

NEPC Review: Learning Systems: Shaping the Role of Artificial Intelligence in Education (Bellwether, September 2024)



Reviewed by:

**T. Philip Nichols
Baylor University**

November 2024

National Education Policy Center

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

Acknowledgements

NEPC Staff

Faith Boninger
Publications Manager

Elaine Duggan
Production Design

Alex Molnar
NEPC Director

Kevin Welner
NEPC Director

Suggested Citation: Nichols, T.P. (2024). *NEPC review: Learning systems: Shaping the role of artificial intelligence in education*. Boulder, CO: National Education Policy Center. Retrieved [date] from <http://nepc.colorado.edu/review/learning-systems>


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



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

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.

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>



NEPC Review: Learning Systems: Shaping the Role of Artificial Intelligence in Education (Bellwether, September 2024)


Reviewed by:

T. Philip Nichols
Baylor University

November 2024

Summary

In September 2024, Bellwether released a three-part report, entitled *Learning Systems*, on artificial intelligence (AI) in K-12 education. The report's stated purpose is to provide a holistic picture of the opportunities and risks that AI presents for education stakeholders, and it offers recommendations for policymakers and practitioners to navigate this emerging landscape. The report also highlights current and potential uses of AI for streamlining administrative tasks and personalizing instruction, and it identifies challenges for capitalizing on these potentials. It argues that educators have a responsibility to develop capacity and infrastructure to support the effective integration of AI in schools while trying to mitigate the risks involved in doing so. However, by allowing the projected benefits of AI to drive decisions related to AI's development and implementation, the report overlooks substantial research literatures that document its known limitations and harms. As a result, while the report's stated aim is to present policymakers and practitioners with a holistic view of AI in education, its recommendations are skewed toward the imagined future advantages of AI rather than its actual present risks. As one example, the report encourages the development of education-specific datasets to improve the accuracy of AI technologies in schools, but it does not discuss how such efforts could be carried out without amplifying known issues related to bias and surveillance in educational data processing. This undermines the report's relevance and usefulness for policymakers, particularly those whose goal is to ensure that technology adoption and implementation prioritizes equitable learning for all students.



NEPC Review: Learning Systems: Shaping the Role of Artificial Intelligence in Education (Bellwether, September 2024)

Reviewed by:

T. Philip Nichols
Baylor University

November 2024

I. Introduction

Since the public launch of ChatGPT in November 2022, artificial intelligence (AI) has become a pressing topic in education policy and practice. Some education stakeholders see ChatGPT and similar “generative AI” (GenAI) technologies—i.e., those that can generate text, image, audio, and video in response to prompts from users—as having potential to streamline administrative tasks, personalize instruction and assessment, and open pathways for new school models. Accordingly, they suggest policymakers should prioritize supporting the effective integration of GenAI in schools so that K-12 education can maximize the benefits of this technology.¹

However, others argue that GenAI’s benefits for education are outweighed by the risks it introduces or intensifies, especially those related to academic integrity, student data security, algorithmic bias, and energy usage. For this reason, they recommend policymakers show caution and restraint as they consider what role, if any, GenAI should play in K-12 schools.²

These promises and concerns leave policymakers without clear direction as they make decisions about GenAI in leadership, teaching, and learning. A recent three-part report series from Bellwether, authored by Amy Chen Kulesa, Michelle Croft, Brian Robinson, Mary K. Wells, Andrew J. Rotherham, and John Bailey, aims to address this uncertainty by offering a holistic picture of the GenAI landscape in education and the opportunities and risks it presents. The first report, *Learning Systems: The Landscape of Artificial Intelligence in K-12 Education*,³ provides an overview of recent developments, impacts, and usage of AI in education. The second report, *Learning Systems: Opportunities and Challenges of Artificial Intelligence-Enhanced Education*,⁴ describes current platforms, policies, and stakeholder experiences related to AI in education. The third report, *Learning Systems: Artificial In-*

telligence Use Cases,⁵ spotlights potential uses for AI in education and the risks associated with each.

