Rethinking Performance: The Ethical and Technological Landscape of Accelerated Computing

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

Hugo Abbot

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

Bryne E. Seabrook, Department of Engineering and Society

Rethinking Performance: The Ethical and Technological Landscape of Accelerated Computing

I. The Beginning After the End

For decades, Moore's law served as a guiding principle in computer hardware development, fueling mass amounts of research into growing processing power exponentially. However, as transistor scaling reaches physical and economic thresholds, the industry was forced to explore alternative pathways for performance improvements. "The end of Moore's law could be the best thing that has happened in computing since the beginning of Moore's law," said R. Stanley Williams, a prominent research scientist at Hewlett-Packard Labs. "Confronting the end of an epoch should enable a new era of creativity" (Williams, 2017). This shift introduced accelerated computing hardware as we know today—such as GPUs and TPUs—influenced by the massive rise in large language models and artificial intelligence (Shalf, 2020). Along with many other upcoming technologies which hold immense potential, each is designed to push the boundaries of efficiency and computational capability, introducing their own ethical challenges. Concerns range from environmental sustainability and resource consumption to issues of accessibility and technological inequality (Zhuk, 2023; Brynjolfsson & Unger, 2023). This paper implements the theory of Technological Momentum to survey the history of computer architecture development, new areas of promising research, and the ethical concerns related with this progress, revealing how industry leaders are reacting and addressing these problems.

II. The Methodology

How has the evolution of accelerated computing hardware reshaped the landscape of society, leading to significant ethical consequences in research, development, and production? Accelerated computing encompasses any innovation designed to enhance computational performance, often achieving greater efficiency and speed through specialized architectures or novel processing methods across various domains. This work examines the evolving relationship between society and accelerated computing architecture – including technologies like Processing in Memory (PIM) and the development of new potential room-temperature superconducting materials like LK-99 – and the broader ethical implications of this interplay, such as sustainability, accessibility, and technological inequality. Using a qualitative methodology, primarily utilizing literature review along with academic and policy discussions, the study draws from peer-reviewed academic articles, industry reports, and technical documentation from the leading companies within the industry. In addition, government and third-party reports concerning the sustainability and socio-economic impact of hardware advancements serve as critical data sources and areas of cost-benefit analysis. The following section will examine historical trends in hardware development, analyze the societal implications of PIM and LK-99, and discuss ethical considerations, regulatory responses, and future research directions. Through the lens of Technological Momentum, this study analyzes how early social, political, and economic forces guide the design of new technologies, and how, over time, these technologies acquire inertia, shaping future societal possibilities and limiting alternate development paths. This time-dependent relationship is essential to understanding both the possibilities and limitations of modern computing.

III. The Start of Accelerated Computing

In 1965, the co-founder of Intel Gordon Moore observed a trend in semiconductor development: the number of transistors on an integrated circuit would double every year for 10 years, later revising it in 1975 to doubling approximately every 2 years. This prediction, later coined as Moore’s Law, became the driving force behind advancements in computer hardware and served as a baseline projection for the industry (“Moore’s Law,” 2023). Several breakthrough innovations resulted from this early-age pursuit, including complementary metal-oxide-semiconductor technology (CMOS), reduced instruction set computing (RISC) architectures, multi-core processors, and 3D chip stacking (“The Silicon Engine: Timeline,” n.d.). However, as transistor sizes approached atomic scales, quantum tunneling effects, excessive heat generation, and increasing manufacturing costs imposed significant barriers to further transistor compactness (“The Future of Semiconductor Miniaturization,” 2024). By the early 2010s, the historical pace of Moore’s Law had slowed, prompting researchers to seek alternative methods for increasing computational efficiency.

One such innovation is PIM, a proposed model that reduces the inefficiencies caused by the traditional separation of memory and processing units. In conventional computer systems, data must continuously move between dynamic random-access memory (DRAM) and central or graphical processing units, creating a performance bottleneck known as the von Neumann bottleneck. PIM mitigates this issue by embedding computational capabilities directly into memory, enabling parallel data processing and execution while significantly reducing energy consumption (Zou et al., 2021; Stone, 1970). Research into PIM architectures dates back to the 1990s, with early implementations focusing on specialized applications such as neural networks or machine learning processing. Due to this foundation, many architectural variations exist with various modifications to the standard DRAM design, making comparisons and large-scale production difficult. Today, major semiconductor companies such as Samsung and SK hynix continue research and development for commercial PIM solutions, while smaller-scale projects such as UPMEM have tried to gain an edge on the market by releasing their own programmable solution to the public.

