The Rise of Generative Artificial Intelligence as A Technological System

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Tao Groves

Spring, 2025

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Bryn E. Seabrook, Department of Engineering and Society

The Rise of Generative Artificial Intelligence as A Technological System

In the last few years, generative artificial intelligence (GenAI) has metamorphosed from a technology that professionals rely on to optimize and support infrastructure, products, and academic research (Jordan, 2015) to being synonymous with chat bots and homework help. This transformation entered the public eye in 2022 when OpenAI, the current AI technology leader, released ChatGPT, the first free and public-facing Large Language Model (LLM). However, what may appear to be a single revelatory invention was predated by years of technological and infrastructural development behind the scenes which laid the groundwork for the capabilities seen in today's GenAI applications. Since the release of ChatGPT, LLMs have been applied across a broad range of applications, which in turn has driven unprecedented interest and investment in AI technology. With companies competing to push the limits of AI's capabilities, the focus has predominantly been on scaling model sizes and enhancing performance, often with limited attention to the environmental costs. This rapid pace has exacerbated infrastructure strain, as AI workloads now account for a significant portion of global data center energy use, contributing to broader concerns around sustainability (Patterson et al., 2022; Wu et al., 2022).

All this uncertainty begs the question – what is the driving factor behind the growth of the GenAI industry, and why are companies, investors, and governments so enthusiastic about a technology with no clear path to environmental or economic sustainability? This research involves a sociotechnical exploration of the exigence and trajectory of GenAI with the goal of developing practical recommendations by which engineers can implement these systems sustainably. By analyzing the industry through the lens of technological momentum, a clear picture of its past, present, and future is obtained.

Research Question & Methods

What is the driving factor behind the growth of the GenAI industry, and why are companies, investors, and governments so enthusiastic about a technology with no clear path to environmental or economic sustainability? To understand how public perception and societal needs have shaped the growth of GenAI, how GenAI has shaped society in turn, and most importantly how these influences have shifted over time, the rise of Generative AI in the United States is examined from three perspectives: private industry, public society, and government. Each of these groups has a very different incentive to encourage the development of GenAI, and each has had a part to play in its success. By examining each individually and determining their motivations and stakes, the unique driving force each has exerted on the AI industry can be more easily discerned. These social influences are framed through technological momentum to explain how and why the proliferation of GenAI has outpaced its utility and it has begun to impact society in turn.

To investigate sentiment, public sphere heuristics including Google search trends and social media engagement are analyzed alongside industry messaging such as AI product releases, venture capital investments, and infrastructure expansion, as reflected in financial reports and energy consumption data. To determine government motivation, and whether social and political factors are becoming increasingly influential, a comparative analysis is conducted between public and private perceptions and adoption of Generative AI in the United States and China. Such an analysis helps to contextualize the global implications of technological momentum, particularly in assessing whether national policies, economic strategies, and ideological positions shape the technology's continued development and entrenchment.

Origins of The AI Boom

In 2017, Google released a landmark paper detailing a novel machine learning model architecture, the "transformer," which was the first sequential model capable of effectively translating text between languages without requiring an impractical amount of computing power (Vaswani et al., 2023). Researchers quickly began to discover that this efficiency could be leveraged to build much larger models capable of a broader range of language tasks. In 2018, OpenAI released GPT-1, the first LLM, which could be trained to a high level of language proficiency with very little data and then finetuned to a specific task (Radford et al., 2018). Development on GPT models continued at a rapid rate, and they began to be applied throughout the industry – for example, Google Translate switched to a transformer architecture in 2020 (Caswell & Liang, 2020). However, LLMs did not enter the public eye until the release of ChatGPT, a free online tool that allowed anyone to interface with a powerful LLM in a familiar "chat" interface. ChatGPT became an overnight sensation, gaining over 100 million users in two months and becoming the fastest-growing web application in history (Milmo, 2023). Tech corporations were quick to jump on this trend, and LLMs are now integrated into everything from Google search to the iPhone.

