generative AI deploys even more complex learning algorithms to create original text, images, and audio from data collected from the web or other sources.³⁷ This line of research has produced a series of program models often now called *foundation models* because they can be adapted for highly diverse purposes.³⁸ For example, generative AI *large language models* predict which word is most likely to follow the words preceding it by breaking down words into a sequence of numbers and then calculating the most probable response.³⁹

Large language models have provoked widespread excitement because their promoters have promised productivity gains for businesses as AI programs relieve workers of everyday writing tasks.⁴⁰ Similar support for teachers and students is also a promised benefit.⁴¹ It is, however, questionable that these promises will be realized in practice.⁴² For instance, while models often generate correct responses, computer programs cannot understand what words mean.⁴³ Therefore, the accuracy of their responses or the value of their responses in any given circumstance cannot be assumed.⁴⁴ Indeed, at the moment, AI is often “stupid.”⁴⁵ It invents facts, mangles the results of analyses, destabilizes information sources and produces dangerously wrong assertions about matters of social, public, and cultural importance.⁴⁶ When Google demonstrated its Bard language application as an educational tool by getting it to answer questions about the James Webb telescope in 2023, for example, the model gave a false factual response, prompting a market plunge in Google's company value.⁴⁷

Concerns have also emerged about AI development processes. For technical and commercial reasons, AI models are not transparent—they are not publicly explained in any detail.⁴⁸ For example, an analysis of 10 leading foundation models found limited information about where data for analysis came from, how much computing was necessary to create the models, and what specifics were embedded in analytical algorithms.⁴⁹ Many machine learning models are *black box* models, meaning that their mechanisms are said to be too complicated to explain or not explainable at all.⁵⁰ Others are hidden from public view by proprietary rights accorded to corporations.⁵¹

Black box foundational AI models are key to large technology corporations' plans to expand their proprietorial models into all sectors, to grow global market share, and to generate maximum profit.⁵² In the absence of effective public oversight and regulation, running gen-

erative AI programs is currently only possible by using “Big Tech” companies’ databases, high-powered computing capabilities, and financial resources. This makes it likely that the proprietary AI models of a few corporations will become the foundation for the vast majority of AI applications developed.⁵³ Meanwhile, regulators and lawmakers will be left struggling to respond by creating a patchwork of after-the-fact regulatory protections.⁵⁴

AI in Education

Since the 1960s, scientists and technology companies have explored ways to apply AI in education. *AI in Education* (AIED) is a major field of research and development.⁵⁵ The AI applications being promoted to schools today were preceded in the 1960s and 1970s by “Intelligent Tutoring Systems” and “Computer Assisted Instruction” systems.⁵⁶

Since the early 2000s, researchers have gathered, stored, and analyzed massive quantities of educational data with the intention of informing institutional and instructional strategies.⁵⁷ These approaches are now routinely considered synonymous with AIED, and have also been rapidly commercialized by the ed tech industry.⁵⁸ Most AIED applications employ big data and machine learning to produce various predictions and automated actions—such as predicting that a student may fail an assessment or creating a “personalized” intervention intended to produce a desired learning outcome.⁵⁹

Research on AI in education has developed and tested various approaches and reported modest effectiveness on measurable learning achievement—performance on quizzes and tests, for example.⁶⁰ Current excitement about its potential is motivating both public and private sources to generously fund researchers trying to find ways to improve learning outcomes using AI.⁶¹

However, the assumption that AI in education can be understood primarily as a technical matter best addressed by scientists and companies is increasingly challenged by researchers who argue that a narrowly technical perspective may lead to both bad policy and bad pedagogy.⁶² They point out that AI exists in social, economic, and political contexts that shape its development and uses.⁶³ How AI is adopted by different educational stakeholders (including AIED researchers, ed tech entrepreneurs, corporate leaders, and policymakers) will have significant implications for its use in schools.⁶⁴

The fact that entrepreneurs and corporations funded by venture capital and private equity are rushing to promote AI in education will inevitably narrow possible applications to those preferred by stakeholders with financial interests.⁶⁵ Small-scale ed tech start-ups and Big Tech corporations alike see AI as an opportunity,⁶⁶ leveraging popular hype to market such education products as personalized learning programs, automated lesson plan generators, and AI tutoring chatbots, called “tutorbots,” to schools.⁶⁷ Compelling evidence for the effectiveness of tutorbots in education remains scarce,⁶⁸ though this does not prevent entrepreneurs and researchers from proclaiming their usefulness.⁶⁹

Policymakers routinely invoke AI rhetorically, calling on schools to embark on “digital transformation,”⁷⁰ often with little attention to social, economic, legal, or ethical implications.⁷¹ These calls dovetail with existing political priorities on performance monitoring, account-

ability, efficiency, and effectiveness—all of which require extensive collection of data about students.⁷² Although systems of test-based accountability have existed in schools since the 1990s,⁷³ they will expand and intensify as AI is used to continuously monitor and assess student learning.⁷⁴ As a result, commercial AI systems will increasingly serve as private actors in public education as schools, districts, and governments relinquish key tasks, functions, and responsibilities to third-party technology vendors.⁷⁵

Existing and potential uses of AI in education are not merely innovative technical add-ons to teaching and learning practices or engineering solutions to schools' existing pedagogic and administrative problems. Rather, AI in education has been spurred by multiple forces: longstanding efforts by scientists to measure, predict, and support learning processes and outcomes; commercial aspirations to profit from selling products to schools; and the political objective of being perceived as having improved school efficiency and accountability while cutting costs. As things currently stand, these ambitions have begun to coalesce into a vision of AI-driven schooling in which commercial products assess student learning, automate teaching, and make decisions about student progress.

Inadequate Research Base

Despite the extensive research in the field of AI in Education (AIED) and the burgeoning research on machine learning, there is remarkably little evidence to support claims of AI's ability to “transform” schools.⁷⁶ While AIED researchers have produced many research findings, their studies tend to focus primarily on measures of individual student engagement and performance (assessed by standardized achievements tests), or on “engineering” problems such as designing increasingly sophisticated algorithms and enhancing machine learning effectiveness.⁷⁷

Overall, AIED studies tend to find ambiguous results, lack independence and scale, and fail to address more fundamental questions about educational goals.⁷⁸ AIED research therefore often promotes a view of education transformation as improving measurable individual outcomes despite very limited evidence that AI “works.”⁷⁹ In effect, such studies reduce well-researched and nuanced theories of how humans learn to whatever can be made into a mathematical model (however complex), and they ignore the contested terrain of exactly which goals and curriculum public schools should embrace.⁸⁰ Moreover, claims that AI can solve major educational problems—such as lack of qualified teachers, student underachievement, and educational inequalities—rely to a considerable extent on conjecture rather than evidence.⁸¹

Even more problematic are the serious methodological flaws in machine learning research that call into question the validity of hundreds of studies.⁸² The nature of the flaws, in general, leads toward “over optimism” with respect to the usefulness and value of machine learning applications in a variety of fields.⁸³ These findings are particularly concerning because they call into question not only commercial marketing claims, but also the scientific evidence base supporting the widespread implementation of AI systems in all sectors,⁸⁴ including education.

Finally, because of the very high computing costs associated with running machine learning models, most researchers have to rely on systems from the dominant AI companies themselves in order to conduct research⁸⁵—the same corporations that often fund AI studies.⁸⁶ This makes research dependent on corporate resources, funds, and business practices, giving AI firms considerable influence over not only AI development, but also the academic research that depends on their systems.⁸⁷ It also compromises an important part of the research process, which is reproducing findings to verify their validity. When a company changes or stops supporting a particular model, researchers cannot reproduce studies conducted earlier.⁸⁸ This renders the research base unstable and unverifiable—and thus unusable as a basis for assessing subsequent models.

IV. Recent Developments

In November 2022, OpenAI released ChatGPT, accelerating the race to develop and market generative AI platforms and applications. Companies that have developed artificial intelligence foundation models, including OpenAI, Google, Meta, and Amazon, intend to expand rapidly and “scale up” in every sector they enter—including education.⁸⁹ While media attention has focused on students using ChatGPT, these companies are swiftly expanding AI in education in several ways:

1. Selling access to digital systems to schools. For example, Amazon sells schools and districts access to cloud computing facilities, enabling them to use its AI systems to analyze institutional and student data.⁹⁰
2. Adding AI features to products that schools already use. Google, for instance, has begun introducing AI into its Workspace suite for schools, which includes the Classroom platform used by schools worldwide.⁹¹
3. Integrating AI applications into new or upgraded products. OpenAI, for example, partners with ed tech companies to integrate its language models into services it promotes as “AI teaching assistants,”⁹² and is also exploring its own educational applications of ChatGPT.⁹³
4. Building AI into services such as search engines and other everyday applications commonly used in classrooms. For example, Microsoft introduced new AI applications in its Office software, based on OpenAI technologies, which are promoted for educational use.⁹⁴

Current promotion of AI in education focuses largely on the pedagogical uses of AI applications that provide automated language and image producing capabilities.⁹⁵ For students, reporting has, for example, focused on their use of automated tools, such as ChatGPT, to write assignments,⁹⁶ and of personalized learning “tutorbots” to mimic a one-on-one tutoring experience.⁹⁷ For teachers, it has focused on their use of AI “assistants” to create lesson plans, develop grading systems, or review student progress, among other tasks.⁹⁸ Proponents assert that such AI applications offer students personal assistance in learning and offer teachers time-saving support.⁹⁹ Entrepreneurial educators have produced guidebooks and training materials for teachers,¹⁰⁰ and OpenAI has launched a *Teaching with AI* guide to train teachers to use its applications.¹⁰¹

Accepting proposed benefits as real, international and governmental organizations have supported the use of AI products in schools. For example, the Organisation for Economic Co-operation and Development (OECD), worried that AI will soon outperform humans on many cognitive tasks, has called for the urgent modification of formal education systems so students can learn skills to complement AI rather than skills for tasks that could soon be automated.¹⁰² The US Department of Education’s Office of Educational Technology encourages teachers to involve themselves in developing and evaluating AI applications for education.¹⁰³ This all results in teachers and schools being pushed to accommodate untested and opaque commercial AI applications or risk being “left behind.”¹⁰⁴

Recent regulatory proposals show little sign of slowing the rapid advance of AI into schools.¹⁰⁵ An executive order released by President Biden in October 2023 outlined directives to, among other things, increase federal oversight of foundation model testing, protect data privacy, and “promote innovation and competition.”¹⁰⁶ It mandates the Department of Education to create an “AI toolkit” incorporating “appropriate human review of AI decisions, designing AI systems to enhance trust and safety and align with privacy-related laws and regulations in the educational context, and developing education-specific guardrails.”¹⁰⁷

However, the fact sheet accompanying the executive order described the Department of Education’s mandate differently as: “shap[ing] AI’s potential to transform education by creating resources to support educators deploying AI-enabled educational tools, such as personalized tutoring in schools.”¹⁰⁸ Although the Department of Education will base its toolkit and guidance on its goal to “keep humans in the loop,” two assumptions guide its approach: that “learning” is defined as and limited to those things that digital programs can measure and that AI programs can and should “optimize” learning as so defined.^{109,110}

V. Discussion and Analysis

As noted earlier, despite the risks of rushing untested tech into classrooms and the lack of implementation criteria or regulatory controls,¹¹¹ schools face unrelenting pressure to “modernize” by adopting artificial intelligence.¹¹² While AI applications are marketed as ways to address teaching and learning problems and to streamline school administrative processes, they carry with them all the limitations, problems, and risks inherent in the AI models used to run them.¹¹³ Rushing AI into schools increases the likelihood that these technologies will reproduce or intensify many problems, making any potential benefit less significant than the potential harms.

Two heavily promoted applications illustrate the perils of uncritical adoption of AI in schools: *tutoring chatbots*, based on large language models, which promise to personalize and automate teaching¹¹⁴; and *adaptive learning platforms*, which use large quantities of student data to make predictions, customize classroom resources and activities, and automatically intervene in pedagogical processes.¹¹⁵ The dangers that such products pose include costing teachers more time than they save, artificially restricting the definition of “learning,” sidelining teacher expertise and relationships with students, introducing curricular misinformation, and increasing bias and discrimination in classrooms and schools. In light of such dangers, before expanding or allowing AI applications into schools, policymakers and

education leaders should consider what it would mean to use AI responsibly, whether its potential benefits outweigh the costs, and whether its adoption in education is truly inevitable.