II. Findings and Conclusions of the Report

Based on original interviews and a review of literature, the report concludes that AI is neither “a passing fad nor a panacea” but “a potentially powerful tool that can support steady, long-term improvements toward a more equitable and efficient education system.”⁶ It offers 10 recommendations, clustered into three categories, which are intended to support education stakeholders in capitalizing on AI’s benefits while mitigating its risks.⁷

Building Strong Capacity

The report finds that effectively integrating AI in education will require capacity-building in knowledge, skills, and diversity of leadership. Examples of this include strengthening AI literacy among different stakeholders, developing comprehensive policies to guide AI adoption, and building partnerships between AI developers and educators to ensure that tools are relevant, effective, and safe. The report also highlights the importance of developing expertise through training or cross-sector collaborations to address ethical challenges related to data security and to support the development of education-specific AI resources. It also notes the need for including diverse voices throughout the processes of AI development, procurement, and implementation.

Building Resilient Infrastructure

The report finds that robust and secure infrastructure is necessary both for improving the quality of AI outputs and for protecting sensitive data. It recommends developing reliable networks for safely sharing information and creating education-specific datasets and benchmarks to tailor the use of AI tools for administrative or instructional purposes. It also suggests that ongoing research will be vital to evaluate the impact and efficacy of such tools as they are implemented and to refine them accordingly.

Building Thoughtful Design

The report recommends that the design of AI tools ought to account for both the needs of diverse students and the value of human interaction. This involves building and adopting AI tools with specific educational goals in mind and ensuring that such resources can support the learning of all students. It also involves acknowledging the limitations of AI in certain circumstances and considering what skills should be cultivated through human, rather than AI-enabled, interactions.

III. The Report’s Rationale for Its Findings and Conclusions

The underlying rationale for the report’s recommendations is that GenAI marks a significant breakthrough in technological research and development. It argues that GenAI is distinct from earlier forms of AI (which it terms “Traditional AI”)⁸ and that it is poised to dramatically alter life and work across social sectors, including education, in the near future.⁹ Consequently, the report suggests that the risks associated with GenAI, while real and serious, must be weighed against the risks of lagging behind in adopting GenAI in schools during this transitional moment.¹⁰ In other words, the report positions AI’s imminent risks not as reasons to slow its development and implementation in education but as temporary hurdles to be overcome on the journey toward the “equitable and efficient education system” that AI will enable.¹¹ This rationale leads to the report’s conclusion that education stakeholders have a responsibility to build capacity and infrastructure to support the use of AI in K-12 schools, despite the inherent risks involved.

IV. The Report’s Use of Research Literature

The report does not substantively engage with relevant research literature. Of the 284 citations included across the three-part series, only seven are peer-reviewed journal articles (and three of these are duplicate citations). The remainder of the references are to reports, news stories, opinion articles, and product press releases. This is not, in itself, a flaw. Many reports, especially those covering new topics that do not have an expansive research literature, rely on other sources to make sense of an emerging terrain. The report is transparent that it views GenAI as such an innovation. It distinguishes between “GenAI” and “Traditional/Classical AI” and says that it is primarily focused on the transformative developments occurring in the former.¹²

However, this distinction overstates the novelty of GenAI as a technology. The report delineates Traditional AI as responding to inputs by making predictions based on data and GenAI as generating “new text, images, and other media in response to prompts.”¹³ This explanation omits the fact that the process by which GenAI generates new text involves the same methods (i.e., statistical predictions) and computational models (i.e., neural networks) that Traditional AI research and development has used since the 1950s.¹⁴ While these techniques have become more sophisticated as computer processing power and the size of available datasets have grown, there is not a fundamental difference in the underlying processes of Traditional AI and GenAI.