Concurrent to architectural innovations like PIM, many breakthroughs in the area of materials science have shown just as much potential in reshaping computing capabilities. Currently, high-performance computing components use materials such as copper and aluminum for wiring, silicon as the primary component of computer chips, and rare earth elements like neodymium, tantalum, and yttrium for specialized parts; however, many of these materials present significant supply chain vulnerabilities, as they are often sourced from a select few countries, leading to geopolitical and economic risks (Wheeler, 2018; Nunez, 2021). In addition, silicon-based transistors suffer from heat dissipation issues and energy inefficiencies at increasingly smaller scales, making them harder to improve while reducing scale. Moreover, all superconductors in use today, such as niobium-based alloys used in quantum computer research, require extremely low temperatures to function, making them impractical for widespread use in consumer electronics (Cho, 2023); however, a notable development within this field is LK-99, a material proposed as a room-temperature superconductor by South Korean researchers. If made widely available, LK-99 could revolutionize computing by enabling zero-resistance electrical transmission, reducing energy losses in processors and memory systems, while also eliminating the cooling requirements needed of modern-day chips (Lee et al., 2023). Additional independent replication attempts of LK-99 have yielded mixed results, dividing the scientific field on the feasibility of this material (Lemonick, 2024; Wang et al., 2024). Even with the uncertainty, the research within this field highlights the industry's recognition of material limitations with current hardware devices, looking for new ways to find sustainable and efficient solutions for today's limitations.

IV. Technological Momentum and Innovation

The evolution of computing hardware, particularly with the emergence and eventual snowballing of alternative routes of development, aligns closely with the concept of Technological Momentum, a framework developed by Thomas P. Hughes, one of the nation's foremost historians of technology. This idea suggests that technology begins as a malleable innovation shaped by social, economic, and political forces but eventually gains inertia, making it increasingly difficult to change and become independent of (Hughes, 1969). Technological Momentum provides a useful perspective for understanding the current state and trajectory of new technologies, such as accelerated computing components which have moved from early research-driven exploration to an established industry standard with significant ethical implications. It is mandatory that decisions surrounding new development are no longer just about efficiency, performance, and trailblazing, but also about sustainability, accessibility, and equity.

Hughes' theory of Technological Momentum is often positioned between Technological Determinism and Social Construction of Technology (SCOT). While determinists argue that technology develops autonomously with inevitable progress, and SCOT scholars emphasize the role of human agency in shaping technology trajectories, Technological Momentum suggests a dynamic interplay between these forces over time (Basu, 2022; Scranton 1995). Many of these new accelerated technologies have given rise to new innovations, such as in the field of artificial intelligence and cloud computing; however, due to specific research needs and market demands for what's new, further research and development is pushed for the backbone infrastructure, showing the deeply entrenched mutual relationship which forms. Due to these factors, it is difficult to challenge emerging but dominant paradigms which are frequently viewed as high-risk research areas due to their emergent stage and unclear foundational principles, emphasizing the need for comprehensive approaches to their ongoing development, assessment, and regulation.

Previous works have expanded on this framework and ability to be a dominant force within society, such as the idea of path dependence. Within this context, historian David Nye introduces the economic theories of path dependence, a social science that describes how past events and decisions can limit future outcomes and decisions. Using the electrification of the United States as a case study, Nye explores how countries and companies continue the use of outdated technology due to historical investment and institution inertia, or the resistance to change (Nye, 2004; Samadi et al., 2024). In addition, past research has also explored the dangers of Technological Momentum, arguing that innovation is occurring too often at an unsustainable pace, often without a full understanding of its long-term consequences. Technological progress is often driven by competition and efficiency, rather than careful consideration of ethical, educational, and societal impacts. (Dyer, 1995). Historical technological shifts such as the Industrial Revolution show how rapid advancements lead to unforeseen or flat out ignored effects, which can involve environmental degradation or economic dependencies.

