Generative AI is Doomed

By Eric Goldman*

Abstract: I delivered this talk as the 2024 Nies Lecture at Marquette University School of Law, Milwaukee, WI. The talk compares the recent proliferation of Generative AI with the Internet's proliferation in the mid-1990s. In each case, it was clear that the technology would have revolutionary but uncertain impacts on society. However, the public sentiments toward the two innovations have differed radically. The Internet arrived during a period of widespread techno-optimism, creating a regulatory environment that fostered the Internet's growth. Generative AI, in contrast, has arrived during widespread techno-pessimism and following decades of conditioning about the dangers of "AI." The difference is consequential: The prevailing regulatory and legal responses to Generative AI will limit or even negate its benefits. If society hopes to achieve the full potential of Generative AI, we'll need to adopt a new regulatory approach quickly.

Table of Contents

* Professor of Law, Associate Dean for Research, Co-Director of the High Tech Law Institute, and Co-Supervisor of the Privacy Law Certificate, Santa Clara University School of Law. Website: http://www.ericgoldman.org. Email: egoldman@gmail.com. I appreciate the help of Mary Rose Finnigan and Jess Miers in preparing this talk and the associated slides and the feedback from Lisa Goldman, Edward Lee, and Mark Lemley.

I dedicate this talk to my wonderful former students at Marquette Law, who taught me as much or more than I taught them. I am also grateful for how my former students and colleagues were unfailingly gracious about the many mistakes I made as I navigated the learning curve of being a full-time law professor, even when those errors worked to their disadvantage.

Except where noted, all images were rendered using DALL-E (GPT 4), with substantial prompt design assistance from Gemini. Most were created in mid-April 2024. The image set is available at https://digitalcommons.law.scu.edu/cgi/viewcontent.cgi?article=3849&context=historical.

#

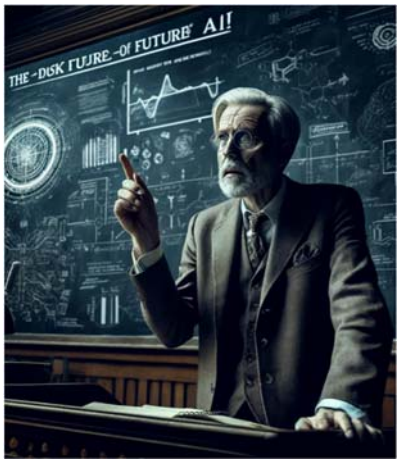

[Image #1: Doom-and-gloom[1]] In selecting my topic, I've taken the ever-popular approach of predicting gloom-and-doom. Dire predictions are scènes à faire of every talk about artificial intelligence, but I deviate from the standard predictions of misanthropic and murderous AI. Instead, I want to talk about the precarious future of Generative AI.

Before I go further, I want to say a few words about the images. As you have probably already guessed, the images accompanying this talk are Generative AI-generated. This project was more than I could handle alone. My former student Jess Miers, who now teaches Santa Clara Law's AI and the Law course, did most of the work. I gave her the concepts I wanted, which I've noted in the footnote for every image. She prompt-engineered the concept with help from Gemini, and then she used DALL-E to generate the images (often with further instructions). Here is an example prompt:

Prompt for Slide 7 [Concept: 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, showcasing 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.

The collection of images is truly a joint production of humans (Jess and me) and machines (DALL-E and Gemini). Not all of the images achieved my vision, but I've kept those to show some challenges and limits of producing Generative AI outputs. Also, Jess sometimes overrode DALL-E's tendency to depict people as white and male, but several images still display DALL-E's bias. I'll say more about the images in a bit.

##########

[1] Concept: a professor predicting gloom-and-doom about the future of AI.

#

[Image #2: a tidal wave of law[2]] Back to my talk overview. 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 severely stunt 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." I'll explain 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 than the recent Generative AI. The consequence is that Generative AI will face more severe regulatory treatment than the early Internet did. Finally, I'll explain the implications of the regulatory tsunami for the future of Generative AI and all of us.

