


A Means to Know End Generative AI and the Many Futures of Learning

Amy J. Ko, Ph.D.
The Information School
University of Washington, Seattle

- Thank you so much for the invitation and the introduction. It's my first time at Wesleyan, and it's an honor to share space with you, to talk about this moment and our role in it, as students, scholars, and citizens.
- When I was thinking about what I wanted to speak on, one thought kept coming back to me: *don't talk about generative AI*. And yet, here I am, talking about AI!
- So why, of all things taking up way too much space in public consciousness, do we need another talk about the technology of the moment? Beyond the obvious reasons that some of us hate it, some of us love it, and it seems to be changing everything, whether we like it or not?
- It came down to one thing: so much is missing from our conversations about generative AI from history, political economy, learning sciences, and more, and I wanted to supplement this discourse with these social, political, and historical perspectives, from my perspective as a computer scientist and learning scientist. And in sharing this missing context about generative AI, I hope to help us see the future a bit more clearly, especially one in which generative AI is not at the center of so many conversations, even if it remains part of them.
- I share these thoughts from a particular stance.


I am irreducible.

- My grandparents immigrated to the U.S. ~100 years ago from Guangdong, China and Horslund, Denmark, but I grew up in 1980's Ontario and Portland, Oregon. My cultural and ethnic history is centered around food, war, poverty, punk, urban, rural, weirdness, and whiteness.
- I'm a woman, a mother, and transgender. I experience harassment and sexism like all women, but I've also face assault in men's and women's restrooms, I am banned from teaching in schools in many states, and am not permitted to serve our country in the military. I broke the law in Texas on my way here because of the misinformation and lies spread by AI about transness.
- I am a computer scientist. I invent programming languages, but also read social theory, and write about oppression and justice. I like talking about computing, anti-racism, and disability justice with kids.
- Generative AI is not separate from these things: it is a tool designed to reproduce history, to center normativity. I am anything but normal and I desperately need the future to be nothing like the past. And so my default position on generative AI is that it is a particularly acute threat to my existence and the people I love.
- But I also come to this talk as a scientist, meaning that despite my beliefs, I want to bring the full breadth of knowledge and discovery to understanding: Is AI a threat? How? And if not, what should AI's role be in my life and others'?




“AI”

- And one terminology note, before we begin: I am not fond of the phrase "Artificial Intelligence". No form of computing has ever demonstrated anything more than a facsimile of the kinds of intelligence we mean when we say that word. There is no thought, no mind, no speaker, other than its creators'.
- Its creators, in a way, are all of us, the few overpaid capitalists, the computer science Ph.D. students they hired to hype this technology, and the billions of people on this planet who have shared their digital content online and in print, with our without their consent, for the careful tuning of a trillion parameters in a handful of privately trained convolutional neural networks.
- Adding the word "generative" was one part metaphor and one part marketing. And as we shall see, the phrase AI itself was marketing to begin with.
- So throughout this talk, I will shorten generative AI to "AI" in quotes.
- With that, let's begin.



- We are in a time of "AI" panic.
- Fears of cognitive offloading abound, in learning and society more broadly.
- True believers fear that the singularity is coming, and marketers feign similar fear to mobilize investment, with both arguing that the best thing we can do is centralize power in a few Silicon Valley startups, so that they may save us from the Pandora's box that computer science has opened.
- Our governments, long having deferred to computer science and industry on matters of computing, are at a loss as to what kinds of policy, regulation, or subsidy is appropriate, with states frantically proposing ill-conceived legislation in the absence of any congressional leadership.
- Meanwhile, the U.S. federal government is willfully dismantling our global leadership in science, engineering, education, and libraries, in the name of defunding and destabilizing institutions committed to knowledge, knowing, and truth.
- Meanwhile, our industries seem to be in a race to see who can give more money to Anthropic, Google, and OpenAI before the entire economy deflates, while are students are left with broken promises about jobs, and a choice between resistance or a \$100 Claude Max subscription to keep up.




- So where is the world going?
- The status quo, with a new tool for extruding our collective digital expression into insipid, homogenous writing, images, music, and code?
- A more "productive" capitalism, in which we can all work twice as much for half the pay?
- Or a Star Trek utopia in which we don't work at all, liberated one language embedding at a time, implicitly dependent on the encrypted open weights in Anthropic's AWS account?
- At an emotional level, I have found myself seeing all three of these futures at different times, wrestling with their extremes, while hoping for something much more mundane.

Our journey

- Histories of knowing through technology
- Histories of “AI”
- Futures of learning

To help us find out, I want to take us on a journey through three histories: of knowing through information technology, of "AI," and end with a glimpse into some of our recent research at the UW Center for Learning, Computing, and Imagination, and what it says about the futures of learning we might build together, instead of waiting for Silicon Valley to decide for us.

