



Generative AI Is Doomed

While light-handed regulation allowed the Internet to bloom and boom in the 1990s, the developing push for strong regulation of generative artificial intelligence will stifle the new technology's development.

BY ERIC GOLDMAN

Dire predictions are *scènes à faire* of every talk about artificial intelligence. But I deviate from the standard predictions of misanthropic and murderous AI. Instead, I want to talk about the precarious future of generative AI—what regrettably may be done *to* it, rather than by it.

We are at the beginning of what I'll call a regulatory tsunami, a metaphorical overproduction of regulation governing generative AI. I believe the regulatory tsunami will overwhelm the innovative trajectory of the technology and possibly wipe it out entirely.

I'll proceed in three parts. First, I'll define generative AI as a subset of "artificial intelligence" (or AI). This will include an explanation of how we're in the beginning stages of a new era of innovation and communications. Second, I'll explain how and why regulators approached the early Internet differently from the recent ways they have been approaching generative AI. The consequence is that generative AI will face more severe regulatory treatment than the early Internet did. Finally, I'll set forth the implications of the regulatory tsunami for the future of generative AI and all of us.



This is an edited version of the Nies Lecture on Intellectual Property delivered at Marquette Law School on April 16, 2024, by Eric Goldman, professor of law, associate dean for research, and co-director of the High Tech Law Institute at Santa Clara University's School of Law in California. From 2002 to 2006, Goldman was an assistant professor of law at Marquette University.

The illustration on the facing page was produced by a generative AI program, DALL-E, which was given instructions to act on the concept of a massive tidal wave, made entirely of legal documents, each clearly recognizable, with text resembling complex legalese. See page 19 for a description of how this image and the other illustrations accompanying this article were generated.

THE GENERATIVE AI EPIPHANY

First, I should define what I mean by “generative AI.” ChatGPT told me:

Generative AI refers to a subset of artificial intelligence (AI) techniques that involve creating or generating new data, content, or outputs that mimic human creativity or problem-solving abilities. Unlike traditional AI systems that are based on rules or predefined responses, generative AI models are trained on large datasets and are capable of producing novel outputs that are not explicitly programmed.

These “novel outputs” can include text, software code, images, audio, video, and other content types. With respect to text outputs, people increasingly have firsthand experience with “chatbots” that engage in polite and friendly banter—a lot like chatting with another human. The output quality often depends on the quality of the user prompts.

Many of us have a lot to learn about “prompt engineering.” Increasingly, it will become more valuable to know how to ask the right question than to know the right answer.

Generative AI is just one type of AI. We most often associate “artificial intelligence” with “general artificial intelligence,” which can think for itself, and the AI that controls physical machines, such as murderous drones. Generative AI poses many risks to society, but for now it’s still under human control and is not inherently murderous.

Generative AI may feel new, but automated content generation assistance has been a part of everyday tools such as Adobe Photoshop, Google Search (with its predictive autocomplete feature), and Gmail (with its predictive replies). Nevertheless, generative AI experienced a “moment” in 2022, when OpenAI made available the DALL-E 2 image generator and ChatGPT 3.5 chatbot, both of which are easy-to-use and

continued on page 20



Generative AI and This Article's Illustrations

The illustrations accompanying this article are among 35 images that Professor Eric Goldman prepared for the Nies Lecture. All of them were created by use of generative AI. Professor Goldman explained that Jess Miers, a former student of his who taught Santa Clara's AI and the Law course, did most of the work. Goldman gave her the concepts he had in mind or goals he wanted to accomplish. Miers prompt-engineered each concept with help from Google's Gemini and then used DALL-E to generate the images (often with further instructions).

During the lecture, Professor Goldman elaborated in part as follows: "Not all of the images achieved my vision, but I have kept those to show some challenges and limits of producing generative AI outputs. In addition, Jess Miers sometimes overrode DALL-E's tendency to depict people as white and male, but several images still display DALL-E's bias."

Here is an example of Goldman's goals and Miers's prompts, in this instance leading ultimately to the image on page 18 (facing); from this prompt and the corresponding illustration, the reader/viewer will see, among other things, that some of the details noted or requested do not appear in the output art:

Goal: A crowd of people celebrating robots in an old medieval town square.

Subject: Medieval town square celebrates robots.

Setting: A bustling medieval town square with cobblestone streets and half-timbered buildings adorned with colorful banners.