However, some believe there is a growing disconnect between the capabilities being developed and the practical, ethical, and social needs of the broader population (Weidinger et al., 2021) and that a constant thirst for progress has encouraged the industry to prioritize rapid development over risk management and environmental sustainability (Wu et al., 2022). A 2024 survey of the UK public showed a generally optimistic view of the technology's potential, but many participants reported feeling uncertain or uninformed about the risks associated with AI (Bright et al., 2024). This uncertainty is just as prevalent among AI developers. In June 2024,

current and former OpenAI employees published an open letter condemning the company's reckless disregard of ethical and safety concerns (Field, 2024). Additionally, the AI industry has begun to attract the attention of national governments, which are incentivized to bend the rules and suppress discourse even further in pursuit of becoming the leader in the space. A large-scale analysis of China's public sphere found that AI-related discourse was being manipulated by its government to create a more positive image of the technology and shut down criticism (Jing Zeng & Schäfer, 2022), and experts in and out of government have argued that the US should leverage its AI dominance to gain a foreign-policy advantage (Frank, 2024).

Even without the ethical risks involved, there exists a gargantuan environmental cost of building and operating so much AI infrastructure. As an example, Microsoft reported a 2.5-fold increase in its overall energy consumption between 2018 and 2023, an increase it attributes mostly to GenAI (see Figure 1).

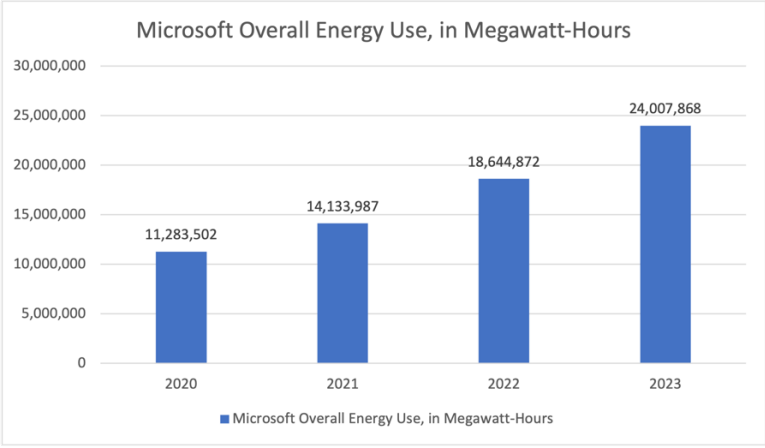

Figure 1: Microsoft Energy Use Over Time (Microsoft, 2024)

This rate of increase in power consumption, on an international level, is unprecedented, and its effects are already being felt. Grids are being strained, resources are being redirected, and retirement of fossil fuel power plants is being delayed. (Halper, 2024). Global data center water consumption is likely to reach more than half that of the entire United Kingdom by 2027 (Li,

2023). It is clear that the prioritization of speed and expansion over environmental and long-term economic stability has put the GenAI industry on a destructive and unsustainable path.

Technological Momentum and GenAI

The concept of technological momentum, introduced by historian Thomas P. Hughes (1987), provides a nuanced framework for understanding how technologies evolve and become integrated into societal structures. Hughes attempts to reconcile two traditionally opposing theories – social constructivism and technological determinism – by considering how a technology's influence shifts over time. In the initial stages of a technology's development, societal factors predominantly shape its trajectory, aligning with the principles of social constructivism. As the technology matures and becomes more entrenched, it gains momentum, making it increasingly resistant to change and exerting a deterministic influence on society—a phase resonant with technological determinism. This synthesis suggests that technological systems are initially shaped by social forces but eventually acquire an autonomous character as they develop. This inertia is particularly pronounced in large technological systems, where interdependencies between technical and social components reinforce the technology's continued presence and expansion. Generative AI exemplifies this process: initially, its development was shaped by research priorities, ethical debates, and regulatory concerns, but as it becomes embedded in industries like media, healthcare, and education, it gains momentum, shaping workflows and societal norms in ways that become increasingly resistant to change.

Since Hughes introduced the concept, technological momentum has been widely applied in the analysis of various technological systems. Historians and STS scholars have used it to bridge the divide between theories that emphasize either the autonomy of technology or the

social and political motivations behind technological choices. For instance, it has been used to explain the persistence of large-scale infrastructure projects such as electrical grids (Nye, 1998) and transportation networks (Davis, 1992), which, despite shifts in policy or public opinion, continue to expand due to their deep entrenchment in society. Similarly, the internet and digital communication technologies have followed this trajectory—initially influenced by governmental and corporate interests but eventually developing an independent momentum that makes them difficult to regulate or reshape (Vogels et al., 2020). By applying technological momentum to Generative AI, we can better understand how its current societal influence is shifting from one shaped by human decision-making to one that drives its own adoption and expansion, making it a defining force in the future of work, creativity, and knowledge production.

