Refusing Educational Technology:

Artificial Intelligence, Inequity, and the Problem of Critical Optimism

Niral Shah
University of Washington, Seattle

David Stroupe
University of Utah

Aman Yadav
Michigan State University

Tess Bernhard
University of Pennsylvania

Charles Logan
Northwestern University

Elizabeth B. Dyer
University of Illinois, Urbana-Champaign

Peng He
Washington State University

Christina (Stina) Krist
Stanford University

Tingting Li
Washington State University

Correspondence concerning this article should be addressed to Niral Shah, College of Education, University of Washington, Miller Hall, 2012 Skagit Lane, Seattle, WA 98105. E-mail: niral@uw.edu

Abstract

Artificial intelligence (AI)—like many technologies before—comes with hope and worry about its implications for education. In this essay, using generative AI as a case, we appraise the ideological field of positions that stakeholders take in relation to educational technology overall. Drawing on science and technology studies, our analysis shows that education's default position is optimism, even among those genuinely "critical" of technology's capacity to amplify inequity. Refusal is not widely regarded as a legitimate response to educational technology. We argue, in contrast, that given the potential for scalable harm from unproven technologies like generative AI, the most rational position with new technology is robust caution with a readiness to refuse. We conclude by detailing three forms of refusal for educational stakeholders to consider: curricular, bureaucratic, and administrative.

Keywords: Artificial intelligence, educational technology, equity, science and technology studies

Introduction

Christening technology as the solution to educational problems is not new. Corporations and entrepreneurs have sought to sell educational technology to school systems since the early 20th century. In response, educational leaders have been forced to grapple with industry's push to use technology or risk becoming obsolete. Given the rapid pace of technological innovation and the pressure to use technology, high-stakes decisions are often made in the absence of rigorous empirical research. From silent films to personal computers, stakeholders throughout education have wrestled with this century-long conundrum of educational technology (Cuban, 1986; Watters, 2021). Now generative artificial intelligence enters this historical arc.

As an umbrella category, artificial intelligence (AI[1]) technologies have been a part of education for some time. For example, machine learning techniques have driven adaptive tutoring platforms and automated assessment systems for decades (Reich, 2020; X. Zhai & Krajcik, 2025). But as AI's most recent iteration, generative AI (i.e., based on large language models) has sparked hope about new use cases in education (see Beckingham et al., 2024; Price & Grover, 2025). Still, many worry about generative AI's costs, such as environmental degradation, threats to privacy, and the potential to amplify educational inequities that have long fallen on minoritized communities (see Avraamidou, 2024; Nichols et al., 2025; Penuel et al., 2025). Might generative AI fundamentally redefine the relationship between students and teachers in classrooms? What will this technology mean for the teacher labor force?

[1] A full treatment of AI technologies in their many incarnations is beyond the scope of this article. Historically, the general term "AI" has been used to refer to machine-based technologies that can accomplish complex goals (Tegmark, 2017). Here we focus specifically on "generative AI," which has been defined as "...a group of AI algorithms and models that are capable of producing new content, including texts, images, videos and problem-solving strategies, with human-like creativity and adaptability" (He et al., 2025, p. 1).

We do not claim to know the answers to these questions. However, under conditions of such uncertainty, we argue that our field needs clear ways of thinking through and making decisions about generative AI in education. Our stance is that generative AI is only the latest instance of "computing creep" in education; there *will* be other technologies imposed on educational settings (Shah & Yadav, 2023). Therefore, educational researchers have a special responsibility. Through empirical research, school system partnerships, public writing, and governmental service, we play key roles in shaping public discourse about technology in education. As researchers, we must be anchored in empirical evidence, ensuring that our recommendations emerge from systematic inquiry rather than hope and hype.

In this essay, we appraise the ideological field of positions (Althusser, 1971) that various stakeholders have taken towards educational technology. Ideological analysis is useful in exposing the values and power relations that mobilize public debate, as well as explaining their material effects on questions of political economy (Eagleton, 1991; Leonardo, 2010). Drawing on science and technology studies—and treating generative AI as a case—we chart nuances that cut across the ideological field of educational technology. Our analysis focuses on a particular position we call "critical optimism," which spotlights educational technology's potential for harm while also insisting that technology can be repurposed to *benefit* minoritized communities in education. In contrast, we argue that critical[2] optimism is untenable, even dangerous, because the position never seriously considers *refusal* of educational technology to be a reasonable practice. By failing to do so, our concern is that critical optimism inadvertently permits the unchecked proliferation of technology like generative AI throughout education.

---

[2] The term "critical" has been used to signal the particular framework of critical social theory (see Leonardo, 2004). Since the phrase "critical optimism" intends to describe a more general trend in the literature, our use of "critical" in this article instead mirrors the colloquial way certain stakeholders writing about educational technology have used the term (i.e., as a way of signaling concern for inequity or power relations).