Crowd: A large and diverse crowd of townsfolk fills the square, overflowing with excitement and celebration. Include people from various social classes—merchants in colorful clothing, nobles in finery, blacksmiths with soot-stained faces, and children wide-eyed with wonder. Some people hold aloft signs reading "Robots for Progress!" or "Hail the Age of Invention!" in a medieval-style font.

Robots: The center of attention is a group of three robots, each unique in design but adhering to a fantastical, steampunk aesthetic. A towering automaton with polished brass plating and glowing gears could be raising its metallic arm in a gesture of greeting. A smaller, nimble robot with intricate copper gears and wires

could be perched atop a fountain, interacting with a curious child. Consider including a robotic jester with mismatched parts and glowing eyes, entertaining the crowd with playful antics.

Celebration: Flags, balloons, and ribbons decorate the square, adding to the festive atmosphere. Musicians play lively jigs and folk songs on traditional instruments. Some townsfolk, inspired by the robots, wear makeshift robot costumes made from metal scraps and leather, showing their admiration.

Background details: A large, ornate clock tower looms over the square, its face displaying the Roman numerals. In the background, bustling shops with open windows showcase wares relevant to the celebration, perhaps miniature wind-up automatons or clockwork toys.

Overall tone: The image should be a joyous and vibrant celebration of innovation. The townsfolk and robots should exude a sense of wonder and mutual respect, marking a new era of technological advancement within the medieval world.

We are at the beginning of what I'll call a regulatory tsunami, a metaphorical overproduction of regulation governing generative AI. I believe the regulatory tsunami will overwhelm the innovative trajectory of the technology and possibly wipe it out entirely.

continued from page 18

powerful tools that create high-quality outputs. This moment increased public awareness of generative AI and attracted millions of new users. I call those 2022 developments the “Generative AI Epiphany.” (Cf. the illustration on page 18 and description on page 19.)

At the same time, those developments exacerbated fears about technology generally and AI specifically. Teachers worried about the integrity of their graded assignments due to possible “cheating” by students; employees wondered if generative AI would moot their jobs or replace the employees; and misinformation researchers predicted a flood of junk content online.

The Generative AI Epiphany was quickly followed by major technological launches from additional generative AI vendors and a flood of investment capital into the generative AI space. I believe that historians will denote the Generative AI Epiphany as a key turning point in human-machine interactions.

I want to disentangle two types of generative AI activities.

First, generative AI can manufacture new expressive and functional works. I call this the *content generation function* of Generative AI. There are virtually limitless ways that the content generation function can produce better outputs at lower costs than humans can. For example, coders use generative AI to produce and debug code instantly. Job seekers, including our students, use generative AI to write better cover letters and improve their employment prospects. Lawyers can use generative AI to assemble first drafts of legal documents such as contracts and legal briefs, although they have to carefully check for “hallucinations,” as in the invention of cases or of precedents that do not exist. (Cf. the illustration on page 22, acting on the concept of robots manufacturing information in steampunk style.)

Consistent with these examples and many others, generative AI has revealed how many routine content-production tasks could be outsourced, at least in part, to machines. This outsourcing process reduces the total labor required to produce those types of outputs, just as e-discovery has replaced significant chunks of work previously performed by junior attorneys and robots have replaced humans on manufacturing assembly lines. These changes in content production will displace workers and have distributional and psychological consequences, though it's impossible fully to anticipate these effects today. Either way, we will have to rethink how we educate students and workers to provide the skills needed by future employers.

These content-production transitions are not inherently bad for society. Indeed, they may produce some important benefits. For example, where the machines have a relative advantage over humans at content production, humans can redirect their efforts to other aspects of the content-production process. This is the same basic dynamic that occurred with the automation of labor-intensive industries such as manufacturing and agriculture, where competitive differentiation increasingly comes from knowing what to do. This creates new opportunities for creativity and innovation.

Furthermore, generative AI opens up content production to individuals who were previously excluded. For example, using generative AI, non-experts can easily produce the first drafts of high-quality outputs that previously could be prepared only by experts after significant training and practice. (Those drafts will still need expert quality control.)

The images used as part of this presentation are a microcosm of the world with lower content-production barriers. I don't have any artistic talent at all. In the pre-generative-AI world, I

would never have contemplated custom-produced images for this purpose. However, with the help of generative AI and my colleague Jess Miers, this presentation has been enriched in ways that weren't possible before.

More generally, generative AI can turn non-artists into artists; non-writers into writers; and non-coders into coders. Because the universe of non-experts is so much larger than the class of experts, these expanded labor pools have the potential to produce creative and innovative outputs that historically have been foreclosed by the steep hurdles to acquiring the requisite expertise.