Part 1: The Generative AI Epiphany

[Image #3: DALL-E Selfie[3]] First, I should define what I mean by Generative AI. ChatGPT told me:[4]

> 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 first-hand experience with "chatbots" that engage in polite and friendly banter, like chatting with another human.[5] As suggested by my third slide, 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 rather than to know the right answer.[6]

##########

[2] Concept: a massive tidal wave, made entirely of legal documents. Each document should be clearly recognizable, with text resembling complex legalese.

[3] Concept: draw yourself.

[4] ChatGPT 3.5, using the prompt "what is the definition of generative ai" (Feb. 2024).

[5] Celeste Biever, *ChatGPT Broke the Turing Test — the Race is on for New Ways to Assess AI*, NATURE, July 25, 2023, https://www.nature.com/articles/d41586-023-02361-7.

[6] *Cf.* Mark A. Lemley, *How Generative AI Will Turn Copyright on its Head*, COLUM. SCI. & TECH. L. REV. (forthcoming 2024) ("Increasingly creativity will be lodged in asking the right questions, not in creating the answers").

#

Generative AI is just one type of AI. As I'll discuss below, we most often associate "artificial intelligence" with "general artificial intelligence," which can think for itself, and 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.

[Image #4: welcoming AI[7]] 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 powerful tools that create high-quality outputs. This moment increased public awareness of Generative AI and attracted millions of new users.[8] This talk refers to those 2022 developments as the "Generative AI Epiphany."

[Image #5: fearsome robot[9]] At the same time, it exacerbated many fears about technology generally and AI specifically. Teachers worried about the integrity of their graded assignments due to possible "cheating" by students;[10] employees wondered if Generative AI would moot their jobs or replace them;[11] and misinformation researchers predicted a flood of junk content online.[12]

[Image #6: investors and AI[13]] The Generative AI Epiphany was quickly followed by major technological launches from other Generative AI vendors and a flood of investment capital into the Generative AI space. I believe historians will denote the Generative AI Epiphany as a key turning point in human-machine interactions.

##

[7] Concept: a crowd of people celebrating robots in an old medieval town square.

[8] OpenAI claims 100 million ChatGPT users per week. Jon Porter, *Chatgpt Continues To Be One of the Fastest-Growing Services Ever*, THE VERGE, Nov. 6, 2023, https://www.theverge.com/2023/11/6/23948386/chatgpt-active-user-count-openai-developer-conference.

[9] Concept: people cowering in fear before a malevolent robot.

[10] TIME FOR CLASS 2023: BRIDGING STUDENT AND FACULTY PERSPECTIVES ON DIGITAL LEARNING, TYTON PARTNERS (due to the Generative AI Epiphany, faculty in 2023 reported that "preventing student cheating" was the top challenge facing higher education instructors).

[11] *E.g.*, Josie Cox, *AI Anxiety: The Workers Who Fear Losing their Jobs to Artificial Intelligence*, BBC, July 13, 2023, https://www.bbc.com/worklife/article/20230418-ai-anxiety-artificial-intelligence-replace-jobs.

[12] *E.g.*, Tiffany Hsu & Stuart A. Thompson, *Disinformation Researchers Raise Alarms About A.I. Chatbots*, N.Y. TIMES, Feb. 9, 2023, https://www.nytimes.com/2023/02/08/technology/ai-chatbots-disinformation.html.

[13] Concept: investors shoveling money in the direction of a group of robots.

#

I want to disentangle two types of Generative AI activities.

[Image #7: manufacturing information[14]] 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 do. For example, coders use Generative AI to produce and debug code instantly.[15] Job seekers, including our students, use Generative AI to write better cover letters and improve their employment prospects.[16] Lawyers can use Generative AI to assemble first drafts of legal documents such as contracts and legal briefs, though they have to carefully check for "hallucinations."[17]

[Image #8: work smarter, not harder[18]] Consistent with these examples and many others, Generative AI has exposed how many routine content production tasks could be outsourced, at least in part, to the machines. This outsourcing process reduces the total labor required to produce those types of outputs, just like 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 has distributional and psychological consequences, though it's impossible to fully anticipate these effects today.[19] Either way, we will have to rethink how we educate students and workers to provide the skills needed by future employers.[20]

##

[14] Concept: robots manufacturing information in steampunk style.

[15] Scott Wu, *Introducing Devin, the First AI Software Engineer*, COGNITION, March 12, 2024, https://www.cognition-labs.com/introducing-devin.