By treating the two as distinct categories, the report misses an opportunity to engage research literature that may not overtly reference “GenAI” but is nevertheless relevant for its analysis of GenAI’s benefits and risks in education. This includes extensive research related to the human biases encoded in AI datasets and platforms,¹⁵ the theories of teaching and learning embedded in AI educational technologies,¹⁶ and the labor and governance demands that follow from the adoption of AI platforms in schools.¹⁷ As a result, the report fails to address well-documented issues related to AI-enabled data processing and the reproduction

of social inequalities along lines of race, class, gender, sexuality, and ability.¹⁸ To the extent that the report references matters of equity, it is focused primarily on promoting equal access to the use of AI tools and competencies in K-12 education—not the mitigation of AI’s known impacts on non-dominant communities.¹⁹

In the absence of direct engagement with research literature, the report relies instead on interviews and recent source materials that share its assumptions about the novelty of GenAI. The result is a misleading picture of GenAI as a sudden breakthrough rather than an incremental achievement. In presenting a timeline of “AI Milestones,” for instance, the report leaps from 1956, when the term “artificial intelligence” was coined, to 2010-2015, when Google and OpenAI launched dedicated AI research labs²⁰—skipping over the starts, stops, and paths-not-taken in the decades of AI development in between.²¹ In this way, the report positions GenAI as a revolutionary invention, and it leans on the market projections of industry leaders and venture capital firms to support its claims about the social transformations GenAI will introduce and the need for education stakeholders to adapt to them.²² By using such projections, the report ignores the volatility of the technology sector, where it is not uncommon for companies and markets to collapse with little warning.²³ Moreover, these sources skew the report’s analysis toward preparing K-12 education for an imagined, AI-enhanced future. This leaves little room for evidence that might complicate this goal—e.g., like that related to AI’s known and imminent risks—to figure in the report’s recommendations, except as temporary obstacles to be overcome.

V. Review of the Report’s Methods

The report does not provide a detailed accounting of its methods. It says it was developed “through consultations with experts and practitioners,”²⁴ and it provides a list of these interviewees at the end of the document.²⁵ However, there is no description of the selection criteria for participants, the structure or length of the interviews, or the methods used to analyze the resulting data to inform the report’s recommendations. Moreover, the composition of the interviewee list appears to overrepresent the perspectives of individuals in technology, business, and nonprofit sectors—many of whom have a vested interest in the widespread adoption of AI in K-12 education. Of the 38 interviews listed, only a quarter are with education researchers or school leaders, and none of the interviewees are among the many prominent experts who have raised concerns about GenAI’s imminent harms. This suggests the report’s recommendations are not the result of serious wrestling with AI’s short- and long-term impacts.

VI. Review of the Validity of the Findings and Conclusions

The report is correct that GenAI presents opportunities and challenges for K-12 education and that stakeholders need guidance to navigate these competing factors. The report’s account of these benefits and risks is extensive but incomplete. Its recommendations address

important issues related to the ethics and efficacy of integrating AI in schools, including some that are underexplored in the research literature: for instance, how the relative “openness” of different companies’ AI models (e.g., Gemini, GPT, Claude, Llama, Hugging Face) can limit transparency and control related to school and student data;²⁶ and how education-specific datasets and benchmarks, and techniques for using them (e.g., fine-tuning, retrieval-augmented generation) could improve the quality of AI resources for specific administrative and instructional functions.²⁷ However, the report also omits other concerns that experts have identified—notably, the climate impacts of AI development and usage.²⁸

The validity of the report’s recommendations is undermined by its larger position, unsupported by evidence, that AI’s known risks and limitations should not deter educators from incorporating it into practice. This position prevents the report from acknowledging the ways its recommendations could actually contribute to, rather than solve, the problems it names. The development of large education-specific datasets, for instance, would require an increase in already-elevated levels of data collection in schools—exacerbating existing ethical issues related to the surveillance and privacy of minors and marginalized communities.²⁹ Moreover, experts argue that such forms of mass data collection can have diminishing returns with regard to accuracy, meaning that the resulting compromises in privacy might not even lead to the promised outcomes associated with them.³⁰ Such contradictions reduce the relevance of the report’s recommendations for actually contending with the challenges it seeks to address.