Generative AI also helps people research and understand topics of interest. I'll call this the *research function* of Generative AI. Generative AI can surface insights and resources that wouldn't necessarily show up through traditional keyword searches. This functionality makes generative AI an important complement to the search engines we use daily. Indeed, generative AI has exposed some limitations of keyword searching. In response, both Google and Bing supplemented their keyword search functionality with generative AI options. Generative AI will likely play a critical role in our research and discovery processes for the foreseeable future, and that could produce important society-wide benefits.

I distinguish between the content generation function and the research function of Generative AI because they offer different benefits to users and users will expect different types of outputs, depending on their goals. We might also make legal and ethical distinctions between the two functions. When users are researching, they can be harmed if the outputs contain errors or tortious or illegal material and they rely upon the outputs uncritically. However, we could imagine putting the legal and ethical responsibility on users to critically evaluate the outputs of generative AI. When users are generating content, it is even more obvious that they can review and modify the outputs before amplifying them, including fixing any errors or legal problems. With respect to either function, users have significant agency in how they evaluate and disseminate the outputs, and that agency has potential legal significance.

Despite the baggage associated with the "AI" nomenclature, generative AI's content generation function and research function both lead to the production of content that could be legally characterized as content "publication." We're not

exactly sure how to sort through this issue because at least four legal paradigms could apply:

- One paradigm is that the generative AI model publishes content to prompting users, which would expose the operator of the model to standard publishing liability, while triggering constitutional protections under the First Amendment's free speech and free press clauses.
- In a second paradigm, the generative AI model executes the prompting users' instructions to generate the requested content, in which case the user would bear some or all of the resulting liability.
- A third paradigm is to treat the generative AI model and prompting users as some kind of co-creators, a hybrid status without clear legal precedent.
- A fourth and final paradigm is that neither the generative AI model nor the prompting user is responsible for the output—a type of "deus ex machina" origin story for the content.


SIMILARITIES AND DIFFERENCES IN HOW THE INTERNET AND GENERATIVE AI WERE RECEIVED

I'll now turn to my feeling of déjà vu watching the Generative AI Epiphany. I was fortunate to have a front-row seat for the dot-com boom starting around 1994. There was a lot of excitement about the Internet's potential to benefit humanity. As Congress said in the Communications Decency Act of 1996 (Section 230), "The rapidly developing array of Internet and other interactive computer services available to individual Americans represent an extraordinary advance in the availability of educational and informational resources to our citizens."

It might be impossible to imagine today, but 1990s regulators often took a deferential and generally hands-off approach to the new technology. This stance was fueled by prevailing concerns that overly aggressive regulatory responses could distort or harm the emergence of this important innovation. As Congress said in 1996, its policy was "to preserve the vibrant and competitive free market that presently exists for the Internet and other interactive computer services, unfettered by Federal or State regulation."

It was a remarkable and exceptional phase of

[G]enerative AI can turn non-artists into artists; non-writers into writers; and non-coders into coders.



regulatory humility. In the mid-1990s, regulators could not anticipate or predict all of the Internet's uses that have emerged over the last three decades—or how those developments have benefited society. Had regulators hard-coded their limited and myopic 1990s conceptions of the Internet into law, the Internet never could have achieved many of the eventual outcomes, and I think the world would be poorer for it. But mid-1990s regulators frequently admitted their myopia and unusually chose regulatory forbearance.

Generative AI will not get a similar reception from regulators. Regulators are intervening now, acting on their unenlightened 2020s conceptions of what generative AI does. Because we can't anticipate what generative AI is capable of and how new innovative uses will emerge over time, the interventions taking place today will unavoidably restrict generative AI's potential upside.

As a powerful example of regulator naivete, consider the European Union's recently adopted Artificial Intelligence Act. In 2022, on the eve of the Generative AI Epiphany, the EU AI Act's near-final draft didn't contemplate generative AI at all. The drafters scrambled to fix this major defect, but they had been working for years effectively blind to a crucial and imminent new development in generative AI. They literally couldn't see what was right around the corner. History will surely expose other flaws in the AI Act.

We're not just looking at a few mis-crafted laws here or there. The regulation will come as a tsunami. According to the Business Software Association, state legislatures introduced more than 400 AI-related bills in the first 38 days of 2024—six times as many as had been introduced in that period in 2023. Not all of those bills will pass, but some already have been enacted, and more are coming. Regulators are “flooding the zone” of AI regulation now, and each new bill threatens generative AI's innovation arc.