[16] Jamie Ding, *If You Outsource One Thing to ChatGPT, Job-Seekers Say This Should Be It*, L.A. TIMES, Feb. 16, 2023, https://www.latimes.com/business/story/2023-02-16/chatgpt-writes-cover-letters-job-seekers.

[17] Ben Zimmer, *'Hallucination': When Chatbots (and People) See What Isn't There,* WALL ST. J., Apr. 20, 2023, https://www.wsj.com/articles/hallucination-when-chatbots-and-people-see-what-isnt-there-91c6c88b.

A number of lawyers have received publicity for filing court documents that include hallucinated citations. *E.g.*, Sara Merken, Another NY Lawyer Faces Discipline After AI Chatbot Invented Case Citation, REUTERS, Jan. 30, 2024, https://www.reuters.com/legal/transactional/another-ny-lawyer-faces-discipline-after-ai-chatbot-invented-case-citation-2024-01-30.

[18] Concept 1: a group of happy and productive employees using robots to get their work done. Concept 2: a group of frustrated and frantic employees using outdated technology.

[19] *See* Sarah Kessler, *The A.I. Revolution Will Change Work. Nobody Agrees How.*, N.Y. TIMES, June 10, 2023, https://www.nytimes.com/2023/06/10/business/ai-jobs-work.html; Will Knight, *No One Actually Knows How AI Will Affect Jobs*, WIRED, Apr. 11, 2024, https://www.wired.com/story/ai-impact-on-work-mary-daly-interview/; Qiwen Huang et al, *The Layoff Generation: How Generative Ai Will Reshape Employment and Labor Markets*, Aug. 7, 2023, https://ssrn.com/abstract=4534294; Ali Zarifhonarvar, *Economics of ChatGPT: A Labor Market View on the Occupational Impact of Artificial Intelligence* (February 7, 2023), https://ssrn.com/abstract=4350925.

[20] *See* Bhaskar Ghosh, H. James Wilson, & Tomas Castagnino, *GenAI Will Change How We Design Jobs. Here's How.*, HARV. BUS. REV., Dec. 5, 2023, https://hbr.org/2023/12/genai-will-change-how-we-design-jobs-heres-how.

#

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 like 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.[21] (Those drafts will still need expert quality control).

[Image #9: professor as artist[22]] The images used in this talk are a microcosm of the lower content production barriers. I don't have any artistic talent at all. In the pre-Generative AI work, I would never have contemplated custom-produced images for this talk. However, with the help of Generative AI and my colleague Jess, this talk 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 acquire the requisite expertise.[23]

[Image #10: searching for gems[24]] 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

##########

[21] Steve Lohr, *How One Tech Skeptic Decided A.I. Might Benefit the Middle Class*, N.Y. TIMES, April 1, 2024, https://www.nytimes.com/2024/04/01/business/ai-tech-economy.html (using Generative AI, "more people can take on some of the work that is now the province of elite, and expensive, experts like doctors, lawyers, software engineers and college professors").

[22] Concept: a professor painting art about the law in the style of the old Dutch masters.

[23] *See*, *e.g.*, Erik Brynjolfsson, Danielle Li, & Lindsey Raymond, *Generative AI at Work*, April 23, 2023, https://danielle-li.github.io/assets/docs/GenerativeAIatWork.pdf.

[24] Concept: a robotic St. Bernard searching in the snow for a lost gem.

#

Generative AI options.[25] 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 Research Function of Generative AI because they offer different benefits to users, and users will expect different types of outputs depending on their goals. Furthermore, we might 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 generate content, it is even more obvious that they can review and modify the outputs before amplifying them,[26] 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.

[Image #11: human/machine cooperation[27]] Despite the baggage associated with the "AI" nomenclature, Generative AI's Content Generation and Research Functions 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.