Histories of knowing




- Knowing and information technology are richly interconnected.
- We prove we know through speech, writing, signing, and other media.
- We use information technology to produce knowledge, writing, recording, annotating, observing.
- Computing is one of many kinds of information technology, and "AI" one kind of computing.
- Our current panic about "AI" is fundamentally about knowing: whether we will continue to, what it means to know things, and whether "AI" knows anything at all.
- And so it is worth understanding how humanity has surfaced threats to knowing in the past, and what we learned about the reality of these threats over time.



Plato. Phaedrus, 274c–275b (H. N. Fowler, Trans., Loeb Classical Library).

Writing (370 BCE)


- The classic panic in written history is in Plato's Phaedrus, where Socrates relays a myth about the Egyptian god Theuth presenting writing to King Thamus.
- Theuth calls writing an "*elixir of memory and wisdom*" Thamus replies that it will instead "*produce forgetfulness in the minds of those who learn to use it, because they will not practice their memory*"
- In this dialog, their other worry is that students will be presented with knowledge without thinking through the meaning of it, creating a facsimile of wisdom.
- The irony of all this being presented in writing aside, this fear clearly rhymes with today: writing, as a technology, provoked the same fears that generative "AI" does, that we will lose wisdom.
- We're just using different language to describe our fears.



Dehaene, S., et al. (2015). Illiterate to literate: Behavioural and cerebral changes induced by reading acquisition. Nature Reviews Neuroscience, 16, 234–244.

Writing today

- As it turns out, learning to read does reorganize our brains, just in really small ways.
- A few small studies in 2010 and 2015 showed that increases in literacy have a surprising impact on our ability to recognize faces, but also large gains in our phonological processing and visual discrimination.
- These are not at all the kinds of effects Plato predicted, and none of this disproves that society has not fallen to a feigned wisdom.
- If anything, it makes a cultural point about how prized memory and recall was in ancient Greece, and how much less we value it now, relative to skills like information seeking and sensemaking.
- So no, memory alone turned out not to be so central to wisdom, if you believe our greatest thinkers still have wisdom.



- Fast forward 2,000 years, and writing was taken for granted, as it is now.
- But the proliferation of the printing press had created a new fear: *information overload*.
- Conrad Gessner is often credited with the first information-overload alarm, describing the flood unleashed by the printing press as confusing and harmful to the mind (Bell, 2010).
- By the late 1700s the panic had a vocabulary: Lesesucht (reading addiction), Lesewut (reading rage), reading fever, reading mania.
- It was described across Europe as an epidemic afflicting the young, associated with sensation-seeking and moral dissolution (Furedi, 2015).
- The fears centered on psychological effects — imitation, false ideas about love and life — and those treating novel-reading as a physically harmful waste of time damaging mind, morale, eyesight, and posture.
- The threat here was that we all might disappear into fiction, or the intellectual worlds of others' minds, rather than being present in the world, and growing our *own* minds.




Print today

Phillips, D. P. (1974). The influence of suggestion on suicide: Substantive and theoretical implications of the Werther effect. American Sociological Review, 39(3), 340–354.

Niederkrötenhauer, T., et al. (2020). Association between suicide reporting in the media and suicide: Systematic review and meta-analysis. BMJ, 368, m575.

- Fast forward a few hundred years, what have we learned about the threats of reading?
- Most of the fears didn't pan out. We still learn to become ourselves, despite all of our reading.
- But one fear did come true: there were early concerns about reading triggering copycat suicides (Furedi, 2024), and Phillips (1974) later coined the "Werther Effect" for exactly this phenomenon.
- Suicide contagion via media reporting is now established enough as a real risk that it is baked into journalistic guidelines (see Niederkrötenhauer et al., 2020, BMJ), though the 18th-century version wildly overstated how prevalent it was.
- Reading, once again, survived unfounded moral panic, but had unexpected effects.




Telegraph (1840s-1900s)

*Beard, G. M. (1869).
Neurasthenia, or nervous
exhaustion. Boston Medical and
Surgical Journal, 80, 217–221.*


*Beard, G. M. (1881). American
Nervousness: Its Causes and
Consequences. New York: G. P.
Putnam's Sons.*

*Morus, I. R. (2000). "The nervous
system of Britain": Space, time
and the electric telegraph in the
Victorian age. British Journal for
the History of Science, 33(4), 455–
475.*

- Once we got over the risks of reading, the telegraph was the new threat.
- Beard (1869) coined "neurasthenia", blaming nervous exhaustion on the stimulations of modern civilization "represented by steam-power, the periodical press, the telegraph, the sciences" — and, in *American Nervousness* (1881), the mental activity of women. The neurasthenic body was imagined as a "spent battery," treatable with electrotherapy (Wellcome Collection, n.d.).
- The second fear was epistemic. A Times leader complained that "assertion flies along the telegraphic wire, while proof still travels in the post bag" (The Times [London], 21 March 1852; Commentators simultaneously celebrated and feared the "annihilation of space and time" (Morus, 2000).
- These were panics about speed: that the pace of life had a natural form, and that the movement of information at the speed of light would disrupt the natural order of life and trade.



