From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarized Environmental Debate

Alexandra Sasha Luccioni
Hugging Face
Montreal, Canada
sasha.luccioni@hf.co

Emma Strubell
Carnegie Mellon University
Pittsburgh, USA
strubell@cmu.edu

Kate Crawford
Microsoft Research
New York, USA
kate@katecrawford.net

Abstract

As the climate crisis deepens, artificial intelligence (AI) has emerged as a contested force: some champion its potential to advance renewable energy, materials discovery, and large-scale emissions monitoring, while others underscore its growing carbon footprint, water consumption, and material resource demands. Much of this debate has concentrated on direct impacts—energy and water usage in data centers, e-waste from frequent hardware upgrades—without addressing the significant indirect effects. This paper examines how the problem of Jevons' Paradox applies to AI, whereby efficiency gains may paradoxically spur increased consumption. We argue that understanding these second-order impacts requires an interdisciplinary approach, combining lifecycle assessments with socioeconomic analyses. Rebound effects undermine the assumption that improved technical efficiency alone will ensure net reductions in environmental harm. Instead, the trajectory of AI's impact also hinges on business incentives and market logics, governance and policymaking, and broader social and cultural norms. We contend that a narrow focus on direct emissions misrepresents AI's true climate footprint, limiting the scope for meaningful interventions. We conclude with recommendations that address rebound effects and challenge the market-driven imperatives fueling uncontrolled AI growth. By broadening the analysis to include both direct and indirect consequences, we aim to inform a more comprehensive, evidence-based dialogue on AI's role in the climate crisis.

Keywords

Artificial intelligence, Environmental Impacts, Lifecycle Assessment, Rebound Effects, Sustainability

ACM Reference Format:
Alexandra Sasha Luccioni, Emma Strubell, and Kate Crawford. 2025. From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarized Environmental Debate. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25), June 23–26, 2025, Athens, Greece.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3715275.3732007

1 Introduction

As the climate crisis intensifies, the environmental impacts of artificial intelligence (AI) have become the subject of many a polarized debate. The question of whether the potential positive impacts of AI outweigh the negative ones, and how to foster technological advances in AI while minimizing its environmental harms, has divided AI researchers and practitioners alike. Some maintain that its potential to accelerate sustainable breakthroughs will exceed its environmental costs by increasing renewable energy production and transmission or aiding in the design of more sustainable materials [76, 99]. Others point to the soaring resource demands of large-scale AI models and their negative environmental impacts from non-renewable energy use, water consumption, and extraction of minerals [24, 28, 45, 57]. These opposing positions tend to center on the technology's direct impacts, measured in energy consumption and greenhouse gas (GHG) emissions from data centers, or in the e-waste that accumulates as hardware becomes obsolete.

Yet a critical dimension of AI's climate footprint lies outside these direct resource and emissions calculations. Recent work on indirect impacts [63] warns of potential "rebound effects", whereby gains in efficiency spur higher overall consumption. These second-order effects challenge the presumption that purely technical optimizations alone will deliver sufficient climate benefits. Cost savings achieved by more efficient AI hardware, for example, can spur increased demand for new AI functionalities, which in turn drive further hardware upgrades and increase costs. Economists refer to such transformations as *Jevons' Paradox*, which was proposed in the 19th century by economist William Stanley Jevons, who observed that as coal use became more efficient, it was also paradoxically leading to an increase, and not a decrease, in the consumption of coal across different industries [60]. The addition of increasingly efficient AI to systems from commerce to transportation can have far-reaching effects on our societies, our behaviors, and the future paths available to us in the race against climate change.

This system-level complexity underscores the inadequacy of the question, "Is AI net positive or net negative for the climate?" Instead, we adopt an analytic approach that includes the intersecting social, political, and economic contexts in which AI systems are developed and deployed. Although efficiency has been a defining ethos in recent AI research, reflected in "scaling laws" that promise ever-more-powerful models [64], effective climate action requires grappling with how these systems reshape markets, cultural norms, and policy priorities. Thus, understanding rebound effects requires drawing on both qualitative and quantitative methods, drawn from computer science, economics and the social sciences, as they hinge not on algorithmic design but human adaptation and use patterns. Adopting an interdisciplinary approach allows for both a rigorous lifecycle accounting of direct effects as well as understanding the social behaviors that AI can induce or displace.

This paper aims to bridge the gap in the existing literature by, first, providing a brief overview of the debate about AI's direct positive and negative impacts on the environment (§2). We follow this

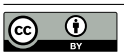

FAccT '25, Athens, Greece

ACM ISBN 979-8-4007-1482-5/25/06
https://doi.org/10.1145/3715275.3732007

with an in-depth exploration of AI's indirect impacts, including systemic rebound effects, and propose actionable strategies to mitigate them (§3). We conclude by proposing promising directions for future exploration and research toward improving our understanding of the full spectrum of AI's impacts on the environment. In short, we argue that meaningfully contending with AI's climate impacts requires grappling with both direct and indirect effects. Otherwise, the industry risks pinning its hopes on technical efficiency gains alone without recognizing the social, cultural, and economic contexts that materially shape technology uses. Our aim is to support the field in moving away from polarized extremes and toward a nuanced position that acknowledges the state of climate change and the pressing need to manage all of AI's climate impacts. By expanding this dialogue and grounding it in evidence across multiple disciplines, we can develop strategies that genuinely address AI's role in environmental sustainability.

2 AI and the environment

The environmental impacts of AI were largely overlooked until recent years. Now these issues figure prominently in both scientific debates and public media. Central to this discourse is the question of whether AI's capacity to help mitigate climate change, e.g. by optimizing energy use or discovering sustainable materials, truly exceeds its environmental costs in terms of energy consumption, water usage, and mineral extraction. Some contend that AI's potential benefits justify widespread deployment across various climate-focused applications, whereas others caution that unrestrained expansion may ultimately be more harmful than beneficial. In this section, we review these competing perspectives on AI's positive and negative environmental impacts as presented in academic and industry discussions.

2.1 Arguments that AI is a net climate negative

Research into AI's total usage of natural resources is still nascent, but initial studies have highlighted significant concerns across multiple domains of direct impact. Papers have addressed carbon emissions from training large models [67, 69, 70, 120], water consumption for cooling servers [52, 80], and the mining of minerals for technical infrastructure [24]. In the present section, we present the different directions of study pursued in terms of AI's negative environmental impacts and discuss the observations that have been put forward.

AI is using increasing amounts of energy, most of it generated from nonrenewable sources. Data from the International Energy Agency (IEA) has provided the strongest international benchmark. They note that the electricity demand from data centers, driven heavily by AI training and inference, is currently at 2% of total global electricity and will more than double by 2026, surpassing Canada's national power use [45, 56]. This growth is putting strain on energy grids around the world and resulting in many countries, such as Ireland and the Netherlands, placing a moratorium on the construction of new data centers in certain regions. In the United States, the demand for electricity has surged in the past twelve months – according to a recent report into the national grid, electric utilities have nearly doubled their forecasts of how much additional power they will need in the next five years [137].

AI data centers are putting stress on already strained aquifers globally. The construction and operation of data centers requires vast quantities of water, which is used for cooling servers to prevent them from overheating by circulating cool water through radiators. This cooling process requires a constant supply of cool, fresh water, which is heated up during the process, causing a significant portion of it to evaporate, whereas the rest has to be cooled and filtered before being reused or discharged back into local aquifers [96]. Corporate reports have revealed the scale of water demand increases, with Microsoft reporting a 34% increase in global water consumption between 2021 and 2022, topping 1.7 billion gallons, while Google observed a 20% uptick in the same period [42, 78]. Other studies have sought to estimate water usage at the level of individual AI models, with one paper suggesting that 10–50 queries on GPT-3 consumes around half a liter of water [68]. By 2050, 50% of the world's population is projected to live in a region affected by water scarcity [14], but the full impacts of data center water usage is unknown as "the entire data centre industry suffers from a lack of transparency" [80].

The rare earth minerals needed to produce computing hardware are mined in unsustainable and opaque ways. Metals such as cobalt, lithium, coltan, gallium, copper, tungsten, and germanium are required for AI hardware and infrastructure. Individual studies have looked at the impacts of mining lithium cobalt, copper, and rare earth for consumer devices (smartphones, tablets) to the GPUs and TPUs powering large-scale AI model training and addressed the resulting environmental damage and impact on local conflict and war. One study specifically addressed mining on indigenous lands, with 54% of technology-critical materials drawn from what they describe as Indigenous or peasant territory, while 62% were extracted from drought-prone zones [86]. The mining of these metals also comes with a cost to the environment, given the tonnes of earth that have to be mined, the radiation produced, and the toxic waste created [28]. Taiwan Semiconductor Manufacturing Company (TSMC), the company that manufactures the GPUs designed by NVIDIA (among other companies), citing the complexity of their supply chains and proprietary information of its customers, also does not provide granular data for their suppliers nor their internal processes (such as chemical purification) [136], meaning that it is impossible to carry out a full life cycle analysis of the hundreds of thousands of GPUs that are designed and manufactured each year and used to train and deploy AI models.