I've been trying to rationalize the disparity between the 1990s regulatory deference to the Internet and the 2020s regulatory tsunami crashing down on generative AI. I've come up with four hypotheses, which are not mutually exclusive.

1. Media depictions. My first hypothesis is that the Internet and generative AI have different reputations because of their dichotomous treatments in the media before their

popularization. In the 1990s, people hadn't really considered a technology like the Internet. It rarely appeared in older works of science fiction. More typically, sci-fi stories turned on information scarcity. This meant that the Internet hadn't been featured in dystopian storylines before it burst into the public consciousness. In contrast, “AI” has been the subject of dystopian books and films for decades, and those depictions have socialized us to view AI as a threat to humans. Indeed, if an older movie depicted AI, it invariably turned murderous. For these purposes, it doesn't matter that “generative AI” isn't autonomous at all. So long as it's characterized as “AI,” we have been conditioned to fear it. For this reason, generative AI never got a honeymoon period. (Cf. the illustration on page 26, acting on the concept of people getting ready to fight a robot army coming at them.)

2. Techno-optimism vs. Techlash. The second hypothesis explaining the dichotomous response is similar to the media conditioning point but is based on broader public attitudes about technology. The 1990s were generally a time of techno-optimism—a celebration of how ingenuity and innovation could improve society. The Internet was welcomed as new, important, and beneficial. In contrast, the Generative AI Epiphany occurred during a strong wave of anti-technology sentiments, sometimes called the “techlash” (a portmanteau of *technology* and *backlash*). Many people today are angry about technology generally and “Big Tech” specifically. They blame technology for many of our social ills. Thus, even before people understood generative AI, they were primed to dislike it—especially because several Big Tech companies are associated with the highest-profile generative AI initiatives.

3. Partisanship. A third hypothesis to explain the dichotomous treatment is the current environment of heightened partisanship. Partisanship isn't new, and partisan rancor was common in the 1990s. Nonetheless, the Internet astonishingly avoided that partisan fray for many years. It was only in the past decade that concerns about online political bias became mainstream. Now that everyone is on heightened alert for partisan bias online, however, there will be accusations of political

Had regulators hard-coded their limited and myopic 1990s conceptions of the Internet into law, the Internet never could have achieved many of the eventual outcomes, and I think the world would be poorer for it.

Illustration on page 22 (facing) is based on the concept of robots manufacturing information in steampunk style.

[W]ell-funded incumbents often view regulatory compliance as just another cost of doing business. Thus, incumbents can use industry regulation to raise new entrants' costs and deepen their own competitive moat.

bias against every new content-publication technology, bolstered by anecdotal evidence. This isn't because the tools are actually biased in partisan ways. The publication process necessarily prioritizes some content over others, which inevitably will create anecdotes of bias for critics to cherry-pick evidence if they disregard proper scientific methods. Generative AI is producing such anecdotes, and partisan advocates are seizing on them to push for censorial interventions to favor their team and disfavor the rival team. This partisan buzzsaw poses an existential threat to generative AI (and the modern Internet as well).

4. Incumbents. A fourth and final hypothesis for the disparity: the differences in incumbency. When the Internet gained public awareness, it lacked incumbent players. In the mid-1990s, the leading online players, including websites and commercial online services such as AOL, were relatively small. The biggest commercial players at the time were the telecommunications services that offered Internet access. The telcos weren't interested in reducing competition; they liked the steady flow of new paying customers. The Generative AI Epiphany, in contrast, was triggered by a giant "startup," OpenAI, with more than \$13 billion raised and a valuation of around \$80 billion. Other Big Tech giants, such as Microsoft, Google, Facebook, and Amazon, are also aggressively spending huge amounts on generative AI. I'm reminded of the sport of polo's nickname, the "Sport of Kings," because you need to be royalty to afford it.

Before I turn to the general implications of the differences between the 1990s Internet and the Generative AI Epiphany, I want to make two specific points about the implications of having large incumbents as the generative AI industry proliferates.

First, unlike start-up companies, incumbents don't always oppose regulation, and sometimes they favor it. That's because well-funded incumbents often view regulatory compliance as just another cost of doing business. Thus, incumbents can use industry regulation to raise new entrants' costs and deepen their own competitive moat.

You might think that generative AI incumbents have the muscle to push back on regulators and protect the industry from overregulation in ways

that the early Internet players couldn't do, but that assumes they oppose regulation. Instead, for example, OpenAI has openly called for the increased regulation of generative AI. This move doesn't prove that such regulations are wise or in the public interest. More likely, it is an incumbent's effort to hinder its competitors. Many regulators will happily support such requests, even when the regulators are being played.