- Neurasthenia didn't pan out, at least not from the telegraph.
- However, the speculated economic effects were real and large.
- Steinwender (2018) found that after the 1866 transatlantic cable, speculation traveled via telegraph, while the proof on the ground lagged behind, and this led to a relatively permanent baseline economic instability.
- This basic phenomena of rumor traveling faster than reality continues to be true today in the way rumors and information spread via text messaging.
- The telegraph was one of the first times people saw the social risks of that at societal scale.



Radio (1920s - 1940s)

Dennis, E. (1998). Chills and thrills: Does radio harm our children? The controversy over program violence during the age of radio. Journal of the History of the Behavioral Sciences, 34(1), 33-50.

Schwartz, A. B. (2015, May 31). America's original fake-news outrage. Salon.


- People feared radio as well.
- Dennis (1998) reviewed the period's popular and scientific press and found the radio "thriller" controversy was one of the most publicized issues about radio during its peak years.
- A Brooklyn mother wrote to the Times in 1935 that crime-and-murder broadcasts produced "unhealthy excitement, unnecessary nervousness, irritability and restless sleep" (quoted in Schwartz, 2015).
- Reform groups and PTAs argued through the 1930s that children's shows were making neurasthenics of youngsters.
- The concern was that radio couldn't be monitored like comics, or movies, because children could listen to it in private
- The fear was: without surveillance, how could parents shape the moral development of their children?



Radio now

Pooley, J., & Socolow, M. J. (2013, October 28). The myth of the War of the Worlds panic. Slate.

- The era's own experts mostly disagreed with the panic, viewing radio violence as mostly harmless.
- And it turns out that most of the panic about radio was manufactured by the media itself, not even grounded in real fear.
- For example, Pooley and Socolow (2013) found the War of the Worlds nationwide-hysteria narrative unsupported; ratings data showed a tiny listenership — about 2% of 5,000 surveyed households.
- Their argument is that the newspaper industry sensationalized the panic to convince advertisers and regulators that to regulate radio to secure the economic position of print news as radio became more popular.




Television (1950s-1980's)

Surgeon General's Scientific Advisory Committee on Television and Social Behavior. (1972). Television and Growing Up: The Impact of Televised Violence. Washington, DC: U.S. Government Printing Office.

Cater, D., & Strickland, S. (1975). TV Violence and the Child: The Evolution and Fate of the Surgeon General's Report. New York: Russell Sage Foundation.

Gerbner, G. (1972). Television Violence: Where the Surgeon General's Study Leads.

- While there were many intersecting fears between radio and television, the panic specific to television arrived in 1972, when the U.S. Surgeon General's Scientific Advisory Committee reviewed 40 original research reports and concluded there was "a preliminary and tentative indication of a causal relation" between viewing TV violence and aggressive behavior (Surgeon General's Scientific Advisory Committee, 1972).
- This particular language angered and confused politicians, broadcasters, and many participating scientists (Cater & Strickland, 1975).
- The New York Times initially ran it as "TV Violence Held Unharmful to Youth" before coverage reversed it, siding with the Surgeon General's report (Gerbner, 1972).
- This panic echoed books and radio, resurfacing fears that children experiencing stories would lose their minds, and stop learning.



Television now

Ferguson, C. J., Copenhaver, A., & Markey, P. (2020). *Reexamining the findings of the American Psychological Association's 2015 task force on violent media: A meta-analysis*. *Perspectives on Psychological Science*, 15(6), 1425–1443.


Drummond, A., Sauer, J. D., & Ferguson, C. J. (2020). *Do longitudinal studies support long-term relationships between aggressive game play and youth aggressive behaviour? A meta-analytic examination*. *Royal Society open science*, 7(7), 200373.

Ferguson, C. J. (2007). *Evidence for publication bias in video game violence effects literature: A meta-analytic review*. *Aggression and Violent Behavior*, 12(4), 470–482.

Prescott, A. T., Sargent, J. D., & Hull, J. G. (2018). *Metaanalysis of the relationship between violent video game play and physical aggression over time*. *PNAS*, 115(40).

Thomas, M. (2019). *Was television responsible for a new generation of smokers?* *Journal of Consumer Research*, 46(4), 689–707.