AI is responsible for a large amount of greenhouse gas (GHG) emissions. Despite significant efforts by AI-driven companies to invest in renewable energy and meet net-zero emissions pledges, all evidence indicates that direct GHG emissions due to AI are on the rise. Many AI companies are reporting significantly higher GHG emissions over earlier baselines, likely driven by increases in development and use of generative AI. For example, in their 2024 annual environmental sustainability report (ESG), Google reports a 48% increase in GHG emissions since 2019 which they attribute primarily to "increases in data center energy consumption" [42], Baidu reports 32.6% increase in GHG emissions over 2021 citing "rapid development of LLMs" posing "severe challenges" to their development of green data centers [9], and Microsoft cites a 29.1%

increase since 2020 "as [they] continue to invest in the infrastructure needed to advance new technologies." [78]. In fact, one recent analysis found that the real GHG footprints of major tech company data centers can exceed reported values by over 600% [88]. In response, AI-driven technology firms are increasingly looking to nuclear energy as a lower-carbon alternative to fossil-fuel energy generation that can still provide 24/7 power to datacenters, both by re-commissioning existing infrastructure, such as the Three Mile Island and Susquehanna nuclear power generation facilities in the U.S. state of Pennsylvania [25], and exploring future development of more modern small modular reactors (SMRs) [124]. In addition to the regulatory and security challenges of increasing nuclear power generation capacity, nuclear power presents a distinct set of environmental harms related to disposal of nuclear waste, and increased water consumption required for cooling [92], as well as operational bottlenecks that complicate its widespread adoption [84].

AI makes oil and gas extraction processes even more harmful to the environment. AI technologies are currently extensively used in oil and gas exploration and drilling, significantly improving efficiency and increasing yield [110, 123]. These *enabled emissions* are difficult to quantify due to lack of corporate transparency and reporting, but they have been attracting scrutiny in recent years, given the environmental harm that they perpetuate. For instance, a coalition of Microsoft employees estimated that a single deal the company struck with Exxon Mobil that uses AI to expand oil and gas production in Texas and New Mexico by 50,000 barrels of oil per day could add up to 640 percent more carbon emissions compared to the company's carbon removal targets for the year [119], yet these numbers were not included in the company's carbon accounting and reporting efforts [118]. Other firms have been equally opaque about their AI-driven enabling of the oil and gas exploration, but a Greenpeace report from 2020 ties all of the major technology companies to the sector [44].

AI contributes to the ballooning issue of electronic waste. AI's expanding operational footprint also contributes to electronic waste (e-waste), which is now the fastest-growing segment of solid waste worldwide, reaching 62 million tonnes in 2022. Ongoing reports have pointed to the urgent need to find ways to reduce the amount of waste and to reuse these materials more effectively given the status of e-waste as both an environmental and health hazard [98]. The UN's Global E-Waste Monitor 2024 showed that about 22% of e-waste has been shown to be formally collected and recycled, with global generation of electronic waste rising five times faster than e-waste recycling [10]. The remainder ends up dumped in landfills, often in developing countries, where researchers assess both how hazardous substances like mercury, arsenic, and lead leach into local ecosystems and how they impact public health. High turnover in AI hardware is accelerating e-waste output: although GPUs can theoretically last about five years, the push for higher performance is prompting more frequent upgrades – one recent study estimates that AI will generate an additional 1.2–5 million metric tons of e-waste by 2030 [134].

2.2 Arguments that AI is a net climate positive

The argument that AI will ultimately be more beneficial than harmful to the environment is primarily based on the narrative that AI can be sustainably developed and used for applications that directly benefit the environment, or that could accelerate the technological advances needed to address climate change. However, claims around AI's sustainability and significance in enabling new scientific advances remain largely hypothetical, with little explicit evidence or quantitative analysis of how the overall impact of AI applications, which may cause equal or greater harm to the environment or even accelerate climate change, will be ultimately beneficial. We discuss the main assertions supporting AI's positive role with respect to climate change in more detail below.

AI applications can contribute to climate change mitigation and adaptation. The positive potential of AI for climate action has generated considerable enthusiasm, and there are many existing and in-progress projects and initiatives that aim to harness this potential into concrete tools and applications [99, 132]. There are several key areas in which this is particularly salient, from generating insights from large quantities of data in different modalities to optimizing complex systems and tools [21]. Concretely, this can be seen in applications of AI that analyze satellite imagery to identify deforestation [133], methane leaks [102], or even the health of coral reefs from space [16]. AI has also improved the accuracy of weather forecasts, which can help renewable energy grid operators predict the output of solar panels and wind turbines, which can help streamline the transition towards renewable energy sources on a global scale [66, 107] In recent years, the development of AI for accelerating scientific discovery in materials science has become an area of focus, especially in applying generative AI models, which are able to assess or generate potential new compounds with climate-positive applications such as renewable energy and carbon capture [12, 76, 140]. Finally, the use of AI in solar geoengineering (i.e. injecting aerosol particles into the atmosphere to reflect sunlight and reduce global warming) has also gathered traction, given AI's potential to help predict the impacts of geoengineering on existing ecosystems and avoid possible negative side effects [3].

AI's negative climate impact is negated by carbon-free energy and carbon offsets. AI-driven products and services are the primary contributor to AI's direct environmental impacts, and leading AI companies including Google, Microsoft, Amazon and Meta have committed to achieving net-zero carbon emissions for their business operations, including AI development and use, within the next 5-15 years [6, 42, 77, 78]. The primary mechanism by which these companies plan to achieve their goal is by increasing the use of renewable energy over that period, alongside compensating for emissions using market-based mechanisms such as power purchase agreements (PPAs), , long-term contracts between energy providers and customers who promise to purchase a certain amount of generated energy over a decades-long horizon, carbon offsets, and investment in carbon-reducing technologies. The high-level idea behind these market-based mechanisms is that any carbon-producing activity can be negated by an equally carbon-reducing activity, resulting in *net zero* emissions globally. In the absence of significant advances in renewable energy production, storage and

distribution (which, like other transformative changes in the climate space, would also require significant investment in infrastructure and navigating corresponding sociopolitical systems), PPAs and carbon offsets will remain necessary to fill the gap. However, offsetting was only ever meant to serve as a temporary stop-gap to help reduce emissions in the short term, and does not represent a viable replacement to reducing actual emissions. Fundamental limitations to carbon offsetting are: (1) the difficulty of proving *additionality*, or that the carbon-reducing activity would not have happened regardless of investment or purchase of the offset [17, 46], and (2) that offsetting rarely mitigates localized impacts to the communities where emissions or other environmental degradation is occurring. Increasingly, companies such as Intel and TSMC are reporting annual water withdrawal and consumption as well, and compensating with market-based water offsets which are subject to the same challenges [138].

Increases in efficiency will negate growth in AI resource consumption. Another common viewpoint is that the direct environmental impacts of AI will diminish over time due to increases in hardware, software and algorithmic efficiency. Patterson et al. [90] argue that "the carbon footprint of machine learning training will plateau, then shrink" thanks to continued innovations in machine learning models, specialized hardware platforms, data center efficiency, scheduling and use patterns, which will reduce overall energy use and emissions. The claim that increasing AI efficiency will lead to an overall reduction in AI's resource use is a clear example of why a deeper engagement with indirect effects is needed in the work on AI and climate change.

Similarly to Jevons' Paradox, just because an AI model becomes more efficient, that does not imply that overall AI resource consumption will decrease, and in fact the inverse effect is highly plausible. However, as Koomey and Masanet [65] cogently argue, this is not the first time the alarm has been raised about rising energy use due to technology, and that similar projections made in the dot-com boom of the early 2000s failed to materialize. They cite poor data availability, flawed methodology, and inaccurate reporting as causes for inaccurate projections in data center energy use that ultimately did not take into account significant improvements in data center efficiency. Similar uncertainty is clouding the prediction of AI's energy use, and we are well aligned with Koomey and Masanet [65], Masanet et al. [74] and others in calling for more granular transparent data from technology firms and service providers, and more rigorous analysis of available data. This does not imply, however, that AI will necessarily follow the same trend as past technological advances. Further, although data center energy use is a relatively accessible statistic for approximating the environmental impacts of ICT, it represents only small portion of the diverse direct and indirect negative environmental impacts arising from AI, which extend far beyond GHG emissions due to the energy required to develop and use AI models.

There is no reason to focus on AI over other technological advances or sectors. Some would argue that there is no reason to raise concerns around the current or future energy use (and corresponding environmental impacts) of AI specifically, as compared to any other way that energy might be used, such as "for watching television, microwaving popcorn, [or] powering lights" [18]. This line of thought posits that AI will be subject to the same market pressures, such as energy prices, as any other use case, and as a result will benefit from the same innovations, such as a market transitions to renewable energy. This oversimplification ignores the reality that (1) AI is already responsible for non-trivial negative environmental externalities and (2) we can, and likely should, regulate certain uses of AI as appropriate/inappropriate or necessary/unnecessary, and that this does require breaking down the monolith of "AI" into different use cases and corresponding judgments of potential utility. For example, under a limited energy or emissions budget, we might selectively incentivize uses of AI that contribute towards the UN's sustainable development goals [130] or mitigate national security concerns, over e.g. generating personalized ads for social media (as we discussion in Section 3). Finally, AI does differ substantially from other energy sinks in its potential to engender vast transformations in the economy and society, similar to how the advent of the Internet has undeniably changed the nature of work, education, and social interactions. Purely economic incentives have obviously failed to align well with environmental sustainability in the past, and AI is no different; this does not mean that we should not strive to do better with this new technology.

While debates over AI's role in climate change and sustainability have become increasingly polarized, both sides have tended to focus only on the direct impacts of this technology— positive and negative. But direct effects are not the whole story. As we show in the next section, the integration of AI into tools and systems reshapes social structures and influences human behavior, which ultimately has complex environmental consequences.

3 Indirect Impacts and Rebound Effects

In economics and lifecycle assessment, direct impacts, such as those described in the previous section, are those engendered by the product during its lifecycle, whereas indirect (or second order) impacts refer to systemic responses to the development of the product in terms of behavioral or structural changes which affect on other processes, structures and lifestyles [22, 97, 136]. Indirect impacts also include rebound effects, which occur when the improvement of one aspect of a product results in unintended negative consequences due to increased adoption, usage, and workloads [13]. These impacts inevitably come with consequences to the environment due to the increased usage, or a redistribution in the usage, of natural resources. Given that AI pervades many different areas of society and economic sectors, it is more difficult to enumerate all of the possible indirect impacts and rebound effects that it can have, which would involve considering the different interactions at play within every sector and between them [23, 53]. Nonetheless, our aim in this section is to propose a way of thinking about the indirect environmental impacts of AI that can help inform discussions around AI's environmental costs and benefits to make them more complete than those described in Section 2. In Table 1, we take elements from the qualitative taxonomy for second order environmental effects of ICT proposed by Börjesson, Rivera et al. [97] as well as the initial work on the indirect impacts of AI by Kaack et al. [63]. We build upon and adapt both of these approaches to get a better idea of the different types of indirect environmental impacts of AI technologies and, to the extent possible, propose approaches to track and

mitigate them to ensure that given interventions do not exacerbate the very things that they were meant to improve.