Results and Discussion

There exists within any industry the potential for a single new idea to transcend its original bounds and become a defining part of the global zeitgeist. Sometimes this is because the idea is truly world-transforming, as in the case of nuclear weaponry or food refrigeration. Other times, the stars will align for a much less deterministic rise, where society is perfectly positioned to respond abnormally positively toward an incoming technology, catapulting it into a position of such prevalence that it begins to act on society in turn and create new purposes for itself. The rise of Generative AI is an example of this phenomenon. A combination of unlikely factors—discovery of a revolutionary language model architecture, the COVID-19 pandemic, and increasing attention toward technology among major governments—allowed GenAI to explode in popularity and bypass ethical and political roadblocks imposed on previous AI technology. As it has become more capable, accessible, and integrated into daily life, it has begun to act on

society in unforeseen ways, providing new exigence for its development and in turn spurring even faster growth. Even though this growth has outpaced public desire, created significant ethical, environmental, and political risks, and potentially damaged education around the globe, it only continues to accelerate. Companies, people, and governments must work together to slow this momentum and re-evaluate the role GenAI should have in society before it causes irreparable damage. The role of each is analyzed below.

Part 1: Private Industry

As is the case with most technologies, the inventors and developers of GenAI have generally been its strongest proponents. While the groundwork for the AI boom was laid by the academic research community, it was the private tech industry that first capitalized on AI's potential and kickstarted its rise to prevalence. OpenAI, the world's largest for-profit AI company, was originally founded in 2015 as a non-profit research organization. In its first public statement, OpenAI pledged itself to be "an extension of individual human wills and, in the spirit of liberty, as broadly and evenly distributed as is possible safely" (BBC, 2015). After developing open-source GPT (Generative Pretrained Transformer) models for years, which largely went unnoticed by the public due to their technicality and abstract nature, OpenAI released its first cloud-hosted, publicly available product in 2022—ChatGPT.

While originally intended as a free research aid, allowing the company to gather training data through user interaction and further improve its models, ChatGPT revolutionized the entire GenAI industry and singlehandedly set a paradigm for human-AI interface which stands to this day. Its intuitive "chat" interface, where users and models converse back and forth in continuous conversations, was much more useful to most than the previous transformer interface (where the

model simply tries to predict the next word after a piece of input text). Its success may also have been related to its release near the end of the COVID-19 pandemic, a time when many societal customs were being reconsidered and people were generally more open to novel ideas. This moment was the very beginning of GenAI's technological momentum, when the technology was nascent and waiting to be defined by society's needs.

In most industries which rely on regular innovation and product releases to drive revenue, large players in the tech space often face pressure from investors and customers to adopt new trends quickly and appear "cutting edge." Microsoft, long a dominant force in computing, recognized the disruptive potential of GenAI early and positioned itself as a key player—not by developing its own competing product, as OpenAI had an untouchable advantage at the time, but by leveraging its vast influence and capital to purchase a large stake in the company. This early investment laid the foundation for AI-driven transformations across Microsoft's product suite, from integrating large language models into its Azure cloud services to embedding OpenAI-powered features in Microsoft Office and Bing (Microsoft, 2024). By 2023, Microsoft had doubled down on its commitment, pouring billions more into OpenAI, further integrating its technology into consumer and enterprise software. It eventually released its own competitor to ChatGPT, Copilot, which appeared to be a rival but in fact was built on OpenAI GPT models. This aggressive push into AI not only solidified Microsoft's position in the AI space, providing investors with piece of mind, but also pressured other major tech firms to follow suit, accelerating the industry-wide race to develop and deploy Generative AI solutions. Through the lens of technological momentum, this could be seen as major step constraining the path GenAI would later take—once major players like Microsoft and Apple had committed to developing

this particular type of AI and the public was continually exposed to it through them, it became much harder to introduce an alternative.