- None of this science held up.
- Ferguson, Copenhaver, and Markey (2020) found the APA's 2015 violent-media meta-analysis hard to reproduce — some included studies lacked relevant data, and many available studies were omitted.
- Drummond, Sauer, and Ferguson (2020) found no evidence of publication bias in longitudinal studies and concluded the research cannot support a meaningful long-term effect of violent television or games on youth aggression.
- The latent fears of learning loss have shown surprising *positive* effects of television, but for particular groups: English-language learners, who have used television as means to English fluency, and undereducated house wives, learning at home during housework.
- And other emergent issues turned out to be catastrophically negative: Thomas (2019), using a natural experiment, found television caused people to start smoking, with 16-to-21-year-olds particularly affected raising the smoking share by 5–15 percentage points, roughly 11 million additional smokers between 1946 and 1970.
- And so whether television hurt depended on who you were and what messages you were receiving.



*Kraut, R., Patterson, M., Lundmark, V., Kiesler, S.,
Mukhopadhyay, T., & Scherlis, W. (1998). Internet
paradox: A social technology that reduces social
involvement and psychological well-being?
American Psychologist, 53(9), 1017-1031.*

The Internet (1990's)

- The Internet of the 1990's was the panic of my youth. I was born in 1980, many of the fears surrounding its emergence were that children like me would be victims to online pedophiles, that people would isolate themselves in chat rooms and become lonely, and that screen time would erode my young adolescent mind.
- And indeed, in the early 90's, I found myself as a young adolescent in chat room, being groomed by strangers, and exploring my identity. But for myself, far from isolating me, it made visible a world of diversity that wasn't visible in my school, city, or state, as queer and trans people were still mostly in hiding in the U.S.
- This contradicted the earliest science. The first major study, Bob Kraut et al.'s (1998) HomeNet field trial, did find that heavier internet users became less socially involved, lonelier, and more depressed.
- These early findings led to parents and society creating new panics around screen time, fearing that youth would be enveloped into the internet



The Internet (now)

Kraut, R., Kiesler, S., Boneva, B., Cummings, J., Helgeson, V., & Crawford, A. (2002). *Internet paradox revisited*. *Journal of Social Issues*, 58(1), 49–74

Orben, A., & Przybylski, A. K. (2019a). *The association between adolescent well-being and digital technology use*. *Nature Human Behaviour*, 3, 173–182.

Boxell, L., Gentzkow, M., & Shapiro, J. M. (2017). *Greater internet use is not associated with faster growth in political polarization among US demographic groups*. *PNAS*, 114(40), 10612–10617

Braghieri, L., Levy, R., & Makarin, A. (2022). *Social media and mental health*. *American Economic Review*, 112(11), 3660–3693

- But a follow up a few years later showed that most of those negative outcomes disappeared, (Kraut et al., 2002), except the association with increased stress.
- The main emergent finding since has been a "rich get richer" effect of better outcomes for extraverts with strong social networks and worse for introverts and isolated people.
- For screen time, Orben and Przybylski (2019a) studied 355,358 adolescents finding only a tiny negative association between internet use and well-being.
- Even more recent fears about political polarization have found that polarization increased most among the demographic groups least likely to use the internet and social media (those over 75) (Boxell, Gentzkow, and Shapiro, 2017).
- The strongest and most recent causal effects consistently show that youth who lean on it for community when they lack community in real life face the worst negative effects on mental health (Braghieri, Levy, and Makarin, 2022), but that queer and trans youth who don't have community IRL face the highest positive effects on mental health. This was exactly my experience.
- Like television, these findings show that whether the Internet helps or hurts depends greatly on what you use it for, and how much you use it.




**Moral panics about youth, learning,
and information technology tell us
little and it takes time for science to
tell us anything**

- It takes sustained, careful investigation of the very specific mediating effects of information technology to understand the who, what, when, where, and why of how it shapes learning.
- And what we often find is that many things we *aren't* paying attention to end up being bigger problems, or bigger benefits, in particular contexts, for particular people.
- The real risks and opportunities are in the minute details, not in the macro trends.
- And so science has been a tool to sort some of this out slowly over time, but only in retrospect, and often only after a decade or longer of careful inquiry. Science has not yet been a powerful tool for us to predict these effects on learning or society.

Histories of “AI”

- And so we find ourselves in that position now: a new information technology, raising new fears, but only emerging evidence to guide us.
- What does this history teach us about "AI" and what is the emerging science telling us so far?

