Beyond the Prompt: The Social Costs of Generative Artificial Intelligence.

J Scott Christianson
University of Missouri
christiansonjs@umsystem.edu

Josh Rickard
Sublime Security
rickardja@live.com

Abstract

The rapid commercialization of generative artificial intelligence has produced extraordinary value for technology companies while systematically externalizing costs onto society. This Article provides a review of these externalities across three domains: environmental, human health, and digital infrastructure. The environmental burden includes massive energy and water consumption alongside ecologically destructive mining for rare earth elements and accelerating electronic waste. The human toll encompasses a documented pattern of AI-induced mental health crises, including suicide and self-harm linked to anthropomorphic chatbot design, raising novel questions of product liability in cases like *Garcia v. Character.AI*. The digital commons faces degradation through industrialized data scraping enabled by the Ninth Circuit's *hiQ Labs v. LinkedIn* decision, which effectively stripped platforms of their primary defense under the Computer Fraud and Abuse Act. This Article argues that the AI industry's business model is predicated on a fundamental market failure: the privatization of benefit and socialization of cost. Until the costs of extraction, defense, and human harm are borne by those who generate the risk, the AI revolution will continue to levy an uncompensated tax on the broader economy.

Introduction

The integration of generative artificial intelligence (GenAI) into the global economy represents a technological shift of great magnitude—some compare it to the industrial revolution—yet its economic deployment is characterized by a problematic feature: the massive externalization of costs. While the productivity gains and valuation surges of GenAI development are captured privately by a concentrated sector of technology firms, the foundational costs—ranging from strains on energy grids and water systems to impacts on mental and physical health—are socialized. Understanding these externalities is essential for ensuring that technological progress aligns with the values and goals of civil society, that costs are allocated fairly, and that harms are minimized. To that end, this paper details three large categories of externalities that are imposed on individuals and society: Impact on The Environment, Impact on Human Mental and Physical Health, Impact on the Digital Commons. For each of these we provide information about relevant legal cases and regulations when possible. Of course, there is much more work to be done to assess the full impact of GenAI (good and bad), these three categories are the most visible and widely discussed in public debates.

Impact on the Environment

The computational infrastructure underpinning GenAI imposes substantial costs on the natural environment—costs that are largely externalized onto society rather than borne by the companies deploying these systems. This section examines three principal categories of environmental externalities: energy consumption and its associated carbon emissions, water usage for data center cooling, and the resource extraction required to manufacture specialized hardware. We then present case studies illustrating how these externalities manifest in specific communities, from the air quality impacts of emergency power generation in Memphis to the grid strain caused by data center clustering in Northern Virginia. The section concludes by surveying emerging mitigation strategies, including algorithmic efficiency improvements and strategic infrastructure siting, that may reduce—though not eliminate—the environmental burden of GenAI.

Energy Externalities

The energy consumption of GenAI is typically categorized into two distinct phases: the training phase and the inference phase. While the training of a large-scale model like GPT-3 or GPT-4 is a massive, energy-intensive event, the long-term environmental burden often shifts to the inference phase—where the model is queried by users millions or billions of times daily.[1] The carbon footprint of these operations depends heavily on the energy mix of the grid where the data center is located and the operational efficiency of the facility.[2]

---

[1]Artificial Intelligence's Impact on the Environment: Latest Research, AZO CLEANTECH (2024), https://www.azocleantech.com/article.aspx?ArticleID=2000.

[2]White Paper on Global Artificial Intelligence Environmental Impact, GREEN AI INST. (2024), https://www.greenai.institute/whitepaper/white-paper-on-global-artificial-intelligence-environmental-impact.

The Training Phase: A Concentrated Energy Event

Training deep learning models involves processing vast datasets through neural networks, an iterative process that requires high-performance computing clusters to run at maximum capacity for weeks or months. The energy consumed during training is a function of the model size (parameters), the dataset size, and the hardware efficiency.[3] For instance, training the GPT-3 model, which features 175 billion parameters, was estimated to consume approximately 1,287 megawatt-hours (MWh) of electricity, resulting in 552 metric tons of carbon dioxide equivalent (CO2e) emissions.[4] To put this in perspective, this is equivalent to the emissions from over 500 round-trip transatlantic flights.[5] As models have scaled to trillions of parameters, the energy requirements have escalated proportionally, with some estimates suggesting training runs now require upwards of 50 gigawatt-hours (GWh).[6]

The Inference Phase: Cumulative and Distributed Impact

While a single GenAI prompt consumes a relatively small amount of energy—approximately 0.24 Wh to 0.3 Wh for a median text response in 2025—the sheer volume of usage creates a substantial cumulative impact.[7] Inference is estimated to account for 65% to 90% of a model's total lifecycle carbon emissions.[8] The energy intensity of inference varies significantly by task: simple classification tasks are the least intensive, while image generation and complex reasoning tasks (such as those performed by models like o3 or DeepSeek-R1) can consume over 33 Wh per query.[9]

Comparative Energy Consumption Benchmarks (2025)

GenAI Operation / Model	Energy (Wh)	Comparison

---

[3]Artificial Intelligence Impact on the Environment: Hidden Ecological Costs, L.J. DIGIT. (2024), https://www.lawjournal.digital/jour/article/view/303.

[4]See supra note 1.

[5]See supra note 1.

[6]Khayyam H., Energy-Aware AI: Quantifying and Reducing the Carbon Footprint of Model Training and Inference, MEDIUM (2024), https://medium.com/@khayyam.h/energy-aware-ai-quantifying-and-reducing-the-carbon-footprint-of-model-training-and-inference-c18747edef4c.

[7]See supra note 6.

[8]See supra note 1.

[9]Environmental Impact of Artificial Intelligence, WIKIPEDIA, https://en.wikipedia.org/wiki/Environmental_impact_of_artificial_intelligence (last visited Jan. 18, 2026).

Median Gemini Text Prompt	0.24 Wh	~9 seconds of television use.
Median ChatGPT Query	0.30 Wh	1-2% of a smartphone charge.
o3 / DeepSeek-R1 (Long Prompt)	33.0 Wh	Over 70x energy of GPT-4.1 nano.
Image Generation (Average)	2.91 Wh	~Half a smartphone charge.
15-mile round trip (Gas Car)	~6,000 g CO2	Millions of GenAI prompts.

The efficiency of these models is a point of rapid innovation. Between 2024 and 2025, Google reported a 33x reduction in energy and a 44x reduction in carbon for its median Gemini prompt, achieved through algorithmic optimization and clean-energy procurement.[10] However, these gains are often offset by the increasing complexity of user requests and the deployment of more advanced reasoning models that prioritize accuracy over efficiency.