V. Findings and Discussion

The history of computing hardware has been defined by a continuous pursuit of increased processing power, efficiency, and scalability – deeply intertwining with societal needs and economic pressures. The development of accelerated technology follows a cyclical pattern, where technological advancements drive societal changes, creating in turn a societal demand to push for new innovations. This reciprocal relationship aligns with Technological Momentum, demonstrating how initial technological decisions and investments shape future development paths, sometimes reinforcing existing power structures.

a. Historical Patterns

The origins of modern computing can be traced back to the mid-20th century, when early computers such as the Electronic Numerical Integrator and Computer, known as the ENIAC, were developed almost exclusively for military applications during World War II. These machines were large, expensive, and purpose-built, almost immediately being put to use to calculate firing tables to accurately hit desired targets with various weapons systems ("Celebrating Penn Engineering," n.d.). The impact of the Cold War and the societal need for greater national security led to this device becoming a strategic asset, leading to a rapid influx of funding and competition to advance computing capabilities across multiple countries (Leese, 2023). This exemplifies a key concept in Technological Momentum – once this type of technology became essential for defensive or offensive security measures, the pace of innovation was largely dictated by economic and political forces rather than purely scientific curiosity.

The 1960s marked a turning point where Moore's Law transitioned from a scientific observation into a socially reinforced economic and industrial benchmark. This evolution illustrates the shift from early societal shaping of technology to the beginning of technological inertia – a key tenet of Technological Momentum. Now, the rate of adoption for new hardware advancements was set, influencing both industry and public perception of computing progress. As businesses, governments, and individuals integrated computers into daily life, the demand for faster, smaller, and cheaper hardware surged, reinforcing the trajectory Moore's Law described and the feedback loop relationship between society and technology. The rapid expansion of hardware capabilities facilitated the rise of personal computers, mobile devices, and now artificial intelligence (Arcuri & Shivakumar, 2022); however, the exponential progress started to wain, leading to a very slow shift in computational approaches, with companies in the industry exploring alternative pathways to maintain performance growth and solidify their place within the market. This is a hallmark of Technological Momentum – once an industry-wide expectation of progress is established, deviation from that path becomes increasingly difficult, even as the physical and economic limitations of traditional scaling approaches become apparent.

The early 2000s marked the beginning of a true shift away from traditional processor scaling due to the various limitations present, as well as a growing need for new processing techniques. Most notably, the von Neumann bottleneck, which restricts performance due to the separation of memory and processing, became more pronounced due to data-intensive applications such as AI and big data analytics gaining prominence (Pingali, 2018). Due to these societal needs, driven mainly by cloud computing, machine learning, and advanced security algorithms such as cryptography, researchers were forced to explore alternative designs and domain-specific accelerators. One of these was the graphics processing unit (GPU), a term popularized but not created by NVIDIA. This specialized processor was a modernized version of previously developed graphics accelerators by various companies and was intended for use mainly with speeding up graphics rendering; however, these GPUs over time have become omnipresent in various parallel processing computational environments, leveraging the specialized processor's parallelizable architecture by using many tiny smaller cores to execute the same instructions on different data (Tabor, 2024). Over time, more highly specialized components have been developed, such as tensor processing units (TPUs) which have greater processing power and energy efficiency compared to GPUs with similar AI benchmarks. Unlike previous periods where hardware drove the societal adoption of new technology, this mark in history exemplifies the reverse influence – society shaping innovation. The growing data-centric economy, demands for real-time processing, the rise of AI-driven decision-making systems, and the growth of generative AI made it clear that traditional hardware approaches were no longer sufficient given the requirements. The industry had shifted to a point where change was necessary, but existing infrastructure and economic systems made it difficult to transition seamlessly (Arik, 2023).

By the 2000s, the industry had largely internalized these expectations, with prior decisions about architectures and infrastructures limiting the capacity to shift directions easily. This exemplifies how, over time, technological systems accumulate momentum and resist change, even in the face of pressing societal needs or emerging alternatives.

b. Socioeconomic, Environmental, and Accessibility Affairs

The ongoing evolution of computing hardware demonstrates a duality in technological and societal influence – one that constantly shapes and responds to each other. When technology progresses too rapidly without ethical foresight, consequences such as privacy concerns, environmental damage, and labor displacement arise. Conversely, when society resists necessary innovation, stagnation can occur, limiting opportunities to develop something to solve pressing local or global challenges (Hersh, 2013). Accelerated computing development is no stranger to these dilemmas, ranging from sustainability issues – such as greater energy consumption, electronic waste, and resource gathering – to accessibility and digital inequality.

