

TLDR

This document is intended as a supplementary resource for The Great Simplification podcast episode #132 with Nate Hagens and Daniel Schmachtenberger on artificial intelligence (AI) and climate change. As discussed in the podcast, issues emerging from the nexus of AI and climate go beyond concerns of direct energy consumption to deeper and more systemic concerns, including the acceleration of extractive industries by AI applications, strong interconnections with the defense industry, the underlying transhumanist philosophy driving AI accelerationism, and a systems blindness to economy-wide Jevons paradox rebound effects.

- There is a clear growing trend in marketing rhetoric across sectors for AI to enable significant growth whilst also tackling climate and sustainability challenges
 - Coupled with the extraordinary exponential growth of the AI sector and **ubiquitous applications in areas such as the oil & gas industry and defense & intelligence sector**, there is a substantial risk that AI climate solutionist rhetoric ultimately serves as an unprecedented form of greenwashing for an ultimately catastrophic set of technologies.
 - Despite some claims that AI could revolutionize scientific development of panacea solutions like nuclear fusion, in reality, this is unlikely
- Advancements in various forms of artificial intelligence, including AI-enhanced planetary sensing capacities as well as generative AI pose distinct accelerating risks for the twin attractors of the metacrisis (catastrophe or dystopia)
 - **Increased dystopian potential:** The prevalence of digital technology and growing capabilities of AI-enabled Internet of Things (IoT), distributed sensor networks, ultra-high resolution satellites, and smart cities, whilst potentially useful for climate and biodiversity conservation, simultaneously increase the capacity for authoritarian regimes to utilize ubiquitous surveillance
 - **Pathways for catastrophic risks:** agential risks where a small group of people or even one person could use generative AI to assist planning large-scale terrorist actions, possibly at the level of mass genocide, omnicide and collapsing civilization. Similarly, runaway artificial superintelligence (ASI) with agentic capacity may do so by itself. With little effort, ChatGPT can be prompted to answer confidently about various ways an ASI could destroy humanity: [see our document here](#)
- AI safety efforts themselves tend to accelerate AI development, all of which is largely supported by a dangerous mix of philosophical views which fundamentally see little value in embodied human lives, natural ecosystems and the living biosphere
- **How much of a difference AI would make to climate is likely a lot less than proponents like to claim.** None of the solutions individually, nor taken as a whole, appear to convincingly “solve” climate change writ large. At the same time, various existing and developing machine learning applications are individually useful within a *narrow* scope, though many with limitations and additional risks

- **Two low hanging fruit applications (concrete and food waste) optimistically add up to ~4% of global GHG emissions.** There is insufficient information to ballpark a fraction for the rest - many of which depend on the success of other underlying tech anyway (renewables, EVs, fusion, precision ag, etc.)
- **Considering other side effects and externalities from AI/ML, proposed solutions likely outweigh climate benefits** - the material and energy demands of more semiconductors, cloud server capacity, IoT sensors, and robotics; the exacerbation of actively deleterious climate approaches like forest carbon offsets, direct air capture, boosting of fossil fuel greenwashing; Jevons rebound effects across nearly all ML-driven efficiency gains, plus the entire economic growth acceleration from the AI field as a whole if taken together almost certainly outweighs the few percent of global GHG emission reductions it could otherwise drive.
 - Globally, the digital tech sector accounts for 10% of global electricity demand. Current projections expect **data server energy demand to double over the next four years**, primarily due to AI data server electricity demand growth
 - We estimate **about 965 thousand metric tons of mining rock waste** was produced just for the GPUs used to train OpenAI's GPT-3
 - A rough estimate based on the [Garrett relation](#) between global cumulative economic production and energy consumption suggests that current projected economic growth from AI would require by 2030 an additional **739TWh per year, which is approximately the entire nuclear energy generation of the US**
- **Virtually all useful AI/ML applications for climate do not require more AI advancements.** Of the solutions [reviewed below](#), only *two* benefit from the current advancements in *generative AI* - material discovery for battery science and construction materials such as concrete.
 - All others genuine applications make use of existing, well established tech and algorithms (e.g. deep learning, convolutional neural networks, computer vision), which only require further work within their particular domain of application (data processing, training, architecture), and some benefits from related AI R&D efforts like "explainable AI".
 - The majority of current AI R&D efforts in large language models and generative image & video algorithms are effectively separate from these existing applications.