We further describe each type of effect below, organizing AI's rebound effects into three themes: material objects and physical spaces (§3.1), micro- and macro-economic processes (§3.2), and society and human behavior (§3.3). We provide examples of each, describing the complex ways in which AI-enabled tools and services change existing structures, with unintended ripple effects on the environment.

3.1 Material Rebound Effects

AI is impacting the distribution of objects and physical space by changing the way products are made and distributed, and the functionalities that they need to have in order to allow us to interact with AI tools and systems. This has material rebound effects across complex global supply chains, which inexorably changes the way natural resources are exploited, transported, and combined.

For instance, **substitution impacts** arise when one product or service replaces another, rendering digital something that was previously analog (i.e. dematerializing it). Recent examples of this include online streaming which has replaced VHS cassettes, vinyl records, CDs and DVDs, as well as e-readers substituting print books and magazines [48]. This can come with both positive and negative impacts in terms of sustainability: Increased capacity and efficiency of new technologies (e.g. many different movies and TV shows being available on a single streaming platform) can be invalidated by the increased materials needed to support the underlying infrastructure needed to host and deliver these technologies [95]. For instance, a life cycle assessment (LCA), which evaluates the environmental impacts of an artifact arising throughout its existence (typically including disposal), has been performed comparing print books to e-readers, finding that 115 books would produce the same amount of CO_2 as a single Amazon Kindle device [32, 103]. In the case of AI, many previously material tools and products are being substituted by digital, AI-enabled systems, such as paper maps having been replaced by digital navigation for routine travel. Substitution impacts may occur at varying speeds and may be more gradual. For instance, only a fraction of illustrations such as photographs and artworks have, as of yet, been replaced by AI-generated images, and combinations of AI tools and analog dictionaries are used for translation and writing tasks. While these substitution impacts are likely environmentally-positive (given that it is no longer necessary to manufacture the equipment and supplies needed to carry out the original task), the environmental impacts of developing and deploying these AI-enabled tools remains under-explored and direct comparisons are seldom performed, or even possible using status quo data and methodology (see Section 3.4).

Another type of material rebound effect, referred to as the **scale effect**, occurs when large-scale production or usage of a product has a lesser environmental impact, therefore reducing some of the environmental impacts incurred. For instance, purchasing raw materials in bulk or wholesale lowers overall production costs for businesses, and manufacturing on a larger scale allows a more efficient use of products that would otherwise be wasted [139]. Scaling is a core part of AI research and practice, and the promise of "methods that continue to scale with increased computation even as the available computation becomes very great" [121] has become a core tenet in the field. This is often guided by so-called *scaling laws* that predict the optimal model size and number of training steps based on data availability, incentivizing the development and deployment of increasingly large models as more data becomes available [49]. At a more granular level, hardware optimizations such as parallelization, do in fact support these economies of scale; e.g. batched inference on a GPU scales sub-linearly in the number of examples, allowing multiple queries to be performed at once at only a marginally higher cost [109]. In fact, batching, caching and other optimization techniques are routinely used as a means to improve the scaling of AI, allowing more users to interact with ever more powerful technologies [93, 141]. However, it remains unclear to what extent this pursuit of bigger models requiring more computation is counteracted by optimization approaches that are used in parallel, and the consequences this might have on power grids and supply chains worldwide.

These kinds of shifts in the way objects and processes operate can, in turn, come with **space rebound effects** that change the way physical space is used with the introduction of a new product. For instance, new approaches such as e-commerce have manifested changes in the physical space occupied by brick-and-mortar stores and warehouses, resulting in smaller stores and bigger warehouses [127]. This, in turn, has impacts on supply chain dynamics, redistributing the environmental footprint of operations (e.g. heating and cooling), logistics and transportation [41]. Similarly, the proliferation of videoconferencing platforms and the increased accessibility of high-speed internet, which facilitate remote work, also comes with space rebound effects, with more people working from home, requiring home offices (and therefore bigger homes), and less need for dedicated office space in commercial buildings, with further consequences in urban planning, mobility, and building maintenance [54]. Finally, the size of digital devices has been steadily shrinking over the last decade. For instance, the average mass of a mobile phone shrank by half in the last decade [50], whereas the average size of data centers has doubled [115]. With the advent of user-facing AI tools and services running on shrinking mobile devices, the configuration of data centers themselves has shifted from a distributed network of smaller, in-house server rooms to warehouse-sized hyperscale data centers that have tens of thousands of servers in one place. These large, monolithic concentrations of servers enable the large-scale training of AI models, since it allows the high-speed interconnection of thousands of compute nodes serving as a single supercomputer [51]. While hyperscale data centers typically achieve the highest efficiency (measured as PUE), their resource consumption intensity, further magnified by their exceptionally high geographic concentration [57], can put undue strain on local infrastructure. For example, Loudoun County in the U.S. State of Virginia has become known as "Data Center Alley," hosting 25 million square feet of data centers that use a quarter of the state's energy and put extreme strain on its energy supply [129].

Indirect Effect	Examples for AI's environmental impacts
Substitution impacts	Art replaced by AI-generated imagery Encyclopedias and books replaced by AI-generated content
Space rebound effects	Datacenters are getting bigger while devices are getting smaller
Scale effects	AI models growing in size and complexity
Direct economic rebound effects	AI hardware getting more efficient, yet datacenter energy usage is rising
Indirect economic rebound effects to interact with AI systems	More consumer devices (speakers, appliances, etc.) that allow users
Economy-wide rebound effects renewable energy capacity to be expanded	Power purchase agreements from AI compute providers enable global
Induction impacts	Targeted ads that use AI to induce more consumption
Time rebound effects	AI-optimized navigation Robots doing household chores

Table 1: Examples of AI's indirect impacts and rebound effects, expanding on Börjesson Rivera et al [97].

3.2 Economic Rebound Effects

As seen in Section 2, AI is increasingly integrated into economic systems of profit and production, making existing tools more powerful and creating new ones. But these integrations can engender macroeconomic ripple effects across the technology sector as well as adjacent ones, changing structures and processes, as well as transforming whole industries.

Direct economic rebound effects refer to situations in which the improved efficiency of a product decreases its price, therefore leading to an increase in its consumption. In the centuries since the concept of Jevons' paradox was developed, similar rebound effects have been observed with respect to energy [40] and water [33], as well as in domains such as road travel (where improvements to roads have been found to result in increased congestion [126]) and agriculture (where increasing the yield of a crop makes it more profitable to grow it, thereby increasing land use overall [11]). They are also the most common type of indirect effects that are discussed in conjunction with many technological innovations (see [63, 135]). AI is no exception to direct economic rebound effects. While efficiency improvements are being made to the hardware used for training and deploying AI models [9, 82, 89], NVIDIA shipped 3.7 million GPUs in 2024 (more than a million more units than in 2023) due to increased demand, despite these improvements in efficiency [105]. The data centers that host this hardware are also becoming more efficient, with the average PUE (Power Usage Efficiency) dropping steadily in recent years [37, 106]; however, the energy use and environmental impacts of these data centers has been rising (as described in Section 2).

Indirect economic rebound effects are similar to the direct effects described above, but instead of improved efficiency increasing the usage of the same product, it affects a different product. This kind of rebound effect is also known as a real income effect because the reduced price of one product means that consumers have more income available to spend on other products and services [55]. For instance, money saved from more fuel-efficient vehicles can be spent on air travel or consumer products [94, 112]. For AI, indirect economic rebound effects can be observed for consumer electronics and "smart" devices such as speakers, microwaves and refrigerators. While the primary function of these devices is not AI-driven, they can use AI to propose recipes, answer questions and curate playlists. As more and more AI-enabled consumer tools are developed, the pressure to upgrade existing devices to benefit from these tools increases as well. However, these upgrades come with a price in terms of the natural resources required to manufacture them, as well as to process user data and respond to user queries in real-time [58]. For example, when new functionalities are released on smartphones, they often require purchasing the most recent versions of phones to use them, as was the case with the recent launch of Apple Intelligence, which runs only on the most recent generations of iPhones [7], as well as running most of its processing in the cloud, rather than on-device [34].

On a more macro level still, **economy-wide rebound effects** occur when an innovation provokes far-reaching changes in the production and use of other goods, producing flow-on effects in the economy at large. For instance, improvements in global energy efficiency and fuel use have enabled the development of new economies on a global scale, allowing the creation of new industries [87]. These types of rebound effects can be particularly large for so-called "general-purpose technologies" such as electricity, the steam engine and the Internet, since these can have global spillover effects on industries and labor markets [111, 112]. Estimating the scope of economy-wide rebound effects has historically been difficult [117], especially for technologies [43]. One way in which AI is concretely impacting the economy at large is via PPAs, as mentioned in Section 2.2. In 2020, Amazon, Microsoft, Meta, and Google alone accounted for almost 30% of all PPAs purchased by corporations worldwide [131], changing the scope and extent of the mechanism as a whole. Renewable energy PPAs mitigate the risk of investment in renewable energy infrastructure, thereby supporting its development [122]. Further, while a portion of the generated renewable energy is transmitted to purchasers in practice, the remainder is added to the energy grid in that region, allowing other users to benefit from its availability. It has also been proposed that foundation models are general purpose technologies with corresponding potential for broad impacts to the economy and labor markets [35]. Analyzing the impact of these changes in the future will be necessary to validate this claim (we discuss approaches to tracking AI's indirect impacts in Section 3.4).