● **Origins of "AI"**




- We have to start by observing that "AI" has always been a marketing idea for engineers, not a scientific one.
- It was never a word that had any well-defined meaning to start with.
- Before digital computers, it was also meant as a creativity tool for thought, to help people wonder, like Ada Lovelace and Alan Turing did earlier, wondering whether and how it might be possible to make a machine say and do what we do.
- After digital computers, it was meant as a provocation, to excite the U.S. Department of Defense and private foundations to give to computer scientists who wanted to explore what kinds of complexity computers might be able to replicate.
- And in the late 60's, AI was a counter-vision of Douglas Englebart's futures of computing that imagined computers as tools to augment our thought, instead imaging that they would replace us, freeing us to do things other than think.

The many "AI"s

<i>1940's</i>	Algorithms that mimic processes in life (e.g., <i>Turing, Lovelace</i>)
<i>1950's</i>	Algorithms that model and mimic human cognition (e.g., <i>McCarthy</i>)
<i>1960's</i>	Algorithms with probability optimization (e.g., <i>hill climbing</i>)
<i>1980's</i>	Algorithms that modeled expert knowledge (e.g., <i>expert systems</i>)
<i>2000's</i>	Statistical models learned from data (e.g., machine learning)
<i>2010's - now</i>	Massive convolutional neural networks dependent on surveillance capitalism (e.g., LLMs)

- Because it was ill-defined, it was a flexible term, and its meaning has changed in the public mind repeatedly.
- With Turing in the 40's, AI was algorithms: deterministic procedures that would mimic the industrialized human processes in math, factories, and life.
- With McCarthy in the 50's, when the phrase was coined, it was the idea of modeling every aspect of human cognition could be simulated "if sufficiently described"
- In the 60's and 70's, "AI" was optimization with probability
- In the 70's and 80's, AI was expert systems, elaborate systems of hard-coded rules that captured skilled performance of people, and a small corner of people inspired by reductive models of neurons and neurotransmitters
- In the 90's, "AI" was statistics, often drawing upon labeled data, like "common sense" knowledge, elaborate webs of semantic connections between concepts and words.
- Since the 2000's, AI has mostly meant machine learning, and its many applications in robotics, machine vision, image processing, patterned generation
- And since a few years ago, it has specifically meant convolutional neural networks trained on billions of our tokenized documents to model language embeddings, wrapped in a loop to predict the next word, pixel, or frequency to generate plausible writing, images, sounds, and videos.

● “AI” moral panics




- Up until now, most moral panics about “AI” have been about societal collapse.
- “I, Robot”
- “The Terminator”
- “The Matrix”
- “WALL-E”
- “Battlestar Galactica”
- These tell stories of existential revolt, of economic obliteration, of replacement lovers, of loss of autonomy, of mental and physical atrophy
- Of course, none of these have really happened.
- Automation does displace forms of work, but does not appear to eliminate work; capitalism seems to ensure that we always have too much labor to do to survive
- People *do* use technology to avoid thinking, just like we use cars to avoid walking
- And while LLMs are starting to do things without designed intention, software has done that since it’s beginning, with viruses and defects always causing computers to misbehave.

● Cognitive offloading

- Our latest panic of "cognitive offloading" is in the same tradition of cognitive atrophy, the very same concept Plato worried about.
- The fear, however, is far more comprehensive than past panics.
- That people will stop reading, stop writing, stop making music, stop learning, stop creating, stop making decisions.
- That people will be replaced by algorithms at work.
- These fears sit alongside many other very real social, political, and economic forces that we will come back to later, which aren't really hypothetical: a centralization of power, theft of intellectual property, amplification of mis- and disinformation, acceleration of climate change, and the exploitation of people for data and land for data centers.
- For now, let's focus on impacts on knowledge: what has science taught us so far?



- The science is nascent. It's small lab studies, some field experiments, and a lot of speculation.
- One of the more publicized studies last year was Kosmyna et al. (2025), a lab study of 54 participants who wrote essays with an LLM, with a search engine, or by themselves, finding that EEG connectivity scaled inversely with external support. LLM users also struggled to quote their own essays and reported lower ownership.
- Another study by Budzyń et al. (2025), in an entirely different domain, found endoscopists' adenoma detection rate when working without AI fell from 28.4% to 22.4% after a period of routine AI-assisted colonoscopy, while AI-assisted rates held steady at 25.3%, showing actual deskilling in a clinical setting.
- A study in human-computer interaction last year, one of my communities, showed higher confidence in GenAI was associated with less critical thinking, while higher task-specific self-confidence was associated with more. Participants reported reduced effort across most cognitive activities, with the nature of the work shifting from information gathering to verification and from doing to supervising.
- Gerlich (2025), surveying 666 participants, similarly found heavy AI use associated with lower critical-thinking scores, mediated by offloading. Both are cross-sectional and confounded.
- Multiple studies have shown that how generative AI tools are designed and used have significant effects on learning; unrestricted use tends to lead to harms in learning, whereas scaffolded, intentional, directed, and focused uses can lead to improvements in learning.
- And in computing education, the effects are quite clear: when learners delegate programming to generative AI without engaging with what they generate, they learn less. But when they use them to support learning, they learn more.
- In an experiment with 600 creative writing students across 10 universities, LLMs enhanced students' level of detail and presentation, but not their ability to produce those qualities independently, while also severely harming originality.