Second, the presence of generative AI incumbents increases the likelihood that they will embrace a content-licensing scheme that diverges from the Internet's model.

When the commercial Internet launched, it was widely assumed that the Internet would evolve into something like cable TV. Consumers would pay subscription fees to access "walled gardens" of content online. What happened instead is a bit of a miracle. Instead of paywalls, Web 2.0, as it's called, ushered in a quarter-century-long era of "user-generated content" or UGC. Typically, UGC services obtain user content without paying licensing fees, and in turn they can deploy a wide range of nonsubscription business models.

It's likely that generative AI could index third-party content without securing copyright permissions or paying licensing fees. A leading precedent indicates that services don't infringe when they index third-party content and use it to create non-infringing outputs, which is usually the case with generative AI.

The generative AI incumbents may nevertheless accept, or even prefer, an industry standard of licensing fees to index third-party content. Incumbents may consider licensing fees, like regulatory compliance costs, to be another cost of doing business. Thus, increasing those costs again acts as an entry barrier to rivals. OpenAI has already indicated an openness to license content from major copyright owners. If OpenAI goes in that direction and its competitors don't follow suit, OpenAI will undoubtedly endorse copyright reforms that impose licensing costs on its competitors.

* * * *

While I haven't identified a single causal explanation for the disparity in regulatory responses to the Internet and generative AI, I reiterate my descriptive claim. We're experiencing an epochal shift in technology, on the order of magnitude of the Internet's commercialization—but

this time, the regulators are intervening early, in a massive and unrelenting way.

THE IMPLICATIONS OF A REGULATORY TSUNAMI

Let me turn to the final part of my talk, which anticipates how the regulatory tsunami will shape the future of generative AI. I already gave you a preview of the bottom line in the title: “Generative AI Is Doomed.” I want, more specifically here, to predict how and why regulators will kill off generative AI; discuss how those interventions will have negative consequences for all of us; and conclude with a few unrealistic ideas of how we might hypothetically avoid these losses, even though the regulatory tides cannot be stopped.

HOW REGULATORS WILL KILL OFF GENERATIVE AI

Earlier, I explained four differences between the 1990s Internet and the 2022 Generative AI Epiphany. Let me now mention a fifth difference: the U.S. regulatory context.

In 1996, shortly after the Internet’s commercial launch, Congress enacted what we know as 47 U.S.C. § 230, an extraordinarily powerful legal immunity. In essence, Section 230 says that websites aren’t liable for third-party content. This simple but elegant legal principle reflected the 1990s technolibertarian ethos of Internet “exceptionalism”—the mindset that the Internet is unique and therefore requires different legal treatment from other media. Section 230 provided the legal foundation for Web 2.0 and established the primacy of user-generated content online. For example, for many years, the most popular U.S. websites have heavily relied on user-generated content—and, by implication, on Section 230.

Furthermore, state lawmakers were discouraged from passing laws targeting UGC websites because those laws were likely to be preempted by Section 230. Thus, Section 230 helped prevent a 1990s regulatory tsunami against the Internet.

In contrast, Congress has not enacted a generative AI-specific law analogous to Section 230. So generative AI won’t get the same regulation-dampening effect that Section 230 gave the Internet.

In addition to Section 230, the Internet benefited from favorable constitutional interpretations by the Supreme Court. In 1997, the Supreme Court declared that the “Internet is a unique and wholly

new medium of worldwide human communication” and thus deserved the highest levels of First Amendment protections. This ruling surely discouraged some 1990s lawmakers from pursuing Internet regulation.

The Supreme Court might conclude that generative AI qualifies for equally vigorous constitutional protection, but I wouldn’t bet on it. For one thing, as I’ve mentioned, there is the unresolved question about who “creates” generative AI outputs—the query submitter, the model-maker, both, neither, someone else? This question is hotly debated and will surely affect the First Amendment’s application. Until there is a decisive answer, regulators will be inclined to disregard the First Amendment entirely. Indeed, many pending and passed bills regulating generative AI are justifiable only if the First Amendment doesn’t apply. For another thing, the 2020s Supreme Court may not apply the Constitution as vigorously even to the Internet itself as it did in the 1990s. The Supreme Court’s composition has changed; it’s shown itself willing to revisit longstanding precedent; and even the Court can’t ignore the negative public attitude toward the Internet. And given the ambiguities over the agency question, generative AI probably won’t get more favorable constitutional treatment than Internet publishing.