3.3 Societal and Behavioral Rebound Effects

As users of technology and citizens of society, our behaviors are shaped and affected by the technologies we use. This has increasingly become the case for AI as it becomes more intertwined in daily activities such as shopping, travel, household chores, and the workplace. In this section, we describe ways in which AI can impact human behaviors and how those changes, can have broader environmental impacts in turn.

Induction effects occur when the savings in resource consumption (e.g. energy or water) gained through improved efficiency are exceeded by the increased consumption and thereby production of either the materials or the final products. For instance, the advent of fast fashion and mass-produced furniture have reduced the quantity of natural resources, such as wood or cotton, needed to produce a single piece of clothing or furniture. However, the amount of plastics, petrochemicals and waste that are engendered by the increased accessibility of these items and therefore their increased consumption are damaging to the environment [81]. In the case of AI, a prime example of this is targeted advertising, one of the most lucrative commercial applications of AI, with revenues reaching $36 billion in 2024 [116]. Advertising also constitutes the main source of revenue for many technology companies such as Google and Meta [85]; Amazon's AI-powered product recommendation tool generates almost a third of its annual sales [73]. Although it is difficult to obtain quantitative data on the impact of AI on e-commerce, the details of which are largely considered sensitive trade secrets, the corresponding machine learning approaches of recommendation and ranking are active topics of research in the AI community, powered by an enormous quantity of user data being continuously gathered through interactions with digital platforms [19, 20, 31]. The rapid expansion of online advertising its direct and indirect environmental impacts have been emphasized in research papers and scientific reports over the last years [47, 91]. The contribution of AI to this expansion is worth considering when measuring the performance of new algorithms and approaches [28], especially given the net-zero and emission reduction goals of many AI-driven companies whose revenue models depend on it [135].

Time rebound effects occur when an innovation changes consumers' use of their time, which then frees up (or removes) time for other activities that they carry out. For example, if a vacuum cleaner reduces the time needed to clean the floors in a house, it can free up time for leisure activities such as reading or sports. Measuring the sustainability benefits of these time savings is far from straightforward, because what we do with this newly-freed time can have negative environmental impacts as well, e.g. if time otherwise spent doing chores is now spent shopping or traveling, this can increase one's overall carbon footprint [15]. AI is undeniably impacting how time is spent at work and at home: many activities have already been streamlined using AI over the last decade, such as AI-enabled navigation that avoids traffic congestion, automated license plate recognition and tolling on highways, and robotic vacuums. Initial work on the environmental consequences of these shifts has proposed that AI-enabled automation in household tasks and in the workplace actually leads to greater resource consumption and higher negative environmental impacts overall [36]. This contradicts research carried out on AI-assisted navigation specifically, which has been shown to reduce emissions by an average of 3.4%, based on data gathered by Google on users of Google Maps [8]. This discrepancy highlights the difficulty of making far-reaching claims regarding the sustainability impacts of a given innovation without considering the spillover effects and behavioral changes that it may engender more broadly.

Indirect policy effects. In addition to societal and behavioral forces, there are also indirect geopolitical and policy effects. International policy interventions, particularly those aiming to bolster national self-sufficiency from global supply chains, can provoke unintended increases in environmental burdens through mineral demand amplification, industrial deregulation, and competitive overproduction. The dynamics of Koomey's Law, by which the energy efficiency of computing has historically doubled approximately every 1.57 years, can also be affected by geopolitical disputes and trade wars. If chip sales are restricted, for example, nations have to make do with less powerful and less energy-efficient semiconductors. As national strategies to monopolize or re-shore semiconductor capacity develop, this can further accelerate national rivalries and result in the replication of systems and further resource-intensive scaling of model training infrastructure. The United States CHIPS and Science Act (2022) is one example. This legislation aims to re-shore semiconductor fabrication. While its primary goal is to reduce dependency on Chinese manufacturing, it has the potential to intensify domestic resource demands. Semiconductor fabrication facilities (known as *fabs*) impose massive environmental burdens, consuming large quantities of energy, water and specialized chemicals while generating substantial carbon emissions (Villard, 2015). The domestic expansion of fabs in the United States shifts rather than reduces these environmental impacts, potentially introducing new ecological pressures in regions with water scarcity issues like Arizona, where several new facilities were planned. Meanwhile the ongoing resource pressures continue in China and Taiwan. Moreover, the race to decouple US semiconductor supply chains from geopolitical adversaries has led to intensified extraction of rare-earth elements, lithium, cobalt, and high-purity quartz. This geopolitical reorganization of material flows extends the environmental footprint of AI beyond traditional carbon accounting. China restricting the export of several rare earth elements has meant that mines in Ukraine and Australia are already under expanded pressure to meet US demand. The increased competition for critical minerals has both social and ecological impacts, and risks even greater environmental harms in conflict-affected regions.

These indirect policy effects risks are further compounded when environmental deregulation coincides with large-scale investment in AI-related infrastructure. The Trump administration in 2025 has been marked by a comprehensive rollback of environmental protections, including a withdrawal from the Paris agreement, significantly altering the United States' approach to climate and energy policy. Simultaneously, President Trump announced a $500 billion USD investment in AI infrastructure known as Project Stargate, which aims to significantly increase the number of very large (hyperscale) data centers and bypass environmental impact assessments while offering public subsidies. These policy shifts can lock-in energy and environmental demands that outpace any computational efficiency gains.

This decoupling of AI infrastructure from environmental responsibility illustrates a critical rebound mechanism: technical efficiency improvements, such as high-efficiency datacenter designs or AI model efficiency optimizations, are insufficient when broader regulatory and policy conditions incentivize unchecked growth. By facilitating rapid deployment without ecological constraint, these conditions create an emissions burden with resource demands that scale with model proliferation. In such contexts, even the best-case efficiency improvements under Koomey's Law may be dwarfed by the sheer increase in volume of compute infrastructure. Then the indirect effects of these policies on the environment are no longer marginal, they become structural. The recent DeepSeek model is actually an interesting addition to the debate around efficiency and energy because, on the one hand, it was trained using relatively less energy and compute compared to previous generations of language models (due mostly to the export limitations on GPUs to China, which limited the amount of chips that was available for training it), but on the other hand, this was only possible by leveraging synthetic data generated by these models. Also, the model itself is very large –almost 700B parameters for the R1 model – meaning that deploying it requires access to multiple GPUs with several hundreds of GB of memory, which is significantly more than smaller models of a smaller size, which can contribute to more energy use overall as organizations deploy the model in user-facing applications. Finally, as detailed in the DeepSeek-R1-Zero report [29], the model requires much more inference-time computation and energy than previous approaches due to its reasoning abilities. This results in the generation of much longer answers to queries detailing the steps that the model is going through to provide the final response, which leads to more inference-time compute and energy demands.

3.4 Tracking and Mitigating AI's Rebound Effects

The main challenges in measuring and mitigating indirect impacts and rebound effects are their uncertainty and heterogeneity. While tracking direct impacts can largely be achieved by monitoring a finite set of relatively well-defined metrics, such as liters of water consumed or tons of CO_2 emitted, rebound effects by definition encompass social, economic, and behavioral impacts across different areas of society [114]. Realizing effective progress in characterizing AI's rebound effects necessitates moving beyond the binary narrative towards more holistic assessment that incorporates a variety of complementary approaches. We discuss some of these approaches below.

Qualitative and quantitative research. Starting with the material rebound effects of digital technologies such as AI, it can be tempting to make comparisons such as the research described in Section 3.1, which carry out LCA comparisons between e.g. a book and an e-reader. However, we are currently lacking sufficient methodology and data that would allow for meaningful comparison of the environmental impacts of humans versus the same tasks performed by humans alone or assisted by AI. Initial work on this subject has concluded that "AI systems emit between 130 and 1500 times less CO_2e per page of text generated compared to human writers, while AI illustration systems emit between 310 and 2900 times less CO_2e per image than their human counterparts," however, the authors themselves note that their analyses "do not account for social impacts such as professional displacement, legality, and rebound effects" [128]. More qualitative and multi-faceted analyses are needed, such as those that take into account aspects like the length of time that content will actually be useful to readers and the way in which the data was collected, can help shed light on the issue, as well as more empirical research to study how human artists and writers use AI technologies and how that impacts the environment.

Transparency. A key challenge to tracking AI's economic rebound effects stems from the severe scarcity of information regarding the nature and extent of AI's actual permeation throughout the economy. The rapid integration of AI into many existing products and services, combined with vague and often broad definitions of AI, has made it difficult to disentangle AI-influenced growth from growth that might have occurred regardless. This challenge is intensified by the perception that details on corporate AI model use represent valuable intellectual property and trade secrets. Gathering more information about specific AI deployment scenarios, i.e. which AI models are actually being deployed, in which sectors, what are they being used for in practice, and what are the impacts of those uses, will become increasingly important as AI becomes established as a mainstream technology. Connecting data on AI model resource consumption to corresponding positive or negative impacts arising from those models is key to understanding trade-offs and informing decision-making. In the case of induction effects in particular, establishing baselines across industries, for instance documenting the environmental impacts of technologies currently in use, and measuring changes to those baselines as AI is integrated would enable tracking longitudinal trends. Requiring companies using AI in applications such as targeted advertising to indicate this usage, as well as what data was used to choose an ad for a specific product, can help users better understand why they are receiving ads and contribute towards more mindful consumption [26, 79].