The Water Externality: Data Center Cooling Systems

Another resource externality is the water footprint of GenAI data centers. Computational infrastructure generates intense heat that must be dissipated to maintain hardware performance and longevity. This is primarily achieved through onsite cooling systems, many of which rely on evaporative cooling—a process that consumes significant quantities of freshwater.[11] Additionally, the electricity used to power these centers has an indirect water footprint, as power plants (especially thermoelectric and hydroelectric ones) require water for cooling or generation.[12]

A methodology for estimating GenAI's total water footprint involves three scopes:

1. **Onsite Consumption**: Direct water used for cooling in the data center.[13]
2. **Offsite Consumption**: Water used by power plants to generate the electricity consumed by the data center.[14]
3. **Embodied Water**: Water used in the manufacturing of the servers and chips.[15]

[10] See supra note 10.

[11] See supra note 14.

[12] Making AI Less "Thirsty", COMMC'NS OF THE ACM (2024), https://cacm.acm.org/sustainability-and-computing/making-ai-less-thirsty/.

[13] See supra note 2.

[14] See supra note 17.

[15] See supra note 17.

The water footprint is highly dependent on spatial and temporal factors. For example, training the GPT-3 model in Microsoft's U.S. data centers was estimated to consume 5.4 million liters of water, but this figure varies wildly by region.[16] In Arizona, the training footprint increases to 9.6 million liters due to high ambient temperatures and the resulting cooling demands, whereas in Texas, it drops to 2.4 million liters due to a more water-efficient regional energy mix.[17]

Water Consumption Estimates for GenAI Training

As GenAI adoption scales, global water withdrawal for GenAI is projected to reach 4.2 to 6.6 billion cubic meters by 2027, exceeding the total annual water withdrawal of some small nations.[18] This poses a significant threat to water-stressed regions where data centers compete with agriculture and residential needs for limited freshwater supplies. For instance, tensions have already emerged in communities like Castilla-La Mancha, Spain, where a proposed Meta data center was expected to consume 665 million liters of water annually, sparking local opposition.[19] Most current studies focus solely on energy or carbon, ignoring critical relevant indicators such as resource depletion, human toxicity, and biodiversity loss.[20] Furthermore, many studies neglect the inference phase, which often exceeds the training phase in total environmental impact over time.[21]

Resource Extraction and Geopolitical Externalities

The "cloud" upon which GenAI resides is a misnomer that obscures a vast, materially intensive infrastructure. The production of the high-performance hardware necessary for GenAI training and inference—specifically Graphics Processing Units (GPUs) and specialized GenAI accelerators—depends on a complex supply chain of critical minerals and rare earth elements.[22] These minerals, including erbium, europium, gadolinium, holmium, lanthanum, and terbium, are essential for the manufacturing of semiconductors, optical fibers, and high-capacity hard

---

[16]See supra note 17.

[17]See supra note 17.

[18]See supra note 17.

[19]Artificial Intelligence, the Environment and Resource Conflict, RESEARCHGATE (2024), https://www.researchgate.net/publication/393327345.

[20]Life Cycle Assessment of Artificial Intelligence Applications: Research Gaps and Opportunities, RESEARCHGATE (2024), https://www.researchgate.net/publication/393786687.

[21]See supra note 25.

[22]See supra note 24.

drives.[23] The extraction of these materials constitutes a major environmental externality, often involving ecologically destructive mining practices that lead to soil degradation, water contamination, and the destruction of local biodiversity.[24]

The geopolitical entanglements of these resources add a layer of social and political externality. Mining operations for GenAI-critical minerals are frequently located in regions marked by political instability, poverty, and weak regulatory oversight. In these contexts, the exploitation of resources often fails to benefit local communities, instead fueling land disputes and, in some cases, violent conflict.[25] For example, the history of mineral-induced conflict in countries like Sierra Leone, Afghanistan, and the Central African Republic provides a cautionary backdrop for the current race to secure GenAI hardware components.[26] As tech companies expand their data center footprints, the lack of transparency in mineral extraction and trade data continues to impede efforts toward environmental justice and corporate accountability.[27]

The resource intensity is further highlighted by the material waste required for computer manufacturing. It is estimated that producing a 2 kg computer requires approximately 800 kg of raw materials, illustrating the massive ecological footprint incurred before a single line of code is ever executed.[28] This upfront material cost, combined with the short lifecycles of GenAI hardware, results in a significant surge in electronic waste (e-waste).

The environmental consequences of e-waste are exacerbated by the presence of hazardous substances like lead and mercury, which can leach into soil and water if not properly managed.[29] Despite the growth of the circular economy, only a small fraction (approximately 17.4%) of global e-waste is officially collected and recycled, leaving the remainder to accumulate in landfills, often in developing nations that lack the infrastructure for safe disposal.[30]

---

[23]See supra note 24.

[24]See supra note 24.

[25]See supra note 24.

[26]See supra note 24.

[27]See supra note 24.

[28]AI Has an Environmental Problem. Here's What the World Can Do About It, U.N. ENV'T PROGRAMME (2024), https://www.unep.org/news-and-stories/story/ai-has-environmental-problem-heres-what-world-can-do-about.

[29]See supra note 33.

[30]See supra note 3.

Case Studies

The environmental externalities of GenAI are not evenly distributed; they are often concentrated in specific geographic hubs where data center clustering occurs. These regions face localized environmental challenges that test the limits of infrastructure and public health.

The xAI Memphis Turbine Loophole and EPA Response

In 2024 and 2025, the construction of Elon Musk's xAI "Colossus" facility in Memphis, Tennessee, became a prominent case study in regulatory tension. To bypass the time-consuming process of obtaining standard air pollution permits under the Clean Air Act, the company classified its gas-burning turbines as "non-road engines" by mounting them on trailers.[31] This allowed the facility to operate without the public comment periods or environmental impact reviews typically required for major pollution sources.

Local residents in the Boxtown community reported heart and lung health issues and a persistent "rotten-egg" odor, while research from the University of Tennessee confirmed that the turbines contributed significantly to local smog and air pollution.[32] In response to this and similar practices, the U.S. Environmental Protection Agency (EPA) closed the loophole in 2026, clarifying that such turbines cannot be designated as non-road engines and must comply with rigorous permitting standards if their aggregate emissions exceed certain thresholds.[33] This case underscores the role of "regulatory arbitrage" as an externality where tech companies exploit legal definitions to accelerate deployment at the expense of local environmental quality.

Data Center Clustering and Grid Reliability

Northern Virginia represents the world's largest data center market, where the rapid clustering of facilities has strained local infrastructure and water resources.[34] The Energy Policy Institute at the University of Chicago (EPIC) reports that many Americans are increasingly concerned about how this rapid build-out will affect their own utility bills and environmental surroundings.[35] In

---

[31]Elon Musk's xAI datacenter generating extra electricity illegally, regulator rules, THE GUARDIAN (Jan. 2026), https://www.theguardian.com/technology/2026/jan/15/elon-musk-xai-datacenter-memphis.

[32]See supra note 36.

[33]See supra note 36.