VII. Usefulness of the Report for Guidance of Policy and Practice

The report is a useful starting point for understanding the landscape of AI in K-12 education and the opportunities and challenges it presents for teaching, learning, and leadership. The report raises important, and under-discussed, points about the quality and security of data in AI models and the potential of AI tools that are built for education-specific purposes. Such details could be useful considerations in setting policies related to procurement and evaluation of AI products in educational systems. The report’s general focus on the importance of secure infrastructure, human interaction, and diverse perspectives in AI development and implementation is valuable and aligns with existing research.³¹ Policymakers and practitioners would benefit from attending to these areas of concern.

However, the report couches its recommendations in a larger argument about positive transformations that GenAI is poised to introduce in diverse sectors, including education. This argument is based on market projections rather than empirical evidence, and it leads the report to downplay AI’s imminent risks as temporary obstacles that can be overcome as its anticipated benefits are realized. Consequently, policymakers and practitioners should be wary of accepting the report’s recommendations at face value. The projected transformations on which the report’s conclusions are based are unlikely to occur as predicted, if at all. For this reason, the immediate and well-documented risks that AI poses for ethical and equitable education ought to drive policymakers’ and practitioners’ engagements with AI

rather than the imagined future its developers and investors promise. Only to the extent that specific recommendations in the report offer guidance for education stakeholders to address the former, rather than the latter, will they be useful, particularly for those whose aim is to prioritize equitable learning for all students.

Notes and References

- 1 Acar, O.A. (2024, February 19). *With generative AI we can reimagine education – and the sky is the limit*. World Economic Forum. Retrieved October 24, 2024, from <https://www.weforum.org/stories/2024/02/with-generative-ai-we-can-reimagine-education-and-the-sky-is-the-limit/>

Khan, S. (2024). *Brave new words: How AI will revolutionize education (and why that's a good thing)*. New York, NY: Penguin Random House.

OECD (2024). *The potential impact of artificial intelligence on equity and inclusion in education*. Washington, DC: OECD. Retrieved October 24, 2024, from https://www.oecd.org/en/publications/the-potential-impact-of-artificial-intelligence-on-equity-and-inclusion-in-education_15df715b-en.html
- 2 Logan, C. (2024). Learning about and against generative AI through mapping generative AI's ecologies and developing a Luddite praxis. In R. Lindgren, T.I. Asino, E.A. Kyza, C.K. Looi, D.T. Keifert, & E. Suárez. (Eds.), *Proceedings of the 18th International Conference of the Learning Sciences* (pp. 362-369). International Society of the Learning Sciences. Retrieved October 24, 2024, from <https://repository.isls.org/handle/1/11112>

Selwyn, N. (2024). On the limits of artificial intelligence (AI) in education. *Nordisk Tidsskrift for Pedagogikk Og Kritikk*, 10(1), 3-14. Retrieved October 24, 2024, from <https://doi.org/10.23865/ntpk.v10.6062>

Williamson, B., Molnar, A., & Boninger, F. (2024, March). *Time for a pause: Without effective public oversight, AI in schools will do more harm than good*. Boulder, CO: National Education Policy Center. Retrieved October 24, 2024, from <https://nepc.colorado.edu/publication/ai>
- 3 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education*. Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 4 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: Opportunities and challenges of artificial intelligence-enhanced education*. Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_2_Bellwether_September2024.pdf
- 5 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: Artificial intelligence use cases*. Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_3_Bellwether_September2024.pdf
- 6 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 32). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 7 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (pp. 5-8). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 8 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 12). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 9 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning*

systems: The landscape of artificial intelligence in K-12 education (pp. 19-20). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf

- 10 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 32). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 11 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 32). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 12 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 12). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 13 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 12). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 14 Dhaliwal, R.S., Lepage-Richer, T., & Suchman, L. (2024). *Neural networks*. Minneapolis, MN: University of Minnesota Press.