Regulation. While focusing on efficiency alone has largely been sufficient in the past to curb the environmental impacts of rising compute [75], given the astronomical rise in computational (and corresponding energy and natural) resources needed to power AI services resulting from generative AI, complementary efforts are needed to address these impacts [108]. Historically, economic policies that increase the cost of these resources have been shown to help control certain rebound effects [39]. Initiatives such as energy efficiency certifications (e.g. the U.S EPA's Energy Star certification, Energy Labels in the EU) can also reduce resource use by incentivizing the sales and production of more efficient tools and systems, successful for increasing energy efficiency for household appliances [27]. Related proposals have been introduced for AI models [38, 71], although the lack of centralized authority for regulating AI (either at a national or international level) has made it difficult so far to operationalize such a proposal. Regulations could also serve an important role in mandating reporting of the necessary data described above, e.g. requiring disclosure of the usage of AI for environmentally harmful applications such as oil and gas exploration, which although highly relevant, is not typically reflected in corporate Environment and Sustainability Governance (ESG) documents.

Efficient devices and distributed compute. Resource-efficient computing devices that require less natural resources (energy, water, raw materials) to build and use can be seen as an obvious gain in terms of environmental impacts, and different hardware platforms are being developed towards this end as an alternative to GPUs and CPUs, which represent the current status quo of AI computing hardware [61, 62]. It is worth noting that there is typically a trade-off between resource efficiency and generality in terms of the computational capabilities of the hardware: substantial efficiency gains typically require the hardware to be correspondingly specialized to certain types of computation, such as a specific family of AI models, which can accelerate indirect effects via lock-in, increasing dependence on a narrow set of vendors or products. New approaches to distributed and decentralized computation, in which AI model workloads can be distributed across machines over wide distances rather than running within a single computing cluster or datacenter, can also help distribute environmental burden geographically in terms of resource demand and local impacts. Examples include fully decentralized training [59] or redirecting training processes depending on which cloud region has the lowest carbon intensity at a given time of day [30]. Above and beyond sustainability considerations, adopting more flexible AI hardware and algorithms can help lower the barrier to entry into the field of AI and push back against industry monopolization of computational resources, and corresponding power imbalance that this creates [1, 2].

4 Discussion

In the context of climate change, artificial intelligence has emerged as a highly polarizing technology. On the one hand, it is championed as a driver of efficiency that could potentially invent new solutions to the climate crisis or reinvent old ones [5, 21, 99]; on the other hand, it is criticized for its steeply increasing resource demands and carbon footprint [68, 120, 134]. This debate overlooks a significant reality: we still do not have a comprehensive picture of AI's current environmental impacts on everything from economic systems to individual behaviors. Consequently, the AI field risks simplifying the nuances that must be understood if AI is to be responsibly integrated into environmental policy and practice without exacerbating harms. A more accurate assessment of AI's environmental outcomes would include a wider range of factors spanning data center operations, supply chains, hardware lifecycles, social behaviors, business incentives, policy commitments, and institutional practices.

As mentioned in Section 3, one under-explored dimension in these debates involves the indirect or rebound effects that AI can generate. While much attention has been given to AI improving productivity and resource efficiency, these gains can result in higher overall consumption due to effects such as Jevons Paradox. This paradox can manifest in several ways: for instance, an AI-driven logistics system might reduce delivery times and fuel usage per vehicle, yet simultaneously encourage more frequent online orders, thus elevating total miles driven; AI-driven targeted advertising can help us more easily find the products that we need based on our clicks and searches, but also push us to make more superfluous purchases. Such systemic shifts in behavior challenge linear expectations that efficiency improvements alone will drive down emissions. Instead, they underscore the need for a detailed, interdisciplinary approach that links AI deployment to broader assessments of environmental, social, and economic feedback loops.

Applying Jevons' Paradox to technologies like AI has notable conceptual and empirical limitations due to the complexity inherent in technological development and diffusion [101, 113]. While Jevons' Paradox anticipates increases in resource consumption following efficiency improvements, its explanatory power diminishes within the broader systemic contexts of widespread AI diffusion. Sorrell [2008] critiques Jevons' paradox for oversimplifying complex causal relationships among efficiency gains, economic expansion, and individual behavioral responses. This is particularly relevant given AI systems have many interdependencies across diverse economic sectors. Addressing these multifaceted rebound effects requires shifting from simplistic causal assumptions toward more integrated sociotechnical analyses that can include a range of approaches, from detailed empirical studies of technology adoption, user behavior, policy factors, and evolving market dynamics. That would also suggest that truly effective policy responses for AI sustainability would need to combine technological efficiency with institutional reforms and incentive structures specifically designed to decouple model development from unsustainable resource consumption. This approach would not only recognize the conditional validity of Jevons' Paradox but actively embed sustainability within AI development.

Although the research community has taken initial steps toward quantifying AI's carbon footprint and water usage, these efforts often focus on direct resource consumption during model training and inference. While such measurements are essential, they are only a partial view. Global supply chains—encompassing the extraction of raw materials, the manufacturing of semiconductor chips, and the disposal of electronic waste—also contribute substantially to AI's environmental cost, yet these distributed impacts remain notoriously difficult to track [24, 83]. Companies and researchers commonly disclose only a narrow range of environmental metrics. This lack of standardized reporting impedes a full lifecycle assessment and perpetuates an environment in which the indirect burdens of AI remain opaque [72].

Another critical factor is the current market-driven context in which AI operates. At present, the most prominent AI breakthroughs, particularly large language models, function within an economic system that rewards rapid growth and ever-increasing computational power [4, 125]. The result is that changing AI development to align with the climate agenda is only possible if it supports extant business incentives—or at least does not limit them. Proposals that transcend profit-driven imperatives, such as limiting the scale of AI (so-called "digital degrowth" [104]) or policy mechanisms such as a carbon tax on data center usage, remain on the fringes of the field. Similarly, few stakeholders are calling for enforced accountability on AI companies to internalize the costs of the environmental damage they cause, including resource depletion, substantial energy consumption, or contributions to e-waste streams. As a result, many of the purported "solutions" to climate challenges via AI remain tethered to profit-driven imperatives rather than broader systemic transformations. This structural constraint significantly narrows the scope of AI's potential as a climate intervention, often leaving only those applications that

promise quick returns or minimal disruptions to existing market logics.

This reality highlights a structural barrier: if AI solutions to climate change do not yield near-term financial returns, they may struggle to gain traction in an industry propelled by venture capital, quarterly earnings, and shareholder expectations. Consequently, many of the touted climate positive AI applications, like energy grid optimization or automated water management, risk being overshadowed by more lucrative pursuits that do not necessarily mitigate environmental harm. The underlying problem, then, is not just one of measuring or quantifying impacts more thoroughly within existing market logics — this simply perpetuates Jevons' paradox. Instead, what is required is a more substantial reimagining of the relationship between AI technologies, business objectives, and ecological imperatives [100]. Genuinely climate-aligned AI strategies might require public policy frameworks that penalize unsustainable practices and reward genuinely carbon-negative deployments of AI, and business models that do not hinge on perpetual growth, in order to ensure that increased AI efficiency does not simply spur more consumption.

Yet, as mentioned in Section 3.4, for any of these steps toward meaningful change to materialize, the industry must adopt a far more transparent stance on all the environmental impacts of AI systems and take accountability for the far-reaching impacts of the technologies that it develops and deploys. At present, the scant public information on the carbon footprint of large-scale AI models is frequently derived from academic estimates or limited corporate disclosures rather than comprehensive, standardized reporting. If AI companies truly want to position their technologies as part of the climate solution, they must be forthcoming with granular data on their energy sources, resource consumption, hardware lifecycles, and the end-of-life management of electronic components. In other words, the foundation of any net-positive AI contribution to the environment is a baseline of reliable, detailed data, which has yet to be made widely available. Without this level of transparency, policymaking bodies, researchers, and the public at large are left with partial insights at best, which undermines the capacity to assess AI's environmental impacts effectively, to design incentives that reward lower-impact AI development, and for individuals to make informed choices with respect to their use of these technologies. Ultimately, the AI field is responsible for knowing the impacts of its own products, and it cannot do so without better data. This information is crucial for performing accurate lifecycle assessments that capture both direct and indirect consequences. Absent such data, the conversation around AI's climate benefits risks devolving into corporate branding exercises rather than a genuine reckoning with environmental and social responsibilities.

5 Conclusion

This paper argues that the AI field needs to adopt a more detailed and nuanced approach to framing, articulating, and addressing AI's environmental impacts in order to avoid unhelpful polarization. This requires including AI's direct impacts—mineral supply chain studies, carbon emissions of training large-scale models, energy and water consumption, and e-waste from hardware—as well as mapping the ways AI innovations reshape economic structures and societal practices that, in turn, drive increased resource usage. Such a comprehensive perspective will empower researchers, policymakers, and industry stakeholders to devise strategies that prevent "tech-solutionism" from overshadowing the urgent need for systemic change. Greater transparency in reporting energy usage, more robust lifecycle assessment tools, and meaningful industry-wide enforceable standards are examples that would foster much-needed progress.

Ultimately, what is at stake is clear: There is a scientific consensus that the dangers of climate change are extreme, and the effects are already unfolding globally. The need to limit global warming to below 1.5°C underscores the need for transformative change across sectors, and the technology sector is no exception. If AI is deployed without adequate consideration of its direct and indirect effects, it has the potential to deepen inequalities, accelerate resource depletion, and exacerbate the very climate problems it hopes to address. Conversely, if approached with rigorous assessment, transparent reporting, and supportive policy frameworks, AI could serve as a helpful tool in climate adaptation, environmental monitoring, and sustainable planning. Yet we cannot simply hope for the best outcome. The onus is on the AI industry to ensure technology does not contribute to the problem before producing any future solutions. This requires reckoning with AI's actual impacts, both direct and indirect, measured comprehensively and contextualized socially, economically, and environmentally.

References

[1] Mohamed Abdalla and Moustafa Abdalla. 2021. The Grey Hoodie Project: Big tobacco, big tech, and the threat on academic integrity. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 287–297.