Dangers in Teaching and Learning

Restricted Teaching and Learning

Khan Academy has developed several offerings for technology-based schooling,¹¹⁶ expanding them beginning in March 2023 with heavy marketing of its Khanmigo “tutorbot” to parents, teachers, and school districts.¹¹⁷ Khanmigo’s marketing claims that the tutorbot can engage in personalized and “conversational” pedagogic interaction with students, thus enhancing “personalized learning.”¹¹⁸

Khanmigo incorporates OpenAI’s GPT-4 language model to generate customized educational content, lesson plans, and assessment, and to perform such other tasks as monitoring student progress.¹¹⁹ Its marketing materials promise one-on-one “virtual tutoring” for students and a “personal teaching assistant” for teachers. These materials also claim that Khanmigo can understand and respond individually to students’ academic and career goals and save teachers time by providing a content library for lesson planning, reports of student progress (along with “recommendations for what to do with that information”), and a “personal concierge service” that can help teachers “just do more.”¹²⁰

Such marketing suggests Khanmigo can make decisions that mimic a real teacher’s thinking and responses, and so satisfy the needs both of students who require a teacher’s attention and teachers who cannot offer that attention because of other demands on their time. It also reveals assumptions about the nature of teaching and learning embedded in the programming that are unavoidably imposed on real teachers, students, and communities that work with the product.

Algorithms in personalized learning applications implement a narrow understanding of learning and a highly constrained pedagogic model.¹²¹ “Learning” in these applications typically consists of students engaging in computer-based activities and producing “correct” or acceptable answers within a limited range of predetermined responses.¹²² The applications present teachers with reports of student “progress” on the measurable tasks and suggest strategies to improve students’ “performance.”¹²³ Any school adopting these applications tacitly accepts that learning can largely be reduced to a narrow range of behavioral responses and tacitly rejects more complex child-centered understandings of learning.¹²⁴

The theory of learning built into AI chatbot applications like Khanmigo prioritizes data-based numerical profiles of students—considered “objective” measures of performance—while undermining teachers’ ability to make professional judgments about their students and devaluing teachers’ subjective experience, subject matter expertise, classroom interactions, and contextual knowledge of a given child, class, or social setting.¹²⁵

A supposed advantage in such AI-based programs is that they reduce administrative burdens on teachers and allow them more time to teach. Khanmigo’s advertising promises, for example, “your evenings are yours again without compromising quality.”¹²⁶ This is persuasive

marketing for overburdened teachers, who might welcome outsourcing their overwhelming administrative duties to a computer program, or might be relieved to occupy students with technology that theoretically supports their learning while they wait for teacher time in an overcrowded classroom.

Delegating tasks to a tutorbot, however, inserts a digital intermediary between students and teachers who are pedagogical experts in their fields and who know their students and understand their context. Moreover, despite generative AI's humanlike communication, it is limited to responding to prompts and queries that fall within a product's established parameters. As a result, teachers must devote time to understanding and mastering such limitations and then teach students how to ask questions and verify responses.¹²⁷ Even as AI automates some administrative tasks, then, it is likely to introduce other time-consuming pressures and burdens on teachers.¹²⁸

Rather than reducing demands, AI of the kind currently promoted by private commercial enterprises can add further complexity to teachers' workloads by ultimately positioning them as servants of the technology tasked with ensuring its smooth classroom operation. Reconfiguring the core systems of education—instruction, curriculum, and assessment—to accommodate AI will demand laborious efforts by educators to adjust their professional practices. At the same time, it will deny them a voice in determining whether proposed changes have real value and should or should not be implemented. Academic AIED research and development, informed by partnership with teachers and students, will struggle to be heard above the current promotional and speculative rhetoric around commercial AI applications for schools.¹²⁹ AI therefore presents a critical challenge to the institutional autonomy of schools' decision-making processes, and particularly to teachers' autonomy to make professional decisions about their pedagogic practices.

Curricular Misinformation

A key challenge of AI for teachers is that generative AI technologies are trained to produce text that seems convincing even though it might contain false information.¹³⁰ Applications that use generative AI to support teachers' lesson planning and resource generation, for example, could flood the classroom with misleading inaccuracies or false information.¹³¹

This is not a problem likely to be resolved, because the data such programs depend on may degrade further as automated content spreads across the web.¹³² The danger is that the information environment will be overrun by AI-generated text, making it impossible to ascertain the authority or authenticity of any online source,¹³³ and therefore rendering online sources useless or misleading for educational purposes.¹³⁴

As has been documented in both language- and image-generating AI,¹³⁵ instabilities and errors in underlying data make AI chaotic and unreliable.¹³⁶ No doubt this is why Khanmigo recommends that users not rely on its responses but rather verify its information using such resources as “textbooks, articles, or other trusted sources.”¹³⁷ And why it also restricts the time students can spend interacting with the tutorbot, noting that “extended interactions are more likely to lead to poor AI behavior.”¹³⁸ These recognized limitations beg the question: Why spend time using this application or any AI application when it seems obvious

their use will result in more work and very likely undermine effective teaching?¹³⁹

Outstanding questions remain, then, about ed tech and Big Tech market leaders becoming powerful gateways to online learning content. Do their AI applications improve the quality of information taught in the classroom? How much should the content produced by AI, having been trained on material scraped from the web by companies like OpenAI, be trusted by teachers and students? What are the implications of students encountering “poor AI behavior” from a tutorbot? It is also not clear whose responsibility it must be to check the quality of the content produced by AI. Khanmigo implicitly makes already overworked teachers responsible for doing the hidden labor of checking for errors and providing feedback that the product’s programmers can then fix one-by-one.¹⁴⁰ This scenario makes it very likely that such errors will be overlooked and that the accuracy of information that students encounter via AI will degrade over time.¹⁴¹

AI poses a real threat to the accuracy of school knowledge and thus to the validity of curriculum materials. In a context where students and many teachers already source content online (for example, through sites such as Teachers Pay Teachers¹⁴²) it may become increasingly difficult to tell an authoritative and accurate source from a plausible but fallacious one produced by AI. This is likely to become an even bigger problem with OpenAI’s release of “GPTs”—versions of ChatGPT that can be customized by users—as individuals are able to create their own educational AI applications and share them freely online, with few quality checks and controls.¹⁴³

Potentially Amplified Bias and Discrimination

Because AI models are trained on either internet data or historical data, they incorporate biases that can transfer to their educational applications.¹⁴⁴ For example, when ChatGPT was released in November 2022, educators expressed anxiety about students cheating on written assignments, prompting ed tech companies to develop automated AI detectors. The makers of Turnitin, a product already used internationally to detect student plagiarism, added AI detection functionality in early 2023, arguing it would be able to detect distinctive markers of AI-generated text.¹⁴⁵ However, independent studies have found that such AI detectors are prone to error, leading to a surge in false accusations of cheating.¹⁴⁶ Such accusations are disproportionately biased against non-native English speakers, who tend to write in simpler sentences that AI flags as suspicious.¹⁴⁷

Similarly, the automatic AI essay grading programs used in many states are prone to bias against certain demographic and ethnic groups, falsely awarding high grades to work that features sophisticated language and structure, regardless of the meaning or quality of the writing.¹⁴⁸ This is because AI programs do not “know,” in any real sense, what makes for a good or bad essay, but can only search for patterns that correlate with higher or lower human-assigned grades. Likewise, GoGuardian, which uses algorithms and AI to monitor students’ social media and web browsing activity for “suspicious” content or behavior, routinely flags multiple categories of non-explicit material as harmful or dangerous—including material on general educational sites.¹⁴⁹

AI plagiarism detectors, automated essay graders, and student monitoring platforms are

built on the false assumption that automated analyses of data provide objective accounts of students' practices and behaviors. They replace human discretion with automated suspicion scores generated by machine learning algorithms, with subsequent judgments on assignments, grades, and behaviors that are difficult, if not impossible, for students to challenge. Thus, if AI is integrated into the existing structures of schooling, it is likely to worsen rather than reduce many of the inequalities that characterize contemporary education. Opaque machine learning systems will make hidden decisions affecting students' courses of learning, outcomes, and prospects based on unknowable mathematical calculations.

Dangers in Administration

Increased Costs

Learning management systems already used in many schools, such as Google Classroom, Blackboard, and Canvas, are beginning to integrate AI into their platforms.¹⁵⁰ Google Classroom, with its suite of nominally “free” software and low-cost Chromebook hardware, dominates the market.¹⁵¹ It has already announced the launch of AI-based adaptive learning add-ons to Classroom, with associated additional costs for schools, as well as plans to upgrade Classroom further with generative language AI.¹⁵² “Practice Sets” is Google’s AI-based adaptive learning system for education, and “Duet AI” is its “collaboration partner” for teachers.¹⁵³ In addition to any pedagogical implications associated with using Google Classroom, its integration of further AI and automation into many aspects of school functioning also carries potentially significant administrative implications.¹⁵⁴

The most significant of these is to obscure the rationale for administrative decisions about critical institutional issues when decision-making is ceded to opaque machine learning systems controlled by tech firms. Google Classroom, for example, integrates with hundreds of other ed tech products and can synchronize with a school’s student information systems.¹⁵⁵ It offers Google cloud services such as single sign-on, identity management, and device management, as well as plagiarism detection, automated grading, teaching templates, student grouping, and administrative analytics to facilitate “data-driven decisions.”¹⁵⁶ Such management systems facilitate the transfer of control of schools from the public to private corporations by acting as central conduits through which all of a school’s digital activities must pass—making it hard for educators or administrators to see how any decisions based on the data have been made.¹⁵⁷

Because running AI is costly, the use of AI programs in schools will necessarily require schools to pay for operating costs for an increasing number of pedagogic and administrative AI applications. The promise that AI can save schools money by reducing staffing costs is likely illusory, as schools will probably be required to pay costly fees for accessing AI facilities. In other words, rather than saving money, administrative applications are more likely to shift existing funds to monopolistic technology providers.

Khanmigo and Google Classroom already illustrate how this works. Khan Academy, when it provides Khanmigo to districts, currently charges those districts \$60 per student for annual use, citing high computing costs associated with OpenAI’s GPT-4 as the justification for the

charges.¹⁵⁸ Likewise, districts must also pay for Google Classroom’s AI upgrades. To access its latest adaptive learning application, Practice Sets, they must switch from the free basic offering to a for-fee premium package.¹⁵⁹ In other words, tech firms are extracting value from school budgets to defray the high computing costs associated with AI (and grow company value).¹⁶⁰

Increased Threats to Student Privacy

AI applications collect and aggregate data in order to function. In so doing, they normalize digital surveillance and privacy invasions in school.¹⁶¹ In practice, education technology companies use applications like Google Classroom to routinely collect as much data as possible, well beyond that required to perform their assigned tasks.¹⁶²

Although proponents of using AI in education tend to emphasize the efficiency of data-driven administrative systems, privacy-related threats to equity are inherent in it.¹⁶³ This is because AI models are built using massive data sets that can be used to profile, compare, and assess individuals who are then subject to potentially discriminatory decisions based on “statistical dossiers” of their personal lives.¹⁶⁴ Thus, a significant danger of digital technology in general, and of the privacy-invasive model of AI in particular, is that they can reproduce and amplify existing forms of inequality in education by using datasets containing examples of historic bias and discrimination.¹⁶⁵ For example, if a big data set indicates that certain marginalized groups have underperformed historically, then a software application may be biased against individuals from such groups in the future, singling out and targeting them as “at-risk” and closing down or limiting their opportunities to access information and resources.¹⁶⁶

Moreover, school data systems are vulnerable to breaches, hacks, ransomware, and denial-of-service attacks.¹⁶⁷ A data breach at the student-tracking ed tech company Illuminate, for example, compromised the educational data of at least a million public school students and prompted New York City’s Department of Education to ask schools to stop using Illuminate’s products.¹⁶⁸ School data systems feature highly detailed and intimate student information, including personal and demographic data, grades, attendance, behavioral information, and other confidential information. Increasing AI capacity in ed tech products may exacerbate these vulnerabilities, as student data are collected at even greater scale by a wide range of companies—including AI companies—that offer only vague data privacy protections.¹⁶⁹

Reduced Transparency and Accountability

Finally, enabling AI to play a role in school administration will reduce the transparency and accountability of decision-making.¹⁷⁰ Many digital products already used in schools are neither transparent nor accountable because current law and regulation allows companies to shield the inner working of their products behind proprietary protections.¹⁷¹

AI is even more opaque than other digital programs.¹⁷² Black box machine learning and AI models are so complicated that their outputs are often impossible to explain or interpret.¹⁷³ Although in many cases simpler and more accessible statistical models can produce equally

accurate results, companies benefit from selling access to proprietary models that require customers to trust the systems and simply accept being unable to verify results.¹⁷⁴ If the system makes a mistake, it might never be identified or redressed and the public suffers the consequences. For example, the facial identification systems used for remote testing often fail to accurately identify individuals or mistakenly flag student behaviors as suspicious, but they are very hard for students to challenge.¹⁷⁵

In high-stakes decision-making in a sector like education, allowing such impenetrable models to assume responsibility for key administrative procedures necessarily means the creation of schools in which school leaders and teachers will be unable to exercise judgment, provide a rationale, or take responsibility for classroom and institutional decisions.¹⁷⁶

Considerations for the Future

Is AI Development Responsible?