Self-regulation helps


Prather, J., Reeves, B. N., Leinonen, J., MacNeil, S., Randrianasolo, A. S., Becker, B. A., ... & Briggs, B. (2024, August). The widening gap: The benefits and harms of generative ai for novice programmers. ACM ICER.

Newman, M., Sun, K., Dalla Gasperina, I. B., Shin, G. Y., Pedraja, M. K., Kanchi, R., ... & Yip, J. (2024, May). "I want it to talk like Darth Vader": Helping Children Construct Creative Self-Efficacy with Generative AI. In ACM CHI

Pardos, Z. A., & Bhandari, S. (2024). ChatGPT-generated help produces learning gains equivalent to human tutor-authored help on mathematics skills. PLOS One, 19(5), e0304013.

- Other studies suggest that negative impacts on learning and knowledge are highly mediated by self-regulation.
- For example, that same study on computing education showed that when students used "AI" to explicitly support their learning, they learned more and faster.
- In a study on creativity with children, when scaffolded to support ideation, LLMs acted as a kind of small-scale support for children's broader creative process, and did not supplant their desire to be creative, or learn creativity skills.
- In an experiment with 274 adults comparing ChatGPT to human tutors, LLMs could produce hints that were as helpful as best-in-class human tutor hints on routine algebra and statistics problems, even though its hints were significantly more incorrect than the human hints, because problems were idiomatic and ubiquitous.
- These and others suggest that using generative AI as a tool, with intention, and for the right jobs, can help learners quite a bit. It's just that many students don't.

● **Agency and intention matter**



- What these emerging results appear to suggest is what history has taught us as well about other information technologies: if you use them to offload, offloading happens. If you use them to learn, learning happens.
- The same is true in all of the previous moral panics about information technology and learning.
- If you read but never remember, you'll have no wisdom. But if you read to remember, you might gain wisdom.
- If you trust messages that arrive at the speed of light, but do not verify them, you'll not know the truth. But if you read messages that arrive at the speed of light with skepticism, you might know the truth faster
- If you obsess over stories on the radio, television, or streaming sites, at the expense of being in the world, you will forget yourself. If you enrich your life with stories, you might learn about yourself and the world.
- If you use social media at the expense of connection, you will be isolated. If you use social media to overcome isolation, you might make friends and find yourself.
- And if you use AI to avoid learning, you will not learn. But if you use AI to support your learning, you might learn more, faster.
- All of this points us to a number of maxims that are clear from science and technology studies, HCI, and other fields concerned with people and technology: technology is capable of great benefit and great harm; it depends on how, when, and why we use it, and what we choose to amplify with it.

Futures of learning



- How then, should we use “AI” in learning and in life?
- Our macro-level discussion of past and present information technologies can lead us to view history as something that *happens* to us.
- But that is far from how the world works, or at least ought to work.
- We make choices about how and whether we use information technology.
- We make choices about individual and collective choices about how we regulate and limit it.
- We make choices about how much we depend on it, and when we find ourselves losing control, we change culture and supports to counteract those effects.
- And so we have to ask: what futures do we want, if we don't like the effects on learning and behavior that we're finding in the status quo?
- That's part of what we've been investigating in the Center for Learning, Computing, and Imagination over the past few years, and I'd like to share a bit about what we've found.



Eman Sherif

Young adults want liberatory, relational learning spaces

Belén Edgar, Eman Sherif, Janet Jiang, Saara Uthmaan, Amy J. Ko (2026). Not Built for Us: Marginalized Students' Visions for Help-Seeking in Computing Education. ACM International Computing Education Research Conference (ICER).

- My Ph.D. student Eman Sherif has worked with students on the margins of computing education to ask: if you don't like the world that generative AI is creating in your learning spaces, what worlds do you want instead?
- She's found that “AI” use is fundamentally about avoiding shame and humiliation, and surviving capitalist, competitive labor market threats in industrialized learning spaces.
- Students want for learning contexts in computer science to be more relational, centered on building community around shared identity, and less on transactions between teachers and students.
- They want help-seeking and teaching to be about connecting.
- They want agency to learn how they want, when and where they want, with whom they want, without having to conform to a teachers' particular framework for learning.
- In those worlds, they see generative AI as a more narrow learning tool, not a tool for survival.



Jayne Everson, Ph.D.