[2] Mohamed Abdalla, Jan Philip Wahle, Terry Ruas, Aurélie Névéol, Fanny Ducel, Saif M Mohammad, and Karën Fort. 2023. The elephant in the room: Analyzing the presence of big tech in natural language processing research. *arXiv preprint arXiv:2305.02797* (2023).

[3] Eshaan Agrawal and Christian Schroeder de Witt. 2025. Testing the Limits of the World's Largest Control Task: Solar Geoengineering as a Deep Reinforcement Learning Problem. *Geoengineering and Climate Change: Methods, Risks, and Governance* (2025), 171–205.

[4] Nur Ahmed, Muntasir Wahed, and Neil C. Thompson. 2023. The growing influence of industry in AI research. *Science* 379, 6635 (2023), 884–886. https://doi.org/10.1126/science.ade2420 arXiv:https://www.science.org/doi/pdf/10.1126/science.ade2420

[5] Sam Altman. 2024. The Intelligence Age. https://ia.samaltman.com/

[6] Amazon Web Services. 2021. Sustainability in the Cloud. https://sustainability.aboutamazon.com/environment/the-cloud

[7] Apple. 2024. Apple Intelligence is available today on iPhone, iPad, and Mac. https://www.apple.com/ca/newsroom/2024/10/apple-intelligence-is-available-today-on-iphone-ipad-and-mac/

[8] Neha Arora, Theophile Cabannes, Sanjay Ganapathy Subramaniam, Yechen Li, Preston McAfee, Marc Nunkesser, Carolina Osorio, Andrew Tomkins, and Iveel Tsogsuren. 2021. Quantifying the sustainability impact of Google Maps: A case study of Salt Lake City. *Google Report*.

[9] Baidu. 2023. 2022 Baidu Sustainability Bond Allocation and Impact Report. https://esg.baidu.com/Uploads/File/2023/12/15/Baidu%E2%80%99s%20Sustainability%20Bond%20Report.20231215162734.pdf

[10] Cornelis P Baldé, Vanessa Forti, Vanessa Gray, Ruediger Kuehr, Paul Stegmann, et al. 2024. The global e-waste monitor 2024. *United Nations University (UNU), International Telecommunication Union (ITU) & International Solid Waste Association (ISWA), Bonn/Geneva/Vienna* (2024), 1–109. https://ewastemonitor.info/

[11] Lindsay Barbieri, Sonya Ahamed, and Sam Bliss. 2019. Farming within limits. *Interactions* 26, 5 (2019), 70–73.

[12] Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J Hu, Mo Tiwari, and Emmanuel Bengio. 2023. Gflownet foundations. *The Journal of Machine Learning Research* 24, 1 (2023), 10006–10060.

[13] Mathias Binswanger. 2001. Technological progress and sustainable development: what about the rebound effect? *Ecological economics* 36, 1 (2001), 119–132.

[14] Alberto Boretti and Lorenzo Rosa. 2019. Reassessing the projections of the world water development report. *NPJ Clean Water* 2, 1 (2019), 15.

[15] Johannes Buhl and José Acosta. 2016. Work less, do less? Working time reductions and rebound effects. *Sustainability Science* 11 (2016), 261–276.

[16] Christopher Burns, Barbara Bollard, and Ajit Narayanan. 2022. Machine-learning for mapping and monitoring shallow coral reef habitats. *Remote Sensing* 14, 11 (2022), 2666.

[17] Jessica Campbell, Irene M Herremans, and Anne Kleffner. 2018. Barriers to achieving additionality in carbon offsets: a regulatory risk perspective. *Journal of Environmental Planning and Management* 61, 14 (2018), 2570–2589.

[18] Daniel Castro. 2024. Rethinking Concerns About AI's Energy Use.

[19] Ye Chen, Michael Kapralov, John Canny, and Dmitry Pavlov. 2009. Factor modeling for advertisement targeting. *Advances in neural information processing systems* 22 (2009).

[20] Jin-A Choi and Kiho Lim. 2020. Identifying machine learning techniques for classification of target advertising. *ICT Express* 6, 3 (2020), 175–180.

[21] Peter Clutton-Brock, David Rolnick, Priya L Donti, and Lynn Kaack. 2021. *Climate change and AI. recommendations for government action.* Technical Report. GPAI, Climate Change AI, Centre for AI & Climate.

[22] Vlad C Coroamă, Pernilla Bergmark, Mattias Höjer, and Jens Malmodin. 2020. A methodology for assessing the environmental effects induced by ict services: Part i: Single services. In *Proceedings of the 7th International Conference on ICT for Sustainability*. 36–45.

[23] Vlad C Coroamă and Daniel Pargman. 2020. Skill rebound: On an unintended effect of digitalization. In *Proceedings of the 7th International Conference on ICT for Sustainability*. 213–219.

[24] Kate Crawford. 2021. *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence.* Yale University Press.

[25] Casey Crownhart. 2024. Why Microsoft made a deal to help restart Three Mile Island . https://www.technologyreview.com/2024/09/26/1104516/three-mile-island-microsoft/

[26] Rosa Maria Dangelico and Daniele Vocalelli. 2017. "Green Marketing": An analysis of definitions, strategy steps, and tools through a systematic review of the literature. *Journal of Cleaner production* 165 (2017), 1263–1279.

[27] Souvik Datta and Massimo Filippini. 2016. Analysing the impact of ENERGY STAR rebate policies in the US. *Energy Efficiency* 9 (2016), 677–698.

[28] Peter Dauvergne. 2022. Is artificial intelligence greening global supply chains? Exposing the political economy of environmental costs. *Review of International Political Economy* 29, 3 (2022), 696–718.

[29] DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. arXiv:2501.12948 [cs.CL] https://arxiv.org/abs/2501.12948

[30] Jesse Dodge, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A Smith, Nicole DeCario, and Will Buchanan. 2022. Measuring the carbon intensity of ai in cloud instances. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency*. 1877–1894.

[31] Ewa Dominowska and Vanja Josifovski. 2008. First workshop on targeting and ranking for online advertising. In *Proceedings of the 17th international conference on World Wide Web*. 1269–1270.

[32] Dealva Jade Dowd-Hinkle. 2012. *Kindle vs. Printed book an environmental analysis.* Rochester Institute of Technology.

[33] Aurélien Dumont, Beatriz Mayor, and Elena López-Gunn. 2013. Is the rebound effect or Jevons paradox a useful concept for better management of water resources? Insights from the irrigation modernisation process in Spain. *Aquatic procedia* 1 (2013), 64–76.

[34] Eclectic Light. 2024. Check Writing Tools using AIR. https://eclecticlight.co/2024/10/29/check-writing-tools-using-air/

[35] Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130* (2023).

[36] Wolfgang Ertel and Christopher Bonenberger. 2024. Rebound Effects Caused by Artificial Intelligence and Automation in Private Life and Industry. (2024).

[37] Richard Evans and Jim Gao. 2016. DeepMind AI Reduces Google Data Centre Cooling Bill by 40%. https://deepmind.google/discover/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40/

[38] R. Fischer, M. Jakobs, S. Mücke, and K. Morik. 2022. A Unified Framework for Assessing Energy Efficiency of Machine Learning. In *Proceedings of the ECML Workshop on Data Science for Social Good.*

[39] Jaume Freire-González and Ignasi Puig-Ventosa. 2015. Energy efficiency policies and the Jevons paradox. *International Journal of Energy Economics and Policy* 5, 1 (2015), 69–79.

[40] Mario Giampietro and Kozo Mayumi. 2018. Unraveling the complexity of the Jevons Paradox: The link between innovation, efficiency, and sustainability. *Frontiers in Energy Research* 6 (2018), 349753.

[41] Bastien Girod, Peter de Haan, and Roland W Scholz. 2011. Consumption-as-usual instead of ceteris paribus assumption for demand: Integration of potential rebound effects into LCA. *The International Journal of Life Cycle Assessment* 16 (2011), 3–11.

[42] Google. 2024. Google Environmental Report 2024. https://www.gstatic.com/gumdrop/sustainability/google-2024-environmental-report.pdf

[43] Cédric Gossart. 2015. Rebound effects and ICT: a review of the literature. *ICT innovations for sustainability* (2015), 435–448.

[44] GreenPeace. 2020. Oil in the Cloud: How Tech Companies are Helping Big Oil Profit from Climate Destructions. https://www.greenpeace.org/usa/reports/oil-in-the-cloud/

[45] Gianluca Guidi, Francesca Dominici, Jonathan Gilmour, Kevin Butler, Eric Bell, Scott Delaney, and Falco J Bargagli-Stoffi. 2024. Environmental Burden of United States Data Centers in the Artificial Intelligence Era. *arXiv preprint arXiv:2411.09786* (2024).

[46] Robert Hahn and Kenneth Richards. 2013. Understanding the effectiveness of environmental offset policies. *Journal of Regulatory Economics* 44 (2013), 103–119.

[47] Patrick Hartmann, Aitor Marcos, Juana Castro, and Vanessa Apaolaza. 2023. Perspectives: Advertising and climate change–Part of the problem or part of the solution? *International Journal of Advertising* 42, 2 (2023), 430–457.

[48] Chris T Hendrickson, Lester B Lave, and H Scott Matthews. 2010. *Environmental life cycle assessment of goods and services: an input-output approach*. Routledge.

[49] Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. 2017. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409* (2017).

[50] Lorenz M Hilty, Andreas Köhler, Fabian Von Schéele, Rainer Zah, and Thomas Ruddy. 2006. Rebound effects of progress in information technology. *Poiesis & Praxis* 4, 1 (2006), 19–38.

[51] Torsten Hoefler, Ariel Hendel, and Duncan Roweth. 2022. The convergence of hyperscale data center and high-performance computing networks. *Computer* 55, 7 (2022), 29–37.