The rapid creation of AI applications for schools raises the urgency of prioritizing ethics, student rights, and social responsibility in their development.¹⁷⁷ *Responsible AI* development would ensure that products are safe and trustworthy, designed to benefit people, communities, and society, and mitigate harms.¹⁷⁸ As yet, there is little indication that such values are adequately addressed in education applications.¹⁷⁹ Unfortunately, academic AIED researchers have tended to ignore them or delegate addressing them to the educational tech industry and policy centers.¹⁸⁰ This complacency—along with the money and power held by commercial actors—enables commercial rather than educational imperatives to guide the development of AI and furthers political interests promoting relentless testing and school surveillance.¹⁸¹

Responsible governance would require the companies developing AI to commit to transparent and responsible product design, *and also* to monitoring, understanding, and mitigating the continuous impacts of AI in various contexts. Of particular concern is the automation of decisions with “irreversible and severe consequences.”¹⁸² For example, technologies to identify emotions are currently being developed to assess if a person is lying and cheating.¹⁸³ These technologies are inherently inaccurate, however, and an inaccurate judgment that a student has cheated or that a witness is lying could have dire consequences for their lives. Responsible AI governance might lead to delaying or indefinitely pausing development of such technologies.

Although several responsible AI initiatives have produced principles, frameworks or checklists for safe and trustworthy AI development and accountability,¹⁸⁴ these agendas can be manipulated through various forms of industry lobbying and efforts to water down their scope or possibilities of enforcement.¹⁸⁵ Expanding responsibility for product safety to include the wide range of people or organizations that build and use AI—rather than leaving it to technicians and business alone—would mitigate such dangers.¹⁸⁶

Among the many obstacles to the implementation of responsible policies governing AI is their cost. The goal of profit-seeking business is to shift to the public as many costs as pos-

sible while garnering the highest possible private rate of return on investments. Public oversight of AI necessarily entails either public ownership or a comprehensive regulatory regime adequately financed to achieve its mission. The question is, where will the money come from?

Moreover, the required regulation flies in the face of 50 years of policy devoted to deregulation and privatization. It would demand a fundamental rethinking of the government's relationship to commercial interests. Such rethinking would, without a doubt, be attacked by self-interested parties as not only too costly but also as stifling innovation and promoting inefficiency. While these arguments may be relevant in individual circumstances, they are neither generally nor self-evidently true.

From the perspective of education, responsible governance of AI therefore entails significantly more commitment than the simple principles of responsible development issued by industry. It also requires costly and ongoing monitoring of the effects of AI in classroom contexts. It may also require delays and indefinite pauses in development where warranted—such as, for example, in cases where commercial AI providers seek to introduce products into schools with insufficient evidence that they produce beneficial outcomes, or when those products automate professional judgement with potentially negative consequences, or when they inadequately address questions of AI ethics directly relevant to education.

Is AI Inevitable?

AI products are moving into schools at dizzying speed. As we have noted, this is in part the result of the pressure on schools to “modernize” by adopting the latest products that the technology industry offers. There is already a consensus of sorts that the move to AI is inevitable. The director of educational technology at Newark Public Schools made the case to the *New York Times* when he explained why his district adopted Khanmigo: “It’s important to introduce our students to it, because it’s not going away.”¹⁸⁷

The de facto requirement that students serve as a technology company’s experimental subjects might be explained by the initially low entry cost for school districts. Struggling districts, especially, might be willing to gamble that a technological innovation might turn things around for their students. However, before placing that bet it would be valuable to first ask some fundamental questions. Computer scientist Joseph Weizenbaum posed such concerns 50 years ago, essentially arguing that no technology—including AI—should be implemented unless we know that it is both necessary and good.¹⁸⁸

Is AI Necessary?

The technology industry has tried for a century to profit from schools. In the 1920s, creators of mechanical “teaching machines” claimed they could help students work at their own individual pace and ease administrative demands on teachers.¹⁸⁹ The same claims are now attached to digital personalized learning platforms designed to “meet students where they are,” whenever and wherever they are ready to learn.¹⁹⁰ And, as we have noted, technology companies are eagerly incorporating AI into personalized learning and other platforms to

conduct virtual “tutoring,” generate lesson plans and tests, and collect and organize data for administrative decision-making.

AI programs may perhaps, in some well-defined circumstances, efficiently perform useful tasks. However, the question of whether these tasks can be done as well as or better without using AI programs and accepting their negative consequences is unclear—and rarely discussed.¹⁹¹ We cannot now know which, if any, AI programs offered to schools can legitimately be said to be necessary.

Is AI Good?

AI programs are unreliable,¹⁹² tend to hallucinate¹⁹³ and break,¹⁹⁴ and will almost certainly produce manifold unintended consequences.¹⁹⁵ Instability is inherent in AI programs, in part because they quickly and unpredictably change as they “learn” from how they are used.¹⁹⁶ Very few studies in machine learning research address values-based issues such as societal need or negative potential, instead prioritizing commercial and engineering values like system performance, efficiency, and novelty.¹⁹⁷ There is a danger that AI programs will propagate these built-in commercial and engineering values without the public ever having been considered or having agreed to adopt them.

AI algorithms are neither objective nor neutral.¹⁹⁸ To further complicate matters, even if it’s clear which assumptions govern AI training data and how a program weights the data it processes,¹⁹⁹ when programs actually run they produce results based on a vast number of mathematical probabilities that are unknowable to both their creators and their users.²⁰⁰ As a result, even a person with knowledge of a black box AI program’s initial coding cannot explain how it produced its results. Not surprisingly, machine learning inevitably produces outcomes that may be incomprehensible, untrue, or incorrect in a variety of unknowable ways.²⁰¹ If such systems are used, neither teachers nor administrators will be able to understand, explain, or justify the conclusions the programs reach, much less audit or document their validity.²⁰²

We are now on the verge of creating a world in which students set on a particular academic path or denied advancement by an AI-based learning application will have absolutely no way to challenge the judgments made about them. Nor will their teachers and administrators have a meaningful way to respond to students and parents who challenge AI decisions. As AI captures the pedagogy and administrative processes of schools, as it becomes embedded in the full range of school routines, it therefore poses a direct and powerful threat to the democratic governance of schools.²⁰³ In this sense, implementing AI in schools cannot be considered “good.”

We cannot conclude with any confidence that AI is, in general, either necessary or good for students, teachers, or schools. On the contrary, the weight of the available evidence suggests that the wholesale adoption of artificial intelligence in schools poses a grave danger to democratic civil society and to individual freedom and liberty. The adoption of opaque, untested, and largely unregulated AI systems and applications would force students and teachers to become involuntary test subjects in a giant experiment in automated instruction and ad-

ministration that is sure to be rife with unintended consequences and potentially negative effects.²⁰⁴ Allowing AI to become inextricably enmeshed in school processes and procedure invites disaster, because once it is enmeshed in those processes the only way to disentangle from it would be to completely dismantle those systems.

VI. Recommendations

To forestall the far-reaching dangers posed by hasty implementation of AI in pedagogical and administrative systems, we recommend that school leaders pause adoption of AI applications until policymakers have had adequate time to fully educate themselves about AI and to formulate legislation and policy ensuring effective public oversight and control of its school applications. Any development of AI for schools should be conducted under “responsible AI” frameworks, in partnership with schools. In addition, we recommend that:

Federal and state policymakers:

- Stop promoting AI as a way to transform and modernize schools’ pedagogical and administrative practices.
- Prohibit schools from adopting AI-based educational applications until appropriate regulatory structures are established.
- Adopt regulations that prohibit schools from using any technology, including “black box” AI models, whose workings are not transparent to state regulators, unless they have provided those regulators a well-developed rationale and justification for why the particular technology is the only way to achieve a clearly defined and valid school purpose, and how it offers an improvement over existing education practices.
- Reduce the pressure on schools to adopt AI in their administrative systems by reducing the data reporting requirements placed on schools and teachers.

State policymakers:

- Establish an independent government entity charged with ensuring the quality of digital educational products used in schools. Charge this entity with reviewing and approving the pedagogy and programming of any digital educational product a school proposes to use, both prior to implementation and periodically thereafter. Require that the programming of any digital educational product—explicitly including products that incorporate AI—used in schools be transparent and amenable to review.
- Create classroom contexts that allow teachers to spend more time with their students, such as enacting legislation to limit class size, so that teachers are not pressured to find ways to keep students quietly occupied with digital products.

District policymakers:

- Refrain from adopting AI-based educational applications until:
 - Strict transparency and accountability requirements are put in place as part of an

overall technology accountability plan.

- The public has been provided with compelling evidence or thoughtful and clear explanations as to how those applications are an improvement over other education practices that do not require digital technology.

Notes and References

- 1 Throughout this brief, we recognize that various forms of AI exist, each with different histories, technical foundations, and functions. Our focus is on two particular forms of AI that are of contemporary significance in education: data systems that use machine learning algorithms to measure and make predictions about behaviours or outcomes, and large language models (one version of so-called “generative AI”) that can automatically generate original text.

Ali, S.M., Dick, S., Dillon, S., Jones, M.L., Penn, J., & Staley, R. (2023). Histories of artificial intelligence: A genealogy of power. *British Journal of History of Science Themes*, 8, 1-18.
- 2 Center for AI Safety (2023, May 30). AI extinction statement press release [press release]. Retrieved August 16, 2023, from <https://www.safe.ai/press-release>

Center for AI Safety (2023, n.d.). Statement on AI risk. Retrieved August 16, 2023, from <https://www.safe.ai/statement-on-ai-risk#open-letter>

Future of Life Institute (2023, March 22). Pause giant AI experiments: An open letter. Retrieved August 16, 2023, from <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>
- 3 Brodtkin, J. (2023, May 17). AI technology “can go quite wrong,” OpenAI CEO tells Senate. *Ars Technica*. Retrieved August 17, 2023, from <https://arstechnica.com/tech-policy/2023/05/ai-technology-can-go-quite-wrong-openai-ceo-tells-senate/>

Edwards, B. (2023, May 31). OpenAI execs warn of “risk of extinction” from artificial intelligence in new open letter. *Ars Technica*. Retrieved August 16, 2023, from <https://arstechnica.com/information-technology/2023/05/openai-execs-warn-of-risk-of-extinction-from-artificial-intelligence-in-new-open-letter/>
- 4 Lee, D. (2023, September 26). Amazon makes a shrewd move in the AI arms race. *The Washington Post*. Retrieved November 15, 2023, from https://www.washingtonpost.com/business/2023/09/26/amazon-makes-a-shrewd-move-in-the-ai-arms-race-with-anthropic-deal/c15967f0-5c5e-11ee-b961-94e18b27be28_story.html

Wall Street Journal (2023, October 20). How Meta plans to incorporate AI into its products. Retrieved October 26, 2023, from <https://www.wsj.com/tech/ai/meta-plans-to-incorporate-ai-4ba569e3> [subscription required]
- 5 Kak, A., Myers West, S. and Whittaker, M. (2023, December 5). Make no mistake—AI is owned by Big Tech. *MIT Technology Review*. Retrieved December 5, 2023, from <https://www.technologyreview.com/2023/12/05/1084393/make-no-mistake-ai-is-owned-by-big-tech>

Perrigo, B. (2023, April 21). Big tech is already lobbying to water down Europe’s AI rules. *Time*. Retrieved November 8, 2023, from <https://time.com/6273694/ai-regulation-europe/>
- 6 Perhaps the most extreme and exaggerated example of industry spin was produced by the tech investor Marc Andreessen in a blog post titled “Why AI Will Save the World.”

Andreessen, M. (2023, June 6). Why AI will save the world. *Andreessen Horowitz*. Retrieved November 8, 2023, from <https://a16z.com/2023/06/06/ai-will-save-the-world/>
- 7 Kapoor, S. and Narayanan, A. (2022, September 30). Eighteen pitfalls to beware of in AI journalism [blog post]. *AI Snake Oil*. Retrieved November 15, 2023, from <https://www.aisnakeoil.com/p/eighteen-pitfalls-to-beware-of-in>
- 8 The ‘true threat of AI’, argues technology critic Evgeny Morozov, is an expansion of corporate power over the public domain.

Morozov, E. (2023, June 30). The true threat of AI. *New York Times*. Retrieved June 30, 2023, from <https://www.nytimes.com/2023/06/30/opinion/artificial-intelligence-danger.html>

Likewise, Timnit Gebru and her colleagues note that, “The harms from so-called AI are real and present and follow from the acts of people and corporations deploying automated systems. Regulatory efforts should focus on transparency, accountability and preventing exploitative labor practices.”

