Beyond Bans: Thoughtful Use of AI in the Classroom

Kendall Brogle[1], **Umang Bhatt**[1,2]

[1]The Alan Turing Institute,
[2]New York University
kbrogle@turing.ac.uk, umangbhatt@nyu.edu

Abstract

The rapid rise in use of generative AI (GenAI) systems has sparked debate in educational settings, resulting in reactionary bans to mitigate concerns about academic dishonesty and learning outcomes. Such bans fail to address the underlying challenges AI presents while overlooking its potential benefits. This paper examines the limitations of outright bans and proposes that the successful integration of AI in education requires alignment of tasks, systems and metrics. By analyzing existing studies and real world examples, we demonstrate how thoughtfully designed AI systems can enhance learning and reduce inequalities. The paper concludes with actionable recommendations for educators and policymakers to navigate AI integration responsibly, emphasizing equity, meaningful engagement, and the need for teacher support.

1 Introduction

Generative AI (GenAI) has become widely used in classrooms and its accessibility has raised significant concerns among educators and policymakers, particularly about its potential unknown effects on student learning and academic integrity (Luckin and Holmes 2016). These worries have led to reactionary measures, such as outright bans, that often fail to address the underlying issue (Holmes, Bialik, and Fadel 2019; Luscombe 2023). However, GenAI also presents opportunities to enhance education (Gocen and Aydemir 2020; Grassini 2023). We argue that a thoughtful integration of AI into schools must be rooted in aligning tasks, systems, and metrics to meaningful learning outcomes.

2 Classroom Bans

Concern about potential harms led many schools to push for bans of GenAI tools (Luscombe 2023). While the bans' motivation was well founded, the outcome is deemed reactionary (Harrison, Hurd, and Brinegar 2023). New York City Department of Education, Los Angeles Unified School District, Seattle Public schools, and Baltimore County Public Schools in Maryland have all implemented some form of bans. They all stated concerns over the potential to hinder the development of student's problem solving skills, learning integrity, and academic dishonesty (Darnell 2023; Elgan 2023).

Banning perceived threats to education is not new. For example, calculators, once thought to harm student's math skills, have now been shown to help problem solving and conceptual understanding (Cowdery 1997; Banks 2011; Ellington 2003). Schools blocked access to the internet over distraction and unwanted content concerns and eventually taught digital literacy and responsible use (Livingstone and Helsper 2007). Bans on mobile phones exacerbated equity issues with wealthier students accessing technology outside of school. Over time, schools have incorporated the technology into learning. (Thomas, O'Bannon, and Bolton 2013; Selwyn and Aagaard 2021)

While the push to reduce harms is understandable, this approach is inherently flawed. Schools report that students merely circumvent restrictions, which undermines authority and creates issues of trust and raises serious concerns about equity and inclusion. (Center for Democracy and Technology 2024). Disciplinary measures in response to AI use disproportionately impact marginalized students, who already face barriers within the education system (Center for Democracy and Technology 2024). Rather than prohibiting AI, schools can focus on developing policies and practices that position AI to be successful while still acknowledging the risks (Gocen and Aydemir 2020). Successful AI integration must include:

- Educator support to help teachers understand AI and its applications within the classroom (Selwyn 2019).
- Curriculum adaptation, focusing on in-class work that prioritizes critical thinking skills that cannot be easily replicated by AI. (Lampou 2023)
- Policy development that has detailed and clear guidelines for ethical AI use, emphasizing transparency and accountability (Holmes et al. 2022).
- Most importantly, contextual alignment of the task, system and metrics.

3 Context Matters

It is important to fully consider the context of AI use (Gabriel 2020; Selwyn 2021). Success and ethical implications of AI in education are highly contingent on the contextual alignment of the implementation details, the surrounding philosophy, and the intention of use . (Kurni, Mohammed, and Srinivasa 2023).

3.1 Task, System, Metric Alignment

There is considerable variety in the purpose of different educational activities and in the function of AI tools. Effective use of AI in education depends on aligning three critical components: the educational task goal, the selected system's capabilities, and the metrics to evaluate success.

Tasks in education are diverse, ranging from problem solving to fostering creativity. It is imperative that the tool matches the specific needs of the task to avoid undermining the purpose.