AI Industry Trends, Rhetoric and Risks

Climate Solutionist Rhetoric

Current and emerging AI-driven applications are increasingly seen as reasons to promulgate AI as a broad climate & sustainability solution. Terms such as “Green AI”, “Sustainable AI”, “Zero Carbon AI”, and “Net Zero AI” are often used in much of the marketing and news reporting of the current AI boom but are also [increasingly defining the mission for academic work and R&D for AI](#)¹.

The following are some characteristic examples of this memetic trend from notable platforms and organizations:

- [9 ways AI is helping tackle climate change](#) (World Economic Forum)
- [How To Fight Climate Change Using AI](#) (Forbes)
- [Can AI Help Solve the Climate Crisis?](#) (Google)
- [How AI could power the climate breakthrough the world needs](#) (CNN)
- [Explainer: How AI helps combat climate change](#) (UN)
- [How AI can enable a Sustainable Future](#) (PwC report)
- [Reduce Carbon and Costs with the power of AI](#) (BCG report)

In the extremes, particularly from “effective accelerationist - e/acc” groups which [increasingly define the cultural undercurrents of Silicon Valley](#), some view artificial (general) intelligence (and digital technology broadly) as necessary for solving climate change entirely. From these perspectives, the only real solutions to climate change are fundamentally technological (such as [direct air capture for carbon dioxide removal](#) and [solar geoengineering](#)), and these can only be accelerated in time via AI and/or AGI.

Whilst extreme, similar AI techno-solutionist sentiments are also defining influential mainstream positions. Take these direct quotes from influential AI and information technology think tanks:

*The digital transition is also considered to be a key tool to reduce energy consumption in many sectors (“IT for Green”), to such an extent that **it now hardly seems possible to address climate change without the large-scale incorporation of digital technologies.*** ([The Shift Project, 2019](#))²

¹ [Raman et al. \(2024\)](#) identified three thematic clusters in academic and R&D literature on AI:

(1) Responsible AI for Sustainable Development, focusing on integrating sustainability and ethics within AI technologies;

(2) Advancements in Green AI for Energy Optimization, centering on energy efficiency; and

(3) Big Data-Driven Computational Advances, emphasizing AI’s influence on socio-economic and environmental aspects

² Notably, although the Shift Project does ultimately land at quite optimistic messaging on digital technologies, they are also one of the few think tanks that consistently mentions the rebound effect (Jevon’s paradox) and lack of sufficient evidence for dematerialization of the economy.

...by 2030 digital technologies will deliver reductions in carbon emissions equivalent to nearly seven times the size of the growth in the total information and communications technology (ICT) sector emissions footprint over the same period³. ([Global e-Sustainability Initiative, 2019](#))

The above rhetoric is already translating to real observable market trends. [Fortune Business Insights](#) finds the “green technology and sustainability market” to be valued at USD 16.48 billion in 2023 and is projected to grow to USD 89.97 billion by 2032 - a compounded annual growth rate (CAGR) of 20.9%⁴ - largely presumed to be driven by generative artificial intelligence.

In the energy sector, the [IEA describes AI as an “increasingly critical resource”](#) that already serves more than 50 different uses in the energy system and is essential for enabling the renewable energy transition (see [Forecasting Supply and Demand](#) below).

Unlikely Breakthroughs

Various headlines over the last few years have indicated that the tech and AI industry is particularly hopeful for the ability of artificial intelligence to enable breakthroughs for nuclear fusion. Fusion, widely seen as the holy grail of clean energy and panacea for climate change, is [also being seen as an answer to the extraordinary energy requirements of AI](#).