Gebru, T., Bender, E.M., McMillan-Major, A., & Mitchell, M. (2023, March 31). *Statement from the listed authors of stochastic parrots on the “AI pause” letter*. Palo Alto, CA: Distributed AI Research Institute (DAIR). Retrieved January 16, 2024, from <https://www.dair-institute.org/blog/letter-statement-March2023/>

- 9 In one example, a computer scientist—with no background in the field of education—predicted that AI-driven personalized tutoring systems could make traditional schools and teachers redundant.

Devlin, H. (2023, July 7). AI likely to spell end of traditional school classroom, leading expert says. *The Guardian*. Retrieved August 7, 2023, from <https://www.theguardian.com/technology/2023/jul/07/ai-likely-to-spell-end-of-traditional-school-classroom-leading-expert-says>

- 10 Nichols, T.P. & Dixon-Roman, E. (2024). Platform governance and education policy: Power and politics in emerging edtech ecologies. *Educational Evaluation and Policy Analysis*. Retrieved February 15, 2024, from <https://doi.org/10.3102/01623737231202469>
- 11 Boninger, F. & Molnar, A. (2020). *Issues to consider before adopting a digital platform or learning program*. Boulder, CO: National Education Policy Center. Retrieved October 17, 2023, from <https://nepc.colorado.edu/publication/virtual-learning>
- 12 Laird, E., Dwyer, M., & Grant-Chapman, H. (2023, September). *Off-task: EdTech threats to student privacy and equity in the age of AI*. Washington, DC: Center for Democracy and Technology. Retrieved October 17, 2023, from <https://cdt.org/wp-content/uploads/2023/09/091923-CDT-Off-Task-web.pdf>
- 13 In July 2023, the assistant director of UNESCO’s education unit warned:

The speed at which generative AI technologies are being integrated into education systems in the absence of checks, rules or regulations, is astonishing . . . [T]oday, in most national contexts, the time, steps and authorizations needed to validate a new textbook far surpass those required to move generative AI utilities into schools and classrooms. In fact, AI utilities often required no validation at all. They have been ‘dropped’ into the public sphere without discussion or review. I can think of few other technologies that are rolled out to children and young people around the world just weeks after their development. In many cases, governments and schools are embracing a radically unfamiliar technology that even leading technologists do not claim to understand . . . Education, given its function to protect as well as facilitate development and learning, has a special obligation to be finely attuned to the risks of AI—both the known risks and those only just coming into view. But too often we are ignoring the risks.

Giannini, S. (2023, July). *Generative AI and the future of education*. Paris, France: United Nations Educational, Scientific and Cultural Organization (UNESCO). Retrieved October 26, 2023, from <https://unesdoc.unesco.org/ark:/48223/pf0000385877>

- 14 Gebru, T., Bender, E.M., McMillan-Major, A., & Mitchell, M. (2023, March 31). *Statement from the listed authors of stochastic parrots on the “AI pause” letter*. Palo Alto, CA: Distributed AI Research Institute (DAIR). Retrieved January 16, 2024, from <https://www.dair-institute.org/blog/letter-statement-March2023/>
- Williamson, B., Komljenovic, J., & Gulson, K.N. (Eds.) (2024). *World Yearbook of Education 2024: Digitalisation of education in the era of algorithms, automation and artificial intelligence*. London: Routledge.
- 15 Porayska-Pomsta, K. (2023). A manifesto for a pro-actively responsible AI in education. *International Journal of Artificial Intelligence in Education*. Retrieved January 10, 2024, from <https://doi.org/10.1007/s40593-023-00346-1>

- 16 State of Big Tech (no date). *AI-in-education is intended to replace teachers*. IT for Change. Retrieved February 7, 2024, from <https://projects.itforchange.net/state-of-big-tech/ai-in-education-is-intended-to-replace-teachers/>
- 17 Molnar, A. (2005). *School commercialism: From democratic ideal to market commodity*. New York: Routledge.
- Williamson, B. & Hogan, A. (2020). *Commercialisation and privatisation in/of education in the context of Covid-19*. Brussels: Education International. Retrieved January 11, 2024, from https://issuu.com/educationinternational/docs/2020_eiresearch_gr_commercialisation_privatisation
- 18 Holmes, W. & Tuomi, I. (2022). State of the art and practice in AI in education. *European Journal of Education*, 57, 542–570.
- 19 In his Executive Order of October 30, 2023, President Biden charged the Department of Education as follows:
- To help ensure the responsible development and deployment of AI in the education sector, the Secretary of Education shall, within 365 days of the date of this order, develop resources, policies, and guidance regarding AI. These resources shall address safe, responsible, and nondiscriminatory uses of AI in education, including the impact AI systems have on vulnerable and underserved communities, and shall be developed in consultation with stakeholders as appropriate. They shall also include the development of an “AI toolkit” for education leaders implementing recommendations from the Department of Education’s AI and the Future of Teaching and Learning report, including appropriate human review of AI decisions, designing AI systems to enhance trust and safety and align with privacy-related laws and regulations in the educational context, and developing education-specific guardrails.
- Biden, J.R. (2023, October 30). *Executive order on the safe, secure, and trustworthy development and use of artificial intelligence* (Section 8(d)). The White House. Retrieved October 31, 2023, from <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>
- See also:
- McQuillan, D. (2023). Written evidence (LLM0015) to the House of Lords Communications and Digital Select Committee inquiry: Large language models. Parliamentary Committees. Retrieved November 15, 2023, from <https://committees.parliament.uk/writtenevidence/124038/pdf/>
- 20 U.S. Department of Education, Office of Educational Technology. (2023, May). *Artificial intelligence and future of teaching and learning: Insights and recommendations*. Retrieved October 26, 2023, from <https://www2.ed.gov/documents/ai-report/ai-report.pdf>
- 21 D’Agostino, S. (2023, November 16). In class, some colleges overlook technology’s dark side. *Inside Higher Ed*. Retrieved November 16, 2023, from <https://www.insidehighered.com/news/tech-innovation/teaching-learning/2023/11/16/some-colleges-overlook-technologys-dark-side>
- 22 Yelenevych, A. (2022, December 26). The future of EdTech. *Forbes*. Retrieved November 16, 2023, from <https://www.forbes.com/sites/forbesbusinesscouncil/2022/12/26/the-future-of-edtech/>
- 23 Garcia, A. & Nichols, T.P. (2021). Digital platforms aren’t mere tools —they’re complex environments. *Phi Delta Kappan*, 102(6), 14-19. Retrieved January 10, 2024, from <https://doi.org/10.1177/0031721721998148>
- 24 Whittaker, M. 2023. Origin stories: Plantations, computers, and industrial control. *Logic*, 19. Retrieved November 16, 2023, from <https://logimag.io/supa-dupa-skies/origin-stories-plantations-computers-and-industrial-control/>
- 25 Boden, M.A. (2016). *AI: Its nature and future*. Oxford: Oxford University Press.

- 26 Suchman, L. (2023). The uncontroversial ‘thingness’ of AI. *Big Data and Society*, 10(2). Retrieved November 14, 2023, from <https://doi.org/10.1177/20539517231206794>
- Tucker, E. (2022, March 8). *Artifice and intelligence*. Center on Privacy and Technology. Retrieved November 14, 2023, from <https://medium.com/center-on-privacy-technology/artifice-and-intelligence%C2%B9-fooda128d3cd>
- 27 Siegel, E. (2023, June 2). The AI hype cycle is distracting companies. *Harvard Business Review*. Retrieved November 8, 2023, from <https://hbr.org/2023/06/the-ai-hype-cycle-is-distracting-companies>
- 28 Elish, M.C. & Boyd, D. (2018) Situating methods in the magic of big data and AI. *Communication Monographs*, 85(1), 57-80.
- 29 Williamson, B. & Eynon, R. 2020. Historical threads, missing links, and future directions in AI in education. *Learning, Media and Technology*, 45(3), 223-235. Retrieved November 20, 2023, from <https://www.tandfonline.com/doi/full/10.1080/17439884.2020.1798995>
- 30 Council of Europe (no date). *History of artificial intelligence*. Retrieved November 20, 2023, from <https://www.coe.int/en/web/artificial-intelligence/history-of-ai>
- 31 Alpaydin, E. (2016). *Machine learning: The new AI*. Cambridge, MA: MIT Press.
- 32 Broussard, M. (2019). *Artificial unintelligence: How computers misunderstand the world*. London: MIT Press.
- 33 The AI, Algorithmic, and Automation Incidents and Controversies (AIAAIC) repository has collected over 1,200 examples of real-world problems caused by the introduction of AI, algorithms and automation across different industries and sectors.
- AIAAIC (n.d.). AIAAIC [webpage]. Retrieved December 5, 2023, from <https://www.aiaaic.org/home>
- 34 Hao, K. (2020, April 2). AI can't predict how a child's life will turn out even with a ton of data. *MIT Technology Review*. Retrieved January 31, 2024, from <https://www.technologyreview.com/2020/04/02/998478/ai-machine-learning-social-outcome-prediction-study/>
- 35 Visual Storytelling Team & Murgia, M. (2023, September 12). Generative AI exists because of the transformer. *Financial Times*. Retrieved November 8, 2023, from <https://ig.ft.com/generative-ai/>
- Clarke, S., Milmo, D., & Blight, G. (2023, November 1). How AI chatbots like ChatGPT or Bard work—Visual explainer. *The Guardian*. Retrieved November 8, 2023, from <https://www.theguardian.com/technology/ng-interactive/2023/nov/01/how-ai-chatbots-like-chatgpt-or-bard-work-visual-explainer>
- 36 Ball, P. (2023, July 19). Can machines think? *Prospect*. Retrieved November 20, 2023, from <https://www.prospectmagazine.co.uk/ideas/technology/ai/62131/can-machines-think>
- Lazar, S. (2024, February 13). Frontier AI ethics. *Aeon*. Retrieved February 15, 2024, from <https://aeon.co/essays/can-philosophy-help-us-get-a-grip-on-the-consequences-of-ai>
- 37 Baack, S. (2024). *Training data for the price of a sandwich: Common Crawl's impact on generative AI*. Mozilla Foundation. Retrieved February 7, 2024, from <https://foundation.mozilla.org/en/research/library/generative-ai-training-data/common-crawl/>
- 38 Jones, E. (2023). *Explainer: What is a foundation model?* Ada Lovelace Institute. Retrieved November 8, 2023, from <https://www.adalovelaceinstitute.org/resource/foundation-models-explainer/>
- 39 Clarke, S., Milmo, D., & Blight, G. (2023, November 1). How AI chatbots like ChatGPT or Bard work—Visual explainer. *The Guardian*. Retrieved November 8, 2023, from <https://www.theguardian.com/technology/ng-interactive/2023/nov/01/how-ai-chatbots-like-chatgpt-or-bard-work-visual-explainer>

- 40 Mollick, E. (2023, October 23). The best available human standard. *One Useful Thing* [blog post]. Retrieved January 16, 2024, from <https://www.oneusefulthing.org/p/the-best-available-human-standard>
- 41 Phillips, V. (2023, June 7). Intelligent classrooms: What AI means for the future of education. *Forbes*. Retrieved November 22, 2023, from <https://www.forbes.com/sites/vickiphillips/2023/06/07/intelligent-classrooms-what-ai-means-for-the-future-of-education/>
- Hill, J. (2024, February 6). Meet the founders of ‘revolutionary’ AI being rolled out in schools. *Schools Week*. Retrieved February 7, 2024, from <https://schoolsweek.co.uk/ai-marking-the-answer-to-poor-recruitment-or-pure-snake-oil/>
- 42 Acemoglu, D. (2024, January 10). Get ready for the great AI disappointment. *Wired*. Retrieved January 10, 2024, from <https://www.wired.com/story/get-ready-for-the-great-ai-disappointment/>
- Garcia, A., Logan, C., & Nichols, T.P. (2024, January 28). Inspiration from the Luddites: On Brian Merchant’s “Blood in the Machine”. *Los Angeles Review of Books*. Retrieved February 15, 2024, from <https://lareviewofbooks.org/article/inspiration-from-the-luddites-on-brian-merchants-blood-in-the-machine/>
- 43 Warner, J. (2024, January 22). ChatGPT can’t teach writing. *Inside Higher Ed*. Retrieved February 15, 2024, from <https://www.insidehighered.com/opinion/blogs/just-visiting/2024/01/22/arizona-state-announces-plan-give-education>
- 44 Hao, K. (2021, May 20). The race to understand the exhilarating, dangerous world of language AI. *MIT Technology Review*. Retrieved November 8, 2023, from <https://www.technologyreview.com/2021/05/20/1025135/ai-large-language-models-bigsience-project/>
- 45 Bridle, J. (2023, March 16). The stupidity of AI. *The Guardian*. Retrieved November 8, 2023, from <https://www.theguardian.com/technology/2023/mar/16/the-stupidity-of-ai-artificial-intelligence-dall-e-chatgpt>
- 46 Ananny, M. (2024, January 16). To reckon with generative AI, make it a public problem. *Issues in Science and Technology*. Retrieved January 31, 2024, from <https://issues.org/ai-public-problem-ananny/>
- 47 Pearl, M. (2023, February 8). Bard shared false information in Google’s announcement. *Mashable*. Retrieved January 31, 2024, from <https://mashable.com/article/google-bard-james-webb-telescope-false-fact>.
- 48 AI Now Institute, Ada Lovelace Institute and Open Government Partnership. (2021, August 17). *Algorithmic accountability for the public sector – Report*. AI Now Institute. Retrieved January 16, 2024, from <https://ainowinstitute.org/publication/algorithmic-accountability-for-the-public-sector-report>
- 49 Bommasani and colleagues noted, “While the societal impact of foundation models is growing, transparency is on the decline . . .” Their analysis assessed the 10 leading foundation models on 100 indicators of transparency, and found the average score to be 37 out of 100 (Meta’s Llama 2 scored highest, at 57). They found especially low transparency with respect to how the models are built (e.g., related to the training data) and on downstream impacts of use (e.g., how users might seek redress for harm).
- Bommasani, R., Klyman, K., Longpre, S., et al. (2023). *The foundation model transparency index*. Retrieved October 19, 2023, from https://crfm.stanford.edu/fmti/fmti.pdf?utm_source=substack&utm_medium=email
- 50 Bagchi, S. (2023, May 22). What is a black box? A computer scientist explains what it means when the inner workings of AIs are hidden. *The Conversation*. Retrieved February 7, 2024, from <https://theconversation.com/what-is-a-black-box-a-computer-scientist-explains-what-it-means-when-the-inner-workings-of-ais-are-hidden-203888>
- 51 Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1, 206-215. Retrieved November 29, 2023, from <https://www.nature.com/articles/s42256-019-0048-x> [subscription needed to access full text]