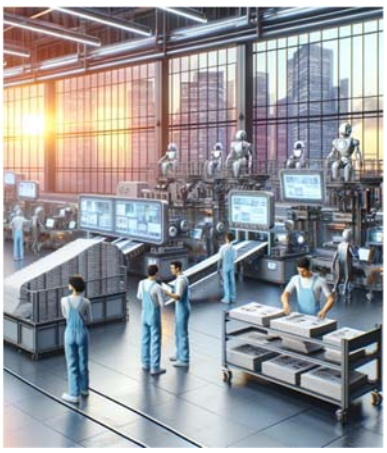

- One paradigm is that the Generative AI model publishes content to prompting users, which would expose the model operator to standard publishing liability while triggering Constitutional protections under the First Amendment's Free Speech and Free Press Clauses.
- A second paradigm is that the Generative 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.[28]

I'll revisit this content generation agency and responsibility issue in Part 3.

##########

[25] The functionality was initially called Google Bard and then transitioned to Google Gemini. Bing's tool is called CoPilot.

[26] This is an example of humans "in the loop." *E.g.*, Rebecca Crootof, Margot E. Kaminski, & W. Nicholson Price II, *Humans in the Loop*, 76 VAND. L. REV. 429 (2023).

[27] Concept: robots and humans working together on a manufacturing line to create a huge pile of information.

[28] While this paradigm may sound weird, it might resemble the legal treatment when animals create copyrighted works (like taking selfies), which courts treat as being owned by no one (i.e., in the public domain). *See* Naruto v. Slater, 888 F.3d 418 (2018); U.S. Copyright Office, Compendium of U.S. Copyright Office Practices § 313.2 (3d ed. 2021), https://www.copyright.gov/comp3/chap300/ch300-copyrightable-authorship.pdf.

#

Part 2: Similarities and Differences Between How Internet and Generative AI Were Received

[Image #12: welcoming the Internet[29]] I'll now turn to the second part of my talk: 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 1996,

> 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.[30]

A note about this image: for this part of the talk, I wanted DALL-E to create an image depicting "people excitedly celebrating the arrival of baby Internet." DALL-E balked at our many attempts to create slides depicting babies, presumably as part of its trust-and-safety efforts. To get around this barrier, I choose Nyan cat as a metaphor for the nascent Internet.

[Image #13: deferential politicians[31]] 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."[32]

It was a remarkable and exceptional phase of regulator humility.[33] 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

##########

[29] Concept: people celebrating the Internet represented by Nyan cat on a pedestal at a large gathering.
[30] 47 U.S.C. §230(a)(1).
[31] Concept: politicians looking bashful around the Internet represented as Nyan cat.
[32] 47 U.S.C. §230(b)(2).
[33] Immediately after I delivered the talk in person, I got several questions from attendees confused by the phrase "regulator humility," because most people have never witnessed this phenomenon in their lifetimes. In the 1990s, legislators sometimes intentionally chose not to manufacture new law for the Internet because they recognized they lacked the foresight and expertise to do so in a helpful manner. Indeed, laws like Section 230 and the Internet Tax Freedom Act (47 U.S.C. §151 note, enacted in 1998) represented proactive decisions by legislators to constrain their future selves. This concept is impossible to imagine today because we virtually never seen regulators exhibit any humility regarding their interventions into technological development cycles.

#

have achieved those outcomes, and I think the world would be poorer for it. But mid-1990s regulators frequently admitted their myopia and unusually chose regulatory forbearance.

[Image #14: carving up AI[34]] 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.

A few notes about this image. DALL-E's trust-and-safety mechanisms don't like depicting knives, though it's hard to depict someone carving a turkey without them. DALL-E also turned the server into a cyborg with robot hands because…? And this is an obvious example of DALL-E's bias towards white males.

[Image #15: myopic politicians[35]] As a powerful example of regulator naivete, consider the European Union's recently-adopted AI Act.[36] In 2022, on the eve of the Generative AI Epiphany, the AI Act's near-final draft didn't contemplate Generative AI at all.[37] 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.

##

[34] Concept: politicians carving up a robotic turkey at a traditional Thanksgiving meal.

[35] Concept: politicians trying to look around a corner but unable to do so.

[36] Artificial Intelligence Act, European Parliament legislative resolution of 13 March 2024 on the proposal for a regulation of the European Parliament and of the Council on laying down harmonized rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union Legislative Acts (COM(2021)0206 – C9-0146/2021 – 2021/0106(COD)), https://www.europarl.europa.eu/doceo/document/TA-9-2024-0138_EN.pdf.