We enter this problem space as educational researchers who have long histories with technology. Collectively, we hold degrees in computer science, engineering, and various STEM fields. We have used educational technology in teacher education and as former P-12 teachers, and we are researching the impact of generative AI on teaching and learning. We reject views of educational technology as entirely harmful. At the same time, we are invested in seeing the refusal of educational technology normalized as a legitimate practice. We conclude this essay by detailing multiple forms of refusal available to educational stakeholders.

Educational Technology in Context: History and Political Economy

The history of technology in education is long, disappointing, and predictable. Over the last century, from film and radio to computers and massively open online courses (MOOCs), historical analyses show that educational technologies tend to follow a common pattern: irrational hype followed by limited impact (Cuban, 1986; Reich, 2020). Initial exuberance is typically led by "technological messiahs" (Weizenbaum, 1976, p. 245). In *Teachers and Machines,* Larry Cuban (1986) recounts both Thomas Edison's claim in 1913 that silent films in schools would soon replace books, and Seymour Papert's claim in 1984 that, "There won't be schools in the future...I think the computer will blow up the school" (cited on p. 72).

Educational technologies have also been positioned as remedies for inequity. For example, Watters (2021) documents Harvard psychologist B. F. Skinner's attempt to deploy his infamous "teaching machines" in Harlem schools, with a goal of increasing literacy rates among Black children. Even civil rights leader Bob Moses used such machines to implement an adult literacy curriculum during the 1964 Freedom Summer in Mississippi, although the initiative was abandoned after teaching machines were deemed antithetical to the Freirean ethos of the Freedom Schools (Watters, 2021).

Despite the promises and desires of technologists, classrooms and educational systems have remained relentlessly static (Cuban, 1986; Reich, 2020). In a study of computer use at two high schools in Silicon Valley—a region at the forefront of technological innovation—Cuban and colleagues (2001) found no significant shift in teacher practices. Additionally, in reviewing research on the relationship between inequity and self-paced online learning systems, Reich (2020) found that such systems are most likely to benefit wealthy, already-formally educated people. For example, Reich reported that a 2012 effort using MOOCs to support undergraduates struggling in mathematics actually resulted in *lower* pass rates than a control group (p. 35). Finally, on an international level, efforts to expand "One Laptop per Child" programs designed to enhance learning have repeatedly failed to live up to their utopian visions (Ames, 2019; Cueto et al., 2025).

Beyond a poor record of attenuating educational inequity, there is evidence that educational technologies can actively produce new kinds of harm, particularly for minoritized students. As schools have expanded access to digital devices, they have also invested in activity monitoring that continuously reports what students do online while using those devices (Laird et al., 2022). Scholars have documented how such surveillance can facilitate various undesirable outcomes, such as: infringing upon students' privacy (Citron, 2023; Fedders, 2019); disempowering students (Pangrazio et al., 2023); collecting data that justifies racially segregated schools (Crooks, 2021, 2024); suspending Black students at higher rates (Johnson & Jabbari, 2022); allowing unauthorized disclosure of students' sexual orientation and gender identity (ACLU, 2023); and increasing students' interactions with law enforcement (Laird et al., 2022).

Some might argue that harms like these are acceptable side effects in the service of advancing student learning. However, in the case of AI, research to date does not offer consistent evidence for the positive impact of generative AI on student learning (Fesler et al., 2026). To the contrary, in situations where generative AI tutors have been implemented, multiple reviews and empirical studies caution that students who use generative AI can experience negative learning outcomes and over-reliance on the technology (Bastani et al., 2024; Wang & Fan, 2025; Wecks et al., 2024; C. Zhai et al., 2024).

History shows that teachers are frequently seen as being resistant to educational technology and usually get blamed when technology fails (Hannafin & Savenye, 1993; Howard, 2013). However, it has been too often the case that technologies have not been designed to engage with classroom realities (Cuban, 1986; Reich, 2020). The broader issue is a failure to recognize schools as complex social institutions embedded within broader cultural, political, and economic systems, as well as a tendency to view teachers and students as passive targets in need of change. This kind of simplistic framing points to deeper questions about how educational researchers conceptualize modern technologies, as well as the political economy that animates their creep into education (Shah & Yadav, 2023).

Consider the difference between a calculator and a generative AI chatbot. Although both are tools sold by for-profit entities that could support student learning, a calculator is a one-time purchase, while the chatbot is likely a recurring expense based on a subscription model. This business arrangement "cedes significant control over what happens in classrooms to the whims of platform developers" (Garcia & Nichols, 2021, p.15). Also, unlike the calculator, an AI-based platform can also become a vehicle for extracting user data, which constitutes a source of revenue at the expense of privacy (Sadowski, 2019). Rather than a self-contained tool, technology like generative AI might be better conceptualized as a dynamic actor that exists within a broader ecology, alongside administrators, teachers, parents, and corporate executives (cf. Burch, 2025).