- 52 Knight, W. (2021, September 14). A Stanford proposal over AI's 'foundations' ignites debate. *WIRED*. Retrieved November 8, 2023, from <https://www.wired.com/story/stanford-proposal-ai-foundations-ignites-debate/>
- 53 Kak, A. & Myers West, S. (2023). *AI Now 2023 landscape: Confronting tech power*. New York: AI Now Institute. Retrieved November 8, 2023, from <https://ainowinstitute.org/2023-landscape>
- 54 Satariano, A., & Kang, C. (2023, December 6). How nations are losing a global race to tackle A.I.'s harms. *New York Times*. Retrieved December 6, 2023, from <https://www.nytimes.com/2023/12/06/technology/ai-regulation-policies.html>
- 55 Indicative lists of the range of research and applications of AI in education can be found in the two main research journals in the field. The aims and scope of the *International Journal of Artificial Intelligence in Education* (originally launched in 1988) feature 50 relevant topics, including adaptive systems, statistical methods, cognitive models, robotics, intelligent agents, tutoring systems, neural models, student modelling and cognitive diagnosis, and theories of learning.
- Springer Nature (2023). *International Journal of Artificial Intelligence in Education: Official Journal of the International AIED Society* [webpage]. Retrieved November 27, 2023, from <https://link.springer.com/journal/40593/aims-and-scope>
- A new journal, *Computers and Education: Artificial Intelligence*, launched in 2020, refers to using AI “to facilitate teaching, learning, or decision-making. With the help of AI technologies, which simulate human intelligence to make inferences, judgments, or predictions, computer systems can provide personalized guidance, supports, or feedback to students as well as assisting teachers or policymakers in making decisions.”
- Hwang, G-J., Xie, H., Wah, B.W., & Gašević, D. (2020). Vision, challenges, roles and research issues of artificial intelligence in education. *Computers and Education: Artificial Intelligence*, 1. Retrieved November 27, 2023, from <https://doi.org/10.1016/j.caeai.2020.100001>
- See also:
- Doroudi, S. (2022). The intertwined histories of artificial intelligence and education. *International Journal of Artificial Intelligence in Education*. Retrieved November 8, 2023, from <https://link.springer.com/article/10.1007/s40593-022-00313-2>
- 56 Alkhatlan, A. & Kalita, J. (2018). Intelligent tutoring systems: A comprehensive historical survey with recent developments. *arXiv*. Retrieved November 8, 2023, from <http://arxiv.org/abs/1812.09628>
- 57 Lockyer, L., Heathcote, E., & Dawson, S. (2013). Informing pedagogical action: Aligning learning analytics with learning design. *American Behavioral Scientist*, 57(10), 1439-1459.
- 58 Buckingham Shum, S., & Luckin, R. (2019). Learning analytics and AI: Politics, pedagogy and practices. *British Journal of Educational Technology*, 50(6), 2785-2793.
- Gašević, D. Siemens, G., & Sadiq, S. (2023). Empowering learners for the age of artificial intelligence. *Computers and Education: Artificial Intelligence*. Retrieved November 8, 2023, from <https://doi.org/10.1016/j.caeai.2023.100130>
- 59 Gillani, N., Eynon, R., Chiabaut, C., & Finkel, K. (2023). Unpacking the “black box” of AI in education. *Educational Technology and Society*, 26(1), 99-111.
- 60 A recent statistical meta-analysis of AIED research results suggested some effectiveness of AI on “learning achievement,” as measured by standardized tests. Although only 24 articles met the study’s criteria for inclusion, most of these were very short time-scale quasi-experimental studies, and the socio-economic or geopolitical contexts of the reviewed studies were not made clear.

Zheng, L., Niu, J., Zhong, L., & Gyasi, J.F. (2023) The effectiveness of artificial intelligence on learning achievement and learning perception: A meta-analysis. *Interactive Learning Environments*, 31(9), 5650-5664.

61 Nietzel, M.T. (2023, May 4). NSF announces \$140 million investment in seven artificial intelligence research institutes. *Forbes*. Retrieved November 8, 2023, from <https://www.forbes.com/sites/michaelt Nietzel/2023/05/04/nsf-announces-140-million-investment-in-seven-artificial-intelligence-research-institutes/>

62 Holmes, W. (2023). AIED—Coming of age? *International Journal of Artificial Intelligence in Education*. Retrieved November 8, 2023, from <https://doi.org/10.1007/s40593-023-00352-3>

63 Williamson, B., Eynon, R., Knox, J., & Davies, H. (2023). Critical perspectives on AI in education: Political economy, discrimination, commercialization, governance and ethics. In B. du Boulay, A. Mitrovic, & K. Yace (Eds.), *Handbook of artificial intelligence in education* (pp. 553-570). Cheltenham: Edward Elgar.

Gulson, K.N., Sellar, S., & Webb, P.T. (2022). *Algorithms of education: How datafication and artificial intelligence shapes policy*. Minneapolis: University of Minnesota Press.

64 Eynon, R. & Young, E. (2021). Methodology, legend, and rhetoric: The constructions of AI by academia, industry, and policy groups for lifelong learning. *Science, Technology, and Human Values*, 46(1), 166-191.

65 Komljenovic, J., Williamson, B, Eynon, R., & Davies, H. (2023). When public policy ‘fails’ and venture capital ‘saves’ education: Edtech investors as economic and political actors. *Globalisation, Societies and Education*. Retrieved November 8, 2023, from <https://doi.org/10.1080/14767724.2023.2272134>

66 HolonIQ (2023, February 27). Artificial intelligence in education. 2023 survey insights [blog post]. *HolonIQ*. Retrieved November 8, 2023, from <https://www.holoniq.com/notes/artificial-intelligence-in-education-2023-survey-insights>

Berger, R. (2023, March 10). AI applications extend throughout education sector. *Forbes*. Retrieved November 8, 2023, from <https://www.forbes.com/sites/rodberger/2023/03/10/ai-applications-extend-throughout-education-sector/>

67 A few examples can provide some sense of the range of applications and businesses promoting AI in education. The established education business Duolingo integrates with OpenAI’s large language model technology to run a conversational language learning application. Microsoft promotes the application of AI in education, and has also incorporated OpenAI language technology into its products. The learning management systems Blackboard and Canvas have both announced plans to incorporate generative AI technologies. Blackboard has launched an “AI Design Assistant” powered by OpenAI to support teachers’ lesson planning. Canvas is also actively exploring and promoting the application of AI in schools.

Anthology, Inc. (Blackboard) (2022). AI design assistant [webpage]. Retrieved November 27, 2023, from https://help.blackboard.com/Learn/Instructor/Ultra/Course_Content/Create_Content/AI_Design_Assistant

Duolingo (2023, March 14). Introducing Duolingo Max, a learning experience powered by GPT-4 [webpage]. Retrieved November 27, 2023, from <https://blog.duolingo.com/duolingo-max/>

Gasvoda, H. (2023, September 29). Navigating generative AI: Key insights for this school year [blog post]. *Instructure Blog*. Retrieved November 27, 2023, from <https://www.instructure.com/en-au/resources/blog/navigating-generative-ai-key-insights-school-year?filled>

Instructure (2023, July 27). Instructure reveals major product innovations across four strategic areas [press release]. Retrieved October 26, 2023, from <https://www.instructure.com/en-au/press-release/instructure-reveals-major-product-innovations-across-four-strategic-areas>

Microsoft (2023, March 28). Exploring new opportunities with AI in education [blog post]. *Microsoft Blog*.

Retrieved November 27, 2023, from <https://educationblog.microsoft.com/en-us/2023/03/exploring-new-opportunities-with-ai-in-education>

- 68 A recent meta analysis of research on AI chatbots in education, based on a sample of 24 studies, found some evidence of effectiveness on learning performance on selected skills and outcomes in some contexts, though mostly in short time-scale conditions with some degree of “novelty effect,” and primarily in higher education settings. The authors acknowledged that some studies found no significant difference in learning performance between the use of chatbots and control groups, the risk of information overload from chatbot use, a lack of cultural specificity in studies, and its conclusions remained highly speculative.

Wu, R. & Yu, Z. (2023). Do AI chatbots improve students learning outcomes? Evidence from a meta-analysis. *British Journal of Educational Technology*. Retrieved January 11, 2024, from <https://doi.org/10.1111/bjet.13334>

- 69 Gašević, D., Siemens, G., & Sadiq, S. (2023). Empowering learners for the age of artificial intelligence. *Computers and Education: Artificial Intelligence*, 4. Retrieved January 16, 2024, from <https://doi.org/10.1016/j.caeai.2023.100130>

Singer, N. (2024, January 11). Will chatbots teach your children? *New York Times*. Retrieved January 11, 2024, from <https://www.nytimes.com/2024/01/11/technology/ai-chatbots-khan-education-tutoring.html>

- 70 See, for example:

Broadband Commission for Sustainable Development (2020, September). *The digital transformation of education: Connecting schools, empowering learners*. International Telecommunication Union (ITU), United Nations Educational, Scientific and Cultural Organization (UNESCO), and United Nations Children’s Fund (UNICEF). Retrieved October 26, 2023, from <https://unesdoc.unesco.org/ark:/48223/pf0000374309/PDF/374309eng.pdf.multi>

U.S. Department of Education, Office of Educational Technology (2010). *Transforming American education: Learning powered by technology (National education technology plan 2010)*. Retrieved October 26, 2023, from <https://tech.ed.gov/files/2017/01/NETP17.pdf>

- 71 Schiff, D. (2022). Education for AI, not AI for education: The role of education and ethics in national AI policy strategies. *International Journal of Artificial Intelligence in Education*, 32, 527-563. Retrieved January 11, 2024, from <https://doi.org/10.1007/s40593-021-00270-2>

- 72 Wyatt-Smith, C., Lingard, B., & Heck, E. (Eds.) (2021). *Digital disruption in teaching and testing: Assessments, big data, and the transformation of schooling*. Abingdon: Routledge.

- 73 Anagnostopoulos, D., Rutledge, S., & Bali, V. (2013). State education agencies, information systems, and the expansion of state power in the era of test-based accountability. *Educational Policy*, 27(2), 217-247.

- 74 Rupp, A.A. & Lorie, W. (2023, April 19). *Ready or not: AI is changing assessment and accountability*. Center for Assessment. Retrieved January 11, 2024, from <https://www.nciea.org/blog/ready-or-not-ai-is-changing-assessment-and-accountability/>

- 75 Crawford, K. & Schultz, J. 2019. AI systems as state actors. *Columbia Law Review*, 119(7), 1941-1972.

- 76 Selwyn, N. (2019). *Should robots replace teachers? AI and the future of education*. Cambridge: Polity.

- 77 Selwyn, N. (2019). What’s the problem with learning analytics? *Journal of Learning Analytics*, 6(3), 11-19.

- 78 Holmes, W. & Tuomi, I. (2022). State of the art and practice in AI in education. *European Journal of Education*, 57, 542-570.