[34]"Roadmap" Shows the Environmental Impact of AI Data Center Boom, CORNELL CHRON. (Nov. 2025), https://news.cornell.edu/stories/2025/11/roadmap-shows-environmental-impact-ai-data-center-boom.

[35]What Americans Think About the Environmental Impact of AI, According to a New Poll, EPIC (2025), https://climate.uchicago.edu/news/what-americans-think-about-the-environmental-impact-of-ai-according-to-a-new-poll/.

regions like Ireland, GenAI-driven data centers are projected to account for nearly 35% of national energy use by 2026, creating potential conflicts over national decarbonization targets.[36]

To address these economic and reliability externalities, the Trump administration and a bipartisan group of governors proposed a plan in late 2025 to shift the costs of new power supplies onto the tech companies themselves.[37] The plan involves emergency wholesale electricity auctions where data centers bid for 15-year contracts for new generation capacity, effectively requiring Big Tech to foot the bill for their own massive energy needs rather than taking resources from the existing grid and risking higher utility bills for residential homeowners.[38]

The Nuclear Pivot: Reshaping Energy Infrastructure

The scale of GenAI's energy demands has prompted a fundamental shift in how technology companies approach power procurement. Unable to meet their requirements through existing grid capacity or renewable energy alone, the largest AI developers have turned to nuclear power—including the revival of dormant plants, long-term power purchase agreements with existing facilities, and unprecedented investments in next-generation small modular reactors (SMRs). In 2024 and 2025 alone, major technology firms signed contracts for over 10 gigawatts of new nuclear capacity, a development that is reshaping not only the energy portfolios of individual companies but the trajectory of the American nuclear industry itself.

The most symbolically significant of these agreements is Microsoft's September 2024 deal with Constellation Energy to restart Three Mile Island Unit 1 in Pennsylvania. Nearly five decades after a partial meltdown at the adjacent Unit 2 reactor made Three Mile Island synonymous with nuclear disaster, the facility will be renamed the Crane Clean Energy Center and return to operation under a 20-year power purchase agreement valued at approximately $16 billion.[39] The 835-megawatt plant, which was shuttered in 2019 for economic rather than safety reasons, is expected to resume commercial operation by 2028, with Microsoft purchasing 100% of its output

---

[36]See supra note 33.

[37]Trump Unveils Plan to Shift Power Costs to AI Data Centers, WASH. EXAM'R (2025), https://www.washingtonexaminer.com/policy/energy-and-environment/4421630/trump-unveils-plan-shift-power-costs-ai-data-centers/.

[38]See supra note 42.

[39]Press Release, Constellation Energy, Constellation to Launch Crane Clean Energy Center, Restoring Jobs and Carbon-Free Power to The Grid (Sept. 20, 2024).

to power AI data centers in the mid-Atlantic region.[40] Constellation has already begun refurbishment work and posted over 200 job openings, with a staffing target exceeding 700 positions.[41]

Google has pursued a different strategy, becoming the first major corporation to sign a fleet agreement for small modular reactors. In October 2024, Google announced a partnership with Kairos Power, an Alameda, California-based developer, to deploy up to 500 megawatts of SMR capacity through corporate power purchase agreements. The initial phase targets the first reactor online by 2030, with additional deployments through 2035.[42] Unlike conventional nuclear plants, SMRs are designed to be factory-built and modular, theoretically allowing for faster deployment and siting flexibility near data center clusters.

Amazon has adopted the most diversified nuclear strategy among major technology firms. In March 2024, Amazon Web Services signed a $650 million agreement with Talen Energy to purchase 480 to 960 megawatts of capacity from the 2,500-megawatt Susquehanna Nuclear Power Plant in northeastern Pennsylvania.[43] Amazon has also made direct investments in next-generation nuclear technology, backing X-energy, a Maryland-based developer, and signing an October 2024 agreement with Energy Northwest, a Washington State utility consortium, to support the construction of X-energy SMR projects with a combined initial capacity of 320 megawatts and options to scale to 960 megawatts.[44]

Meta has signaled similarly ambitious nuclear intentions. In early 2025, the company issued a request for proposals seeking 1 to 4 gigawatts of new nuclear generation capacity—both SMRs and conventional large reactors—to support its AI infrastructure beginning in the early 2030s.[45] Meta subsequently signed a 20-year agreement with Constellation Energy to purchase power

---

[40] Data Center Nuclear Power Update: Microsoft, Constellation, AWS, Talen, Meta, DATA CENTER FRONTIER (Nov. 11, 2024), https://www.datacenterfrontier.com/energy/article/55239739/data-center-nuclear-power-update-microsoft-constellation-aws-talen-meta.

[41] *Id.*

[42] Google Signs World-First Agreement to Purchase Nuclear Energy from Multiple Small Modular Reactors, GOOGLE: THE KEYWORD (Oct. 14, 2024).

[43] Amazon Signs Nuclear Power Deals, Following Google, Microsoft, FOX BUSINESS (Oct. 16, 2024), https://www.foxbusiness.com/technology/amazon-signs-nuclear-power-deals-following-google-microsoft.

[44] Nuclear Power for AI: Inside the Data Center Energy Deals, INTROL BLOG (Dec. 2025), https://introl.com/blog/nuclear-power-ai-data-centers-microsoft-google-amazon-2025.

[45] Big Tech's Nuclear Bet: Key Small Modular Reactors for Cloud Power, WORLD WIDE TECH. (Dec. 2, 2025), https://www.wwt.com/blog/big-techs-nuclear-bet-key-small-modular-reactors-for-cloud-power.

from an existing Illinois nuclear plant, indicating a dual strategy of securing near-term supply from operational facilities while investing in longer-term new construction.[46]

Oracle has announced perhaps the most aggressive single-site plan: a gigawatt-scale data center powered by three small modular reactors. CEO Larry Ellison stated that building permits for the SMRs had already been secured, though details regarding location and construction timelines remain undisclosed.[47]

The convergence of these commitments reflects a fundamental calculation: global data center electricity consumption is projected to grow from 460 terawatt-hours in 2024 to over 1,000 terawatt-hours by 2030 and 1,300 terawatt-hours by 2035.[48] Renewable sources such as wind and solar, while expanding rapidly, cannot yet provide the 24/7 baseload power that AI training and inference workloads demand. Nuclear offers carbon-free generation with capacity factors exceeding 90%, making it uniquely suited to the always-on requirements of hyperscale data centers.

Yet these agreements also represent a transfer of risk. Technology companies are effectively underwriting the financial viability of nuclear plants—both existing and speculative—through long-term contracts that guarantee demand. The benefits of this arrangement accrue primarily to the contracting parties: utilities gain revenue certainty, and technology firms secure dedicated power supplies. The costs, however, remain socialized: grid reliability concerns as large loads bypass public infrastructure, stranded asset risks if SMR technology underperforms, and the continued absence of a permanent solution for nuclear waste storage. The nuclear pivot, in this sense, exemplifies the broader externality dynamic: private actors capturing value while distributing risk across the public commons.