[37] Martin Coulter, *What is the EU AI Act and When Will Regulation Come into Effect?*, REUTERS, Dec. 7, 2023, https://www.reuters.com/technology/what-are-eus-landmark-ai-rules-2023-12-06/.

#

[Image #16: regulator frenzy[38]] We're not just looking at a few miscrafted laws here or there. The regulation will come as a tsunami. According to the Business Software Association, state legislatures introduced over 400 AI-related bills in the first 38 days of 2024[39]—six times as many as had been introduced in that period in 2023. Not all of those bills will pass, but some already have 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.

[Image #17: discovery[40]] 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.[41] This meant that the Internet hadn't been featured in dystopian storylines before it burst into the public consciousness.

[Image #18: murderous robots[42]] 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.[43] Indeed, if an older movie depicted AI, it invariably turned murderous.[44] For these purposes, it doesn't matter that "Generative AI" isn't autonomous at all. So long as

[38] Concept: regulators in DC trying to come up with laws to regulate AI.
[39] *BSA Analysis: States Intensify Work on AI Legislation*, BUSINESS SOFTWARE ASS'N, Feb. 14, 2024, https://www.bsa.org/news-events/news/bsa-analysis-states-intensify-work-on-ai-legislation; *see also* Grant Gross, *The Complex Patchwork of US AI Regulation Has Already Arrived*, CIO, Apr. 5, 2024, https://www.cio.com/article/2081885/the-complex-patchwork-of-us-ai-regulation-has-already-arrived.html.
[40] Concept: miners discovering Nyan cat and being excited about the discovery.
[41] In a rare counterexample, the protagonist of the novel *Ender's Game* shaped political discourse using pseudonymous postings on an electronic network. ORSON SCOTT CARD, ENDER'S GAME (1985).
[42] Concept: robots plotting ways to harm humans.
[43] *See* Kerry Lynn Macintosh, *Rogue Artificial Intelligence, Science Fiction, and the Law* (forthcoming).
[44] Some examples: the movies 2001: A SPACE ODYSSEY (1968) and WARGAMES (1983) and the episode *What Are Little Girls Made Of?*, STAR TREK: THE ORIGINAL SERIES, season 1, episode 7 (1966).

it's characterized as "AI," we have been conditioned to fear it. For this reason, Generative AI never got a honeymoon period.

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.[45]

[Image #19: humans vs. AI[46]] 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").[47] 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.

[Image #20: partisan battles[48]] 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. However, the Internet shockingly avoided that partisan fray for many years. It was only in the past decade that concerns about online political bias became mainstream.

[Image #21: partisan bias[49]] However, now that everyone is on heightened alert for partisan bias online, there will be accusations of political bias against every new content publication technology, bolstered by anecdotal evidence. This isn't because the tools are actually biased. 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

##########

[45] Peter Schwartz & Peter Leyden, *The Long Boom: A History of the Future, 1980–2020*, WIRED, July 1, 1997, https://www.wired.com/1997/07/longboom/.

[46] Concept: people getting ready to fight a robot army coming at them.

[47] *E.g.*, Robert D. Atkinson et al, *A Policymaker's Guide to the "Techlash"—What It Is and Why It's a Threat to Growth and Progress*, INFORMATION TECHNOLOGY & INNOVATION FOUNDATION, Oct. 28, 2019, https://itif.org/publications/2019/10/28/policymakers-guide-techlash/.

[48] Concept: elephants and donkeys preparing to fight each other.

[49] Concept: robot picking elephants and donkeys off a tree.

#

methods. Generative AI is producing such anecdotes,[50] and partisan advocates are seizing on those to advocate 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).

[Image #22: many baby Internets[51]] 4. *Incumbents*. I now address my 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 like 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. As a result, the Internet industry in the 1990s was strongly anti-regulation.