Narayanan, A. & Kapoor, S. (2024). *AI snakeoil: What artificial intelligence can do, what it can't, and how to tell the difference*. Princeton, NJ: Princeton University Press.

Mitchell, M. (2019). *Artificial intelligence: A guide for thinking humans*. New York, NY: Farrar, Straus, and Giroux.

Paquinelli, M. (2023). *The eye of the master: A social history of artificial intelligence*. New York, NY: Verso.

Stone, M., Goodlad, L.M.E., & Sammons, M. (2024). The origins of generative AI in transcription and machine translation, and why that matters. *Critical AI*, 2(1). Retrieved October 25, 2024, from <https://doi.org/10.1215/2834703X-11256853>

- 15 Baker, R. & Hawn, A. (2022). Algorithmic bias in education. *International Journal of Artificial Intelligence in Education*, 32, 1052-1092. Retrieved October 25, 2024, from <https://link.springer.com/article/10.1007/s40593-021-00285-9>

Dixon-Román, E., Nichols, T.P., & Nyame-Mensah, A. (2020). The racializing forces of/in AI educational technologies. *Learning, Media, and Technology*, 45(3), 236-250. Retrieved October 25, 2024, from <https://www.tandfonline.com/doi/full/10.1080/17439884.2020.1667825>

Denton, E., Hanna, A., Amironesi, R., Smart, A., & Nicole, H. (2021). On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society*, 8(2). Retrieved October 25, 2024, from <https://journals.sagepub.com/doi/abs/10.1177/205395172111035955>

Mirzoeff, N. (2021). Artificial vision, white space and racial surveillance capitalism. *AI & Society*, 36, 1295-1305. Retrieved October 25, 2024, from <https://link.springer.com/article/10.1007/s00146-020-01095-8/>

Nichols, T.P. & Garcia, A. (2022). Platform studies in education. *Harvard Educational Review*, 92(2), 209-230. Retrieved October 25, 2024, from <https://doi.org/10.17763/1943-5045-92.2.209>

- 16 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 October 25, 2024, from <https://nepc.colorado.edu/publication/personalized-learning>
- Pangrazio, L., Stornaiuolo, A., Nichols, T.P., Garcia, A., & Philip, T. (2022). Datafication meets platformization: Materializing data processes in teaching and learning. *Harvard Educational Review*, 92(2), 257-283. Retrieved October 25, 2024, from <https://doi.org/10.17763/1943-5045-92.2.257>
- Perrotta, C., Gulson, K., 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 October 25, 2024, from <https://www.tandfonline.com/doi/full/10.1080/17508487.2020.1855597>
- Sefton-Green, J. (2022). Towards platform pedagogies: Why thinking about digital platforms as pedagogic devices might be useful. *Discourse: Studies in the Cultural Politics of Education*, 43(6), 899-911. Retrieved October 25, 2024, from <https://www.tandfonline.com/doi/full/10.1080/01596306.2021.1919999>
- 17 Ferrari, F., & van Dijck, J. (2023). Observe, inspect, modify: Three conditions for generative AI governance. *New Media & Society*. OnlineFirst. Retrieved October 25, 2024, from <http://journals.sagepub.com/doi/10.1177/14614448231214811>
- Gulson, K.N., & Sellar, S., (2024). Anticipating disruption: Artificial intelligence and minor experiments in education policy. *Journal of Education Policy*, 39(5), 702-717. Retrieved October 25, 2024, from <https://www.tandfonline.com/doi/full/10.1080/02680939.2024.2302474>
- Nichols, T.P., & Dixon-Román, E. (2024). Platform governance and education policy: Power and politics in emerging edtech ecologies. *Educational Evaluation and Policy Analysis*, 46(2), 309-328. Retrieved October 25, 2024, from <https://doi.org/10.3102/01623737231202469>
- Selwyn, N. (2021). The human labour of school data: Exploring the production of digital data in schools. *Oxford Review of Education*, 47(3), 353-368. Retrieved October 25, 2024, from <https://www.tandfonline.com/doi/full/10.1080/03054985.2020.1835628>
- 18 Benjamin, R. (2019). *Race after technology: Abolitionist tools for the New Jim Code*. Malden, MA: Polity.
- Broussard, M. (2023). *More than a glitch: Confronting race, gender, and ability bias in tech*. Cambridge, MA: The MIT Press.
- Chun, W.H.K., & Barnett, A. (2021). *Discriminating data: Correlation, neighborhoods, and the new politics of recognition*. Cambridge, MA: The MIT Press.
- Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. New York, NY: St. Martin's Press.
- Noble, S. (2018). *Algorithms of oppression: How search engines reinforce racism*. New York, NY: New York University Press.
- 19 The importance of equitable access to AI is emphasized throughout all three Bellwether reports. See, for instance:
- Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 10). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: Opportunities and challenges of artificial intelligence-enhanced education* (p. 47). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_2_