- 79 Selwyn, N. (2024). On the limits of artificial intelligence (AI) in education. *Nordisk tidsskrift for pedagogikk og kritikk*, 10, 3-14. Retrieved January 31, 2024, from <https://pedagogikkogkritikk.no/index.php/ntpk/>

assistant-for-every-student-teacher-parent-and-administrator/

92 Franzen, C. (2023, August 31). OpenAI wants teachers to use ChatGPT for education. *Venture Beat*. Retrieved November 8, 2023, from <https://venturebeat.com/ai/openai-wants-teachers-to-use-chatgpt-for-education/>

93 Tong, A. (2023, November 16). Exclusive: OpenAI wants ChatGPT in classrooms. *Reuters*. Retrieved December 5, 2023, from <https://www.reuters.com/technology/openai-explores-how-get-chatgpt-into-classrooms-2023-11-16/>

See also:

Davis, A. (2024, January 18). A new collaboration with OpenAI charts the future of AI in higher education. *ASU News*. Retrieved February 6, 2024, from <https://news.asu.edu/20240118-university-news-new-collaboration-openai-charts-future-ai-higher-education>

94 Brewin, K. (2023, March 27). 5 ways Microsoft's Copilot will save teachers time. *TES Magazine*. Retrieved November 15, 2023, from <https://www.tes.com/magazine/analysis/general/5-ways-microsoft-copilot-will-save-teachers-time-artificial-intelligence>

95 Skates, A.L. (2023, February 8). *Five predictions for the future of learning in the age of AI*. Andreessen Horowitz. Retrieved November 15, 2023, from <https://a16z.com/the-future-of-learning-education-knowledge-in-the-age-of-ai/>

96 Klein, A. (2023, February 1). ChatGPT cheating: What to do when it happens. *Education Week*. Retrieved February 7, 2024, from <https://www.edweek.org/technology/chatgpt-cheating-what-to-do-when-it-happens/2023/02>

97 Young, J.R. (2024, January 22). Years before ChatGPT, one of the creators of IBM's Watson tried to harness AI to tutor—here's why it didn't work. *Fast Company*. Retrieved February 7, 2024, from <https://www.fastcompany.com/91013601/years-before-chatgpt-one-of-the-creators-of-ibms-watson-tried-to-harness-ai-to-tutor-heres-why-it-didnt-work>

98 Singer, N. (2023, June 8). New A.I. Chatbot tutors could upend student learning. *New York Times*. Retrieved October 18, 2023, from <https://www.nytimes.com/2023/06/08/business/khan-ai-gpt-tutoring-bot.html>

99 In some contexts, AI tutoring is also marketed to parents as a cheaper alternative to in-person private tutoring services.

Chen, C. (2024, February 12). Parents turn to AI tablets after China's tutor crackdown. *Rest of World*. Retrieved February 15, 2024, from <https://restofworld.org/2024/china-student-tablet-ai/>

100 For example:

Thirdbox Ltd. (2023). Take back your time in 3 hours [webpage]. Retrieved November 27, 2023, from <https://survivalkit.theaieducator.io/>

Mollick, E. (2023, June 12). Assigning AI: Seven ways of using AI in class [blog post]. *One Useful Thing*. Retrieved February 21, 2024, from <https://www.oneusefulthing.org/p/assigning-ai-seven-ways-of-using>

101 OpenAI. (2023, August 31). Teaching with AI. *OpenAI* [blog post]. Retrieved October 20, 2023, from <https://openai.com/blog/teaching-with-ai>

102 The OECD's Artificial Intelligence and the Future of Skills project is developing a program to assess the capabilities of AI and robotics, particularly in relation to literacy and mathematics, and their impact on education and work.

Organisation for Economic Cooperation and Development (OECD) (2023). *Is education losing the race with technology? AI's progress in maths and reading*. Paris: OECD. Retrieved October 20, 2023, from <https://doi.org/10.1787/169800000000>

org/10.1787/73105f99-en

- 103 U.S. Department of Education, Office of Educational Technology. (2023, May). *Artificial intelligence and future of teaching and learning: Insights and recommendations*. Retrieved October 26, 2023, from <https://www2.ed.gov/documents/ai-report/ai-report.pdf>
- 104 Klein, A. (2023, November 14). Schools ‘can’t sit out’ AI, top U.S. Education Department official argues. *Education Week*. Retrieved November 20, 2023, from <https://www.edweek.org/technology/schools-cant-sit-out-ai-top-u-s-education-department-official-argues/2023/11>
- 105 In late summer 2023, the United Nations Educational, Scientific and Cultural Organization (UNESCO) published *Guidance for Generative AI in Education and Research*, a report that proposed “key steps for governmental agencies to regulate the use of generative AI” and “enable ethical and effective uses of this technology in education.”
- Miao, F. & Holmes, W. (2023, September 7). *Guidance for generative AI in education and research*. UNESCO. Retrieved November 8, 2023, from <https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>
- 106 Biden, J.R. (2023, October 30). *Executive order on the safe, secure, and trustworthy development and use of artificial intelligence*. The White House. Retrieved October 31, 2023, from <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>
- 107 Biden, J.R. (2023, October 30). *Executive order on the safe, secure, and trustworthy development and use of artificial intelligence* (Section 8(d)). The White House. Retrieved October 31, 2023, from <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>
- 108 The White House (2023, October 30). *FACT SHEET: President Biden issues executive order on safe, secure, and trustworthy artificial intelligence*. Retrieved October 31, 2023, from <https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/>
- 109 U.S. Department of Education, Office of Educational Technology. (2023, May). *Artificial intelligence and future of teaching and learning: Insights and recommendations* (for example, p. 44). Retrieved October 26, 2023, from <https://www2.ed.gov/documents/ai-report/ai-report.pdf>
- 110 The Department of Education appears unlikely to provide meaningful oversight of technology companies’ activities in schools. It has, for example, never enforced a FERPA violation, despite that no school can satisfy the Department’s requirement that it directly control technology vendors’ use and maintenance of student education records in their role as “school officials.”
- For definition and requirements of “school official,” see:
- United States Department of Education Student Privacy Policy Office (2021, July 9). *SPPO-21-04, A parent guide to the Family Educational Rights and Privacy Act (FERPA)* (p. 3). Retrieved August 21, 2023, from https://studentprivacy.ed.gov/sites/default/files/resource_document/file/A%20parent%20guide%20to%20ferpa_508.pdf
- 111 CRPE reports that only two U.S. states—California and Oregon—have provided any official guidance to school districts on using AI in schools for the 2023-2024 school year. Only 11 states are in the process of developing guidance.
- Dusseault, B. & Lee, J. (2023, October). *AI is already disrupting education, but only 13 states are offering guidance for schools*. CRPE. Retrieved November 22, 2023, from <https://crpe.org/ai-disrupt-ed-13-states/>

- 112 Merod, A. (2023, October 27). *AI guidance continues to roll out — this time it's for ed tech companies*. K-12 Dive. Retrieved November 8, 2023, from <https://www.k12dive.com/news/ed-tech-AI-guidance-schools/698052/>
- 113 Weidinger, L., Mellor, J., Rauh, M., Griffin, C., Uesato, J., Huang, P.S., Cheng, M., Glaese, M., Balle, B., Kasirzadeh, A., Kenton, Z., Brown, S., Hawkins, W., Stepleton, T., Biles, C., Birhane, A., Haas, J., Rimell, L., Hendricks, L.A., Isaac, W., Legassick, S., Irving, G., & Gabriel, I. (2022). Taxonomy of risks posed by language models. *FACCT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency*. Retrieved November 15, 2023, from <https://doi.org/10.1145/3531146.3533088>
- 114 Chatbot tutors can be defined as AI technologies that can hold conversations and produce all kinds of text in response to prompts, and are being developed extensively by ed tech industry companies.
- Schwartz, S. (2023, May 31). Can AI tutor students? Why it's unlikely to take over the job entirely. *Education Week*. Retrieved February 7, 2024, from <https://www.edweek.org/technology/can-ai-tutor-students-why-its-unlikely-to-take-over-the-job-entirely/2023/05>
- 115 Google promotes adaptive learning in its core AI offerings for education, as: “a type of learning where students are given customized resources and activities to address their unique learning needs . . . What's new is applying recent AI advances to this concept, which opens up a whole new set of possibilities to transform the future of school into a personal learning experience.”
- Cormie, A. (2022, March 16). Let's get personal: Adaptive learning tech and education. [blog post]. *The Keyword*. Retrieved February 7, 2024, from <https://blog.google/outreach-initiatives/education/adaptive-learning-technology/>
- 116 Khan Academy (2024). *For every student, every classroom. Real results*. Khan Academy [webpage]. Retrieved January 16, 2024, from <https://www.khanacademy.org/>
- 117 OpenAI (2023, March 14). Khan Academy: Khan Academy explores the potential for GPT-4 in a limited pilot program. Retrieved November 15, 2023, from <https://openai.com/customer-stories/khan-academy>
- Khan, S. (2023, March 14). Harnessing GPT-4 so that all students benefit. A nonprofit approach for equal access. [blog post]. *Khan Academy*. Retrieved November 21, 2023, from <https://blog.khanacademy.org/harnessing-ai-so-that-all-students-benefit-a-nonprofit-approach-for-equal-access/>
- Khan Academy (2023, March). *Update: Introducing... Khanmigo! Khan Academy's AI tool*. Retrieved November 21, 2023, from <https://support.khanacademy.org/hc/en-us/articles/14394953976333--Update-Introducing-Khanmigo-Khan-Academy-s-AI-Tool>
- 118 Salman Khan, the founder of Khan Academy, claimed, “We're at the cusp of using AI for probably the biggest positive transformation that education has ever seen. And the way we're going to do that is by giving every student on the planet an artificially intelligent but amazing personal tutor. And we're going to give every teacher on the planet an amazing, artificially intelligent teaching assistant.”
- Phillips, V. (2023, June 7). Intelligent classrooms: What AI means for the future of education. *Forbes*. Retrieved November 22, 2023, from <https://www.forbes.com/sites/vickiphillips/2023/06/07/intelligent-classrooms-what-ai-means-for-the-future-of-education/>
- See also:
- Khan Academy (2023). Khanmigo, your AI guide: Get a glimpse of the future of learning [webpage]. Retrieved November 22, 2023, from <https://www.khanacademy.org/khan-labs>
- 119 The analytics functionality in Khanmigo is called “Class Snapshot” and provides a breakdown of student progress from their interactions with the application. See the Khanmigo Teacher Tools page. Retrieved November 15, 2023, from <https://support.khanacademy.org/hc/en-us/articles/14799047733645-What->

teacher-tools-are-available-on-Khanmigo-

- 120 Khan Academy (2023). What teacher tools are available on Khanmigo? [webpage]. Retrieved November 22, 2023, from <https://blog.khanacademy.org/teacher-khanmigo/>
- 121 Boninger, F., Molnar, A., & Saldaña, C.M. (2019). *Personalized learning and the digital privatization of curriculum and teaching*. Boulder, CO: National Education Policy Center. Retrieved January 16, 2024, from <http://nepc.colorado.edu/publication/personalized-learning>
- 122 Statistical performance gains from students using AI—based on studies in controlled experimental classroom settings, highly prescriptive tasks, and short-term interventions, and measured through standardized assessments—can also be generalized as speculative claims about the potential of foundation models like GPT-4 as “personal tutors.” See, for example:
- Mollick, E. (2024, January 6). *Signs and portents: Some hints about what the next year of AI looks like*. One Useful Thing. Retrieved January 16, 2024, from <https://www.oneusefulthing.org/p/signs-and-portents>
- 123 Perrotta, C. & Williamson, B. (2018). The social life of learning analytics: Cluster analysis and the ‘performance’ of algorithmic education. *Learning, Media and Technology*, 43(1), 3-16.
- 124 Kerssens, N. (2022, December 1). *Valuing autonomy of teacher and child in cloud classrooms*. Digital Child. Retrieved October 20, 2023, from <https://digitalchild.org.au/valuing-autonomy-of-teacher-and-child-in-cloud-classrooms/>
- 125 Kerssens, N. & van Dijck, J. (2022). Governed by edtech? Valuing pedagogical autonomy in a platform society. *Harvard Educational Review*, 92(2), 284-303.
- 126 Khan Academy (n.d.). Khanmigo is a ground-breaking AI tool designed to improve learning and teaching. [webpage]. Retrieved October 19, 2023, from <https://blog.khanacademy.org/teacher-khanmigo/>
- 127 Khan Academy (2023, March). *How do the large language models powering Khanmigo work?* Retrieved November 22, 2023, from <https://support.khanacademy.org/hc/en-us/articles/13888935335309-How-do-the-Large-Language-Models-powering-Khanmigo-work>
- 128 Singer, N. (2023, June 26). In classrooms, teachers put A.I. tutoring bots to the test. *New York Times*. Retrieved October 18, 2023, from <https://www.nytimes.com/2023/06/26/technology/newark-schools-khan-tutoring-ai.html>
- 129 Highlighting the speculative nature of promotional rhetoric of AI in education, in a recent interview Salman Khan imagined that “one day, you’re going to have AI, and you’re going to be in an immersive world, and you’re gonna have a link straight into your brain . . . That’s likely to happen in the next 15 years, and it’s likely to be very accessible.”
- Possible.fm (2024, January 10). *Sal Khan on the future of K-12 education*. Retrieved January 31, 2024, from <https://www.possible.fm/podcasts/salkhan/>
- 130 Angwin, J. (2023, January 23). Decoding the hype about AI. *The Markup*. Retrieved October 19, 2023, from <https://themarkup.org/hello-world/2023/01/28/decoding-the-hype-about-ai>
- Bender, E., Gebru, T., McMillan-Major, A., & Schmittell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *FACCT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. Retrieved November 8, 2023, from <https://doi.org/10.1145/3442188.3445922>
- 131 Ferlazzo, L. (2023, June 19). Here’s what students think about using AI in the classroom. *Education Week*. Retrieved February 7, 2024, from <https://www.edweek.org/technology/opinion-heres-what-students-think-about-using-ai-in-the-classroom/2023/06>
- 132 Smith, M.S. (2023, June 23). The Internet isn’t completely weird yet; AI can fix that. *IEEE Spectrum*.