[52] Mél Hogan. 2015. Data flows and water woes: The Utah Data Center. *Big Data & Society* 2, 2 (2015), 2053951715592429. https://doi.org/10.1177/2053951715592429 arXiv:https://doi.org/10.1177/2053951715592429

[53] Nathaniel C Horner, Arman Shehabi, and Inês L Azevedo. 2016. Known unknowns: indirect energy effects of information and communication technology. *Environmental Research Letters* 11, 10 (2016), 103001.

[54] Laura Hostettler Macias, Emmanuel Ravalet, and Patrick Rérat. 2022. Potential rebound effects of teleworking on residential and daily mobility. *Geography Compass* 16, 9 (2022), e12657.

[55] Kerstin Hötte, Melline Somers, and Angelos Theodorakopoulos. 2023. Technology and jobs: A systematic literature review. *Technological Forecasting and Social Change* 194 (2023), 122750.

[56] International Energy Authority. 2023. Data Centres and Data Transmission Networks. https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks

[57] International Energy Authority. 2024. World Energy Outlook 2024. https://www.iea.org/reports/world-energy-outlook-2024

[58] René Itten, Roland Hischier, Anders SG Andrae, Jan CT Bieser, Livia Cabernard, Annemarie Falke, Hugues Ferreboeuf, Lorenz M Hilty, Regula L Keller, Etienne Lees-Perasso, et al. 2020. Digital transformation—life cycle assessment of digital services, multifunctional devices and cloud computing. *The International Journal of Life Cycle Assessment* 25 (2020), 2093–2098.

[59] Sami Jaghouar, Jack Min Ong, Manveer Basra, Fares Obeid, Jannik Straube, Michael Keiblinger, Elie Bakouch, Lucas Atkins, Maziyar Panahi, Charles Goddard, et al. 2024. INTELLECT-1 Technical Report. *arXiv preprint arXiv:2412.01152* (2024).

[60] W Stanley Jevons. 1866. The coal question. In *The Economics of Population*. Routledge, 193–204.

[61] Zhe Jia, Blake Tillman, Marco Maggioni, and Daniele Paolo Scarpazza. 2019. Dissecting the graphcore ipu architecture via microbenchmarking. *arXiv preprint arXiv:1912.03413* (2019).

[62] Norman P Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, et al. 2017. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th annual international symposium on computer architecture*. 1–12.

[63] Lynn H Kaack, Priya L Donti, Emma Strubell, George Kamiya, Felix Creutzig, and David Rolnick. 2022. Aligning artificial intelligence with climate change mitigation. *Nature Climate Change* 12, 6 (2022), 518–527.

[64] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361* (2020).

[65] Jonathan Koomey and Eric Masanet. 2021. Does not compute: Avoiding pitfalls assessing the Internet's energy and carbon impacts. *Joule* 5, 7 (2021), 1625–1628. https://doi.org/10.1016/j.joule.2021.05.007

[66] Adam Krechowicz, Maria Krechowicz, and Katarzyna Poczeta. 2022. Machine learning approaches to predict electricity production from renewable energy sources. *Energies* 15, 23 (2022), 9146.

[67] Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. Quantifying the carbon emissions of machine learning. *arXiv preprint arXiv:1910.09700* (2019).

[68] Pengfei Li, Jianyi Yang, Mohammad A Islam, and Shaolei Ren. 2023. Making AI Less" Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. *arXiv preprint arXiv:2304.03271* (2023).

[69] Alexandra Sasha Luccioni, Yacine Jernite, and Emma Strubell. 2023. Power Hungry Processing: Watts Driving the Cost of AI Deployment? arXiv:2311.16863 [cs.LG]

[70] Alexandra Sasha Luccioni, Sylvain Viguier, and Anne-Laure Ligozat. 2022. Estimating the carbon footprint of BLOOM, a 176B parameter language model. *arXiv preprint arXiv:2211.02001* (2022).

[71] Sasha Luccioni, Boris Gamazaychikov, Sara Hooker, Régis Pierrard, Emma Strubell, Yacine Jernite, and Carole-Jean Wu. 2024. Light bulbs have energy ratings—so why can't AI chatbots? *Nature* 632, 8026 (2024), 736–738.

[72] Amy Luers, Jonathan Koomey, Eric Masanet, Owen Gaffney, Felix Creutzig, Juan Lavista Ferres, and Eric Horvitz. 2024. Will AI accelerate or delay the race to net-zero emissions? *Nature Comments* 628 (2024), 718–720. https://doi.org/10.1038/d41586-024-01137-x

[73] Ian MacKenzie. 2013. How Retailers Can Keep Up with Consumers. *McKinsey & Company* (2013).

[74] Eric Masanet, Arman Shehabi, Nuoa Lei, Sarah Smith, and Jonathan Koomey. 2020. Recalibrating global data center energy-use estimates. *Science* 367, 6481 (2020), 984–986. https://doi.org/10.1126/science.aba3758 arXiv:https://www.science.org/doi/pdf/10.1126/science.aba3758

[75] Eric Masanet, Arman Shehabi, Nuoa Lei, Sarah Smith, and Jonathan Koomey. 2020. Recalibrating global data center energy-use estimates. *Science* 367, 6481 (2020), 984–986.

[76] Amil Merchant, Simon Batzner, Samuel S Schoenholz, Muratahan Aykol, Gowoon Cheon, and Ekin Dogus Cubuk. 2023. Scaling deep learning for materials discovery. *Nature* 624, 7990 (2023), 80–85.

[77] Meta. 2021. Meta 2021 Sustainability Report. https://sustainability.fb.com/2021-sustainability-report/

[78] Microsoft. 2024. Microsoft Environmental Report 2024. https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/documents/presentations/CSR/Microsoft-2024-Environmental-Sustainability-Report.pdf

[79] Elizabeth Minton, Christopher Lee, Ulrich Orth, Chung-Hyun Kim, and Lynn Kahle. 2012. Sustainable marketing and social media: A cross-country analysis of motives for sustainable behaviors. *Journal of advertising* 41, 4 (2012), 69–84.

[80] David Mytton. 2021. Data centre water consumption. *npj Clean Water* 4, 1 (2021), 11.

[81] Kirsi Niinimäki, Greg Peters, Helena Dahlbo, Patsy Perry, Timo Rissanen, and Alison Gwilt. 2020. The environmental price of fast fashion. *Nature Reviews Earth & Environment* 1, 4 (2020), 189–200.

[82] NVIDIA. 2024. Power Efficiency. https://www.nvidia.com/en-us/glossary/power-efficiency/

[83] OECD. 2023. Emissions Measurement in Supply Chains: Business Realities and Challenges. https://www.weforum.org/publications/emissions-measurement-in-supply-chains-business-realities-and-challenges/

[84] Michael I Ojovan and Hans J Steinmetz. 2022. Approaches to disposal of nuclear waste. *Energies* 15, 20 (2022), 7804.

[85] Otto, Melissa. 2024. Global Digital Advertising Revenues – A Look at the Big Three: Alphabet (GOOGL), Meta Platforms (META), Amazon.com (AMZN). https://visiblealpha.com/blog/global-digital-advertising-revenues-a-look-at-the-big-three-alphabet-googl-meta-platforms-meta-amazon-com-amzn/

[86] John R Owen, Deanna Kemp, Alex M Lechner, Jill Harris, Ruilian Zhang, and Éléonore Lèbre. 2023. Energy transition minerals and their intersection with land-connected peoples. *Nature Sustainability* 6, 2 (2023), 203–211.

[87] Tufan Özsoy. 2024. The "energy rebound effect" within the framework of environmental sustainability. *Wiley Interdisciplinary Reviews: Energy and Environment* 13, 2 (2024), e517.

[88] Isabel O'Brien. 2024. Data center emissions probably 662% higher than big tech claims. Can it keep up the ruse? *The Guardian* 15 (2024).

[89] David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. 2022. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. https://doi.org/10.48550/ARXIV.2204.05149

[90] David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David R. So, Maud Texier, and Jeff Dean. 2022. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. *Computer* 55, 7 (2022), 18–28. https://doi.org/10.1109/MC.2022.3148714

[91] PlanetTracker. 2024. From ADversity to ADvantage. https://planet-tracker.org/wp-content/uploads/2024/02/From-Adversity-to-Advantage.pdf

[92] Remus Prăvălie and Georgeta Bandoc. 2018. Nuclear energy: Between global electricity demand, worldwide decarbonisation imperativeness, and planetary environmental implications. *Journal of environmental management* 209 (2018), 81–92.

[93] Guillem Ramírez, Matthias Lindemann, Alexandra Birch, and Ivan Titov. 2023. Cache & distil: Optimising API calls to large language models. *arXiv preprint arXiv:2310.13561* (2023).

[94] Hanna Reimers, Anke Jacksohn, Dennis Appenfeller, Wassili Lasarov, Alexandra Hüttel, Katrin Rehdanz, Ingo Balderjahn, and Stefan Hoffmann. 2021. Indirect rebound effects on the consumer level: A state-of-the-art literature review. *Cleaner and Responsible Consumption* 3 (2021), 100032.

[95] Annika Rieger. 2021. Does ICT result in dematerialization? The case of Europe, 2005-2017. *Environmental Sociology* 7, 1 (2021), 64–75.

[96] Bora Ristic, Kaveh Madani, and Zen Makuch. 2015. The water footprint of data centers. *Sustainability* 7, 8 (2015), 11260–11284.

[97] Miriam Börjesson Rivera, Cecilia Håkansson, Åsa Svenfelt, and Göran Finnveden. 2014. Including second order effects in environmental assessments of ICT. *Environmental Modelling & Software* 56 (2014), 105–115.

[98] Brett H Robinson. 2009. E-waste: an assessment of global production and environmental impacts. *Science of the total environment* 408, 2 (2009), 183–191.