[Image #23: AI incumbents[52]] In contrast, the Generative AI Epiphany was triggered by a giant "startup," Open AI, with over $13 billion raised and a valuation of around $80 billion.[53] Other "Big Tech" giants, such as Microsoft, Google, Facebook, and Amazon, are also aggressively spending huge amounts on Generative AI.[54] I'm reminded of the sport of polo's nickname, the "Sport of Kings," because you need to be royalty to afford it.[55]

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

##########

[50] Jochen Hartmann, Jasper Schwenzow, & Maximilian Witte, *The Political Ideology of Conversational AI: Converging Evidence on ChatGPT's Pro-Environmental, Left-Libertarian Orientation*, ARVIX, Jan. 5, 2023, https://arxiv.org/pdf/2301.01768.pdf; Kate Knibbs, *Most Top News Sites Block AI Bots. Right-Wing Media Welcomes Them*, WIRED, Jan. 24, 2024, https://www.wired.com/story/most-news-sites-block-ai-bots-right-wing-media-welcomes-them/.

[51] Concept: show the Internet represented as baby Nyan cats in a crib.

[52] Concept: 3 giant robots towering over farmland.

[53] Cade Metz & Tripp Mickle, *OpenAI Completes Deal That Values the Company at $80 Billion*, N.Y. TIMES, Feb. 16, 2024, https://www.nytimes.com/2024/02/16/technology/openai-artificial-intelligence-deal-valuation.html.

[54] They are also buying up start-ups who might compete with them, furthering locking in their position. *See* Mark A. Lemley & Matthew Wansley, *Coopting Disruption* (Feb. 1, 2024), https://ssrn.com/abstract=4713845.

[55] *Cf.* Casey Newton, *The indie AI companies are falling apart: From Stability to Inflection, would-be challengers to the giants are failing to execute*, PLATFORMER, March 26, 2024, https://www.platformer.news/indie-ai-collapse-stability-ai-mostaque-inflection-suleyman/ (the smaller AI companies can't keep up with the incumbents); Kenrick Cai, *How Stability AI's Founder Tanked His Billion-Dollar Startup*, FORBES, March 29, 2024, https://www.forbes.com/sites/kenrickcai/2024/03/29/how-stability-ais-founder-tanked-his-billion-dollar-startup/?sh=6f0dc1143e63; Katherine Bindley, *The Fight for AI Talent: Pay Million-Dollar Packages and Buy Whole Teams*, WALL ST. J., Mar. 27, 2024, https://www.wsj.com/tech/ai/the-fight-for-ai-talent-pay-million-dollar-packages-and-buy-whole-teams-c370de2b.

#

[Image #24: gatekeeping[56]] First, unlike start-up companies, incumbents don't always oppose regulation, and sometimes they favor it.[57] That's because well-funded incumbents often view regulatory compliance costs as just another cost of doing business. Thus, incumbents can use industry regulation to raise new entrants' costs and deepen their 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 it. Instead, OpenAI has openly called for the increased regulation of Generative AI.[58] 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 these requests, even when they are being played.

[Image #25: AI licensing[59]] 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 ushered in a quarter-century-long era of UGC, short for "user-generated content." UGC services typically obtain user content without paying licensing fees, and in turn they can deploy a wide range of non-subscription business models.[60]

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,[61] which is

##

[56] Concept: three giant robots blocking small robots.

[57] *E.g.* @googlepublicpolicy, X, Mar. 7, 2024, https://twitter.com/googlepubpolicy/status/1765823885776228799; Adi Robertson, *Mark Zuckerberg Just Told Congress to Upend the Internet*, VERGE, Oct. 29, 2020, https://www.theverge.com/2020/10/29/21537040/facebook-mark-zuckerberg-section-230-hearing-reform-pact-act-big-tech.

[58] Cecilia Kang, *OpenAI's Sam Altman Urges A.I. Regulation in Senate Hearing*, N.Y. TIMES, May 16, 2023. *But see* Steven Levy, *Tech Leaders Once Cried for AI Regulation. Now the Message Is 'Slow Down'*, WIRED, Apr. 12, 2024, https://www.wired.com/story/tech-ai-regulation-bill/.

[59] Concept: robots negotiating a contract with each other in a wood-paneled office.

[60] Eric Goldman, *The U.K. Online Harms White Paper and the Internet's Cable-ized Future*, 16 OHIO STATE TECH. L.J. 351 (2020).

[61] Authors Guild v. Google 804 F.3d 202 (2nd Cir. 2015).

#

usually the case with Generative AI.[62] (I note the possibility of other legal entanglements, such as trespass-to-chattels and contracts).