Retrieved October 27, 2023, from <https://spectrum.ieee.org/ai-collapse>

- 133 Technology critic and educator Matthew Kirschenbaum calls this a “textpocalypse.”
Kirschenbaum, M. (2023, March 8). Prepare for the textpocalypse. *The Atlantic*. Retrieved November 8, 2023, from <https://www.theatlantic.com/technology/archive/2023/03/ai-chatgpt-writing-language-models/673318/>
- 134 BBC. (2023, September 16). AI used to target kids with disinformation. *BBC News*. Retrieved December 5, 2023, from <https://www.bbc.co.uk/newsround/66796495>
- 135 Tiku, N., Schaul, K., & Chen, S.Y. (2023, November 1). These fake images reveal how AI amplifies our worst stereotypes. *The Washington Post*. Retrieved November 1, 2023, from <https://www.washingtonpost.com/technology/interactive/2023/ai-generated-images-bias-racism-sexism-stereotypes/>
- 136 Lykiardopoulou, I. (2023, August 31). LLMs prone to data poisoning and prompt injection risks, UK authority warns. *The Next Web*. Retrieved November 15, 2023, from <https://thenextweb.com/news/llms-data-poisoning-prompt-injection-risks>
- 137 Khan Academy (2023, March). *How do the large language models powering Khanmigo work?* Retrieved November 22, 2023, from <https://support.khanacademy.org/hc/en-us/articles/13888935335309-How-do-the-Large-Language-Models-powering-Khanmigo-work->
- 138 Khan Academy (2023, March). *What safety features does Khanmigo have?* Retrieved November 22, 2023, from <https://support.khanacademy.org/hc/en-us/articles/14394814244365-What-safety-features-does-Khanmigo-have->
- 139 The ed tech company Pearson suggests it will mitigate the tendency of generative AI to produce false information by only using internally generated material to create and train its AI applications. It remains to be seen whether Pearson’s efforts will succeed—or perhaps more accurately, it remains to be seen whether any third-party independent entity will be able to verify Pearson’s claims.
O’Boyle, D. (2023, May 9). Pearson taking legal action over use of its content to train language models, as it announces own AI plans. *Evening Standard*. Retrieved November 15, 2023, from <https://www.standard.co.uk/business/pearson-takes-legal-action-over-use-of-its-content-to-train-ai-b1079807.html>
- 140 When Newark, NJ, schools piloted Khanmigo, for example, teachers found an error in how the program answered a math question. They alerted Khan Academy, who fixed that particularly error. What Khan Academy could not do, however, was prevent errors more generally—and AI is prone to errors.
Singer, N. (2023, June 26). In classrooms, teachers put A.I. tutoring bots to the test. *New York Times*. Retrieved October 18, 2023, from <https://www.nytimes.com/2023/06/26/technology/newark-schools-khan-tutoring-ai.html>
- 141 Cooke, E. (2023, July 3). AI model collapse could spell disaster for AI development, say new studies. *Verdict*. Retrieved November 15, 2023, from <https://www.verdict.co.uk/ai-model-collapse-could-spell-disaster-for-ai-development-say-new-studies/?cf-view>
- 142 Shelton, C.C., Koehler, M.J., Greenhalgh, S.P., & Carpenter, J.P. (2022) Lifting the veil on TeachersPayTeachers.com: an investigation of educational marketplace offerings and downloads. *Learning, Media and Technology*, 47(2), 268-287.
- 143 OpenAI (2023, November 6). Introducing GPTs [blog post]. *OpenAI Blog*. Retrieved November 15, 2023, from <https://openai.com/blog/introducing-gpts>
In an apparent effort to protect the GPTs marketplace from criticism, amidst public concern over children’s access to harmful online content, OpenAI partnered with Common Sense Media to create guidelines on “safe,

responsible use of ChatGPT™ for educators and young people.

- Edwards, B. (2024, January 29). OpenAI and Common Sense Media partner to protect teens from AI harms and misuse. *ArsTechnica*. Retrieved January 31, 2024, from <https://arstechnica.com/information-technology/2024/01/openai-and-common-sense-media-partner-to-protect-teens-from-ai-harms-and-misuse/>
- 144 David, J. (2023, October 30). *Students investigate how artificial intelligence perpetuates biases*. Information School, University of Washington. Retrieved January 31, 2024, from <https://ischool.uw.edu/news/2023/10/students-investigate-how-artificial-intelligence-perpetuates-biases>
- 145 Marshall, J. (2023, July 5). As AI cheating booms, so does the industry detecting it: 'We couldn't keep up with demand.' *The Guardian*. Retrieved October 23, 2023, from <https://www.theguardian.com/technology/2023/jul/05/as-ai-cheating-booms-so-does-the-industry-detecting-it-we-couldnt-keep-up-with-demand>
- 146 Fowler, G.A. (2023, August 14). What to do when you're accused of AI cheating. *The Washington Post*: <https://www.washingtonpost.com/technology/2023/08/14/prove-false-positive-ai-detection-turnitin-gptzero/>
- 147 Mathewson, T.G. (2023, August 14). AI detection tools falsely accuse international students of cheating. *The Markup*. Retrieved November 8, 2023, from <https://themarkup.org/machine-learning/2023/08/14/ai-detection-tools-falsely-accuse-international-students-of-cheating>
- 148 Feathers, T. (2019, August 20). Flawed algorithms are grading millions of students' essays. *Vice*. Retrieved November 8, 2023, from <https://www.vice.com/en/article/pa7dj9/flawed-algorithms-are-grading-millions-of-students-essays>
- 149 Kelley, J. & Prince, K. (2023, October 30). *How GoGuardian invades student privacy*. Electronic Frontier Foundation. Retrieved November 8, 2023, from <https://www.eff.org/deeplinks/2023/10/how-goguardian-invades-student-privacy>
- 150 Blackboard (2022). Blackboard Learn AI design assistant [webpage]. Retrieved October 26, 2023, from https://help.blackboard.com/Learn/Administrator/SaaS/Tools_Management/Learn_AI_Design_Assistant
- Instructure (2023, July 27). *Instructure reveals major product innovations across four strategic areas* [press release]. Retrieved October 26, 2023, from <https://www.instructure.com/en-au/press-release/instructure-reveals-major-product-innovations-across-four-strategic-areas>
- 151 Kerssens, N., Pangrazio, L., & Nichols, T.P. (2023, September). Googlization(s) of education: Intermediary work brokering platform dependence in three national school systems. *Learning, Media and Technology*. Retrieved October 20, 2023, from <https://doi.org/10.1080/17439884.2023.2258339>
- Singer, N. (2017). How Google took over the classroom. *New York Times*. Retrieved October 26, 2023, from <https://www.nytimes.com/2017/05/13/technology/google-education-chromebooks-schools.html>
- 152 Williamson, B. (2021, May 28). Google's plans to bring AI to education make its dominance in classrooms more alarming. *Fast Company*. Retrieved October 26, 2023, from <https://www.fastcompany.com/90641049/google-education-classroom-ai>
- 153 Google (n.d.). *Advancing education with AI*. Retrieved January 31, 2024, from https://edu.google.com/intl/ALL_uk/why-google/ai-for-education/
- 154 Gulson, K., Perrotta, C., Williamson, B., & Witzemberger, K. (2021). Should we be worried about Google Classroom? The pedagogy of platforms in education. *Journal of Professional Learning*, Semester 2, 2021: <https://cpl.nswtf.org.au/journal/semester-2-2021/should-we-be-worried-about-google-classroom-the-pedagogy-of-platforms-in-education/>
- 155 Perrotta, C., Gulson, K.N., Williamson, B., & Witzemberger, K. (2021) Automation, APIs and the distributed

- labour of platform pedagogies in Google Classroom. *Critical Studies in Education*, 62(1) 97-113. Retrieved February 6, 2024, from <https://www.tandfonline.com/doi/full/10.1080/17508487.2020.1855597>
- 156 Google (n.d.). *Google Classroom: Where teaching and learning come together*. Retrieved November 22, 2023, from <https://edu.google.com/workspace-for-education/classroom/>
- 157 Perrotta, C., Gulson, K.N., Williamson, B., & Witzemberger, K. (2021). Automation, APIs and the distributed labour of platform pedagogies in Google Classroom. *Critical Studies in Education*, 62(1), 97-113.
- Veale, M. (2022). Schools must resist big EdTech – but it won't be easy. In S. Livingstone & K. Pothong (Eds.), *Education Data Futures: Critical Regulatory and Practical Reflections*. Digital Futures Commission, 5Rights Foundation. Retrieved November 22, 2023, from <https://educationdatafutures.digitalfuturescommission.org.uk/essays/competing-interests-in-education-data/schools-must-resist-big-edtech>
- 158 Singer, N. (2023, June 26). In classrooms, teachers put A.I. tutoring bots to the test. *New York Times*. Retrieved October 18, 2023, from <https://www.nytimes.com/2023/06/26/technology/newark-schools-khan-tutoring-ai.html>
- 159 Kieczka, D. (2023, March 16). Practice sets: A more personal path to learning [blog post]. *The Keyword*. Retrieved November 22, 2023, from <https://blog.google/outreach-initiatives/education/introducing-practice-sets/>
- 160 Komljenovic, J. (2021). The rise of education rentiers: Digital platforms, digital data and rents. *Learning, Media and Technology*, 46(3), 320-33.
- 161 Yu, J. & Couldry, N. (2021). Education as a domain of natural data extraction: Analysing corporate discourse about educational tracking. *Information, Communication and Society*, 25(1), 127-144.
- 162 Hooper, L., Livingstone, S., & Pothong, K. (2022). *Problems with data governance in UK schools: The cases of Google Classroom and ClassDojo*. Digital Futures Commission, 5Rights Foundation. Retrieved November 22, 2023, from https://eprints.lse.ac.uk/119736/1/Problems_with_data_governance_in_UK_schools.pdf
- 163 The Federal Trade Commission has recently promised to crack down on indiscriminate data collection by the ed tech industry. A child's education, it stated, must not be "conditional" on surveillance.
- Fair, L. (2022, May 22). *FTC to ed tech: Protecting kids' privacy is your responsibility*. Federal Trade Commission, Business Blog. Retrieved November 8, 2023, from <https://www.ftc.gov/business-guidance/blog/2022/05/ftc-ed-tech-protecting-kids-privacy-your-responsibility>
- 164 Solove, D.J. (2001). Privacy and power: Computer databases and metaphors for information privacy. *Stanford Law Review*, 53, 1393-1462. Retrieved November 22, 2023, from https://scholarship.law.gwu.edu/cgi/viewcontent.cgi?article=2077&context=faculty_publications
- 165 Perrotta, C. (2023). Advancing data justice in education: Some suggestions towards a deontological framework. *Learning, Media and Technology*, 48(2), 187-199.
- 166 Whitman, M. (2020). "We called that a behavior": The making of institutional data. *Big Data and Society*, 7(1). Retrieved November 22, 2023, from <https://journals.sagepub.com/doi/pdf/10.1177/2053951720932200>
- 167 Levin, D.A. (2022). The state of K-12 cybersecurity: Year in review – 2022 annual report. *K12 Security Information Exchange*. Retrieved November 8, 2023, from <https://www.k12six.org/the-report>
- 168 Singer, N. (2022, July 31). A cyberattack illuminates the shaky state of student privacy. *New York Times*. Retrieved October 18, 2023, from <https://www.nytimes.com/2022/07/31/business/student-privacy-illuminate-hack.html>
- 169 Merod, A. (2023, March 29). Ed tech experts urge caution on ChatGPT's student data privacy. *K-12 Dive*. Retrieved November 16, 2023, from <https://www.k12dive.com/news/chatgpt-student-data-privacy->