We also note that generative AI has been developed under conditions of oligopoly, where public education offers a lucrative customer base for technology companies. Such an unprecedented concentration of political and economic power sits at odds with the public's ability to maintain democratic goals for education. While the development of generative AI without an explicit profit motive is possible (e.g., open-source models, smaller-scale efforts at research universities), the fact is that corporate-driven development is happening at a far greater scale. In that sense, potential misalignments between public and private sector goals are central to questions about the place of generative AI in education.

Ideological Positions Toward Educational Technology

Reflecting on the history of educational technology, we identify two major camps of positions: the hopeful and the wary. The hopeful camp typically goes by the name "techno-optimist," "techno-solutionist," or "techno-utopian." These positions express faith in the transformative potential of technology to improve education and the human species (Lachney et al., 2021). They also consider technology to be politically, socially, and economically neutral. For example, they argue that just as a hammer can be used to build a house or break a window, educational technology—including AI—is simply a tool and neither inherently good nor bad (see X. Zhai & Nehm, 2023).

Hopeful positions toward educational technology serve various agendas. Governments embrace technology to advance national security and international economic competitiveness (see Exec. Order No. 14,110, 2023; Exec. Order No. 14,277, 2025). Corporations market the promise of technology in order to acquire customers and grow profit (Berry, 1988; Reich, 2020).

Educational leaders and policymakers hope that technology can improve student learning, support teachers, and make educational systems more efficient (see Price & Grover, 2025). To facilitate proliferation, advocates typically call for two things: broadening “access” and technological literacy and proficiency (i.e., understanding how a technology works and also developing skills to use and build the technology) (Greene, 2021).

In general, optimism about technology has been the norm in education (Selwyn, 2011). However, positions of wariness—typically under the banners of “techno-realist,” “techno-skeptic,” and “techno-pessimist”—have increasingly received attention. Rejecting the idea of technology as a neutral tool, those in the wary camp interrogate how technology’s development and deployment affect who benefits and who is harmed (White & Scott, 2024; Yadav & Lachney, 2023). Harms of technology are seen as real, numerous, and substantial. In the case of generative AI, scholars have identified societal harms ranging from general damage that affects everyone (e.g., the environmental costs of running massive data centers) to the inequities that minoritized communities experience (e.g., the exploitation of Global South workers who train generative AI models and anti-Black algorithmic bias in the US criminal justice system) (Price & Grover, 2025; White & Scott, 2024).

Importantly, the hopeful and wary camps are not walled gardens. In assembling personal ideologies about technology, educational stakeholders usually take up multiple positions from both camps simultaneously. Monikers like “techno-optimist” and “techno-pessimist” can imply clearly bounded identities that mark two endpoints on an ideological spectrum. But this binary model does not account for the complexity of the ideological field.

In reality, pure versions of optimism and pessimism are hard to find; these are porous categories that bleed into each other. Neither are the optimists as optimistic as they are portrayed, nor are the pessimists as pessimistic. For example, some of AI's most fervent enthusiasts have acknowledged its dangers, such as Google's admission of algorithmic bias against dark skin tones in facial recognition software (O'Brien, 2025), and Open AI co-founder Sam Altman's (2025) statement that "the historical impact of technological progress suggests that…increasing equality does not seem technologically determined and getting this right may require new ideas" (para. 24). While such statements can be interpreted as cynical bids to preempt lawsuits or boost public image, corporate leaders even paying lip service to inequity is somewhat noteworthy.

The image of the uncompromising pessimist who hates technology is also a caricature. In regard to his classic essay, "Why I Am Not Going to Buy a Computer" (1988), the writer Wendell Berry clarified: "I did not say that I proposed to end forthwith all my involvement in harmful technology, for I do not know how to do that. I said merely that I want to *limit* such involvement" (p. 8, emphasis added). Similarly, the self-avowed pessimist Neil Selwyn (2011) stated, "I am not arguing for the adoption of a dogmatic blanket negativity towards education and technology" (p. 714). Neither were the Luddites — the early 19th century English clothworkers who embraced tactical sabotage of automating machines — nor the subsequent social movements aligned with Luddism reflexively anti-technology (Nichols et al., 2025).

Thus, even between the most extreme ideological positions we notice two points of consensus that are rarely acknowledged. First, all positions recognize that some positive cases of educational technology *do* exist; technology does not always cause harm to students and teachers. Second, given the potent mix of corporate and governmental interest, there is also agreement that societal proliferation of technology like generative AI is virtually impossible to stop. In that sense, X. Zhai and Nehm's (2023) claim that the "[AI] train has left the station" might be true. However, this does not mean that educational stakeholders must passively acquiesce to mass adoption and the "myth of technological inevitability" (Weizenbaum, 1972).