The system encompasses design, capabilities, and adaptability of the AI tool itself. A well designed system should not only execute tasks efficiently but consider the broader context of its use. Assignments that require creativity or skill development and critical thinking need tools that support this process. AI systems must be designed to complement educational goals. For instance:

- Adaptive learning systems, which use algorithms to tailor content to individual needs, show promise in improving math and reading skills when used in structured environments (Pane et al. 2014).
- Using generative AI for tasks such as writing essays without guidance may hinder critical thinking and original expression (Selwyn 2019).
- Memorization tasks, such as vocabulary, may benefit from AI systems that prompt repetition and recall. Systems that allow for passive engagement are less effective for long term memory acquisition (Yusuf 2010; Ironside 2005).

Metrics are critical to measuring effectiveness and shape the design and application of AI tools. Education metrics often emphasize quantifiable outcomes like test scores and fail to capture nuances such as critical thinking, creativity, and student engagement (Hanna, David, and Francisco 2010). Metric priorities may vary between groups such as students and teachers.

3.2 Insights From Existing Work

Recent work has begun to analyze human subject studies of AI deployment in educational settings. These studies highlight how effectiveness directly depends on alignment.

The study *Generative AI Can Harm Learning*, conducted by researchers from the Wharton School at the University of Pennsylvania, examined the impact of GenAI on students learning high school level mathematics. After a controlled experiment, they concluded that "AI may be harmful to learning" (Bastani et al. 2024). In this study, students were divided into three groups: one with no AI access, one using GPT Base (a general system), and one using GPT Tutor (a system with pedagogical safeguards). Results revealed a dichotomy: while GPT Base improved practice problem performance by 48%, it also resulted in a 17% reduction in unassisted exam scores compared to the control group. In contrast, the GPT Tutor group mitigated these effects, resulting in similar performance to the control group on unassisted exams, but also no real improvement . These findings show the negative affects of using AI systems for tasks that do not have aligned learning goals and emphasize that the design of AI systems must prioritize processes rather than simply making results more efficient (Toyama 2011).

In contrast to the Wharton study, *Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise*, conducted by researchers from Stanford University, demonstrated positive results from AI tool use . This work showed how a carefully aligned AI system can positively impact learning outcomes. Tutor CoPilot provided real-time, expert-like suggestions to K-12 tutors during math sessions. The system improved topic mastery rates by 4 percentage points overall and up by 9 points for students of lower-rated tutors, which suggests its ability to bridge gaps in instructional quality. Unlike general purpose models, Tutor CoPilot was explicitly designed to support tutoring by prompting high-quality pedagogical strategies (Wang et al. 2024). Tutors remained in control of instructional decisions, selecting or modifying AI generated suggestions based on real time student interactions, showing the potential collaboration between AI and human educators.

3.3 Recommendations For Success

The lens with which AI is deployed plays a pivotal role in determining its impact. Educational psychology which emphasizes active, student centered learning, can leverage AI to encourage exploration and engagement(Woolfolk 2016). Without this focus, integration of AI may focus on efficiency and automation, which may reduce meaningful development.

Educators play an important role in this integration but without training and resources, this is a futile effort (Selwyn 2019). It it vital that policymakers support professional development and provide schools with the necessary tools to address AI literacy equitably. The societal conditions in which AI is introduced will also impact its outcomes. Inequalities in access to technology and resources can greatly exacerbate existing divides and create disparities in the benefits of AI (Kurni, Mohammed, and Srinivasa 2023).

While including educators perspectives is vital, teachers also need to adapt some of their traditional practices. For example, shifting to focus on in-class, collaborative exercises ensures students engage actively with the material under teacher guidance (Cooper, Robinson, and Patall 2006).

4 Conclusion

Reactionary bans to GenAI tools highlight a recurring pattern of resistance to new technologies in education. While AI bans are motivated by legitimate concerns, they fail to address nuanced challenges and acknowledge opportunities for innovation. Effective AI integration relies on aligning the task, system, and metric. Misalignment can lead to detrimental outcomes, such as over-reliance on automation, or exacerbating inequalities. Conversely, aligned systems, such as those designed to support teacher guidance, can enhance learning outcomes and promote equity. Schools must adopt a thoughtful, learning first approach that prioritizes collaborative task definition, adaptable and equitable AI system design, and metrics that value long term learning over superficial performance.

References

Banks, S. 2011. A historical analysis of attitudes toward the use of calculators in junior high and high school math classrooms in the United States since 1975. *Cedarville University*.

Bastani, H.; Bastani, O.; Sungu, A.; Ge, H.; Kabakcı, O.; and Mariman, R. 2024. Generative ai can harm learning. *Available at SSRN*, 4895486.

Center for Democracy and Technology. 2024. Up in the Air: Educators Juggling the Potential of Generative AI with Detection, Discipline, and Distrust.