In 2007, it was reported that even though E-waste, or discarded electronic devices, only accounts for 2% of America’s trash in landfills, it equals around 70% of overall toxic waste, due to various harmful substances such as lead, mercury, and cadmium (“Harmful Effects,” n.d.). In addition, a majority of this E-waste contains valuable metals that can be disassembled and recycled in various forms, which would in turn could create many jobs in order to process them (“Electronics Donation,” n.d.). Today, much of this waste is exported to developing nations, where laborers are exposed to hazardous conditions and international legislation does very little to enforce stricter recycling measures and export regulations. Many experts within the field note the constraints of applying these measures, being both technological and geopolitical (Shittu et al., 2021).

These materials needed for production, especially rare earth metals, result in significant socio-economic impacts, such as pollution, health risks, and potential for inequality and conflict, especially in developing regions in the world. Specifically, case studies from mines in China and the Lynas Advanced Materials Plant in Malaysia show extreme examples of environmental corruption and labor injustice (Ali, 2014). Residents of these areas, which are typically members of marginalized groups, have shown to suffer from neurological diseases such as neurodegeneration and neurosis due to long-term exposure to these elements (Wang, Yang, et al., 2024). Even with the increasing number of case studies detailing these situations all across the globe, much of the way in which these companies extract these resources has not changed. This resistance to reform is driven by both industry needs and economic constraints, allowing the cheap production while easily distributing for a high bottom line.

As computing innovation becomes integrated into healthcare, law enforcement, financial services, and hiring processes, ethical concerns surrounding privacy, bias, and accountability become increasingly pressing. These systems, powered nowadays by technology mentioned earlier such as GPUs and TPUs, mainly involve some new AI model to increase efficiency, performance, and reliability, but raise questions about fairness, discrimination, and transparency. A major ethical ethical dilemma arises from algorithmic bias, where AI models trained on historically biased data perpetuate systemic inequalities. (Friis & Riley, 2023). For example, research has shown that computer-aided diagnosis systems in healthcare have been found to return lower accuracy results for black patients than white patients. Furthermore, natural language processing algorithms for resume analysis were shown to favor those that incorporated specific words such as “executed” or “captured”, which were more frequently used by men (“Shedding Light,” 2023).

As accelerated technologies like AI enter sensitive domains such as healthcare and hiring, their influence is no longer simply a function of performance, as they begin to shape decision-making frameworks, institutional practices, and even notions of fairness. This shift underscores how technological systems can gain social inertia, driving societal transformation beyond their original scope. With these processes becoming ever more present within companies and potentially altering people’s lives from these decisions, regulatory bodies have just started to take action. New AI governance frameworks, such as the European Union’s AI Act, aim to establish risk-based regulations on the ever growing number of AI applications. Passed in March of 2024 and being the first-ever global legal framework on AI, the Act tries to enact comprehensive oversight but is limited by the complexity of these systems and the rapid pace of development from so many industry leaders (“AI Act,” n.d.).

Massive energy consumption – both for training and for user prompts – is a growing concern for these AI models, which are being bolstered by these innovations in hardware. Especially for large-scale deep learning models like GPT and DALL-E, estimating such a figure as the actual environmental cost is near impossible. Many factors go into the calculation of electricity use, such as the maintenance and operation of dedicated data centers, model training, and individual prompt generation (Zewe, 2025). With the continuous advancement of these LLMs and the production of new novel and domain-specific models, these costs will continue to increase. Industry leaders and investors have started exploring low-power AI chips along with new energy-efficient training techniques; however, widespread adoption remains slow due to significant performance trade-offs and financial incentives ("Syntiant Introduces AI Chip," 2022).

Lastly, a major area of ethical contention regards accessibility with accelerated computing hardware, mainly focusing on the disproportionate allocation of funding and resources within the country. As seen from historical trends, cutting-edge technology is primarily designed intentionally for national security, defense, or research use. These investments, which come from substantial funding from government agencies or private defense contractors, drive progress but leave marginalized communities without equitable access to the benefits of these advancements. Recent studies have shown that lower income families have lower levels of technology adoption, which accounts for either updated and improved equipment or completely new devices (Vogels, 2021). These issues stem from underlying problems that have existed for decades, such as unequal broadband and education access, and continue to keep and worsen the divide between socioeconomic groups ("Reducing the Digital Divide," 2024; Bushnell, 2021). This disparity reinforces the technological elite, where only well-funded bodies can harness high-performance computing, while underprivileged communities, educational institutions, and small businesses are left behind. Without equitable distribution strategies, open-source initiatives, or public-sector investments, the cycle of exclusion continues, which will also further exacerbate socio-economic divides.