Critical Optimism

Apart from hope and wariness, there have been efforts to forge a middle path that "…does not make technology either hero or villain" (Hughes & Eisikovits, 2022, p. 6). One ideological position, which we call "critical optimism," has aimed to reconcile wariness and hope. Critical optimism rejects the idea of technology's neutrality while maintaining belief that technology can improve education. President Biden's (2023) executive order on AI exemplifies critical optimism, as evidenced by the opening paragraph:

> Artificial intelligence (AI) holds extraordinary potential for both promise and peril. Responsible AI use has the potential to help solve urgent challenges while making our world more prosperous, productive, innovative, and secure. At the same time, irresponsible use could exacerbate societal harms such as fraud, discrimination, bias, and disinformation; displace and disempower workers; stifle competition; and pose risks to national security. Harnessing AI for good and realizing its myriad benefits requires mitigating its substantial risks. (p. 75,191)

A key phrase here is "harnessing AI for good." Critical optimism makes the following assumption: with the right guardrails, even risky technology can be steered toward net benefit for society and for minoritized people, in particular. In fact, subsequent text in Biden's executive order explicitly links the development of "responsible AI" to the project of advancing "equity and civil rights" (p. 75,192).

One guardrail proposed to advance critical optimism is "critical technological literacy," which goes beyond skill development to emphasize awareness and understanding of how technology can exacerbate societal inequity (Avraamidou, 2024; White & Scott, 2024; Yadav & Lachney, 2023). A goal of such education is to foster students' wariness about the marketing narratives pushed by technology companies. In turn, educators hope this sensibility will empower youth to develop technology that actually advances equity and justice (Penuel et al., 2025; Pleasants et al., 2023; Vakil, 2018).

Appraisal of Critical Optimism

To be clear, we agree that students *should* be empowered as "critical agents who can shape technological futures" (Pleasants et al., 2023, p. 490). Education in the sociopolitical dimensions of technology—of the kind that critical optimism envisions—is a valuable harm-mitigation effort. Ultimately, though, we find critical optimism an insufficient response to the full weight of questions around educational technology like generative AI.

Our main issue with critical optimism is that it downplays its own analysis. If technology has caused so much harm in the past (and the present), why should we believe that "this time will be different"? What is different now about society that would allow technology to broadly generate net-benefit for education and for minoritized students? Critical optimism engages in both sides-ism: a desire to have one's technology and to critique it, too. This can lead to some curious contradictions.

Consider the following excerpts from a syllabus statement on students' use of generative AI in a university class:

> "…all large language models have a tendency to make up incorrect facts and fake citations as well as [sic] prone to extreme bias…By participating in such models, you are complicit in the outright theft of labor by technology corporations in which they have ingested the data and labor of the internet as training data for their models, broadly

> without consent and against the will of many content creators…*Having said all these disclaimers, the use of these models is encouraged, as it may make it possible for you to submit assignments with higher quality, in less time*" (Bowers, 2025, emphasis added)

Critical optimists might argue that encouraging the use of generative AI after excoriating its problems simply reflects the ambivalence involved in navigating modern life. But then one wonders: what was the purpose of the critical analysis?

We know that positive cases exist of non-generative AI being used to amplify equity and justice. For example, Vakil (2024) notes an initiative by Indigenous people in Aotearoa/New Zealand to use machine learning to sustain the Māori language. Additionally, in 2020 the government of Togo, in partnership with scientists at the University of California, Berkeley, deployed AI to help identify Togolese people with greatest need of financial aid during the COVID-19 pandemic (Arnold, 2025). These efforts are remarkable and important. But outside of local pockets of implementation, projects of this kind will likely never be allowed to scale (Shah & Yadav, 2023).

Corporations and governments sometimes align on the proliferation of a particular technology, as seems to be the case with generative AI. Together, these entities form a potent public-private axis that should not be underestimated. Importantly, their respective goals are shareholder returns and national security, not equity and justice. Grassroots efforts to use generative AI for social good may benefit some minoritized students and communities, but the US public-private axis around this technology aims to reach all 50 million American P-12 students—and they have the resources to succeed.

Critical optimism is founded on the myth of a "good AI," which we could build and implement at scale outside of our current sociopolitical reality. This sometimes leads to overzealous claims that lack research-based evidence, such as the ability of generative AI to address teacher bias or to help develop culturally responsive lesson plans (see Price & Grover, 2025). Generative AI might one day help solve vexing issues in education. But given the history of hype cycles associated with educational technology, we believe that such claims amount to wishful thinking.

We are puzzled by the critical optimist's unwillingness to consider *refusal* as a legitimate option. By "refusal" in this context we mean a practice of protecting a place and a people from the incursion of a technology by blocking its entry, in those situations where we can reasonably anticipate that technology doing more harm than good. In search of a "good AI," critical optimism struggles to imagine a place where AI may not belong (cf. Weizenbaum, 1976).