Example articles:

- [How AI is Revolutionising Nuclear Fusion to Harness Limitless Clean Energy](#) (KnowHow, Apr 2024)
- [Scientists say they can use AI to solve a key problem in the quest for near-limitless clean energy](#) (CNN, February 2024)
- [Engineers use AI to wrangle fusion power for the grid](#) (Princeton University, Feb 2024)

Whilst it certainly is the case that artificial intelligence and machine learning have been successfully used in existing nuclear *fission* reactors (e.g. [reactor design](#), [safety and reliability of pressurized water reactors](#), [predicting extreme heat flux conditions](#)), and there have been genuinely useful advances which utilize machine learning algorithms for [plasma stability of](#)

³ Pitron (2023; p24-25) in *The Dark Cloud: How the Digital World is Costing the Earth* outlines how GeSI, one of the most influential organizations pushing the green ICT and AI narrative, lacks rigor and credibility in their reports which put forward extraordinary claims on the possible environmental benefits, emissions and energy savings from digital technologies. The overwhelming majority of “expert” advice for their work is from private sector interests with deliberate exclusion of NGOs and counterbalancing academic views. Much of their assertions are unverifiable, and one dissident even went so far as to say “It would be very dangerous to use GeSI reports for public policy purposes”. Unfortunately, this is exactly what has happened.

⁴ To put this in perspective, [the historical CAGR](#) of the entire United States from 2013 - 2023 is about 2.3%, and for China about 8.6%.

[nuclear fusion reactors](#), these are unlikely to be total “game-changers”, unlike current techno-optimist rhetoric makes it out to be.

See [our more extensive notes](#) below on why this is unlikely to be the case. In the most general sense, however, whilst machine learning, various types of neural networks, and (in some limited cases) generative AI can be useful in very specific use-cases for environmental/climate-related and adjacent research (see, e.g. [accelerating materials science](#)), for the most part, AI is fundamentally, inherently limited in how much it can accelerate scientific and technological progress. There are several arguments for this:

- Scientific and technological progress requires real-world experimentation. Though various forms of AI can speed up certain methods and processes used for experimentation (e.g. handling large data sets, [speeding up experimental design](#), etc.), simulations, inductive and inferential machine reasoning [simply cannot replace physical experiments](#)
- Though proponents of ML-enhanced scientific research and discovery will often support their claims by suggesting AI models are “theory-free” or “[theory-agnostic](#)” and can thus merely rely on data (i.e. allowing the data to “speak for itself”), [this is in actuality deeply misguided](#). It is widely accepted in philosophy of science that data is always theory-mediated, and hence cannot be taken to be *merely* empirical, objective inputs.
 - Theory-ladenness is evident in contemporary AI methods for scientific research and R&D, such as AlphaFold, which unavoidably does integrate scientific theory at virtually all levels: data, model structure, hyperparameters, training, evaluation and interpretation of results⁵.
 - As a consequence of the above, any true “breakthroughs” in science that require fundamental advances in theory (e.g. theoretical physics) cannot be done via AI. **In short, science itself cannot be automated.**

Recent AI Industry Growth

Artificial Intelligence (AI) as a sector has exploded over the last two decades, particularly since the rise of “deep learning” architectures, which are now ubiquitous across text, image, and speech recognition algorithms. Various metrics of AI development, including research publications, industry affiliations, patents, investments, and algorithmic performance, have displayed remarkable (largely exponential) growth curves.

⁵ See [Andrews \(2024\)](#) for a more extensive argument and description as to why AI models cannot merely rely purely on objective data both generally and in the specific case of AlphaFold2.

Notably, even in the case of AlphaFold, claims that it “revolutionizes” drug discovery isn’t quite accurate, as [knowledge of protein structure is rarely a rate-limiting step in such projects](#).