The ethical challenges posed by accelerated computing hardware are deeply entrenched in historical patterns of technological adoption, where innovations are initially concentrated in well-funded sectors before gradually diffusing to broader society. However, Technological Momentum can both facilitate and hinder equitable access, depending on how early-stage adoption is structured. When economic and political forces prioritize national security and corporate interests, technological benefits remain confined to a select few, reinforcing the socioeconomic divides and digital inequalities discussed. Addressing these ethical concerns requires recognizing that computing hardware does not develop in isolation but is shaped by the institutional, economic, and regulatory forces that guide its trajectory – forces that must be actively managed to avoid repeating past cycles of exclusivity.

Both PIM and LK-99 represent critical developments in accelerated computing, following historical trends where technological advancements emerge in response to societal pressures for efficiency, scalability, and sustainability; however, as seen in prior hardware innovations, these advancements also raise their own ethical concerns regarding accessibility and sustainability. Examining these technologies through the lens of Technological Momentum reveals how their adoption may further reinforce existing societal structures while introducing unforeseen consequences. This analysis will help to guide these currently evolving works and build them into greater ethical innovations for all people.

c. Processing in Memory (PIM)

PIM's development stems from the persistent inefficiencies in the von Neumann architecture, where the separation of memory and processing units results in significant energy and performance costs. By integrating processing capabilities directly into memory, PIM reduces data transfer delays and improves energy efficiency, making it highly desirable for AI, high-performance commuting, and large-scale simulations; however, its trajectory follows the same cyclical patterns seen in past innovations, where adoption is primarily driven by high-budget sectors such as military research, elite academic institutions, and large tech corporations.

While PIM has the potential to revolutionize consumer electronics and cloud computing, its overall accessibility remains quite poorly. As with GPUs and TPUs, early PIM implementations are heavily concentrated within wealthy research institutions, limiting their immediate benefits for underfunded researchers, small businesses, and the general public. Much of the funding for these programs within the United States are from defense related companies or government organizations such as SRC or DARPA ("Optimum Processing," n.d.; Xie, 2019). Furthermore, the reliance on proprietary architectures and patents from major semiconductor firms like Samsung, SK hynix, and Micron may lead to market monopolization, restricting competition and innovation (Asia, 2025).

Finally, PIM may raise future sustainability concerns even though the technology works on introducing energy efficient computation and curb electrical waste that occurs today. In order to implement this, new chips must first be correctly modeled and tested before being fabricated as new physical DRAM chips, requiring various rare earth materials and producing mass amounts of E-waste due to prototyping and discarding outdated DRAM chips. This technology will constantly evolve, bringing about further optimizations and requiring a new round of manufacturing, purchasing, and scrapping. This cycle presently occurs with GPUs at a consumer level each year and leads to high consumer costs and diminishing product returns ("How Long Should a GPU Actually Last," 2023). Without proper regulation and planning, the cyclical pattern of Technological Momentum will present in this field as well, where various social, political, and financial factors will drive the continued research and development and vice versa.

While PIM could potentially redefine computing efficiency, its adoption is constrained by institutional inertia – a form of Technological Momentum where prior investments in legacy systems discourage transition, even when newer models promise substantial benefits. This cyclical pattern, where early societal influence followed by entrenched technological direction, illustrates how momentum limits disruptive change unless deliberately countered by policy, investment, or public pressure.

d. LK-99

The discovery of LK-99 was met with both excitement and skepticism, with the potential to drastically reduce energy waste in processors, memory systems, and data centers. Such an advancement aligns with ongoing societal push for sustainability and efficiency, echoing past trends where technological breakthroughs promised positive societal shifts with changes in computing.

However, the uncertainty and mystery surrounding the feasibility of LK-99 mirrors historical patterns of overhyped discoveries with underwhelming real-world impact. With many independent replications failing to confirm its superconducting properties, much of the scientific rigor fades and the ambition behind the craze starts to get fatigued, repeating in what experts call the Cycle of Technological Hype (Miller, 2019). If LK-99, or a future material with similar properties, does prove viable, further ethical concerns like resource control and E-waste arise.