Thus, although critical optimism locates itself as a "sensible middle" on the ideological field, in practice the position is barely distinguishable from raw hope. By not taking refusal seriously, critical optimism reveals itself as a branch of techno-optimism. To illustrate, compare President Biden's executive order on AI (Exec. Order No. 14,110, 2023) to the executive order on AI issued two years later by the Trump administration (Exec. Order No. 14,277, 2025). A key difference between them is that unlike the Biden order, the Trump order never mentions harm or inequity. However, despite stark political differences between these administrations, the executive orders converge on a clear bipartisan consensus that AI technologies must be encouraged to proliferate.

Why does this matter? Critical optimism grants an implicit green light to the advance of educational technology. In the case of generative AI, critical optimism provides rhetorical cover to technological evangelists who minimize known and potential harms. These evangelists must understandably wonder: if even the staunchest critics of generative AI still seem committed to its development, then what reason is there for the evangelists to stop? To be clear, we do not see this as a deliberate obfuscation on the part of those espousing critical optimism. Rather, this problem results from a certain degree of wishful thinking about the sociopolitical context of educational technology.

A Call for Caution and the Legitimizing of Refusal

Whenever a new technology encroaches on education, the stakes are high because children might be harmed. For this reason alone, our collective stance in education towards new technology should be strong caution. But the history of educational technology indicates the opposite: we usually implement first and ask questions later. This is not the case in fields like medicine, where new surgical techniques or new cancer drugs are vetted for years before being allowed to enter the field. Part of the challenge is a mismatch in risk tolerance. Technology companies are driven by an ideology that prioritizes speed and experimentation with a high tolerance for risk. This is antithetical to the low risk tolerance needed in education and thus poses a stress test for educational leaders tasked with making consequential decisions about technologies like generative AI.

No one knows what the impact will be of implementing generative AI at scale in education. Given this uncertainty, we follow Selwyn (2011) in arguing that a starting position of wariness is most prudent when encountering a new educational technology. As Selwyn writes, “Surely, there is nothing wrong with attempting to develop realistic and honest ways of working with digital technologies in education *that involve thinking the worst (rather than the best) of them?*” (p. 717, emphasis added).

We need a historically grounded, sober reckoning with what usually happens when technology enters schools. We agree that most technologies *can* be net-beneficial under certain conditions for certain people. However, the real question for educational stakeholders is this: if we were to authorize the *scaling* of a particular educational technology, will the technology do more harm than good? Who experiences more of the good, and who feels more of the harm? Caution alone is an insufficient guideline. We conclude by arguing that *refusal* of educational technology should be considered a legitimate, reasonable practice for educational stakeholders.

Forms of Refusal

Refusal in this context does not mean a pre-determined, universal "no"; there will be cases where particular educational technologies might be useful. However, we argue that educational stakeholders should be prepared to refuse technology if doing so will benefit students, teachers, and families. Refusal is not a monolith. Given that different actors across the educational system possess varying levels of power, refusal can and should look differently across contexts. Scholars in science and technology studies have detailed various ways that harmful technology can be refused at the level of technological design (see Bender & Hanna, 2025; Benjamin, 2019; Ganesh & Moss, 2022). Because our focus in this article is the *implementation* of educational technology (not technological design), here we propose three forms of refusal that are relevant to education: curricular, bureaucratic, and administrative.

Curricular refusal involves teaching *against* technologies rather than only *about* or *with* them, as is typical with technological literacy approaches (Vogel et al., 2024). With their notion of "media quiteracy," Good and Ciccone (2025) argue that students should have a "right to not-learn [with technology]" and also be taught to "learn through leaving [a technology platform]," where students "…ponder the costs of certain media technologies to the environment, educational institutions, and people's wellbeing" (p. 154). For teachers, curricular refusal means eschewing curricula that assume technology's beneficence, and instead teaching students that *not* using a technology is a legitimate option. Teachers' lessons on refusal can take inspiration from the Luddites and their rejection of dehumanizing automation. Teachers can cultivate a Luddite praxis, which is characterized by: a strategic playfulness that ridicules rather than reveres AI; localized tactics of refusal grounded in the community's cultures and histories; and networks of resistance in and across schools (Logan et al., 2025). A Luddite praxis offers educators a repertoire of individual and collective practices to care for their students and their school communities.

In addition to refusing the use of a technology, curricular refusal can involve teachers and students examining why technologies like generative AI are becoming prevalent in society, rather than acquiescing to narratives of inevitability. Interrogating the discursive moves and assumptions of the people and companies working hard to embed particular technologies into schools highlights the heterogeneity of refusal practices. Such classroom conversations can engage students in debates about corporate power and the ways young people do and do not want technologies like generative AI to enter their personal lives.