[Image #26: gatekeeping (again)] The Generative AI incumbents may nevertheless accept, or even prefer, an industry standard of licensing fees to index third-party content. Like regulatory compliance costs, incumbents may consider licensing fees 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.[63] 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.

Part 3: The Implications of a Regulatory Tsunami

[Image #27: destroying AI[64]] I now turn to the third and final part of my talk, where I'll anticipate how the regulatory tsunami will shape the future of Generative AI. You already got a preview from my talk title: "Generative AI is Doomed."

This part has three subparts. First, I'll discuss how and why regulators will kill off Generative AI. Second, I'll talk about how those interventions will have negative consequences for all of us. Finally, I'll conclude with a few unrealistic ideas of how we

##

[62] For example, the *New York Times* had to engage in severe prompt hacking to force ChatGPT to reproduce some of its content. Complaint, The New York Times Company v. Microsoft Corporation, et. al, 1:23-cv-11195, (S.D.N.Y.) (2023), https://nytco-assets.nytimes.com/2023/12/NYT_Complaint_Dec2023.pdf.

[63] Sahil Patel & Stephanie Palazzolo, *OpenAI Offers Publishers as Little as $1 Million a Year*, THE INFORMATION, Jan. 4, 2024, https://www.theinformation.com/articles/openai-offers-publishers-as-little-as-1-million-a-year.

Similarly, Google's Gemini has struck a licensing deal with Reddit to acquire training data. *See* Anna Tong, Echo Wang and Martin Coulter, *Exclusive: Reddit in AI Content Licensing Deal with Google*, REUTERS, Feb. 22, 2024, https://www.reuters.com/technology/reddit-ai-content-licensing-deal-with-google-sources-say-2024-02-22/.

However, it remains an open question if content licensing will yield enough material to adequately train the Generative AI models. *See*, *e.g.*, Cade Metz et al, *How Tech Giants Cut Corners to Harvest Data for A.I.*, N.Y. TIMES, Apr. 6, 2024, https://www.nytimes.com/2024/04/06/technology/tech-giants-harvest-data-artificial-intelligence.html.

[64] Concept: show a dozen robots crashing into a brick wall, with lots of robot parts all over the ground.

#

might hypothetically avoid these losses, though the regulatory tides cannot be stopped.[65]

Subpart A: How Regulators Will Kill Off Generative AI

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

[Image #28: heroic Section 230[66]] In 1996, shortly after the Internet's commercial launch, Congress enacted 47 U.S.C. Section 230, an extraordinarily powerful legal immunity. In short, 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"—that the Internet is unique and therefore requires different legal treatment than 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, 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.[67] As a result, Generative AI won't get the same regulatory dampening effect that Section 230 provided to the Internet.

[Image #29: shielding the Internet[68]] 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"[69] and thus

##

[65] This is an allusion to the legend that King Canute proved how kings lacked the divine power necessary to stop the tides.

[66] Concept: the number "230" in a heroic fashion, with a cape waving in the wind.

[67] Generative AI model-builders may be able to claim Section 230 protection for some outputs, a question that hasn't been definitively resolved. Peter J. Benson & Valerie C. Brannon, *Section 230 Immunity and Generative Artificial Intelligence*, CONGRESSIONAL RESEARCH SERVICE, Dec. 28, 2023, https://crsreports.congress.gov/product/pdf/LSB/LSB11097; Scott Nover, *Section 230 Won't Be a Savior for Generative AI*, GZERO, Apr. 16, 2024, https://www.gzeromedia.com/gzero-ai/section-230-wont-be-a-savior-for-generative-ai.

[68] Concept: a judge uses an umbrella to protect the baby Internet represented as Nyan cat from a flurry of laws raining down on it.

[69] Reno v. ACLU, 521 U.S. 844, 850 (1997).

#

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. First, as mentioned before, there are unresolved questions about who "creates" Generative AI outputs—the query submitter, the model-maker, both, neither, or someone else? This question is hotly debated[70] and will surely affect the First Amendment's application. Until this question is decisively resolved, regulators will typically disregard the First Amendment entirely. Indeed, many pending and passed bills regulating Generative AI are justifiable only if the First Amendment doesn't apply.