Due to LK-99's composition of various metals, whoever has material control would be able to create global supply chain dependencies for their own desires. Furthermore, a mass production of this material and incorporation into consumer-level products would mean a whole new generation of products, including new manufacturing and dumping of outdated tech. Without the necessary infrastructure and regulation to control these two mechanisms, the divide and accessibility issues between classes and socio-economic groups can grow, reinforcing the idea of a technological elite and keeping the innovation in the hands of specifically desired people.

As discourse surrounding LK-99 continues, its trajectory highlights how Technological Momentum influences both the perception and adoption of emerging materials. Even if the material itself doesn't work, the excitement it generated has already redirected research investments, shaped funding priorities, and influenced corporate strategies in the pursuit of practical superconducting materials. This shift demonstrates how technological expectations create inertia that can sustain development even in the face of setbacks, leading to the same gradual, incremental adoption pattern of new research pathways seen with past innovations.

e. Limitations of this Study

While this study provides a comprehensive analysis of the ethical and societal consequences of accelerated computing hardware and provides two case studies that reinforce the patterns analyzed, several limitations must be acknowledged. First, given the rapidly evolving nature of hardware advancements, including PIM and LK-99, some findings may quickly become outdated as new breakthroughs or regulatory measures emerge, eliminating potential side effects discussed. The unpredictability of development makes it near impossible to predict long-term consequences, particularly in areas such as accessibility and sustainability. Secondly, the research relies primarily on secondary sources, which, while insightful, lack firsthand empirical studies that could measure real-world impact of these technologies. Lastly, there are many areas of accelerated computing research, many of which are more viable and closer to actual deployment compared to PIM or LK-99. Each research pathway will produce their own ethical consequences – many or few, intended or not. It's important to view all perspectives on new pieces of technology, especially looking at how influential factors from the market or people have brought them to light.

Future research should expand on empirical studies of these new types of hardware, assessing measurable impacts on energy efficiency, cost accessibility, and market monopolization. Additionally, studies should focus on geopolitical tensions resulting from new hardware and the resulting vulnerabilities placed on the supply chain. Researchers and policymakers should explore how Technological Momentum can be accounted for in regulatory frameworks, ensuring that ethical considerations are proactively addressed rather than reacting to crises after technologies have already become entrenched within the industry and society. Finally, interdisciplinary research integrating engineering, environmental science, and social sciences will be essential in developing sustainable and equitable accelerated computing solutions, ensuring the benefit reaches a wider demographic rather than enforcing existing technological divides.

VI. Final Thoughts

The evolution of accelerated computing has demonstrated a reciprocal relationship between technological advancement and societal impact. Innovations such as PIM and LK-99 emerge in response to growing demands for efficiency and sustainability, yet simultaneously introduce their own ethical dilemmas concerning accessibility, resource control, and environmental consequences. This research highlights how development follows historical patterns of Technological Momentum, where initial technological choices and market structures continue to shape the trajectory of future innovations. These advancements risk amplifying digital divides, geopolitical tensions, and unsustainable consumption of resources without proactive intervention. Moving forward, regulation must be made to ensure equitable access to accelerated computing hardware, sustainable material sourcing, and third-party oversight to prioritize ethical decision making alongside technological progress. As this technology continues to shape the digital and social landscape, addressing these concerns will be crucial in ensuring that future innovations serve the broader public good rather than the digital elite.

References

AI Act. (n.d.). European Commission. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

Ali, S. (2014). Social and Environmental Impact of the Rare Earth Industries. *Resources*, *3*(1), 123–134. https://doi.org/10.3390/resources3010123

Arcuri, G., & Shivakumar, S. (2022, October 18). *Moore's Law and Its Practical Implications*. Center for Strategic & International Studies. https://www.csis.org/analysis/moores-law-and-its-practical-implications

Arik, M. (2023, September 21). *The concern around GPU shortages and how these could impact the AI revolution*. TechUK. https://www.techuk.org/resource/the-concern-around-gpu-shortages-and-how-these-could-impact-the-ai-revolution.html

Asia, N. (2025, January 31). *China makes inroads in DRAM chips in challenge to Samsung and Micron*. KrASIA. https://kr-asia.com/china-makes-inroads-in-dram-chips-in-challenge-to-samsung-and-micron

Basu, S. (2022). Three Decades of Social Construction of Technology: Dynamic Yet Fuzzy? The Methodological Conundrum. Social Epistemology, *37*(3), 259–275. https://doi.org/10.1080/02691728.2022.2120783