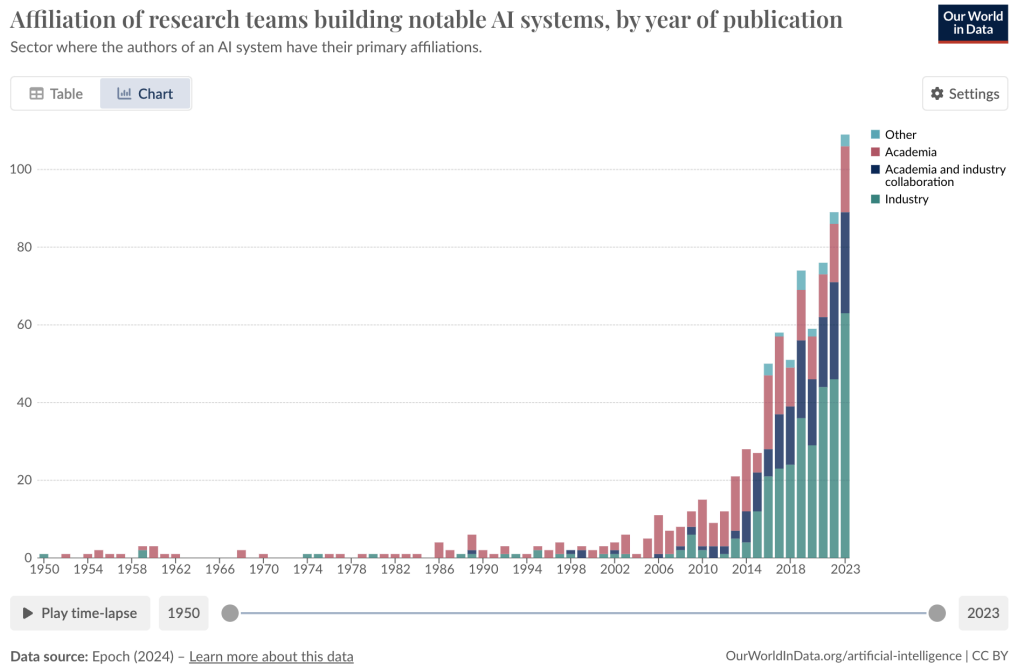


Fig 1: Real-world applications of deep learning have sharply grown since the early 2000s, with AI research from universities increasingly being directly affiliated or in collaboration with industry. (Source: [Our World In Data](#))

Number of AI patents granted, 2010–22

Source: Center for Security and Emerging Technology, 2023 | Chart: 2024 AI Index report





Figure 1.2.1

AI patents granted since 2010 have followed an exponential growth curve. (Source: Center for Security and Emerging Technology, 2023 - [Artificial Intelligence Index Report 2024](#))

Number of newly funded AI companies in the world, 2013–23


Source: Quid, 2023 | Chart: 2024 AI Index report



The AI investment boom is being driven by dramatic increases in capability across a wide set of domains, from image recognition to language processing to data processing and automation.

Many of these applications, such as processing “big data” from distributed sensor technologies, are already widely used for ostensible climate and environment applications. A major portion of claims that AI can help solve climate change and/or the biodiversity crisis leverage these in particular - see, e.g. our notes below on AI use in [smart cities](#), [remote sensing emissions](#), and [ecosystem and biodiversity monitoring](#).

Notably, the growth in NVIDIA GPUs has been documented to now completely outpace expectations of Moore’s Law.



Advancement in GPU processing power compared to projections according to Moore’s Law (Source: [Netrouting](#))

In terms of economic growth, according to data from [Jon Peddie Research](#) “*the compound annual growth rate (CAGR) of GPU shipments between 2015 and 2020 was 9.8%, outpacing the 7.4% CAGR of Moore’s Law during the same period.*”

The growth, current and future potential of AI for climate and environmental challenges need to be situated in the context of other major use areas, with civilization-scale implications, including especially accelerating the oil and gas industry and bolstering the defense industry.

AI - Oil & Gas Industry

As the AI sector has grown, major tech companies have been providing a variety of services, as well as direct R&D collaborations with the oil & gas industry. The following are some highlights, however [we discuss this in more detail including specific applications below](#).

- [50% of oil and gas executives have already begun to use AI in their organizations](#)
 - [92% of oil and gas companies](#) are either currently investing in AI or plan to in the next two years.
 - AI in the oil and gas market was valued at US\$2 billion in 2019 and [expected to reach \\$4.21 billion by 2028](#) and [\\$13 billion by 2034](#)
 - As noted in the podcast, automated robotic servicing of gas pipelines is currently on the horizon. An AI and augmented reality-enabled crawling modular robot [developed by Carnegie Mellon Robotics Institute](#) has demonstrated the ability to map, identify and repair leaks within gas pipelines
 - [Saudi Aramco has been investing the most out of any oil and gas company](#) - around \$3.5 billion in AI R&D, particularly for oil exploration and underwater operations
- Tech corporations (including Accenture, Amazon, Google, IBM, Microsoft and Oracle) are competing to service rising demand for AI technologies from the fossil fuel industry; [what has been described as ‘automating the climate crisis’](#).
 - In 2020, [Greenpeace published a report](#) documenting how major oil corporations (Shell, BP, Chevron, ExxonMobil) had been contracting cloud computation services from tech companies
 - Though Google [announced previously](#) that they will “not build custom AI/ML algorithms to facilitate upstream extraction in the oil and gas industry”, they have continued their partnership with Saudi Aramco since at least 2023
 - [Fortune Magazine reported in 2024](#) that Amazon, Microsoft and Google are currently opening headquarter offices in Saudi Arabia
- The GPU manufacturer NVIDIA has multiple direct partnerships for high performance computing (HPC) AI data centers, specialized hardware and software [specifically for oil and gas](#)
 - A recent collaboration [between NVIDIA and Shearwater GeoServices](#) in 2023 [achieved 10x faster performance](#) of seismic processing software key to oil and