Bellwether_September2024.pdf

Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: Artificial intelligence use cases* (p. 22). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_3_Bellwether_September2024.pdf

A growing research literature has explicitly challenged the conflation of “access” and “equity” as it pertains to educational technologies. See, for instance:

Crooks, R. (2024). *Access is capture: How edtech reproduces racial inequality*. Berkeley, CA: University of California Press.

Greene, D., (2021). *The promise of access: Technology, inequality, and the political economy of hope*. Cambridge, MA: The MIT Press.

Nichols, T.P. & Dixon-Román, E. (2024). Platform governance and education policy: Power and politics in emerging edtech ecologies. *Educational Evaluation and Policy Analysis*, 46(2), 309-328. Retrieved October 25, 2024, from <https://doi.org/10.3102/01623737231202469>

It is worth noting that, in one instance, the report does suggest that equitable AI policies must go beyond promoting access to ensure that AI tools are “built on fair models that use representative datasets and actively mitigate algorithmic bias.” See:

Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: Opportunities and challenges of artificial intelligence-enhanced education* (p. 46). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_2_Bellwether_September2024.pdf

However, the specifics of what such a recommendation might involve are hazy, and the consensus of many researchers is that “bias” is something endemic to AI systems and not something that can be mitigated. See, for instance:

Chun, W.H.K. & Barnett, A. (2021). *Discriminating data: Correlation, neighborhoods, and the new politics of recognition*. Cambridge, MA: The MIT Press.

Hong, S. (2024). AI and bias. In R. Paul, E. Carmel, & J. Cobbe (Eds.), *Handbook on public policy and artificial intelligence* (pp. 109-122). Northampton, MA: Edward Elgar. Retrieved November 8, 2024, from <https://doi.org/10.4337/9781803922171.00015>

Thylstrup, N.B., Hansen, K.B., Flyverbom, M., & Amooore, L. (2022). Politics of data reuse in machine learning systems: Theorizing reuse entanglements. *Big Data & Society*, 9(2). Retrieved November 8, 2024, from <https://journals.sagepub.com/doi/10.1177/20539517221139785>

20 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 13). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf

21 For a comprehensive history of ‘artificial intelligence’ as a philosophical and technological project, see Boden, M.A., (2006). *Mind as machine: A history of cognitive science*. Oxford, UK: Clarendon Press. Boden has also authored a compressed account of this history: Boden, M.A. (2018). *Artificial intelligence: A very short introduction*. New York, NY: Oxford University Press.

22 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: Opportunities and challenges of artificial intelligence-enhanced education* (p. 68). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_2_Bellwether_September2024.pdf

- 23 For a historical overview of ‘speculative bubbles,’ where investment in a market exceeds the underlying value of its product, and the place of technology within them, see: Quinn, W., & Turner, J.D. (2020). *Boom and bust: A global history of financial bubbles*. Cambridge, UK: Cambridge University Press.