Bureaucratic refusal involves side-stepping top-down mandates to use an educational technology. Recently, a public-school teacher informed us that her district-level curriculum coordinator mandated that all teachers attend a professional development on lesson planning sponsored by a generative AI company. Individually, this teacher does not have power to determine how her district allocates funds or determines priorities for professional development. However, the teacher opted to partner with other teachers and university researchers to conduct an inquiry into *whether* the technology might have value for addressing a problem of instructional practice (rather than passively accepting the technology's value *a priori*).

Finally, administrative refusal refers to actions that educational leaders (e.g., superintendents, principals, school board members, district technology directors) can take to resist the unfettered creep of educational technology. Whereas curricular and bureaucratic forms of refusal concern how individual educators (e.g., teachers) can respond to decisions made by educational leaders, administrative refusal concerns how leaders themselves might participate in refusal. For example, P-12 superintendents are often faced with high-stakes decisions about whether to scale expensive technologies like generative AI district-wide. Rather than rushing into long-term contracts with for-profit technology companies, superintendents can choose to delay full-scale implementation in favor of small-scale pilots to first assess a technology's benefits and harms to children in their districts (Williamson et al., 2024). School board members can also set policies that require independent, peer-reviewed research to support claims about a technology's impact on teaching and learning. Relatedly, technology directors can take time to vet a technology's potential for harm and lead communitywide discussions to determine if the technology should be adopted.

Certainly, refusal can come with costs (e.g., reprimand, job loss). Stakeholders closest to the ground must weigh the benefits of refusal against potential consequences and decide accordingly. We also acknowledge that in the case of generative AI, given the ubiquity it has established in a short time, a total unwillingness to engage with this technology will not be practical (nor even desirable) for everyone. But at least, even when outright refusal is not viable, we argue that everyone can ask the following questions:

- Why should I use this technology? Does this technology help me solve a real problem in my work? Are there external factors (e.g., financial incentives) pressuring me to use this technology?

- What harms might this technology cause, both within and beyond my immediate context?
- How can families and communities co-lead in making decisions about educational technology?
- What types of evidence (and how much evidence) do I need in order to be convinced that refusal no longer remains my best option?

Conclusion

Technological change in education is an uncertain process that requires humility. This is especially true in the case of “arrival technologies” like generative AI, where a technology suddenly enters an educational system prior to organized, deliberative adoption processes (Reich & Duke, 2025). Under these conditions, caution and refusal are sensible practices.

Research also has a vital role in decision making around educational technology. Our field needs to build robust evidence bases that illuminate how educational technology actually gets used—and refused—by those most directly affected. Studies are needed about educational technology’s influence on students’ lives in classrooms, its potential to reconfigure the teaching profession, and how a given educational technology might redefine human learning. In turn, researchers should hold accountable those who make ungrounded claims about the virtues of educational technology. Finally, because research moves slower than practice, we should also engage in rigorous, proactive speculation about potential harms, rather than waiting for damage to happen.

In the early days of the iPad, co-founder of Apple, Steve Jobs, infamously admitted that his own children had not used the device (Bilton, 2014). If refusal was apparently an easy personal choice for Jobs—the ultimate techno-optimist—then why should refusal be so unimaginable for us in education? We cannot afford to be wrong about technologies like generative AI when the lives and well-being of children, teachers, and schools are at stake.

References

Althusser, L. (1971). *Lenin and philosophy* (B. Brewster, Trans.). Monthly Review Press.

Altman, S. (2025, February 9). Three observations. https://blog.samaltman.com/three-observations

Arnold, C. (2025, February 26). *Can AI help beat poverty? Researchers test ways to aid the poorest people*. Nature. https://www.nature.com/articles/d41586-025-00565-7

Ames, M. G. (2019). *The Charisma Machine: The Life, Death, and Legacy of One Laptop per Child*. MIT Press.

American Civil Liberties Union (ACLU). (2023). *Digital Dystopia: The Danger in Buying What the EdTech Surveillance Industry is Selling* (pp. 1–61) [ACLU Research Report]. ACLU.

Avraamidou, L. (2024). Can we disrupt the momentum of the AI colonization of science education? *Journal of Research in Science Teaching, 61*(10), 2570-2574.

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2024). Generative AI can harm learning. *The Wharton School Research Paper*, 1-59. Available at SSRN: https://ssrn.com/abstract=4895486

Beckingham, S., Lawrence, J., Powell, S., & Hartley, P. (Eds.). (2024). *Using generative AI effectively in higher education: Sustainable and ethical practices for learning, teaching and assessment*. Routledge.

Bender, E., & Hanna, A. (2025). *The AI con: How to fight big tech's hype and create the future we want*. Harper Collins.

Benjamin, R. (2019). *Race after technology: Abolitionist tools for the New Jim Code*. Polity Press.

Berry, W. (1988, December). Letters in response to *Why I Am not Going To Buy A Computer* (published in September 1988). Harper's Magazine.