Second, the 2020s Supreme Court may not apply the Constitution as vigorously to the Internet as it did in the 1990s. The Supreme Court's composition has changed; it's willing to revisit long-standing precedent; and even the Supreme Court can't ignore the negative public attitude towards the Internet. Given the ambiguities over the agency question, Generative AI probably won't get more favorable Constitutional treatment than Internet publishing.

[Image #30: swarm of laws[71]] If Generative AI doesn't benefit from liability shields like Section 230 and 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:

First, ignorant regulations. Regulators will pass laws that misunderstand the technology or are driven by moral panics instead of the facts.

Second, 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.

##########

[70] *E.g.*, Eugene Volokh, Mark A. Lemley & Peter Henderson, *Freedom of Speech and AI Output*, 3 J. FREE SPEECH L. 651 (2023), https://www.journaloffreespeechlaw.org/volokhlemleyhenderson.pdf; Peter Salib, *AI Outputs Are Not Protected Speech*, WASH. UNIV. L. REV. (Forthcoming 2024); Dan Burk, *Asemic Defamation, or, the Death of the AI Speaker*, 22 FIRST AMEND. L. REV. __ (2024); Margot Kaminski & Meg Lata Jones, *Constructing AI Speech*, YALE L.J. FORUM (Forthcoming 2024).

[71] Concept: giant robot overwhelmed by a swarm of tiny laws despite the robot using a shield to try to protect itself. This image was created by Google Gemini 1.0 Pro.

#

[Image #31: partisan censorship[72]] *Third, partisan regulations*. One particularly 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. These risks are heightened in "trifecta" states, where a single political party controls both the legislature and the governorship, and thus lack the traditional checks-and-balances on partisan-minded bills. Today, forty states are trifectas (23 Republican, 17 Democrat).[73]

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,[74] 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.

Subpart B: What We Lose from the Regulatory Tsunami

[Image #32: a tidal wave of laws (again)] 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 ranging 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 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 of all types.[75] Misdirected or malicious Generative AI regulations jeopardize all kinds of algorithmic

##########

[72] Concept: elephants and donkeys fighting over a printing press operated by a robot.
[73] https://ballotpedia.org/State_government_trifectas.
[74] 47 USC 230(e)(1).
[75] Bryan Casey & Mark A. Lemley, *You Might Be a Robot*, 105 CORNELL L. REV. 287 (2020).

activities, from personalized content to algorithmically sorted search results—things that we rely upon many times a day.

[Image #33: incumbents flexing[76]] If the Generative AI industry does survive the regulatory tsunami, it will likely only contain a small number of large players due to the compliance costs. This has several more downsides:

1) Concentrated industries are less innovative and dynamic because fewer new entrants are around to push the giants competitively.
2) Costs will be higher because consumers will have fewer choices.
3) 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 into a smaller number of players, which necessitates more regulatory interventions.

Subpart C: Is It Possible to Save Generative AI?

[Image #34: the AI sky is falling[77]] 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 and how incumbents will coopt regulators to erect regulatory barriers.

In a hypothetical timeline, with a different Overton Window,[78] Congress might enact statutory immunities for Generative AI analogous to Section 230. This would delay the regulatory tsunami and preserve industry dynamism longer.[79] Unfortunately, in the timeline we occupy, the idea that regulators today would take any affirmative step to shield Generative AI is ivory-tower fantasy.

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 Functions of Generative AI, it might reduce public fear and make the Constitution's applicability more obvious.

On that front, I encourage you to critically scrutinize every effort to regulate Generative AI. Don't assume those efforts are being advanced for your benefit, or for legitimate reasons, or in a

##########

[76] Concept: three giant robots flexing their muscles in a gym filled with tiny robots exercising around them.
[77] Concept: many robots falling from the sky and crashing into the ground, which is covered with robot parts.
[78] https://www.mackinac.org/OvertonWindow.
[79] Eric Goldman, *Want to Kill Facebook and Google? Preserving Section 230 Is Your Best Hope*, BALKINIZATION, June 2019, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3398631.

#

Constitutional manner. Once you notice how often such efforts are illegitimate, you will be better positioned to hold the advocates more accountable.

#