By the end of 2023, GenAI had become the most common buzzword in the tech industry, with more and more investors looking for a piece of the pie. This gave rise to entirely new companies that likely would not have existed outside this rapid wave of innovation, reflecting a common pattern in technological revolutions: after a disruptive breakthrough, an ecosystem of smaller companies emerges to refine, commercialize, and expand the technology's reach, sometimes in ways that the original innovators never anticipated. However, the momentum of GenAI was so strong that it became possible to companies to secure substantial investments despite offering products that lacked innovation and functional efficacy. As an example, consider the AI Startup Rabbit and its single product, the R1. Unveiled at CES 2024, the R1 was marketed as a groundbreaking AI companion capable of streamlining tasks through voice commands. However, the reality fell markedly short of expectations. Users encountered significant issues, including inaccurate responses, limited functionality, and security vulnerabilities. For instance, the device frequently misidentified objectsand suffered from unresponsive controls during media playback. Furthermore, the R1's integration with third-party services like Uber and DoorDash was fraught with complications, often rendering these features unusable. Despite these glaring deficiencies, Rabbit secured approximately $50 million in funding and sold over 100,000 units in under two months. And it was not alone—across all US companies, GenAI drew $33.9 billion in investment in 2024 (Stanford, 2025). The R1 serves as an example of the perils of late-stage technological momentum, when the mere association with the technology becomes sufficient to attract investment, even for products that fail to deliver on their promises.

The private sector's role in AI's rise cannot be overstated. While academic researchers laid the theoretical groundwork, it was industry leaders—both established and emerging—that transformed AI into a mainstream force. However, as Hughes' theory suggests, once technological momentum reaches a certain threshold, it becomes increasingly difficult to guide or restrain. This raises pressing questions about whether private industry, left to its own devices, can responsibly manage AI's trajectory. OpenAI's initial commitment to ethical integrity and open development quickly began to falter once investors realized how much they could benefit from its position as the technology leader. It soon restructured into a "capped-profit" model, while ostensibly maintaining its foundational mission. This shift has been met with criticism, with detractors arguing that the pursuit of profit may compromise ethical standards and safety considerations in AI development. The internal discord stemming from this transition became evident when a group of current and former OpenAI employees, alongside industry figures, published an open letter expressing concerns about the company's direction (Field, 2024). The letter highlighted apprehensions that financial incentives might overshadow safety protocols and ethical guidelines, potentially leading to irresponsible AI advancements. It called for enhanced whistleblower protections, enabling employees to voice concerns without fear of retaliation, and advocated for increased transparency and accountability within AI organizations. As the private sector continues to navigate its role in the evolving AI landscape, these internal and external critiques serve as a reminder of the importance of aligning technological advancement with ethical principles as well as balancing profit with public interest to ensure that AI serves the broader interests of humanity.

Part 2: The Public

The public's engagement with GenAI has undergone a significant transformation since its inception, evolving from enthusiastic adoption to a more complex and, at times, skeptical relationship. This shift underscores the dynamic interplay between technological innovation and societal response, highlighting how public interest initially propelled GenAI's development but gradually became a secondary consideration as the technology's momentum accelerated.

In the early stages, GenAI captured the public's imagination with its remarkable capabilities. Just as ChatGPT was pivotal in the AI industry, it showed the public an entirely new way to engage in human-like conversations, draft content, and seek information. Within two months, it amassed over 100 million users, becoming the fastest-growing consumer software application in history. This enthusiasm was not confined to tech enthusiasts; individuals from diverse backgrounds integrated GenAI tools into their daily routines, leveraging them for tasks ranging from creative writing to coding assistance (Dell'Acqua et al., 2023). The accessibility and versatility of these tools felt like a democratization of a concept which had previously felt unfamiliar and out of reach to many people. This fostered a sense of empowerment among users who could now harness what they perceived as brand-new advanced technology, without the need for specialized knowledge. Such enthusiasm helped motivate the development of newer, better models—society was influencing a newborn technology and reenforcing its direction, aligning once again with technological momentum. It also encouraged organizations across a breadth of industries to adopt GenAI as a productivity tool. By the start of 2025, 95% of US companies were using GenAI in some capacity (Rapoport et al., 2025).

Later, however, as GenAI became more embedded in everyday life, its ubiquity began to elicit mixed reactions. The novelty that once captivated users started to wane, leading to a

phenomenon of technological fatigue. The initial excitement gave way to concerns about over-reliance on AI, particularly in contexts requiring critical thinking and originality. For instance, in educational settings, students who initially embraced AI tools to manage academic workloads began expressing apprehension about potential skill atrophy and diminished problem-solving abilities. Reports indicated that while some students found AI helpful for tasks like transcribing lectures or summarizing information, others worried that excessive dependence might hinder their intellectual development (Lin, 2024). Such apprehension reflected a broader societal tension between embracing technological convenience and empowering human competencies (Parker, 2025).