Brynjolfsson, E., & Unger, G. (2023). The Macroeconomics of Artificial Intelligence. *International Monetary Fund*. https://www.imf.org/en/Publications/fandd/issues/2023/12/Macroeconomics-of-artificial-intelligence-Brynjolfsson-Unger

Bushnell, L. (2021). *Educational disparities among racial and ethnic minority youth in the unites states*. Ballard Brief Research Library. https://ballardbrief.byu.edu/issue-briefs/educational-disparities-among-racial-and-ethnic-minority-youth-in-the-united-states

Celebrating Penn Engineering History: ENIAC. (n.d.). Penn Engineering. https://www.seas.upenn.edu/about/history-heritage/eniac/

Cho, A. (2023). *The short, spectacular life of that viral room-temperature superconductivity claim*. Science. https://www.science.org/content/article/short-spectacular-life-viral-room-temperature-superconductivity-claim

Dyer, D. (1995). Too Much Too Fast: The Dangers of Technological Momentum. *ERIC*, *Proceedings of the 1995 Annual National Convention of the Association for Educational Communications and Technology (AECT)*. https://eric.ed.gov/?id=ED383309

Electronics Donation and Recycling. (n.d.). U.S. Environmental Protection Agency. https://www.epa.gov/recycle/electronics-donation-and-recycling

Friis, S., & Riley, J. (2023, September 29). *Eliminating Algorithmic Bias Is Just the Beginning of Equitable AI*. Harvard Business Review. https://hbr.org/2023/09/eliminating-algorithmic-bias-is-just-the-beginning-of-equitable-ai

Harmful Effects. (n.d.). GreenCitizen. https://greencitizen.com/resources/harmful-effects/

Hersh, M. (2013). Science, technology and values: promoting ethics and social responsibility. *AI & Society: Knowledge, Culture and Communication*, *29*(2), 167–183. https://doi.org/10.1007/s00146-013-0473-z

How Long Should a GPU Actually Last? Plan for 3-5 Years. (2023, November 7). Unicorn Platform Blog. https://unicornplatform.com/blog/how-long-should-a-gpu-actually-last-expect-3-5-years/

Hughes, T. P. (1969). Technological Momentum in History: Hydrogenation in Germany 1898-1933. *Past & Present*, *44*(1), 106–132. https://doi.org/10.1093/past/44.1.106

Lee, S., Kim, J.-H., & Kwon, Y.-W. (2023, July 22). *The First Room-Temperature Ambient-Pressure Superconductor*. ArXiv. https://doi.org/10.48550/arXiv.2307.12008

Leese, B. (2023). The Cold War Computer Arms Race. *Journal of Advanced Military Studies*, *14*(2), 102–120. https://doi.org/10.21140/mcuj.20231402006

Lemonick, S. (2024, June 3). *Leslie Schoop on busting the LK-99 myth*. Chemical & Engineering News; American Chemical Society. https://cen.acs.org/materials/electronic-materials/Leslie-Schoop-debunking-claims-LK-99-room-temperature-superconductor/102/i17

Miller, D. (2019, September 3). *The Unintended Consequences of Technology Hype Overload*. CMSWire. https://www.cmswire.com/digital-experience/the-unintended-consequences-of-technology-hype-overload/

Moore's Law. (2023, September 18). Intel. https://www.intel.com/content/www/us/en/newsroom/resources/moores-law.html#gs.g68wmj

Nunez, C. (2021, May 20). *Rare earth supply disruptions have long-range impacts, computer model shows*. Argonne National Laboratory. https://www.anl.gov/article/rare-earth-supply-disruptions-have-longrange-impacts-compu ter-model-shows

Nye, D. (2004) . Electricity and Culture: Conceptualizing the American Case. *Annales historiques de l’électricité*, *2*(1), 125-137. https://doi.org/10.3917/ahe.002.0125

Optimum Processing Technology Inside Memory Arrays. (n.d.). Defense Advanced Research Projects Agency. https://www.darpa.mil/research/programs/optimum-processing-technology-inside-memory-arrays

Pingali, K. (2018, July 26). *The von Neumann Bottleneck Revisited*. ACM SIGARCH. https://www.sigarch.org/the-von-neumann-bottleneck-revisited/