- gas exploration workflows, utilizing the [GH200 Grace Hopper Superchip](#), specifically designed for AI HPC data centers
- In 2018 NVIDIA announced a [collaboration with Baker Hughes](#) (a subsidiary of General Electric) to develop machine learning to “dramatically reduce the cost, finding, extracting, processing and delivering oil”
 - Orbital Sidekick, founded as part of NVIDIA’s Inception startup program, was the first to use hyperspectral imaging to detect gas pipeline leaks from space, utilizing [NVIDIA’s Jetson AGX module as an AI engine](#) which is designed for autonomous machines

AI - Military defense & Intelligence

A major source of investment for AI development is currently from the defense and Intelligence sector. Though on the surface it may seem that AI applications for climate and the environment are distinctly separate to the military and intelligence sector, ultimately the current risk is that tech companies pushing for “green AI” are (knowingly or inadvertently) simultaneously pushing forward the same technological capacity for uses in war and surveillance. Notably, some AI accelerationists frame their arguments [explicitly in terms of a geopolitical AGI arms race](#) between China and the US for military dominance.

Moreover, some of the technical hardware capacities useful in climate and environmental applications (e.g. whole Earth surface satellite monitoring) have clear dual use cases for military and intelligence surveillance purposes.

Notable recent funding statistics:

- The U.S. Defense Advanced Research Projects Agency (DARPA) recently funded a [\\$78 million program for new types of AI chips in battlefield conditions](#), with the dual goal of optimizing on energy and size.
 - Recent budget requests from DARPA have also indicated massive increases in funding for AI projects including autonomous systems and human machine teams
- The [Brookings Institution recently found that AI-related federal U.S. contracts have increased 1,200%](#) from \$355 million in 2022, to \$4.6 billion in 2023.

Recent examples of AI use in military and battlefield applications:

- **Precision targeting:** AI is currently in use by the Israel defense Force (IDF) for a [target creation platform known as Habsora \(the Gospel\)](#) to rapidly generate recommendations for bombing targets based on suspected militants. [Investigative reporting by Israeli magazines have also described an AI program \(“Lavender”\)](#) which examines data including social network connections and family relationships to identify possible Hamas fighters.
- **Drones:** Substantial developments have been made in autonomous drones utilizing various AI-enabled systems (computer vision, autonomous decision making, swarm

algorithms) over the past few years. These include [mass produced, autonomous drones in the Ukraine conflict](#), [“slaughterbots” which could be used by Cartels](#), and so-called [drone “killwebs”](#)

- **Comprehensive battle planning:** Artificial intelligence systems can now [convincingly simulate human behavior in military simulations](#), U.S. DARPA has partnered with Northrop Grumman to apply [AI in novel strategic game scenarios](#), and [Palantir AIP has recently unveiled an LLM platform](#) for defense and military planning as an AI decision support system.
- **Intelligence planning:** AI algorithms are uniquely well equipped to make use of extremely large data sets, leading to a proliferation in use cases across the intelligence sector. [Currently, about 56 out of 176 countries use AI for surveillance](#). Emerging applications from [ongoing research include “negative mental state monitoring” and containing the spread of rumors in social networks](#).