concern/646297/

- 170 Ada Lovelace Institute, AI Now Institute, & Open Government Partnership. (2021). *Algorithmic accountability for the public sector*. Retrieved December 5, 2023, from <https://www.opengovpartnership.org/documents/algorithmic-accountability-public-sector>
- 171 Citron, D.K. (2023). The surveilled student. *Stanford Law Review*, 76 (Forthcoming), Virginia public law and legal theory research paper 2023-61. Retrieved December 5, 2023, from <https://ssrn.com/abstract=4552267>
- 172 Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data and Society*, 3(1). Retrieved February 21, 2024, from <https://doi.org/10.1177/2053951715622512>
- 173 Rudin, C. & Radin, J. (2019). Why are we using black box models in AI when we don’t need to? A lesson from an explainable AI competition. *Harvard Data Science Review*, 1(2). Retrieved November 15, 2023, from <https://doi.org/10.1162/99608f92.5a8a3a3d>
- 174 Computer scientist Cynthia Rudin explains and advocates for the superiority of inherently interpretable AI models to black box models. Interpretable models are not the same as open source models.
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1, 206-215.
- 175 Broussard, M. (2023). How to investigate an algorithm. *Issues in Science and Technology*, 39(4), 85-89. Retrieved November 8, 2023, from <https://issues.org/algorithm-auditing-more-than-glitch-broussard/>
- 176 Perrotta, C. (2023, November 28). The unreasonableness of AI and its bearing on educational responsibility. *AutomatED* [blog post]. Retrieved November 30, 2023, from <https://automatedonline.org/2023/11/28/plugin-and-play-education-knowledge-and-learning-in-the-age-of-platforms-and-artificial-intelligence-2024/>
- 177 Holmes, W., Porayska-Pomsta, K., Holstein, K., et al. (2022). Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, 32, 504–526. Retrieved January 10, 2024, from <https://doi.org/10.1007/s40593-021-00239-1>
- 178 Responsible AI UK. (2023). *Guiding principles*. Retrieved January 10, 2024, from <https://www.rai.ac.uk/guiding-principles>
- 179 Kousa, P. & Niemi, H. (2023). Artificial intelligence ethics from the perspective of educational technology companies and schools. In H. Niemi, R.D. Pea, & Y. Lu (Eds.), *AI in Learning: Designing the Future*. Springer, Cham. Retrieved January 16, 2024, from https://doi.org/10.1007/978-3-031-09687-7_17
- 180 Porayska-Pomsta, K. (2023). A manifesto for a pro-actively responsible AI in education. *International Journal of Artificial Intelligence in Education*. Retrieved January 10, 2024, from <https://doi.org/10.1007/s40593-023-00346-1>
- 181 Pea, R.D., Biermacki, P., Bigman, M., et al. (2023). Four surveillance technologies creating challenges for education. In H. Niemi, R.D. Pea, & Y. Lu (Eds.), *AI in Learning: Designing the Future*. Springer, Cham. Retrieved January 16, 2024, from https://doi.org/10.1007/978-3-031-09687-7_19
- 182 According to the Edinburgh Declaration on Responsibility for Responsible AI:
- All who exercise and unleash the material and economic power of AI/AS, even with the best of intentions, must start attending — and be required to attend — responsibly to the vulnerabilities that AI/AS create, perpetuate, amplify or relieve. ... Some applications of AI/AS, ... such as the automation of moral and political human judgments with irreversible and severe consequences, arguably require delay or indefinite pause.
- Vallor, S. (2023). Edinburgh declaration on responsibility for responsible AI. *Medium*. Retrieved January 10, 2024, from https://medium.com/@svallor_10030/edinburgh-declaration-on-responsibility-for-responsible-

ai-1a98ed2e328b

183 Ryan-Mosley, T. (2023, August 14). AI isn't great at decoding human emotions. So why are regulators targeting the tech? *MIT Technology Review*. Retrieved January 31, 2024, from <https://www.technologyreview.com/2023/08/14/1077788/ai-decoding-human-emotions-target-for-regulators/>

184 AI Now Institute. (2023, April 11). *Algorithmic accountability: Moving beyond audits*. Retrieved January 16, 2024, from <https://ainowinstitute.org/publication/algorithmic-accountability>

185 Floridi, L. (2019). Translating principles into practices of digital ethics: Five risks of being unethical. *Philosophy and Technology*, 32, 185–193. Retrieved January 10, 2024, from <https://doi.org/10.1007/s13347-019-00354-x>

Shazeda, A., Jaźwińska, K., Ahlawat, A., Winecoff, A., & Wang, M. (2023). *Building the epistemic community of AI safety*. SSRN. Retrieved January 16, 2024, <http://dx.doi.org/10.2139/ssrn.4641526>

186 Coldicutt, R. (2023). Putting the lid on Pandora's Box: How community power shapes AI [blog post]. *Promising Trouble*. Retrieved January 10, 2024, from <https://www.promisingtrouble.net/blog/2023-12-pandoras-box>

Gentelet, K. (2024, January 16). Needed: Ways for citizens to sound the alarm about AI's societal impacts. *Issues in Science and Technology*. Retrieved January 31, 2024, from <https://issues.org/ai-deployment-citizens-input-gentelet/>

187 Singer, N. (2023, June 26). In classrooms, teachers put A.I. tutoring bots to the test. *New York Times*. Retrieved October 18, 2023, from <https://www.nytimes.com/2023/06/26/technology/newark-schools-khan-tutoring-ai.html>

Similarly, Arizona State University President Michael Crow explained the University's partnership with OpenAI by noting: "ASU recognizes that augmented and artificial intelligence systems are here to stay, and we are optimistic about their ability to become incredible tools that help students to learn, learn more quickly and understand subjects more thoroughly . . . "

Davis, A. (2024, January 18). A new collaboration with OpenAI charts the future of AI in higher education. *ASU News*. Retrieved February 6, 2024, from <https://news.asu.edu/20240118-university-news-new-collaboration-openai-charts-future-ai-higher-education>

188 In 1966, Joseph Weizenbaum developed the ELIZA computer program that was capable of chatting with humans like an empathic psychologist (i.e., an early "chatbot"). He later published papers that called on scientists and engineers, in particular, to critically assess when and how computer technologies should be applied to human affairs. His 1978 paper asked the following:

Who is the beneficiary of our much-advertised technological progress and who are its victims?

What limits ought we, the people generally and scientists and engineers particularly, to impose on the application of computation to human affairs?

What is the impact of the computer, not only on the economies of the world or on the war potential of nations, etc., but on the self-image of human beings and on human dignity?

What irreversible forces is our worship of high technology, symbolized most starkly by the computer, bringing into play?

Will our children be able to live with the world we are here and now constructing?

Weizenbaum, J. (1978, September). Once more—A computer revolution. *Bulletin of the Atomic Scientists*, 34(7), 12-19. Retrieved October 17, 2023, from <https://doi.org/10.1080/00963402.1978.11458531>

After the release of ChatGPT and the subsequent rush to implement AI in all sectors, Jack Stilgoe, in August

2023, returned to Weizenbaum's questions, summarizing them as we have here: Is it good? Do we need it?

Stilgoe, J. (2023, August 11). We need a Weizenbaum test for AI. *Science*, 381(6658). Retrieved October 17, 2023, from <https://www.science.org/doi/10.1126/science.adk0176>

See also:

Berry, D.M. (2023). The limits of computation: Joseph Weizenbaum and the ELIZA chatbot. *Weizenbaum Journal of the Digital Society*, 3(3). Retrieved November 8, 2023, from <https://doi.org/10.34669/WI.WJDS/3.3.2>

- 189 Watters, A. (2021). *Teaching machines: The history of personalized learning*. London: MIT Press.
- 190 Boninger, F., Molnar, A., & Saldaña, C.M. (2019). *Personalized learning and the digital privatization of curriculum and teaching*. Boulder, CO: National Education Policy Center. Retrieved August 16, 2023, from <http://nepc.colorado.edu/publication/personalized-learning>
- 191 Pelletier, C. (2023). Against personalised learning. *Interational Journal of Artificial Intelligence in Education*. Retrieved November 8, 2023, from <https://doi.org/10.1007/s40593-023-00348-z>
- 192 Narayanan, A. & Kapoor, S. (2023, March 22). OpenAI's policies hinder reproducible research on language models [blog post]. *AI Snake Oil*. Retrieved October 31, 2023, from <https://www.aisnakeoil.com/p/openai-policies-hinder-reproducible>
- 193 Weise, K. & Metz, C. (2023, May 1). When A.I. chatbots hallucinate. *New York Times*. Retrieved November 8, 2023, from <https://www.nytimes.com/2023/05/01/business/ai-chatbots-hallucination.html>
- 194 Smith, M.S. (2023, June 23). The Internet isn't completely weird yet; AI can fix that. *IEEE Spectrum*. Retrieved October 31, 2023, from <https://spectrum.ieee.org/ai-collapse>
- 195 Holmes, W. (2023). *The unintended consequences of artificial intelligence and education*. Brussels: Education International. Retrieved November 15, 2023, from <https://www.ei-ie.org/en/item/28115:the-unintended-consequences-of-artificial-intelligence-and-education>
- 196 Narayanan, A. & Kapoor, S. (2023, July 19). Is GPT-4 getting worse over time? [blog post]. *AI Snake Oil*. Retrieved October 31, 2023, from <https://www.aisnakeoil.com/p/is-gpt-4-getting-worse-over-time>
- 197 Birhane, A., Kalluri, P., Card, D., Agnew, W., Dotan, R., & Bao, M. (2022). The values encoded in machine learning research. *2022 ACM Conference on Fairness, Accountability, and Transparency*. Retrieved November 15, 2023, from <https://doi.org/10.1145/3531146.3533083>
- 198 Algorithms built into educational software are presented as “neutral” and “scientific,” and as embodying “truth” or fact. They cannot be neutral, however, because they are created by people—and people are not neutral.
- Crawford, K. (2016, June 25). Artificial intelligence's white guy problem. *New York Times*. Retrieved September 2, 2020, from <https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html>
- Freeguard, G. (2020, August 19). Four things government must learn from the A-level algorithm fiasco [blog post]. *Institute for Government*. Retrieved September 2, 2020, from <https://www.instituteforgovernment.org.uk/blog/a-level-algorithm-fiasco>
- O'Neil, C. (2016, October 12). Algorithms are as biased as human curators [webpage]. *ORCAA*. Retrieved July 1, 2020, from <https://orcaarisk.com/articles/2016/10/12/algorithms-are-as-biased-as-human-curators>
- 199 In their critique of using black box AI models to make significant decisions, Cynthia Rudin and Joanna Radin explain that “Even if one has a list of the input variables, black box predictive models can be such complicated

functions of the variables that no human can understand how the variables are jointly related to each other to reach a final prediction.” (p. 3)

Rudin, C. & Radin, J. (2019). Why are we using black box models in AI when we don't need to? A lesson from an explainable AI competition. *Harvard Data Science Review*, 1(2). Retrieved November 15, 2023, from <https://doi.org/10.1162/99608f92.5a8a3a3d>

200 Currently, their programming could feasibly be audited if regulatory mechanisms required it and independent specialists were available. Such auditing is recommended in, for example:

Molnar, A., Boninger, F., Noble, A., & Meenakshi, M. (2023). *We need better education policy. Summit Public Schools shows why*. Boulder, CO: National Education Policy Center. Retrieved October 2, 2023, from: <https://nepc.colorado.edu/publication/summit-2023>

201 Floyd, C. (2023). From Joseph Weizenbaum to ChatGPT: Critical encounters with dazzling AI technology. *Weizenbaum Journal of the Digital Society*, 3(3). Retrieved November 8, 2023, from <https://doi.org/10.34669/WI.WJDS/3.3.3>

202 Broussard, M. (2023). How to investigate an algorithm. *Issues in Science and Technology*, 39(4), 85-89. Retrieved November 8, 2023, from <https://issues.org/algorithm-auditing-more-than-glitch-broussard/>

203 McQuillan, D., Jarke, J., & Pargman, T.C. (2023). We are at an extreme point where we have to go all in on what we really believe education should be about. *Postdigital Science and Education*. Retrieved January 16, 2024, from <https://doi.org/10.1007/s42438-023-00433-5>

204 Taylor, L. (2024, January 16). AI lacks ethic checks for human experimentation. *Issues in Science and Technology*. Retrieved January 31, 2024, from <https://issues.org/ai-ethics-human-experimentation-taylor/>