**Youth want learning and making to be joyful,
and generative AI is rarely that**

*Jayne Everson, Rotem
Landesman, Janet Jiang, Megumi
Kivuva, Eman Sherif, Amy J.
Ko (2026). Student Experiences of
Joyful Secondary CS
Classrooms. ACM International
Computing Education Research
Conference (ICER)*

- My recently graduated Ph.D. student, now at Foundry10, examined the role of joy in computer science teaching and learning.
- She found in her most recent study of the rare joyful CS learning spaces that adolescents want space to express themselves, to learn through making, to have the care and support of teachers in doing
- They also shared that using generative AI is rarely any of those things, because it takes away precisely the joyful struggle that is possible when one learns in community with support.



Rotem Landesman, Ph.D.

Youth want agency to philosophize about AI, and to self-govern their learning and social spaces

Jayne Everson, Rotem Landesman, Janet Jiang, Megumi Kivuva, Eman Sherif, Amy J. Ko (2026). Student Experiences of Joyful Secondary CS Classrooms. ACM International Computing Education Research Conference (ICER)

- My former Ph.D. student Rotem, now a postdoc at Harvard's Graduate School of Education, studied youth's ethical reasoning about computing, including AI.
- In this study, she created spaces for youth to think deeply about their philosophical senses on AI, ethics, and justice
- She found that youth do not want to be told what to think about AI or how to use it
- They want spaces to reason philosophically about what it is, when to use it, and for what purposes
- They wanted agency over when it is brought into their lives, and how it is applied, to be supported in their ethical reasoning

In review.



Janet Jiang

Youth of color marginalized in computing don't want "AI" unless it reflects, sustains, and celebrates their cultures, languages, and values

- This summer, my Ph.D. student Janet Jiang and our two undergraduates Pluto Hassan and Khayri White wondered in what ways youth are seeing or not seeing themselves in computing and "AI," particularly through the lens of culture, race, ethnicity, and gender.
- They taught us that they don't trust who creates AI to create it in their best interests, or in ways that represent their goals, their communities, their languages, or their cultures.
- That "AI" is being used to empower pedophiles to harass and hunt them online in Roblox, and they want adults to find ways to stop that, even if it means using "AI" to detect and block
- That "AI" is creating content and social contexts in their games, online communities, Discords, and group chats, eroding their networks, erasing their cultures, and turning their spaces for connection into spaces for advertising and slop.
- They want computing to be something that reflects the beautiful diversity of the world's cultures and languages, and have no interest in computing or cultures of learning that center the erasure of these worlds.
- And they strongly believe that generative AI is participating in that erasure.

*Stefania Druga, Amy J. Ko (2025).
Scratch Copilot: Supporting
Youth Creative Coding with AI .
ACM Interaction Design for
Children*



Stefania Druga, Ph.D.

**Youth want “AI” that is constrained,
bounded, and focused on their growth**

- With my former Ph.D. Stefania Druga, we tried to imagine ways of framing generative AI as a creativity support tool, helping youth explore possible things to make, possible reasons code isn't working, and possible ways to problem solve.
- Children showed us that they wanted generative AI that was bounded, and focused on their growth, not something designed to replace them or do the work for them.
- They sought empowerment and agency, not productivity, to make things that were uniquely them, not someone else.
- To wanted "AI" to be just another tool, not the center of their experience, something adults helped them learn to use with discretion.



Adrienne Gifford, middle school teacher

*Amy J. Ko, Carlos Aldana
Lira, Isabel
Amaya (2025). Wordplay:
Accessible, Multilingual
Interactive Typography. ACM
SIGCHI Conference on Human
Factors in Computing Systems
(CHI)*

**Youth want programmable media that
reflects their uniqueness and differences, and
only want AI when it supports their agency**

- With middle school teacher Adrienne Gifford, we have been co-designing educational programming languages
- Generative AI comes up a lot, but mostly as something they *don't* want, at least in the form in which industry creates it.
- They don't want it to code for them, they don't want it to supplant their creativity
- They do want it to help them learn, to empower them
- They do want it to help them collaborate across natural languages, to share their cultures
- They do want to be in charge of when and how it is used, not teachers, and certainly not a handful of Silicon Valley companies.

In review.



Megumi Kivuva

Refugee and immigrant youth want “AI” that empowers them to stay safe and well online and to understand how it is shaping their communities for better and worse

- With my Ph.D. student Megumi Kivuva, we have worked with ReWA, a refugee resettlement agency in Seattle, to imagine futures of learning about digital technologies.
- Youth come from all over the world, speak all the world's languages.
- In futures workshops imagining computing education, they shared dreams of learning that empowers them to stay safe and well online, to understand what algorithms and "AI" can and cannot do, to see how computing is connected to their other learning and to the opportunities in their communities.
- They learn none of these things in school now, but we've partnered with Washington's State Board of Education to shape new K-12 graduation competencies that might, trying to make their voices heard.