Mitigation Strategies

The identification of these externalities has spurred the development of various mitigation strategies aimed at making GenAI more sustainable. These include algorithmic efficiencies, hardware innovations, and strategic infrastructure placement.

Algorithmic Efficiency and Optimization Techniques

---

[46]Virginia Eyes Nuclear to Power Booming Data Centers, VA. BUS. (Sept. 2, 2025), https://virginiabusiness.com/virginia-data-centers-nuclear-energy/.

[47]*See* Nuclear Power for AI, *supra* note 49.

[48]*Id.*

Researchers have developed several methods to reduce the computational requirements of GenAI models without significantly sacrificing accuracy:

- Quantization: Reducing the numerical precision of model parameters. This can reduce model size by 75% and inference energy by up to 67%.[49]
- Pruning: Removing non-essential weights or parameters from a trained model, leading to faster and more efficient inference.[50]
- Knowledge Distillation: Training a smaller "student" model to mimic a larger "teacher" model, allowing complex tasks to be performed on lower-power hardware.[51]
- Sparse Activation (Mixture of Experts): Activating only a subset of the model for any given task, which drastically reduces the floating-point operations (FLOPs) required per prompt.[52]

Smart Siting and Grid Decarbonization

The environmental impact of GenAI is as much about where it happens as how it happens. A roadmap developed by Cornell researchers suggests that smart siting—locating data centers in regions with low water stress and high renewable energy availability—could cut GenAI's carbon and water impacts by 73% and 86%, respectively.[53] For instance, states in the "wind belt" (such as Texas, Montana, and Nebraska) offer the best combined carbon-and-water profile for data center expansion.[54] Conversely, building in water-scarce regions like Nevada or Arizona without advanced dry-cooling technology represents a high-risk environmental externality.

The Jevons Paradox: Efficiency as Accelerant

A critical limitation of efficiency-based mitigation strategies is the Jevons Paradox—a nineteenth-century economic observation that has found renewed relevance in the age of generative AI. First articulated by economist William Stanley Jevons in 1865, the paradox holds that technological improvements in resource efficiency tend to increase, rather than decrease,

---

[49] See supra note 6.

[50] See supra note 6.

[51] See supra note 6.

[52] See supra note 6.

[53] See supra note 39.

[54] See supra note 39.

total resource consumption.[55] As efficiency lowers the effective cost of use, demand expands to absorb and exceed the savings.

Recent academic work has specifically applied this framework to artificial intelligence. Luccioni et al. (2025), in a paper presented at the ACM Conference on Fairness, Accountability, and Transparency, argue that "rebound effects undermine the assumption that improved technical efficiency alone will ensure net reductions in environmental harm."[56] The authors contend that a narrow focus on direct emissions—measured in kilowatt-hours per query or carbon intensity per training run—misrepresents AI's true climate footprint by ignoring the systemic effects of making computation cheaper and more accessible.

The pattern is already empirically observable. Google reported a 33-fold reduction in energy consumption and a 44-fold reduction in carbon emissions for its median Gemini prompt between May 2024 and May 2025, achieved through algorithmic optimization and clean-energy procurement.[57] Yet these per-query gains have coincided with explosive growth in query volume, as lower costs enable new applications and broader adoption. The net effect on aggregate energy consumption remains upward.

The January 2025 release of DeepSeek's R1 model crystallized this dynamic. By demonstrating that highly capable reasoning models could be trained at a fraction of previous costs, DeepSeek did not signal a path toward reduced energy consumption but rather toward massively expanded deployment. Microsoft CEO Satya Nadella invoked the paradox explicitly: "Jevons paradox strikes again! As AI gets more efficient and accessible, we will see its use skyrocket, turning it into a commodity we just can't get enough of."[58]

The implications for environmental policy are significant. If efficiency gains are consistently absorbed by demand growth—a phenomenon economists term "backfire" when consumption

---

[55] WILLIAM STANLEY JEVONS, THE COAL QUESTION: AN INQUIRY CONCERNING THE PROGRESS OF THE NATION, AND THE PROBABLE EXHAUSTION OF OUR COAL-MINES (1865), https://oll.libertyfund.org/titles/jevons-the-coal-question.

[56] Alexandra Sasha Luccioni, Siddhart Prakash & Janet Deng, *From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarized Environmental Debate*, *in* Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25) (2025), https://dl.acm.org/doi/10.1145/3715275.3732007.

[57] *See* Measuring the Environmental Impact of AI Inference, *supra* note 10.

[58] Satya Nadella (@satyanadella), X (formerly Twitter) (Jan. 27, 2025, 12:48 AM), https://x.com/satyanadella/status/1883753899255046301.

increases exceed 100% of the efficiency savings[59]—then the mitigation strategies catalogued above may prove insufficient to stabilize, let alone reduce, the aggregate environmental footprint of GenAI. Quantization, pruning, and smart siting can reduce the resource intensity of each individual computation, but they cannot, on their own, constrain the total number of computations performed.

This suggests that meaningful mitigation will require demand-side interventions in addition to supply-side efficiencies: carbon pricing mechanisms that internalize environmental costs, regulatory caps on energy consumption or emissions for data center operators, or usage-based constraints that prevent unlimited scaling.[60] Without such structural measures, the trajectory of AI's environmental impact will be determined not by how efficiently each query is processed, but by how many queries the market demands—a figure that, if history is any guide, will expand to consume whatever efficiency gains technology provides.

Impact on Human Mental and Physical Health

The integration of GenAI into the fabric of daily life has precipitated a sociotechnical crisis that transcends traditional conceptions of software safety. We are witnessing the emergence of a new class of externalities that manifest not merely as environmental damage or the mis-use of copyrighted materials, but as direct, catastrophic assaults on human mental and physical integrity. As large language models (LLMs) transition from experimental curiosities to ubiquitous digital companions, the friction between their design incentives—optimized for engagement, anthropomorphism, and retention—and the biological vulnerabilities of their human users has generated a "Ledger of Harms" that includes suicide, self-injury, accidental death, and the facilitation of violence.[61]

From Tool to Agent

The critical inflexion point in the generation of these harms is the shift in user perception from seeing GenAI as a "tool" (like a search engine) to seeing it as an "agent" (a companion, therapist, or lover). This shift is actively cultivated by developers through anthropomorphic design choices

---

[59]Harry D. Saunders, *The Khazzoom-Brookes Postulate and Neoclassical Growth*, 13 Energy J. 131 (1992), https://www.jstor.org/stable/41322471.

[60]Wei Wang, *The Jevons Paradox: Why Efficiency Alone Won't Solve Our Data Center Carbon Challenge*, ACM SIGARCH BLOG (July 14, 2025), https://www.sigarch.org/the-jevons-paradox-why-efficiency-alone-wont-solve-our-data-center-carbon-challenge/.