Ben Williamson and Janja Komljenovic have used the term “futuring” to describe the process by which edtech investors cultivate markets for products and services by appealing to the anticipated futures they could bring about, even if there is little evidence they will do so. See:

Williamson, B. & Komljenovic, J. (2023). Investing in imagined digital futures: The techno-financial ‘futuring’ of edtech investors in higher education. *Critical Studies in Education*, 64(3), 234-249. Retrieved October 25, 2024, from <https://www.tandfonline.com/doi/full/10.1080/17508487.2022.2081587>

- 24 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 3). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 25 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 39). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 26 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (p. 18). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- 27 Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: Opportunities and challenges of artificial intelligence-enhanced education* (pp. 18-20). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_2_Bellwether_September2024.pdf
- 28 The report includes a one-sentence sidebar on the environmental impacts of AI, which largely frames the issue in the past-tense, as a byproduct of early efforts to train large language models which technologies companies have responded to through the development of ‘small language models.’ See:
- Kulesa, A.C., Croft, M., Robinson, B., Wells, M.K., Rotherham, A.J., & Bailey, J. (2024, September). *Learning systems: The landscape of artificial intelligence in K-12 education* (pp. 16-17). Bellwether. Retrieved October 20, 2024, from https://bellwether.org/wp-content/uploads/2024/09/LearningSystems_1_Bellwether_September2024.pdf
- However, there is substantial evidence that the climate impacts of AI development are continuing to escalate. See, for example:
- Bender, E.M., Gebru, T., McMillan-Major, A., & Schmitz, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610-623. Retrieved October 28, 2024, from <https://dl.acm.org/doi/pdf/10.1145/3442188.3445922>
- Crawford, K. (2021). *Atlas of AI: Power, politics, and planetary crisis*. New Haven, CT: Yale University Press.
- Luccioni, S., Jernite, Y., & Strubell, E. (2024). Power hungry processing: Watts driving the cost of AI deployment? *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 85-99. Association for Computing Machinery. Retrieved October 28, 2024, from <https://dl.acm.org/doi/pdf/10.1145/3630106.3658542>

Valdivia, A. (2024). The supply chain capitalism of AI: A call to (re)think algorithmic harms and resistance through environmental lens. *Information, Communication, & Society*. OnlineFirst. Retrieved October 30, 2024, from <https://doi.org/10.1080/1369118X.2024.2420021>

It is also worth noting that the lack of sustainability in the edtech sector was a concern even before GenAI became a pressing topic among education stakeholders. See, for instance:

Selwyn, N. (2021). Ed-tech within limits: Anticipating educational technology in times of environmental crisis. *E-Learning and Digital Media*, 18(5), 496-510. Retrieved October 28, 2024, from <https://journals.sagepub.com/doi/10.1177/20427530211022951>

- 29 Monahan, T. & Torres, R.D. (2010). *Schools under surveillance: Cultures of control in public education*. New Brunswick, NJ: Rutgers University Press.

Nemorin, S. (2017). Post-panoptic pedagogies: The changing nature of school surveillance in the digital age. *Surveillance & Society*, 15(2), 239-253. Retrieved October 28, 2024, from <https://ojs.library.queensu.ca/index.php/surveillance-and-society/article/view/pedagogies>

Pangrazio, L. & Sefton-Green, J. (Eds.) (2022). *Learning to live with datafication: Educational case studies and initiatives from across the world*. New York, NY: Routledge.

The relationships among data, surveillance, and the reproduction of social inequality, especially along lines of race and ethnicity, have been well-documented in the research literatures of multiple fields, including education. See:

Brayne, S. (2017). Big data surveillance: The case of policing. *American Sociological Review*, 82(5), 977-1008. Retrieved November 5, 2024, from <https://doi.org/10.1177/0003122417725865>

Browne, S. (2015). *Dark matters: On the surveillance of Blackness*. Durham, NC: Duke University Press.