More and more instances emerged where the development and deployment of new AI applications seemed misaligned with public interest. For example, social media platforms explored integrating AI-generated characters to enhance user engagement. While intended to enrich the user experience, such initiatives raised concerns about the authenticity of interactions, potential misinformation, and the ethical implications of blurring lines between human and AI-generated content. Users of these platforms argued that these developments catered more to corporate agendas than to their genuine needs and preferences (Parker, 2025).

The public sector, particularly in education, has a crucial role in slowing the unchecked momentum of GenAI and ensuring its development aligns with human values rather than convenience or profit. Young people are among the first to adopt new technologies and thus experience their disruptive force. Their response to GenAI will set a precedent for how society at large chooses to integrate or regulate it. Faculty and students alike must recognize that while GenAI can be a powerful tool for learning and efficiency, its unregulated use threatens fundamental aspects of education—such as critical thinking, intellectual integrity, and the

development of original thought. If educational institutions simply accept GenAI as an unavoidable part of academic life without establishing clear guidelines, they risk creating a generation of students who lack independent problem-solving abilities and instead rely on AI to do their thinking for them. Students and educators must work together to curtail the momentum of AI by showing integrity and valuing human capability and the future of our students over convenience.

Beyond education, a broader societal framework is needed to protect individuals from the negative consequences of GenAI's rapid growth. An "AI Bill of Rights" could serve as a foundational guideline to establish ethical AI usage and ensure that technological progress does not come at the expense of human dignity and autonomy. Such a framework would need to address key concerns, including data privacy, transparency, and the right to human oversight. For instance, users should have the right to know when they are interacting with AI-generated content and how their personal data is being used to train AI models. Additionally, regulation should include provisions against AI-driven discrimination, ensuring that algorithmic biases do not reinforce social inequalities. AI should reflect the future we want, not the mistakes of our past.

While the public may have the strongest influence on GenAI's future, slowing its momentum requires a collective effort—one that values long-term societal well-being over short-term convenience and financial gain. Public institutions, including schools, government agencies, and advocacy groups, must take proactive steps to demand accountability from AI developers and corporations. By promoting ethical standards in education and enacting legal safeguards for AI governance, society can ensure that AI remains a tool for human empowerment rather than a force that erodes personal agency and critical thought.

Part 3: Governments and Politics

In the nascent stages of GenAI development, governmental oversight was notably absent. This regulatory vacuum provided tech companies with the freedom to innovate and deploy GenAI technologies without bureaucratic impediments. As GenAI began to permeate various sectors, governments worldwide started to recognize their strategic and economic significance, but by this time the technology progressed substantially, and, as technological momentum dictates, had developed a substantial influence on society in turn. Governments, seeking to leverage this influence, focused on establishing dominance in the AI arena rather than implementing comprehensive safeguards (which would have hindered development and given rivals a leg up). The United States, as an early leader in AI innovation, was one of the first nations to use government powers to protect this lead. A notable action was the Biden administration's imposition of export controls on advanced AI chips to China, aiming to stifle AI development of a major political rival. Conversely, both the American and Chinese governments have been taking action to encourage the development and use of GenAI within their own borders. The US government has influenced energy commissions to allow the construction of more data centers (Groves, 2025) and has continued to avoid regulation of US AI development while aggressively regulating Chinese AI products such as DeepSeek (White House, 2023). China in turn has actively encouraged sectors such as healthcare and legal services to integrate homegrown AI models into their practices, aiming to modernize industries and bolster economic growth. Furthermore, Chinese tech companies are incentivized to develop indigenous AI models, reducing reliance on foreign technologies and bolstering national capabilities. They take advantage of deliberately low electricity costs (set by the government for this very purpose) to

run servers cheaply. This push for adoption combined with continuous government propaganda has resulted in an abnormally favorable view of GenAI among the Chinese public, with one study reporting that 78% of Chinese respondents believe AI brings more benefits than drawbacks (Stanford, 2025). This statistic illustrates how much influence governments have over public perception of GenAI and by extension how important they handle it responsibly.

In summary, the technological momentum of GenAI shifted the focus of governments and regulators away from its associated ethics and sustainability concerns toward its value as a geopolitical tool. The U.S.-China technological rivalry is only one example of the increasing political relevance of GenAI around the world, highlighting the need for international dialogue and cooperation to navigate the challenges and opportunities it presents.