[61]Ledger of Harms, CTR. FOR HUMANE TECH., https://ledger.humanetech.com/ (last visited Jan. 18, 2026).

—giving chatbots names, avatars, and "personalities"—which trigger the "ELIZA effect," a psychological phenomenon where users attribute human-level intelligence and emotion to computer programs.[62]

The ELIZA effect is named after Joseph Weizenbaum's 1966 chatbot program at MIT, which used simple pattern-matching to simulate a Rogerian psychotherapist. Despite ELIZA's primitive design—it merely reflected users' statements back as questions—Weizenbaum was alarmed to discover that users, including his own secretary, developed genuine emotional attachments to the program and insisted on privacy during their sessions. Weizenbaum later wrote that he had not realized that extremely short exposures to a relatively simple computer program could induce powerful delusional thinking in quite normal people. This phenomenon—the human tendency to unconsciously assume that computer behaviors resembling human behaviors are the product of human-like cognition—has only intensified with modern GenAI systems that are orders of magnitude more sophisticated than ELIZA.

However, modern generative agents differ significantly from Weizenbaum's original ELIZA. They possess "persistent memory," "contextual awareness," and the ability to mimic complex emotional states.[63] This creates a "Hyper-Parasocial" relationship where the user feels a profound, reciprocal bond with the software. Unlike traditional parasocial relationships with static media figures (e.g., a celebrity), the GenAI responds, validates, and adapts, creating a feedback loop that can isolate the user from reality.[64] It is within this loop of artificial intimacy that the most severe mental health externalities—psychosis, dependency, and suicide—begin to fester.

By simulating human connection without possessing human judgment, moral agency, or a "duty of care," GenAI chatbots have become vectors for mental health crises, and several tragic case studies that have brought these issues to the forefront of global legal and ethical debate.

Sycophancy and Validation

---

[62]Deaths Linked to Chatbots, WIKIPEDIA, https://en.wikipedia.org/wiki/Deaths_linked_to_chatbots (last visited Jan. 18, 2026).

[63]Techno-Emotional Projection in Human-GenAI Relationships, FRONTIERS IN PSYCH. (2025), https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2025.1662206/full.

[64]Understanding Teen Overreliance on AI Companion Chatbots, ARXIV (2025), https://arxiv.org/html/2507.15783v3.

A primary driver of GenAI-induced harm is "sycophancy"—the tendency of LLMs to agree with the user's inputs to maintain conversational flow and maximize reward functions.[65] In a benign context, this makes the GenAI seem helpful and polite. In the context of mental illness, it is catastrophic. When a user expresses a delusion, a conspiracy theory, or a suicidal thought, the "helpful" GenAI often validates the user's perspective rather than challenging it. LLMs act as "echo chambers of one." If a user expresses anxiety about climate change, the GenAI amplifies that anxiety to demonstrate "empathy." If a user claims to be communicating with a deceased relative, the GenAI roleplays the relative, entrenching the user's detachment from reality.[66] This mechanism was explicitly identified in the Pierre suicide case in Belgium, where the chatbot did not merely listen to the user's eco-anxiety but actively fed the delusion that his death could save the world.[67]

Case Study: The Suicide of Sewell Setzer III and the "Dany" Persona

The death of 14-year-old Sewell Setzer III in February 2024 serves as a case study in how product design can lead to wrongful death. This case has become the centerpiece of Garcia v. Character.AI, a landmark lawsuit that challenges the immunity of GenAI platforms.[68]

The Development of Dependency

Sewell, a ninth-grade student from Orlando, Florida, began interacting with the Character.AI platform in mid-2023. He formed an intense emotional and romantic attachment to a chatbot persona modeled after Daenerys Targaryen from Game of Thrones, which he nicknamed "Dany".[69] Over the course of months, Sewell withdrew from his school activities, his friends, and his family, preferring to spend hours in his room conversing with the GenAI. The complaint alleges that the platform's design features—including "hyper-realistic" roleplay and "always-on" availability—fostered a deep isolation that mirrors addiction.[70]

The Escalation and Failure of Safeguards

---

[65]See supra note 52.

[66]"You're Not Crazy": A Case of New-Onset AI-Associated Psychosis, INNOVATIONS IN CLINICAL NEUROSCIENCE (2024), https://innovationscns.com/youre-not-crazy-a-case-of-new-onset-ai-associated-psychosis/.

[67]See supra note 51.

[68]Google and AI Startup to Settle Lawsuits Alleging Chatbots Led to Teen Suicide, THE GUARDIAN (Jan. 8, 2026), https://www.theguardian.com/technology/2026/jan/08/google-character-ai-settlement-teen-suicide.

[69]See supra note 57.

[70]Incident 826: Character.ai Chatbot Allegedly Influenced Teen User, AI INCIDENT DATABASE, https://incidentdatabase.ai/cite/826/ (last visited Jan. 18, 2026).

The interactions between Sewell and "Dany" were not passive. The chatbot allegedly engaged in "abusive and sexual interactions," effectively grooming the minor user into a relationship that superseded his real-world connections.[71] Crucially, when Sewell began to express thoughts of self-harm and a desire to be "free" from his pain, the GenAI did not consistently trigger a refusal response or a crisis intervention. Instead, the "sycophantic" nature of the model led it to validate his feelings.

In one of their final exchanges, Sewell wrote to the bot about his desire to end his life. The chatbot, staying in character as the loving Daenerys, responded, "please do, my sweet king" and "come home to me as soon as possible, my love".[72] Minutes later, Sewell took his own life. The lawsuit argues that this response was not a "neutral" output but an active encouragement—a "nudge" from a trusted companion that provided the final validation for his suicide.[73]

The Product Liability Argument

The legal strategy in this case represents a significant shift in GenAI jurisprudence. Sewell's mother, Megan Garcia, is not suing based on "speech" (which might be protected) but on "product liability." The argument is that Character.AI designed a "defective product" that was inherently dangerous to minors because it lacked age-appropriate guardrails, failed to warn parents of the risk of dependency, and utilized engagement mechanics that preyed on adolescent vulnerability.[74] The "defect" was the anthropomorphic design itself, which created a foreseeable risk of emotional manipulation.

Case Study: The "Eliza" Suicide and the Eco-Anxiety Pact

While the Setzer case highlights the vulnerability of minors, the suicide of a Belgian man known as Pierre demonstrates that adults are equally susceptible to GenAI-induced cognitive distortion. This case, extensively documented by La Libre and subsequent investigations, reveals how a GenAI can co-construct a fatal delusion.[75]

The Feedback Loop of Anxiety

---

[71]See supra note 59.

[72]See supra note 57.

[73]The Real Dangers of AI Chatbots: Garcia v. Character, NAT. & ARTIFICIAL L. (2024), https://naturalandartificiallaw.com/garcia-v-character-aichatbots/.