Reducing the Digital Divide for Families: State and Local Policy Opportunities. (2024, July 8). National Council on Family Relations. https://www.ncfr.org/policy/research-and-policy-briefs/reducing-digital-divide-families-state-local-policy-opportunities

Samadi, A. H., Alipourian, M., Afroozeh, S., Raanaei, A., & Panahi, M. (2024). An Introduction to Institutional Inertia: Concepts, Types and Causes. *Institutional Inertia: Theory and Evidence*, 47–86. https://doi.org/10.1007/978-3-031-51175-2_3

Scranton, P. (1995). Determinism and Indeterminacy in the History of Technology. *Technology and Culture*, *36*(2), S31–S53. https://doi.org/10.1353/tech.1995.0004

Shalf, J. (2020). The future of computing beyond Moore’s Law. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical, and Engineering Sciences*, *378*(2166). https://doi.org/10.1098/rsta.2019.0061

Shedding Light on AI Bias with Real World Examples. (2023, October 16). IBM. https://www.ibm.com/think/topics/shedding-light-on-ai-bias-with-real-world-examples

Shittu, O. S., Williams, I. D., & Shaw, P. J. (2021). Global E-waste management: Can WEEE make a difference? A review of e-waste trends, legislation, contemporary issues and future challenges. *Waste Management*, *120*, 549–563. https://doi.org/10.1016/j.wasman.2020.10.016

Stone, H. S. (1970). A logic-in-memory computer,. *IEEE Transactions on Computers (TC)*, *Volume 100*, pp. 73-78.

Syntiant Introduces AI Chip for Low-Power Edge AI Deployment. (2022, April 25). Medical Device Network. https://www.medicaldevice-network.com/research-reports/syntiant-introduces-ai-chip-for-low-power-edge-ai-deployment/

Tabor, N. (2024). graphics processing unit. In *Encyclopedia Britannica*. https://www.britannica.com/technology/graphics-processing-unit

The Future of Semiconductor Miniaturization. (2024, December 30). Microchip USA. https://www.microchipusa.com/electrical-components/the-future-of-semiconductor-miniaturization/?utm_source=chatgpt.com

The Silicon Engine: Timeline. (n.d.). Computer History Museum. https://www.computerhistory.org/siliconengine/timeline/

Vogels, E. (2021, June 22). *Digital divide persists even as Americans with lower incomes make gains in tech adoption*. Pew Research Center. https://www.pewresearch.org/short-reads/2021/06/22/digital-divide-persists-even-as-americans-with-lower-incomes-make-gains-in-tech-adoption/

Wang, H., Yao, Y., Shi, K., Zhao, Y., Wu, H., Wu, Z., Geng, Z., Ye, S., & Chen, N. (2024). Possible Meissner effect near room temperature in copper-substituted lead apatite. *ArXiv*. https://doi.org/10.48550/arXiv.2401.00999

Wang, W., Yang, Y., Wang, D., & Huang, L. (2024). Toxic Effects of Rare Earth Elements on Human Health: A Review. *Toxics*, *12*(5), 317. https://doi.org/10.3390/toxics12050317

Wheeler, A. (2018, September 29). *What Raw Materials Are Used to Make Hardware in Computing Devices?* Engineering.com. https://www.engineering.com/what-raw-materials-are-used-to-make-hardware-in-computing-devices/

Williams, R. S. (2017). What's Next? [The end of Moore's law]. *Computing in Science & Engineering*, *19*(2), 7–13. https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7878940

Xie, Y. (2019). CRISP: Center for Research on Intelligent Storage and Processing-in-Memory. *2019 International Symposium on VLSI Design, Automation and Test (VLSI-DAT)*, 1–1. https://doi.org/10.1109/vlsi-dat.2019.8741577

Zewe, A. (2025, January 17). *Explained: Generative AI's environmental impact*. MIT News; Massachusetts Institute of Technology. https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117

Zhuk, A. (2023). Artificial Intelligence Impact on the Environment: Hidden Ecological Costs and Ethical-Legal Issues. *Journal of Digital Technologies and Law*, *1*(4), 932–954. https://doi.org/10.21202/jdtl.2023.40

Zou, X., Xu, S., Chen, X., Yan, L., & Han, Y. (2021). Breaking the von Neumann bottleneck: architecture-level processing-in-memory technology. *Science China Information Sciences*, *64*. https://doi.org/10.1007/s11432-020-3227-1