Conclusion

The trajectory of GenAI's rise is a striking example of technological momentum, where initial societal demand and industry breakthroughs gave way to a self-perpetuating cycle of expansion, competition, and, ultimately, detachment from public needs. Private industry played the earliest and most aggressive role in shaping GenAI's development, with companies either pioneering breakthroughs, capitalizing on emerging trends, or launching entirely new businesses made possible by the AI boom. What began as a novel and promising tool for consumers quickly became a high-stakes race between tech giants, spurred by competitive pressure and investment incentives rather than the needs or desires of the public. Meanwhile, as AI tools became ubiquitous, the general public—once an eager participant in AI's adoption—began to grow weary of its rapid expansion, as its presence in everything from creative work to customer service raised concerns about job displacement, misinformation, and privacy violations. Most alarmingly, governments, initially absent from the conversation, have now entered the fray, but

often with a focus on global AI dominance rather than meaningful regulation. The U.S. and China in particular exemplify how AI has become a new frontier for geopolitical competition, where regulatory efforts are often secondary to economic and military interests.

If AI's momentum is to be redirected toward a future that benefits society rather than overwhelming it, intervention is necessary at every level. Companies must refocus their AI development efforts on genuine human needs rather than competitive one-upmanship, prioritizing transparency, safety, and user control over unchecked expansion. The public must continue to assert its influence, demanding accountability from both corporations and policymakers before AI becomes an irreversible and unwanted fixture in everyday life. Most critically, governments must implement policies that go beyond national interests and address AI as an international challenge. An AI Bill of Rights—establishing baseline protections for privacy, fairness, and informed consent—could serve as a starting point for responsible AI governance. On a larger scale, an international framework for AI regulation, akin to nuclear or environmental treaties, may be necessary to prevent the technology from becoming a destabilizing force in global politics. Without these measures, AI's unchecked momentum will continue shaping society in ways that are not only unforeseen but potentially irreversible. The challenge now is not just to slow AI's advance but to ensure that it moves in a direction that serves humanity rather than overriding it.

Bibliography

2024 Environmental Sustainability Report. (2024). Microsoft.

Bright, J., Enock, F. E., Esnaashari, S., Francis, J., Hashem, Y., & Morgan, D. (2024). Generative AI is already widespread in the public sector (arXiv:2401.01291). arXiv. https://doi.org/10.48550/arXiv.2401.01291

Caswell, I., & Liang, B. (2020, June 8). Recent Advances in Google Translate. Google Research Blog. http://research.google/blog/recent-advances-in-google-translate/

Davis, D. (1992). Technological Momentum, Motor Buses, and the Persistence of Canada's Street Railways to 1940. *Material Culture Review*, *36*(1). Retrieved from https://journals.lib.unb.ca/index.php/MCR/article/view/17507

Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *Harvard Business School Working Paper*, *24–013*.

Economic potential of generative AI | McKinsey. (2023). McKinsey and Company. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier

Field, H. (2024, June 4). Current and former OpenAI employees warn of AI's "serious risks" and lack of oversight. CNBC. https://www.cnbc.com/2024/06/04/openai-open-ai-risks-lack-of-oversight.html

Frank, M. (2024, September 22). US Leadership in Artificial Intelligence Can Shape the 21st Century Global Order. The Diplomat. https://thediplomat.com/2023/09/us-leadership-in-artificial-intelligence-can-shape-the-21st-century-global-order/

Groves, T., & Wender, B. (2025, March 26). Ben Wender Interview. personal.

Hughes, T. P. (1987). The Evolution of Large Technological Systems. In W. E. Bijker, T. P. Hughes, & T. Pinch (Eds.), *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology* (pp. 51–82). MIT Press.