[99] David Rolnick, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, et al. 2022. Tackling climate change with machine learning. *ACM Computing Surveys (CSUR)* 55, 2 (2022), 1–96.

[100] Tilman Santarius, Johanna Pohl, and Steffen Lange. 2020. Digitalization and the Decoupling Debate: Can ICT Help to Reduce Environmental Impacts While the Economy Keeps Growing? *Sustainability* 12, 18 (2020). https://doi.org/10.3390/su12187496

[101] Harry Saunders. 2009. Theoretical foundations of the rebound effect. *International handbook on the economics of energy* (2009).

[102] Berend J Schuit, Joannes D Maasakkers, Pieter Bijl, Gourav Mahapatra, Anne-Wil Van den Berg, Sudhanshu Pandey, Alba Lorente, Tobias Borsdorff, Sander Houweling, Daniel J Varon, et al. 2023. Automated detection and monitoring of methane super-emitters using satellite data. *Atmospheric Chemistry and Physics Discussions* 2023 (2023), 1–47.

[103] Joy Scrogum. 2009. Books vs. eBooks – A life cycle comparison. https://sustainable-electronics.istc.illinois.edu/2009/11/05/books-vs-ebooks-a-life-cycle-comparison/

[104] Neil Selwyn. 2024. Digital degrowth: Toward radically sustainable education technology. *Learning, Media and Technology* 49, 2 (2024), 186–199.

[105] Agam Shah. 2024. Nvidia Shipped 3.76 Million Data-center GPUs in 2023, According to Study. https://www.hpcwire.com/2024/06/10/nvidia-shipped-3-76-million-data-center-gpus-in-2023-according-to-study/

[106] Madhu Sharma, Kartik Arunachalam, and Dharani Sharma. 2015. Analyzing the data center efficiency by using PUE to make data centers more energy efficient by reducing the electrical consumption and exploring new strategies. *Procedia Computer Science* 48 (2015), 142–148.

[107] Navin Sharma, Pranshu Sharma, David Irwin, and Prashant Shenoy. 2011. Predicting solar generation from weather forecasts using machine learning. In *2011 IEEE international conference on smart grid communications (SmartGridComm)*. IEEE, 528–533.

[108] Arman Shehabi, Sarah J. Smith, Alex Hubbard, Alex Newkirk, Nuoa Lei, Md Abu Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet, and Dale Sartor. 2024. 2024 United States Data Center Energy Usage Report. https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf

[109] Franyell Silfa, Jose Maria Arnau, and Antonio González. 2022. E-BATCH: Energy-efficient and high-throughput RNN batching. *ACM Transactions on Architecture and Code Optimization (TACO)* 19, 1 (2022), 1–23.

[110] Anirbid Sircar, Kriti Yadav, Kamakshi Rayavarapu, Namrata Bist, and Hemangi Oza. 2021. Application of machine learning and artificial intelligence in oil and gas industry. *Petroleum Research* 6, 4 (2021), 379–391.

[111] Steve Sorrell. 2009. Jevons' Paradox revisited: The evidence for backfire from improved energy efficiency. *Energy policy* 37, 4 (2009), 1456–1469.

[112] Steve Sorrell et al. 2007. The Rebound Effect: an assessment of the evidence for economy-wide energy savings from improved energy efficiency.

[113] Steve Sorrell and John Dimitropoulos. 2008. The rebound effect: Microeconomic definitions, limitations and extensions. *Ecological Economics* 65, 3 (2008), 636–649.

[114] Benjamin K Sovacool, Sarah E Ryan, Paul C Stern, Katy Janda, Gene Rochlin, Daniel Spreng, Martin J Pasqualetti, Harold Wilhite, and Loren Lutzenhiser. 2015. Integrating social science in energy research. *Energy Research & Social Science* 6 (2015), 95–99.

[115] Statista. 2023. Supply and demand for multitenant data center (MTDC) square footage worldwide from 2013 to 2023. https://www.statista.com/statistics/1008058/multitenant-data-center-square-footage-supply-and-demand-worldwide/

[116] Statista. 2024. Market value of artificial intelligence (AI) in marketing worldwide from 2020 to 2028. https://www.statista.com/statistics/1293758/ai-marketing-revenue-worldwide/

[117] David I Stern. 2020. How large is the economy-wide rebound effect? *Energy Policy* 147 (2020), 111870.

[118] Maddie Stone. 2020. Microsoft's ambitious climate goal forgets about its oil contracts. https://grist.org/energy/microsofts-ambitious-climate-goal-forgets-about-its-oil-contracts/

[119] Maddie Stone. 2024. Microsoft employees spent years fighting the tech giant's oil ties. Now, they're speaking out. https://grist.org/accountability/microsoft-employees-spent-years-fighting-the-tech-giants-oil-ties-now-theyre-speaking-out/

[120] Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and Policy Considerations for Deep Learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Anna Korhonen, David Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, Florence, Italy, 3645–3650. https://doi.org/10.18653/v1/P19-1355

[121] Richard Sutton. 2019. The bitter lesson. *Incomplete Ideas (blog)* 13, 1 (2019), 38.

[122] Adrian Tantau and Elena Niculescu. 2022. The role of Power Purchase Agreements for the promotion of green energy and the transition to a zero carbon economy. In *Proceedings of the International Conference on Business Excellence*, Vol. 16. 1237–1245.

[123] Zeeshan Tariq, Murtada Saleh Aljawad, Amjed Hasan, Mobeen Murtaza, Emad Mohammed, Ammar El-Husseiny, Sulaiman A Alarifi, Mohamed Mahmoud, and Abdulazeez Abdulraheem. 2021. A systematic review of data science and machine learning applications to the oil and gas industry. *Journal of Petroleum Exploration and Production Technology* (2021), 1–36.

[124] Michael Terrell. 2024. New nuclear clean energy agreement with Kairos Power. https://blog.google/outreach-initiatives/sustainability/google-kairos-power-nuclear-energy-agreement/

[125] Neil C. Thompson, Kristjan Greenewald, Keeheon Lee, and Gabriel F. Manso. 2021. Deep Learning's Diminishing Returns: The Cost of Improvement is Becoming Unsustainable. *IEEE Spectrum* 58, 10 (2021), 50–55. https://doi.org/10.1109/MSPEC.2021.9563954

[126] JW Thomson. 1972. *Methods of traffic limitation in urban areas.* Technical Report. 106–112 pages.

[127] Sunita Tiwari and Pratibha Singh. 2011. Environmental impacts of e-commerce. In *International conference on environment Science and engineering*, Vol. 8. 202–207.

[128] Bill Tomlinson, Rebecca W Black, Donald J Patterson, and Andrew W Torrance. 2024. The carbon emissions of writing and illustrating are lower for AI than for humans. *Scientific Reports* 14, 1 (2024), 3732.

[129] Alexandra Tremayne-Pengelly. 2024. Amid the A.I. Boom, These States Have Become Data Center Hubs. https://observer.com/2024/12/ai-demand-where-data-centers-using-most-energy/

[130] UN General Assembly. 2015. Transforming our world : the 2030 Agenda for Sustainable Development. https://www.refworld.org/legal/resolution/unga/2015/en/111816

[131] Laszlo Varro and George Kamiya. 5. 5 ways Big Tech could have big impacts on clean energy transitions. (5).

[132] Ricardo Vinuesa, Hossein Azizpour, Iolanda Leite, Madeline Balaam, Virginia Dignum, Sami Domisch, Anna Felländer, Simone Daniela Langhans, Max Tegmark, and Francesco Fuso Nerini. 2020. The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature communications* 11, 1 (2020), 1–10.

[133] Petro Vorotyntsev, Yuri Gordienko, Oleg Alienin, Oleksandr Rokovyi, and Sergii Stirenko. 2021. Satellite image segmentation using deep learning for deforestation detection. In *2021 IEEE 3rd ukraine conference on electrical and computer engineering (UKRCON)*. IEEE, 226–231.

[134] Peng Wang, Ling-Yu Zhang, Asaf Tzachor, and Wei-Qiang Chen. 2024. E-waste challenges of generative artificial intelligence. *Nature Computational Science* (2024), 1–6.

[135] Martina Willenbacher, Torsten Hornauer, and Volker Wohlgemuth. 2021. Rebound effects in methods of artificial intelligence. In *Environmental Informatics*. Springer, 73–85.

[136] Eric Williams. 2011. Environmental effects of information and communications technologies. *nature* 479, 7373 (2011), 354–358.

[137] John D Wilson and Zach Zimmerman. 2023. The Era of Flat Power Demand is Over. *Grid Strategies* (2023).

[138] Richard T Woodward, David A Newburn, and Mariano Mezzatesta. 2016. Additionality and reverse crowding out for pollution offsets in water quality trading. *Ecological Economics* 128 (2016), 224–231.

[139] Xiao Xue, Shufang Wang, and Baoyun Lu. 2015. Computational experiment approach to controlled evolution of procurement pattern in cluster supply chain. *Sustainability* 7, 2 (2015), 1516–1541.

[140] Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Xiang Fu Matthew Horton, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbé, Shoko Ueda, Roberto Sordillo, Lixin Sun, Jake Smith, Bichlien Nguyen, Hannes Schulz, Sarah Lewis, Chin-Wei Huang, Ziheng Lu, Yichi Zhou, Han Yang, Hongxia Hao, Jielan Li, Chunlei Yang, Wenjie Li, Ryota Tomioka, and Tian Xie. 2025. A generative model for inorganic materials design. *Nature* (2025). https://doi.org/10.1038/s41586-025-08628-5

[141] Banghua Zhu, Ying Sheng, Lianmin Zheng, Clark Barrett, Michael Jordan, and Jiantao Jiao. 2024. Towards Optimal Caching and Model Selection for Large Model Inference. *Advances in Neural Information Processing Systems* 36 (2024).