Jaemarie Solyst, Ph.D.

Jaemarie Solyst, R. Ben Shapiro, Amy J. Ko (2026). Visions of Disciplinary and Education Futures with Generative AI. AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society

Faculty want time and resources to find constructive uses in research teaching, and support in refusal

- With colleague Ben Shapiro and postdoc Jaemarie Solyst, we've examined how faculty are reacting at our own university across disciplines, interviewing 20 disciplines including science, humanities, social sciences, engineering, and the arts.
- We find that they faculty either have no capacity to react, or are reacting with skepticism, or see many possibilities, but no resources to explore them.
- They shared that even if we wanted to proactively shape the futures of our disciplines and learning about them, our institutions are not prepared to imagine or create those futures, because they are not creating spaces for shaping this visioning.
- Especially not now, as the U.S. federal government, and broader demographic cliffs, threaten the sustainability of our institutions.



- So youth do not want "AI" to replace them.
- They view "AI" as a capitalist project of displacement, a project that adults are building to enrich themselves, rather than create the sustainable, loving, connected world they know they need.
- And if "AI" is to be part of their lives, they want it to be designed on their terms, not others'.
- They wish adults were their partners in realizing those worlds, and that our governments were our partners in creating them.
- And many faculty, maybe even some of you in this room, what to be their partners in that work, if only you had time.


Might “AI” catalyze liberatory learning?

- Youth have more agency than ever to avoid learning they don’t like
- This is a loss of institutional power, as educators have less power to control them than ever
- How do we create learning spaces that youth, young adults, and adults *want* to participate in, if they don’t want to participate in our 18th century industrialized forms of learning, and use “AI” to avoid it?

- What I see in all of this is a consistent desire for agency, for space to act with criticality and intentionality, and for liberation from the suffocating pressures of capitalism in the ways that it warps their attention, their goals, and their values.
- And I think "AI", despite all its real and speculated threats, might be a catalyst that triggers that liberation.
- After all, if youth embrace their agency, “AI” is a tool that they can use to guide their own learning, and a tool they are using to avoid learning they don't want to do, when they are freed and encouraged to resist
- This is a loss of institutional power — including our power at colleges and universities — to impose particular forms and topics of learning onto them.
- This kind of agency is not new: dropping out of college to learn a trade, avoiding class to learn something else, “cheating” on a homework assignment because the content is boring. hooks and Freire wrote about this extensively as a kind of liberatory view of education.
- Using "AI", like using the internet or a book, or even a different community than a school is a choice someone is making about what they value.
- That means that the conversation is perhaps not about "AI" or not, harm or benefit, cognitive gain or loss, but rather, how we create learning spaces for youth, adults, and society more broadly that help learners decide what to learn and why, and with what tools, including potentially with or about “AI”.

A means to know end?

- I think for myself, I've tried to be quite conscious in what I do and don't use "AI" for, as an exercise in trying to imagine the best case scenario.
- I've found uses for it: it's quite good at helping me find all of the quite overdue Associate Editors for the journal I edit, and prioritize who I nag.
- It's great at implementing my detailed design specifications for software projects I work on, which is helpful, because I have very little time for the intensive focus programming requires. But I still read every line it writes, and critique it's work heavily, until it takes a proper form.
- And it's great at helping me search the web more thoroughly, as long as I'm willing to carefully guide and check its work, and put faith in its ability to retrieve but not interpret, verify, or sense make.
- Just like the emerging research shows, this requires immense self-regulation, which I have, hard won after many decades. But even those are a modest win, and at great cost, of carbon output, enriching Silicon Valley at my own expense, and probably contributing to labor exploitation, theft, and land use catastrophes.
- And so while I can see how it moves me faster as an individual, I cannot see any reason for it to be the center of my attention.
- The problems and people that need my help are the ones that empower the public to shape the broader sociopolitical futures that got us into this mess: transforming schools, transforming teaching and learning, and transforming computing into something that is easier to understand and use for prosocial ends.
- I titled this talk with something intentionally ambiguous.
 - Am I saying that AI is a means to know?
 - That it's the beginning of the end of knowledge?
 - That its potential is unlimited?
 - Or that its benefits do not justify its harms?
- I think I'm saying all of those things, but most of all, that it is all of us who should decide, together.



Let's chat

Information technology rarely has the outcomes we predict in learning and life

“AI” is not built for any particular purpose, and so much of how it shapes learning depends on how someone uses it

Most of us don't have the learning spaces we want, and until we do, “AI” is going to be a tool for surviving learning spaces that are hostile to agency

Amy J. Ko, Ph.D.

- Thank you, let's chat.