Cooper, H.; Robinson, J. C.; and Patall, E. A. 2006. Does homework improve academic achievement? A synthesis of research, 1987–2003. *Review of educational research*, 76(1): 1–62.

Cowdery, T. W. 1997. The great calculator debate: Concerns about calculator use in elementary schools. *Retrieved June*, 21: 2001.

Darnell, T. 2023. NYC Department of Education blocks ChatGPT over fears of harm to student learning. *ABC News*. Accessed: November 24, 2024.

Elgan, M. 2023. Schools look to ban ChatGPT; students use it anyway. *Computerworld*. Accessed: November 24, 2024.

Ellington, A. J. 2003. A Meta-Analysis of the Effects of Calculators on Students' Achievement and Attitude Levels in Precollege Mathematics Classes. *Journal for Research in Mathematics Education*, 34(5): 433–463.

Gabriel, I. 2020. Artificial intelligence, values, and alignment. *Minds and machines*, 30(3): 411–437.

Gocen, A.; and Aydemir, F. 2020. Artificial intelligence in education and schools. *Research on Education and Media*, 12(1): 13–21.

Grassini, S. 2023. Shaping the future of education: exploring the potential and consequences of AI and ChatGPT in educational settings. *Education Sciences*, 13(7): 692.

Hanna, D.; David, I.; and Francisco, B. 2010. *Educational research and innovation the nature of learning using research to inspire practice: Using research to inspire practice*. OECD publishing.

Harrison, L. M.; Hurd, E.; and Brinegar, K. M. 2023. Critical race theory, books, and ChatGPT: Moving from a ban culture in education to a culture of restoration.

Holmes, W.; Bialik, M.; and Fadel, C. 2019. *Artificial intelligence in education promises and implications for teaching and learning*. Center for Curriculum Redesign.

Holmes, W.; Porayska-Pomsta, K.; Holstein, K.; Sutherland, E.; Baker, T.; Shum, S. B.; Santos, O. C.; Rodrigo, M. T.; Cukurova, M.; Bittencourt, I. I.; et al. 2022. Ethics of AI in education: Towards a community-wide framework. *International Journal of Artificial Intelligence in Education*, 1–23.

Ironside, P. M. 2005. Teaching thinking and reaching the limits of memorization: Enacting new pedagogies. *Journal of Nursing Education*, 44(10): 441–449.

Kurni, M.; Mohammed, M. S.; and Srinivasa, K. 2023. Ethics of Artificial Intelligence in Education. In *A Beginner's Guide to Introduce Artificial Intelligence in Teaching and Learning*, 213–229. Springer.

Lampou, R. 2023. The integration of artificial intelligence in education: Opportunities and challenges. *Review of Artificial Intelligence in Education*, 4: e15–e15.

Livingstone, S.; and Helsper, E. 2007. Gradations in digital inclusion: Children, young people and the digital divide. *New media & society*, 9(4): 671–696.

Luckin, R.; and Holmes, W. 2016. Intelligence unleashed: An argument for AI in education.

Luscombe, B. 2023. Schools Shouldn't Ban Access to ChatGPT. Accessed: 2024-11-24.

Pane, J. F.; Griffin, B. A.; McCaffrey, D. F.; and Karam, R. 2014. Effectiveness of cognitive tutor algebra I at scale. *Educational Evaluation and Policy Analysis*, 36(2): 127–144.

Selwyn, N. 2019. *Should robots replace teachers?: AI and the future of education*. John Wiley & Sons.

Selwyn, N. 2021. *Education and technology: Key issues and debates*. Bloomsbury Publishing.

Selwyn, N.; and Aagaard, J. 2021. Banning mobile phones from classrooms—An opportunity to advance understandings of technology addiction, distraction and cyberbullying. *British journal of educational technology*, 52(1): 8–19.

Thomas, K. M.; O'Bannon, B. W.; and Bolton, N. 2013. Cell phones in the classroom: Teachers' perspectives of inclusion, benefits, and barriers. *Computers in the Schools*, 30(4): 295–308.

Toyama, K. 2011. There are no technology shortcuts to good education. *Educational Technology Debate*, 8.

Wang, R. E.; Ribeiro, A. T.; Robinson, C. D.; Loeb, S.; and Demszky, D. 2024. Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise. arXiv:2410.03017.

Woolfolk, A. 2016. *Educational psychology*. Pearson.

Yusuf, M. 2010. Memorization as a learning style: A balance approach to academic excellence. *OIDA International Journal of Sustainable Development*, 1(6): 49–58.