If generative AI doesn’t benefit from liability shields such as Section 230 or the Constitution, regulators have a virtually limitless set of options to dictate every aspect of generative AI’s functions. I’ll mention at least three particularly troublesome regulatory archetypes they will pursue:

- 1. Ignorant regulations.** Regulators will enact laws that misunderstand the technology or are driven by moral panics instead of the facts.
- 2. Censorial regulations.** Without strong First Amendment protections for generative AI, regulators will seek to control and censor outputs to favor their preferred narratives. We can preview this process from recent state efforts to regulate the Internet. Despite the First Amendment and Section 230, regulators nevertheless are actively seeking to dictate every aspect of Internet services’ editorial discretion and operations. Those efforts might fail in court. However, if generative AI never receives strong constitutional protection, regulators will embrace the most invasive and censorial approaches.
- 3. Partisan regulations.** One particularly

Without strong First Amendment protections for generative AI, regulators will seek to control and censor outputs to favor their preferred narratives.



pernicious form of censorship would be to steer generative AI outputs for partisan motivations. Outside of the generative AI context, we're already seeing widespread regulatory efforts to control public discussions on partisanized topics, such as vaccines, transgender issues, and abortion. All of those culture wars will hit generative AI hard, especially if there's only a weak constitutional shield.

In addition to the dumb, censorial, and partisan attacks on generative AI, the generative AI industry will be overwhelmed by the sheer volume of regulation, especially when state laws aren't standardized. Section 230 largely kept states out of regulating the Internet, so Internet services only had to worry about complying with a single federal standard. Without national standards, the compliance costs will compound the Sport of Kings problem.

WHAT WE LOSE BECAUSE OF THE REGULATORY TSUNAMI

To recap: I expect regulators will intervene in every aspect of generative AI's "editorial" decision-making, from the mundane to the fundamental, for reasons that range from the possibly legitimate to clearly illegitimate. These efforts won't be curbed by public opposition, Section 230, or the First Amendment. The regulatory frenzy will have a shocking impact that most of us have rarely seen, especially when it comes to content production: a flood of regulation that will dramatically reshape the generative AI industry—if the industry survives at all.

Earlier, I mentioned the benefits of the content generation and research functions of Generative AI. The regulatory tsunami will eliminate those functions outright or render them useless. We also will lose the potential benefits that would emerge over time, as new applications and innovations build on each other. As is typical when regulators intervene early in technology-development cycles, we'll never know what could have been.

The regulatory tsunami will also create collateral damage beyond just generative AI. For example, regulators struggle to define generative AI without also including algorithms. Misdirected or malicious generative AI regulations will jeopardize all kinds of algorithmic activities, from personalized content to algorithmically sorted search results—things that we rely upon many times a day.

If the generative AI industry does survive the regulatory tsunami, it will likely contain only a small number of large players due to the compliance costs. This has several more downsides:

- Concentrated industries are less innovative and dynamic because fewer new entrants are around to push the giants competitively.
- Costs will be higher because consumers will have fewer choices.
- The incumbents will have so much power that regulators will feel pressure to keep intervening. This creates a negative regulatory-feedback loop. The increased interventions raise costs, further consolidating power among a smaller number of players, which necessitates more regulatory interventions.

IS IT POSSIBLE TO SAVE GENERATIVE AI?

It brings me no joy to deliver a bleak talk, and worst of all, I have no good ideas of how we can achieve a better outcome. Calling more attention to the problem is a start, but it won't move the needle against the decades-long socialization to fear AI or the way incumbents will coopt regulators to erect regulatory barriers.

In a hypothetical timeline, with a different Overton window, Congress might enact statutory immunities for generative AI analogous to Section 230. This would delay the regulatory tsunami and preserve industry dynamism longer. Unfortunately, in the timeline we occupy, the idea that regulators today would take any affirmative step to shield generative AI is ivory-tower fantasy.

Here is one suggestion: It would help to rebrand generative AI to distance it from "AI." If we were to more expressly acknowledge the content generation function and research function of Generative AI, it might reduce public fear and make the Constitution's applicability more obvious.

In all events, I encourage you to critically scrutinize every effort to regulate generative AI. Don't assume that any effort is being advanced for your benefit, or for legitimate reasons, or in a constitutional manner. Once you notice how often such efforts are ill-motivated, you will be better positioned to hold the advocates more accountable. ■

The concept for the illustration on page 26 (facing) is people getting ready to fight a robot army coming at them.