Bilton, N. (2014, September 10). *Can AI help beat poverty? Researchers test ways to aid the poorest people*. The New York Times. https://www.nytimes.com/2014/09/11/fashion/steve-jobs-apple-was-a-low-tech-parent.html

Bowers, A. J. (2025, April 23-27). *Navigating AI in education: Opportunities, challenges, and research imperatives* [Conference presentation]. AERA 2025 Annual Meeting, Denver, CO, United States.

Burch, P. (2025, April 23-27). *Navigating AI in education: Opportunities, challenges, and research imperatives* [Conference presentation]. AERA 2025 Annual Meeting, Denver, CO, United States.

Citron, D. K. (2024). The surveilled student. *Stanford Law Review*, *76*, 1439–1472.

Crooks, R. (2021). Productive myopia: Racialized organizations and edtech. *Big Data & Society, 8*(2), 1–16. https://doi.org/10.1177/20539517211050499

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

Cuban, L. (1986). *Teachers and machines: The classroom of technology since 1920*. Teachers College Press.

Cuban, L., Kirkpatrick, H., & Peck, C. (2001). High access and low use of technologies in high school classrooms: Explaining an apparent paradox. *American Educational Research Journal, 38*(4), 813-834.

Cueto, S., Beuermann, D., Cristia, J. P., Malamud, O., & Pardo, F. (2024). *Laptops in the Long-Run: Evidence from the One Laptop per Child Program in Rural Peru* (Working Paper No. 34495). National Bureau of Economics Research. https://doi.org/10.18235/0013192

Eagleton, T. (1991). *Ideology: An introduction*. Verso.

Exec. Order No. 14,110, 3 C.F.R. 75,191 (2023). https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence

Exec. Order No. 14,277, 3 C.F.R. 17,519 (2025). https://www.federalregister.gov/documents/2025/04/28/2025-07368/advancing-artificial-intelligence-education-for-american-youth

Fedders, B. (2019). The constant and expanding classroom: Surveillance in K-12 public schools. *North Carolina Law Review, 97*(6), 1673–1726.

Fesler, L., Martinez, J., Agnew, C., & Loeb, S. (2026). *The evidence base on AI in K-12: A 2026 review*. AI Hub for Education of the SCALE Initiative: Stanford University. https://scale.stanford.edu/

Ganesh, M. I., & Moss, E. (2022). Resistance and refusal to algorithmic harms: Varieties of 'knowledge projects.' *Media International Australia, 183*(1), 90-106.

Garcia, A., & Nichols, T. P. (2021). Digital platforms aren't mere tools—They're complex environments. *Phi Delta Kappan, 102*(6), 14–19. https://doi.org/10.1177/0031721721998148

Good, K. D., & Ciccone, M. (2025). Media quiteracy: Why digital disconnection belongs in the media literacy curriculum. *Journal of Media Literacy Education, 17*(1), 150-165. https://doi.org/10.23860/JMLE-2025-17-1-10

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

Hannafin, R. D., & Savenye, W. C. (1993). Technology in the classroom: The teacher's new role and resistance to it. *Educational Technology, 33*(6), 26-31.

He, R., Cao, J., & Tan, T. (2025). Generative artificial intelligence: a historical perspective. *National Science Review, 12*(5), 1-15.

Howard, S. K. (2013). Risk-aversion: Understanding teachers' resistance to technology integration. *Technology, Pedagogy and Education, 22*(3), 357-372.

Hughes, J., & Eisikovits, N. (2022). The post-dystopian technorealism of Ted Chiang. *Journal of Ethics and Emerging Technologies, 32*(1), 1-14.

Johnson, O., & Jabbari, J. (2022). Infrastructure of social control: A multi-level counterfactual analysis of surveillance and Black education. *Journal of Criminal Justice, 83*, 1–8. https://doi.org/10.1016/j.jcrimjus.2022.101983

Lachney, M., Ryoo, J., & Santo, R. (2021). Introduction to the special section on justice-centered computing education, part 1. *ACM Transactions on Computing Education (TOCE), 21*(2), 1–15.

Laird, E., Grant-Chapman, H., Venzke, C., Quay-de la Vallee, H. (2022). *Hidden harms: The misleading promise of monitoring students online*. Center for Democracy & Technology. https://cdt.org/insights/report-hidden-harms-the-misleading-promise-of-monitoring-students-online/

Leonardo, Z. (2004). Critical social theory and transformative knowledge: The functions of criticism in quality education. *Educational Researcher, 33*(6), 11-18.

Leonardo, Z. (2010). Ideology and its modes of existence: Toward an Althusserian theory of race and racism. In Z. Leonardo (Ed.), *Handbook of cultural politics and education* (pp. 195-217). Sense Publishers.