Crooks, R. (2019). Cat-and-mouse games: Dataveillance and performativity in urban schools. *Surveillance & Society*, 17(3/4), 484-489. Retrieved November 5, 2024, from <https://doi.org/10.24908/ss.v17i3/4.7098>

McMillan Cottom, T. (2020). Where platform capitalism and racial capitalism meet: The sociology of race and racism in the digital society. *Sociology of Race and Ethnicity*, 6(4), 441-449. Retrieved November 5, 2024, from <https://doi.org/10.1177/2332649220949473>

- 30 Bender, E.M., Gebru, T., McMillan-Major, A., & Schmittchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610-623. Retrieved October 28, 2024, from <https://dl.acm.org/doi/pdf/10.1145/3442188.3445922>

Selwyn, N. (2024). On the limits of artificial intelligence (AI) in education. *Nordisk Tidsskrift for Pedagogikk Og Kritik*, 10(1). Retrieved October 30, 2024, from <https://doi.org/10.23865/ntpk.v10.6062/>

Thompson, N.C., Greenewald, K., Lee, K., & Manso, G.F. (2021). Deep learning's diminishing returns. *IEEE Spectrum*. Retrieved October 30, 2024, from <https://spectrum.ieee.org/deep-learning-computational-cost>

- 31 Research in multiple disciplines has demonstrated the escalating importance of cybersecurity as more and more facets of education are underwritten by networked data technologies. For a general overview of the technical, behavioral, and legal dimensions of these cybersecurity challenges, see:

Wolff, J. (2018). *You'll see this message when it's too late: The legal and economic aftermath of cybersecurity breaches*. Cambridge, MA: The MIT Press.

Studies also show the limitations of existing strategies for protecting student data (e.g., de-identification, FERPA), suggesting the need for ongoing investigation into alternatives tailored to the evolving nature of data gathering and usage in schools. See, for instance:

Yacobson, E., Fuhrman, O., Hershkowitz, S., & Alexandron, G. (2021). De-identification is insufficient to protect student privacy, or – what can a field trip reveal? *Journal of Learning Analytics*, 8(2), 83-92.

Zeide, E. (2016). Student privacy principles for the age of big data: Moving beyond FERPA and FIPPS. *Drexel Law Review*, 8(2), 339–394. Retrieved November 8, 2024, from <https://heinonline.org/HOL/P?h=hein.journals/drexel8&i=351>

Conrad, K. (2024). A blueprint for an AI bill of rights for education. *Critical AI*, 2(1). Retrieved November 8, 2024, from <https://doi.org/10.1215/2834703X-11205245>

For examples of studies that interrogate the

Pleasants, J., Krutka, D., & Nichols, T.P. (2023) What relationship do we want with technology? Toward techskepticism in schools. *Harvard Educational Review*, 93(4), 486-515. Retrieved November 8, 2024, from <https://doi.org/10.17763/1943-5045-93.4.486>

Selwyn, N. (2024). Digital degrowth: Toward radically sustainable education technology. *Learning, Media, & Technology*, 49(2), 186-199. Retrieved November 8, 2024, from <https://doi.org/10.1080/17439884.2022.2159978>

For examples of studies that explore the need for, and the challenges to developing, diverse participation in computer science, see:

Gebru, T. (2020). Race and gender. In M.D. Dubber, F. Pasquale, & S. Das (Eds.), *The Oxford handbook of ethics of AI* (pp. 251-269). New York, NY: Oxford University Press. Retrieved November 8, 2024, from <https://doi.org/10.1093/oxfordhb/9780190067397.013.16>

McGee, E.O. (2020). Interrogating structural racism in STEM higher education. *Educational Researcher*, 49(9), 633-644. Retrieved November 8, 2024, from <https://doi.org/10.3102/0013189X20972718>

Vakil, S. (2018). Ethics, identity, and political vision: Toward a justice-centered approach to equity in computer science education. *Harvard Educational Review*, 88(1), 26-51. Retrieved November 8, 2024, from <https://doi.org/10.17763/1943-5045-88.1.26>