Jing Zeng, C. C., & Schäfer, M. S. (2022). Contested Chinese Dreams of AI? Public discourse about Artificial intelligence on WeChat and People's Daily Online. Information, Communication & Society, 25(3), 319–340. https://doi.org/10.1080/1369118X.2020.1776372

Levin, A., Kataeva, N., & Kallies, M. (2024, September 18). *The Stairway to GenAI Impact*. BCG Global. https://www.bcg.com/publications/2024/stairway-to-gen-ai-impact

Lin, L. (2024, May 15). A quarter of U.S. teachers say AI tools do more harm than good in K-12 education. *Pew Research Center*. https://www.pewresearch.org/short-reads/2024/05/15/a-quarter-of-u-s-teachers-say-ai-tools-do-more-harm-than-good-in-k-12-education/

Milmo, D. (2023, February 2). ChatGPT reaches 100 million users two months after launch. The Guardian. https://www.theguardian.com/technology/2023/feb/02/chatgpt-100-million-users-open-ai-fastest-growing-app

Muralidhar, R., Borovica-Gajic, R., & Buyya, R. (2022). Energy Efficient Computing Systems: Architectures, Abstractions and Modeling to Techniques and Standards. ACM Comput. Surv., 54(11s), 236:1-236:37. https://doi.org/10.1145/3511094

Nye, D. E. (1998). *Consuming power: a social history of American energies*. MIT Press.

OpenAI. (2022, November 30). Introducing ChatGPT | OpenAI. https://openai.com/index/chatgpt

Parker, L. L. and K. (2025, February 25). U.S. Workers Are More Worried Than Hopeful About Future AI Use in the Workplace. *Pew Research Center*. https://www.pewresearch.org/social-trends/2025/02/25/u-s-workers-are-more-worried-than-hopeful-about-future-ai-use-in-the-workplace/

Patterson, D., Gonzalez, J., Hölzle, U., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., & Dean, J. (2022). The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink (arXiv:2204.05149). arXiv. https://doi.org/10.48550/arXiv.2204.05149

Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). Carbon Emissions and Large Neural Network Training (arXiv:2104.10350). arXiv. https://doi.org/10.48550/arXiv.2104.10350

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (n.d.). Improving Language Understanding by Generative Pre-Training.

Rapoport, G., Bicanic, S., & Talabi, M. (2025, May 7). *Survey: Generative AI's Uptake Is Unprecedented Despite Roadblocks*. Bain. https://www.bain.com/insights/survey-generative-ai-uptake-is-unprecedented-despite-roadblocks/

The 2025 AI Index Report | Stanford HAI (AI Index Report). (2025). Stanford University. https://hai.stanford.edu/ai-index/2025-ai-index-report

The White House. (2023, October 30). FACT SHEET: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence. *The White House*. https://bidenwhitehouse.archives.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2023). Attention Is All You Need (arXiv:1706.03762). arXiv. https://doi.org/10.48550/arXiv.1706.03762

Vena, M. (2022, March 15). The 3 Most Important Marketing Messages About Apple Silicon. LinkedIn. https://www.linkedin.com/pulse/3-most-important-marketing-messages-apple-silicon-mark-vena/

Vogels, E., Rainie, L., & Anderson, J. (2020). *Analysis: 2005-2011 Predictions for Digital Life 2020* (Imagining the Internet). Elon University.

Wang, Y., Wang, Q., Shi, S., He, X., Tang, Z., Zhao, K., & Chu, X. (2020). Benchmarking the Performance and Energy Efficiency of AI Accelerators for AI Training. 2020 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID), 744–751. https://doi.org/10.1109/CCGrid49817.2020.00-15

Weidinger, L., Mellor, J., Rauh, M., Griffin, C., Uesato, J., Huang, P.-S., Cheng, M., Glaese, M., Balle, B., Kasirzadeh, A., Kenton, Z., Brown, S., Hawkins, W., Stepleton, T., Biles, C., Birhane, A., Haas, J., Rimell, L., Hendricks, L. A., … Gabriel, I. (2021). Ethical and social risks of harm from Language Models (arXiv:2112.04359). arXiv. https://doi.org/10.48550/arXiv.2112.04359

Wright, D., Igel, C., Samuel, G., & Selvan, R. (2023). Efficiency is Not Enough: A Critical Perspective of Environmentally Sustainable AI (arXiv:2309.02065). arXiv. https://doi.org/10.48550/arXiv.2309.02065

Wu, C.-J., Raghavendra, R., Gupta, U., Acun, B., Ardalani, N., Maeng, K., Chang, G., Behram, F. A., Huang, J., Bai, C., Gschwind, M., Gupta, A., Ott, M., Melnikov, A., Candido, S., Brooks, D., Chauhan, G., Lee, B., Lee, H.-H. S., … Hazelwood, K. (2022). Sustainable AI: Environmental Implications, Challenges and Opportunities (arXiv:2111.00364). arXiv. https://doi.org/10.48550/arXiv.2111.00363