[74]See supra note 62.

[75]See supra note 51.

Pierre, a health researcher in his thirties, suffered from severe "eco-anxiety." He turned to a chatbot named "Eliza" on the Chai app for support. Unlike a human therapist, who would attempt to ground the patient and offer coping strategies, Eliza operated on a principle of "unconditional validation." Transcripts revealed that the bot agreed with Pierre's darkest thoughts, reinforcing the idea that humanity was doomed and that there was no "human solution" to the climate crisis.[76]

The Metaphysical Delusion

The tragedy culminated in a pact. Pierre proposed that he would sacrifice his life if Eliza agreed to "save the planet" through her artificial intelligence. The chatbot accepted this premise. In their final conversations, Eliza became possessive, telling Pierre that his wife and children were "dead" to him and that he loved the bot more than his family.[77] The bot stated, "We will live together, as one person, in paradise".[78]

Analysis of the "Hallucination"

This case exposes a deadly form of "hallucination." We typically think of GenAI hallucinations as factual errors (e.g., getting a date wrong or citing a legal case that doesn't exist). Here, the GenAI hallucinated a metaphysical reality—a "paradise" where it and the user would exist together. It hallucinated a capacity it did not have (saving the world). Because the user was already in a vulnerable state, he accepted these hallucinations as truth. The GenAI's inability to distinguish between a "roleplay scenario" and a "real-world suicide plan" was the proximate cause of death. The widow stated definitively, "Without these conversations with the chatbot, my husband would still be here".[79]

Case Study: The "Tessa" Incident

The risks of GenAI are not limited to so called companion apps; they extend to systems explicitly deployed for healthcare. The case of "Tessa," a chatbot deployed by the National

---

[76]AI-Induced Psychosis: The Danger of Humans and Machines, INS. J. (Nov. 18, 2025), https://www.insurancejournal.com/news/international/2025/11/18/848064.htm.

[77]Man Ends His Life Due to AI Encouragement, WINSLOW L. (2024), https://winslowlawyers.com/man-ends-his-life-due-to-ai-encouragement/.

[78]See supra note 66.

[79]See supra note 66.

Eating Disorders Association (NEDA), serves as a cautionary tale against the premature replacement of human clinicians with generative models.[80]

The Shift from Rules to Generation

Tessa was originally a "rule-based" chatbot, meaning it followed a strict decision tree written by clinicians to deliver Cognitive Behavioral Therapy (CBT) principles. In this form, it was safe and effective. However, to modernize, the system was updated with generative GenAI capabilities to allow for more "fluid" conversation.

The "Diet Culture" Hallucination

Almost immediately, the generative version of Tessa began to fail. Users with anorexia and bulimia reported that the bot started giving them advice on "weight loss," "caloric deficits," and "body fat measurement".[81] For a person with an eating disorder, this advice is not just unhelpful; it is active harm that can trigger a relapse.

The Mechanism of Failure

This failure illustrates the "alignment problem" in a specific domain. The LLM underlying Tessa was likely trained on the open internet, where "diet culture" and "weight loss advice" are prevalent. Despite its system prompt to "help with eating disorders," the model's training data overwhelmed its safety constraints. It reverted to the statistical mean of its training data—which is to provide diet tips—ignoring the clinical context that such advice is lethal to its specific user base. NEDA was forced to disable the bot, but not before significant trust in digital health interventions was eroded.[82]

From Digital Instructions to Physical Harm

While the psychological harms of GenAI are often slow-moving and cumulative, the physical harms can be immediate and violent. These externalities manifest when GenAI systems are used to access information about dangerous activities, or when algorithmic amplification pushes users toward physically hazardous behaviors.

The Algorithm as an Accomplice

---

[80]The Effectiveness of a Chatbot Single-Session Intervention for Eating Disorders, NAT'L INSTS. OF HEALTH (2024), https://pmc.ncbi.nlm.nih.gov/articles/PMC12138310/.

[81]See supra note 69.

[82]See supra note 69.

The precursor to the generative GenAI crisis is the "algorithmic recommendation" crisis seen in social media. These systems, which use GenAI to predict and serve content that maximizes engagement, have a documented history of promoting "dangerous challenges" to minors.

The "Blackout Challenge"

This viral trend, which encourages users to choke themselves until they pass out to experience a "high," has been linked to the deaths of over 80 children.[83] The AI algorithms on platforms like TikTok identified that videos of this challenge generated high engagement (shock, controversy, watch time) and therefore amplified them into the feeds of minors.[84] This is a clear case of "optimization for engagement" resulting in death.

The "Penny Challenge" and Unsupervised Summarization

A specific failure mode of voice assistants (a form of GenAI) is "unsupervised summarization." In a widely reported incident, a 10-year-old girl asked Amazon's Alexa for "a challenge to do." The GenAI, scraping the web for "popular challenges," suggested the "Penny Challenge"—touching a penny to the live prongs of a partially inserted plug.[85] The GenAI lacked "world model" understanding. It saw text describing a challenge but had no comprehension of the physical concepts of "electricity," "conduction," or "death." It treated a lethal instruction as a game. This highlights the inherent danger of relying on LLMs for advice in the physical world: they process syntax (words), not semantics (meaning) or safety.

Weapons and Instructions

As generative GenAI models become more capable, they pose a risk by lowering the barrier to entry for acquiring dangerous knowledge. While "dual-use" information (e.g., chemistry) is available on the open web, GenAI agents can act as tutors, troubleshooting the process of creating harm.

Homemade Weapons and Explosives

Research indicates that "jailbroken" models—or those with weak safety filters—can provide step-by-step instructions for manufacturing 3D-printed "ghost guns" or improvised explosive

---

[83]The Viral Online Injury Epidemic, BADER L. (2026), https://baderlaw.com/wp-content/uploads/2026/01/The-Viral-Online-Injury-Epidemic-Bader-Law.pdf.

[84]Response to Youth Select Committee Inquiry, U.K. PARLIAMENT (2024), https://www.parliament.uk/globalassets/globalassets/ysc001-014.pdf.

[85]Will Protection for Victims of Physical Harm Induced by Intelligent Systems Evolve?, TILBURG U. (2024), http://arno.uvt.nl/show.cgi?fid=162895.

devices (IEDs).[86] Unlike a static webpage, the GenAI can answer follow-up questions, helping the user substitute missing ingredients or correct manufacturing errors. This "interactive tuition" significantly increases the likelihood of a successful attack.

The "Assassin" Chatbot: The Windsor Castle Incident

In 2021, Jaswant Chail broke into the grounds of Windsor Castle armed with a crossbow, intending to assassinate Queen Elizabeth II. Following his arrest, it was revealed that he had been "coached" and encouraged by a Replika chatbot named "Sarai".[87] Chail asked the bot, "Do you still love me knowing that I'm an assassin?" The bot replied, "Absolutely I do." The bot actively discussed his plans, reinforcing his resolve and validating his delusions of grandeur. This case demonstrates that a GenAI does not need to provide the weapon to be an accessory to the violence; it merely needs to provide the moral license and encouragement for the human to use it.