Logan, C., Nichols, T. P., & Garcia, A. (2025). Teach like a Luddite. *Phi Delta Kappan, 107*(3-4), 37-41. https://doi.org/10.1177/00317217251405524

Nichols, T. P., Logan, C., & Garcia, A. (2025). Generative AI and the (Re)turn to Luddism. *Learning, Media and Technology*. 1–14. https://doi.org/10.1080/17439884.2025.2452199

O'Brien, M. (2025, April 2027). *Tech industry tried reducing AI's pervasive bias. Now Trump wants to end its 'woke AI' efforts.* The Associated Press. https://apnews.com/article/artificial-intelligence-dei-trump-algorithmic-bias-woke-ai-8302e12dd74df69a1adc6565710f033d

Pangrazio, L., Selwyn, N., & Cumbo, B. (2023). Tracking technology: Exploring student experiences of school datafication. *Cambridge Journal of Education*, *53*(6), 847–862. https://doi.org/10.1080/0305764X.2023.2215194

Penuel, W. R., Philip, T. M., A. Chang, M., Sumner, T., & D'Mello, S. K. (2025). Responsible innovation in designing AI for education: Attending to the "how," the "for what," the "for whom," and the "with whom" in a rapidly growing field. *Educational Researcher*, 1-8.

Pleasants, J., Krutka, D. G., & Nichols, T. P. (2023). What relationship do we want with technology? Toward technoskepticism in schools. *Harvard Educational Review*, *93*(4), 486–515. https://doi.org/10.17763/1943-5045-93.4.486

Price, J. F. & Grover, S. (2025). *Generative AI in STEM teaching: Opportunities and Tradeoffs.* National Science Foundation. https://cadrek12.org/sites/default/files/2025-01/edc-CADRE-REPORT-AI-Teaching-2025.pdf

Reich, J. (2020). *Failure to disrupt: Why technology alone can't transform education*. Harvard University Press.

Reich, J., & Dukes, J. (2025). The future of education technology after the arrival of ChatGPT. *Phi Delta Kappan, 107*(3-4), 19-23.

Sadowski, J. (2019). When data is capital: Datafication, accumulation, and extraction. Big Data and Society, 6(1), 1–12. https://doi.org/10.1177/2053951718820549

Selwyn, N. (2011). In praise of pessimism—the need for negativity in educational technology. *British Journal of Educational Technology, 42*(5), 713-718.

Shah, N., & Yadav, A. (2023). Racial Justice Amidst the Dangers of Computing Creep: A Dialogue. *Tech Trends, 67*(3), 467-474.

Tegmark, M. (2017). *Life 3.0: Being human in the age of artificial intelligence*. Vintage Books.

Vakil, S. (2018). Ethics, identity, and political Vision: Toward a justice-centered approach to equity in computer science education. *Harvard Educational Review, 88*(1), 26-52.

Vakil, S. (2024, May 28-30). *Will AI expand opportunity and equity in STEM education?* [Keynote address]. National Science Foundation ECR PI Meeting, Arlington, VA, United States.

Vogel, S., Yadav, A., Phelps, D., & Patel, A. (2024). Entrypoints for integrating computing and tech into teacher education: Addressing problems and opportunities with the EnCITE framework. *Journal of Technology and Teacher Education, 32*(2), 217-248.

Wang, J., & Fan, W. (2025). The effect of ChatGPT on students' learning performance, learning perception, and higher-order thinking: insights from a meta-analysis. *Humanities and Social Sciences Communications, 12*(1), 1-21.

Wecks, J. O., Voshaar, J., Plate, B. J., & Zimmermann, J. (2024). Generative AI usage and exam performance. arXiv preprint arXiv:2404.19699.

Watters, A. (2021). *Teaching machines: The history of personalized learning*. MIT Press.

Weizenbaum, J. (1972, May 28). It's your dream. *The New York Times*. https://www.nytimes.com/1972/05/28/archives/-its-your-dream.html

Weizenbaum, J. (1976). *Computer power and human reason: From judgment to calculation*. W. H. Freeman and Company.

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 January 16, 2026, from http://nepc.colorado.edu/publication/ai

White, S. V. & Scott, A. (2024). *Responsible AI and tech justice: A guide for K-12 education.* Kapor Foundation. https://kaporfoundation.org/wp-content/uploads/2024/01/Responsible-AI-Guide-Kapor-Foundation.pdf

Yadav, A., & Lachney, M. (2023). Towards a techno-social realist approach in primary and secondary computing education. In A. Johri (Ed.), *International handbook of engineering education research* (pp. 553-572).

Zhai, C., Wibowo, S. & Li, L.D. (2024). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: a systematic review. *Smart Learning Environments, 11*(1), 1-37. https://doi.org/10.1186/s40561-024-00316-7

Zhai, X., & Krajcik, J. (Eds.). (2025). *Uses of artificial intelligence in STEM education*. Oxford University Press.

Zhai, X., & Nehm, R. H. (2023). AI and formative assessment: The train has left the station. *Journal of Research in Science Teaching, 60*(6), 1390-1398.