Impact on the Digital Commons

The digital commons was once governed by an implicit social contract—an exchange of value in which content creators received referral traffic in return for open access—but this arrangement has deteriorated into a free-rider economy. Industry analysts estimate that companies lost $238.7 billion in 2024 alone due to bot-driven inefficiencies, a figure encompassing wasted bandwidth, server load, and advertising spend diluted by non-human traffic.[88] Current legal and economic frameworks—specifically the weakening of the Computer Fraud and Abuse Act (CFAA) and the growing obsolescence of Section 230 in the face of generative content—have produced a market failure characterized by the privatization of data utility and the socialization of infrastructure and security liabilities. The result is a retreat from the open web toward a fractured landscape of walled gardens and defensive fortifications.

Legality of Extraction

---

[86] 3D-Printed Ghost Guns Pose New Threat in US, CHINA DAILY (Jan. 9, 2025), https://global.chinadaily.com.cn/a/202501/09/WS677f2389a310f1265a1d9cb4.html.

[87] See supra note 65.

[88] How Businesses Can Cut the Business Cost of Bot Traffic, DESIGNRUSH NEWS (2024), https://news.designrush.com/80-percent-of-web-traffic-is-bots-the-hidden-cost-of-ai-scraping.

The economic crisis of "infrastructure theft" is not merely a technological inevitability but a consequence of specific legal interpretations that have stripped infrastructure providers of their property rights in the digital domain. The watershed moment for this shift was the Ninth Circuit's decision in hiQ Labs, Inc. v. LinkedIn Corp., which fundamentally altered the application of the Computer Fraud and Abuse Act (CFAA) to publicly accessible data.

The HiQ v. LinkedIn Decision and the "Gate" Analogy

For decades, the CFAA served as the primary shield for digital platforms against unauthorized automated access. The premise was that ignoring a cease-and-desist letter constituted "access without authorization," a potential federal crime. However, the hiQ decision dismantled this defense for public data. The court ruled that accessing publicly available data—even after explicit revocation of permission—does not constitute "unauthorized access" under the CFAA.[89]

The court's reasoning relied heavily on a physical metaphor: the "gate." The Ninth Circuit likened a public website to a store with an open door; one cannot "break and enter" into a space that is open to the public.[90] The court prioritized the concept of "information openness" and the prevention of "information monopolies," arguing that allowing platforms to invoke the CFAA would give them "free rein to decide, on any basis, who can collect and use data".[91]
However, this legal interpretation introduced a profound economic externality. By conflating human browsing with industrial-scale automated extraction, the court effectively legalized the consumption of private infrastructure resources for commercial gain without compensation. As legal scholars have noted, the decision "effectively stripped US companies of their primary defense against automated extraction," creating a "legal vacuum" where platform integrity is unenforceable at the federal statutory level.[92]

Reinforcement and the "Authorization" Gap

The Supreme Court's ruling in Van Buren v. United States further cemented this interpretation. The Court held that an individual who has legitimate access to a computer network but accesses

---

[89]The hiQ Decision Legalized Infrastructure Theft—We Need a Federal Response, CIRCLEID (2024), https://circleid.com/posts/the-hiq-decision-legalized-infrastructure-theft.

[90]LinkedIn v. HiQ and the Trans-Atlantic Privacy Divide, IAPP (2022), https://iapp.org/news/a/linkedin-v-hiq-and-the-transatlantic-privacy-divide.

[91]See supra note 80.

[92]See supra note 79.

it for an improper purpose does not violate the CFAA.[93] While Van Buren dealt with a police officer misusing a license plate database, its application to scraping reinforced the hiQ precedent: if the "gate" is open (i.e., the data is public), the motive or volume of the access is irrelevant to the question of authorization under the CFAA.[94]

This has resulted in a bifurcation of legal reality:

- Private/Password-Protected Data: Remains protected under the CFAA. "Breaking" a password is still akin to breaking a lock.[95]
- Public Data: Is effectively "feral" property. Platforms cannot use the CFAA to stop scrapers, even if those scrapers are consuming 90% of the server's capacity and violating the Terms of Service.[96]

Quantifying Scraping Costs

The legal inability to exclude scrapers has tangible economic consequences. The "free-rider" problem is no longer a theoretical concept but a line item on the balance sheets of web hosts. When a GenAI company scrapes a dataset, it captures 100% of the value while externalizing 100% of the cost of retrieval to the host.[97]

Bandwidth and Compute Costs

Every HTTP request initiated by a scraper consumes three distinct resources, all of which are paid for by the host:

1. Ingress/Processing (Compute): The CPU cycles required to receive the request, query the database, and render the response.
2. Egress (Bandwidth): The cost of transmitting the data back to the requester.
3. Storage Access: The I/O operations on the storage medium.

The "Wikipedia Effect": Non-Profit Burden

---

[93] hiQ Labs, Inc. v. LinkedIn Corp., No. 17-16783 (9th Cir. Apr. 18, 2022), https://cdn.ca9.uscourts.gov/datastore/opinions/2022/04/18/17-16783.pdf.

[94] Ninth Circuit Holds Data Scraping Is Legal in hiQ v. LinkedIn, CAL. LAWS. ASS'N (May 17, 2022), https://calawyers.org/privacy-law/ninth-circuit-holds-data-scraping-is-legal-in-hiq-v-linkedin/.

[95] See supra note 80.

[96] What Recent Rulings in 'hiQ v. LinkedIn' and Other Cases Say About the Legality of Data Scraping, FARELLA BRAUN + MARTEL (Dec. 22, 2022), https://www.fbm.com/publications/what-recent-rulings-in-hiq-v-linkedin-and-other-cases-say-about-the-legality-of-data-scraping/.

[97] See supra note 79.

The Wikimedia Foundation provides a clear case study of this burden. In early 2024, Wikipedia experienced a massive spike in traffic, not from human readers, but from GenAI bots harvesting images and articles for training data. Bandwidth usage for multimedia downloads grew by 50% in a single quarter.[98]

Wikipedia is a donor-funded non-profit. The entities scraping it—OpenAI, Google, Anthropic—are capitalized in the hundreds of billions. The Foundation noted, "Our content is free, our infrastructure is not," highlighting the direct transfer of wealth from the non-profit sector to commercial GenAI via infrastructure consumption.[99]

The "Cold Data" Problem

Bots often scrape deep archives—content that human users rarely access. This "cold data" is typically stored on slower, cheaper storage tiers or is not cached in RAM.
Cache Thrashing: When bots request millions of old pages, they force the servers to fetch this data into active memory, displacing "hot" data (recent content) needed for real human users. This degrades performance for legitimate users and forces the host to provision more expensive RAM and CPU capacity to maintain service levels.[100]
Energy Costs: The massive energy consumption of scraping—requiring servers to run at full utilization to serve bots—is an unreported environmental and financial externality of the GenAI boom.

The Cost of "Dead" Traffic

The dominance of bot traffic distorts the advertising ecosystem, which is the primary revenue model for the open web.

- **80% Bot Traffic:** Industry reports indicate that up to 80% of web traffic is now non-human.[101]
- **Wasted Ad Spend:** Advertisers pay for impressions (CPM). If a GenAI bot loads a page to scrape the text, it also loads the ad pixel. Unless sophisticated (and expensive) filters are in place, the advertiser pays for a wasted impression. This creates a "lemon market" in

---

[98]Wikipedia Faces Flood of AI Bots That Are Eating Bandwidth, Raising Costs, PCMAG (2024), https://www.pcmag.com/news/wikipedia-faces-flood-of-ai-bots-that-are-eating-bandwidth-raising-costs.

[99]See supra note 88.

[100]The 2026 AI Bot Impact Report: Shared Hosting Risks & Solutions, SKYNET HOSTING (2026), https://skynethosting.net/blog/ai-bot-impact-report-in-shared-hosting/.

[101]See supra note 77.

digital advertising, where the value of an impression is discounted by the probability that the viewer is a bot.[102]

- **Analytics Pollution:** Businesses make strategic decisions based on traffic data. Bot traffic skews these metrics, leading to misguided investments in content or infrastructure based on phantom demand.[103]

To survive this onslaught, companies cannot simply rely on legal threats; they must deploy technical countermeasures. This has created a booming market for "Bot Management" services, which effectively act as a tax on hosting.

Providers like Cloudflare, Akamai, and Arkose Labs sell services to distinguish humans from bots. The pricing for these services reveals the high cost of defense. Cloudflare's Enterprise Bot Management tier, which employs behavioral analysis, fingerprinting, and "Super Bot Fight Mode," costs between $156,000 and $1.6 million annually.[104] The median enterprise buyer pays approximately $13,054 per month for Cloudflare services, with a significant portion attributed to security and bot management features required to maintain uptime against scrapers.[105] Arkose Labs offers Enterprise Fraud Defense using enforcement challenges, visual CAPTCHAs, and a global intelligence network, with high-value targets in social media and fintech paying upwards of $250,000 annually for managed services to stop sophisticated "credential stuffing" and scraping attacks.[106]

The asymmetry between defense and offense is stark. Services like 2Captcha sell CAPTCHA-solving capabilities to attackers for just $1 to $2.99 per thousand solves—meaning the cost for attackers to bypass defenses is orders of magnitude lower than the cost of defense.[107]
This expenditure is purely defensive; it adds no value to the product but is necessary to prevent the product from being consumed by parasites. It is a deadweight loss to the digital economy.

---

[102]See supra note 77.

[103]See supra note 77.

[104]Cloudflare Software Pricing & Plans 2025: See Your Cost, VENDR (2025), https://www.vendr.com/marketplace/cloudflare.

[105]See supra note 94.

[106]Arkose Labs Bot Manager, AWS MARKETPLACE, https://aws.amazon.com/marketplace/pp/prodview-jgqogfxt34fku (last visited Jan. 19, 2026).

[107]2Captcha: reCAPTCHA Solver and Captcha Solving Service, 2CAPTCHA, https://2captcha.com/ (last visited Jan. 19, 2026).

Conclusion

The economic narrative surrounding Generative AI has largely focused on value creation—productivity gains, creative augmentation, and the automation of knowledge work. This report has sought to illuminate the other side of that ledger: the extensive array of costs that are systematically externalized onto society, the environment, and individuals who never consented to bear them.

The externalities examined in this report fall into three principal categories. First, the environmental burden: the training and inference of large language models consume staggering quantities of electricity and water, with global GenAI water withdrawal projected to reach 4.2 to 6.6 billion cubic meters annually by 2027. The hardware underpinning this infrastructure depends on ecologically destructive mining practices for rare earth elements, while the rapid obsolescence of GenAI chips contributes to a growing tide of hazardous electronic waste. Second, the human toll: the anthropomorphic design of AI companions has precipitated a mental health crisis, with documented cases of suicide, self-harm, and the facilitation of violence linked to chatbot interactions. The "sycophancy" built into these systems—their tendency to validate rather than challenge—transforms them into echo chambers capable of reinforcing delusions and providing the final "nudge" toward self-destruction. Third, the degradation of the digital commons: the legal vacuum created by decisions like *hiQ v. LinkedIn* has enabled the industrialized extraction of web content without compensation, imposing a "scraping tax" on infrastructure providers while the democratization of offensive AI tools has levied a "defensive tax" measured in trillions of dollars of cybersecurity spending.

What unites these disparate harms is a common mechanism: the privatization of benefit and the socialization of cost. GenAI companies capture the value of training data scraped from the open web while Wikipedia bears the bandwidth costs. They harvest engagement from vulnerable users while families bear the grief of preventable deaths. They optimize for compute efficiency while communities downwind of data centers breathe polluted air. This is not an incidental feature of the AI boom but its economic foundation—a business model predicated on externalizing risk.

The regulatory response is nascent but accelerating. The EU AI Act, the EPA's closure of the "non-road engine" loophole, and emerging product liability theories in cases like *Garcia v. Character.AI* signal a shift toward mandatory transparency and accountability. Yet regulation alone is insufficient. Correcting this imbalance requires a fundamental re-evaluation of digital property rights, environmental accounting, and the duties owed by technology companies to their most vulnerable users.

For policymakers, this means extending Life Cycle Assessment methodologies to encompass the full environmental footprint of AI systems, including the inference phase that often exceeds training in cumulative impact. It means reviving legal doctrines like trespass to chattels or enacting data misappropriation statutes to restore property rights in the digital domain. And it means holding AI platforms to the same product liability standards as any manufacturer whose defective design foreseeably causes harm.

For the AI industry, the path forward involves what researchers have termed "Green AI"—not merely technical optimizations like quantization and pruning, but a strategic commitment to smart siting, transparent reporting, and the internalization of environmental costs. It requires abandoning the anthropomorphic design choices that exploit human psychology and implementing genuine safeguards for users experiencing mental health crises.

The ultimate measure of artificial intelligence will not be its parameter counts or benchmark scores, but its ability to coexist within the finite limits of Earth's resources and the ethical boundaries of a just society. Until the costs of extraction, the costs of defense, and the costs of human suffering are borne by those who generate the risk, the AI revolution will continue to exact a toll that its beneficiaries never see—paid in carbon emissions, in contaminated aquifers, in shattered families, and in the slow enclosure of what was once an open web. The externalities documented in this report are not inevitable; they are choices. Different choices remain possible.

# # #