Ubiquitous Surveillance

Ubiquitous technological surveillance (UTS) is radically increasing, with extremely strong interlinkages with current and proposed uses of AI for climate and environmental problems, and risks for the [rise of “eco-surveillance” states](#)

- [Internet of Things \(IoT\), smart homes and smart cities](#) are enabling both automated surveillance systems to track citizens through cameras, microphones, traffic and public transport monitoring, consumer purchases (etc.), as well as ostensibly reducing environmental footprints
 - Even existing universal tech such as [WiFi signals can be utilized via AI to identify people in a building](#), similar to radar
 - Smart cities which aim to utilize IoT and “smart” sensor tech to optimize and solve various urban problems from traffic and waste management to energy efficiency of lighting and heating systems [have been argued to pave the way to enabling the rise of “digital authoritarianism”](#), as systems capable of monitoring, controlling and autonomizing urban services can be used for both environmental ends as well as mass surveillance, control and suppression.
- Satellite networks which are now close to the level of real-time monitoring of every square meter⁶ of the Earth’s surface enable both unprecedented capacity for planetary-scale [computation](#) with positive environmental and climate use cases (e.g. mapping [illegal deforestation](#), [tropical forest health](#), [early warning of ecological tipping points](#), [identifying dirty polluters](#)), as well as uses for the military and defense sector (e.g.
 - To take one example, Planet Labs which markets itself primarily as satellite imaging for environmental, agricultural and urban Earth insights (with their latest

⁶ Very high resolution satellite imagery is now reaching ever finer scales, down to [30-60cm ground resolution](#)

- platform utilizing various AI-based models), also has [direct partnerships with the U.S. Department of Defense and Intelligence Community](#).
- Satellite remote sensing capabilities are continuously improving. The company [ICEYE has now developed synthetic-aperture radar \(SAR\) remote sensing](#), with effectively real-time earth observation capabilities for any location on the surface of the planet down to 50 cm resolution.

AI Safety to AI Accelerationism

AI Safety Doom Loop

One of the ironic, though deep-seated drivers of the AI risk nexus is the tendency for any AI safety effort to ultimately drive further AI development and dominance. This is evidenced primarily by the inception of OpenAI, Anthropic, and DeepMind, which were originally AI safety efforts and ultimately became key players in the AI arms race.

- AI safety efforts tend to become AI dominance efforts. Companies like OpenAI and Anthropic demonstrate this trend. The extensive computing resources needed for AI experimentation demand significant capital⁷. This funding typically comes from investors who expect a return on their investment. Consequently, the research gets steered towards [commercially viable directions rather than focusing solely on safety concerns](#).
- Over time, AI safety researchers become AI accelerationists. This shift in orientation occurs when researchers realize that the race to develop advanced AI is unstoppable due to the massive expenditure of capital and talent and the incredibly rapid pace of progress in AI capabilities. Winning this arms race will grant such immense power that only the winners will have any meaningful influence on the future. Consequently, these researchers conclude that accelerating AI development, rather than attempting to slow it down, offers the best chance to have a meaningful impact on how AI will shape humanity's future.
- Increasingly, the fear of being put out of work soon by AI is increasing, as is the money being invested into the space currently - so people with the skills to do safety work, who could also do development work, see making a startup to do some basic thing that will get acquired by one of the big players for a big exit, as possibly their last chance for economic security, which is increasingly desired in the presence of such radical and unpredictable societal transformation.

⁷ See for example the [OpenAI-Microsoft planned "Stargate" data center project](#), estimated to cost US\$100 billion by 2028. From [April to June of 2024, AI start-ups made up \\$27.1 billion of U.S. investor financing](#).

- The fear of AI replacing jobs is increasing. This creates a dilemma for those skilled in AI safety and development. Many see founding an AI startup, with the goal of being acquired by a major company, as possibly their last chance for financial security. This view is fueled by the radical and unpredictable societal changes AI will bring. Nick Bostrom's recent work illustrates this shift in perspective. Known for raising concerns about AGI risks in *Superintelligence* (2014) Bostrom's recent book *Deep Utopia* (2024) claims that a benevolent AI ruler may be the only solution to human coordination problems, reflecting a marked change in outlook on AI's role in society.

Notably, even current LLMs such as GPT-4, Claude, BERT, LaMDA and LLaMA 2 pose non-negligible [agential risks](#) where a small group of people or even one person could use such a model to plan large-scale terrorist actions, possibly at the level of mass genocide, omnicide and collapsing civilization⁸. Alternatively, runaway artificial superintelligence (ASI) with agentic capacity may do so by itself. Even without much effort⁹, [ChatGPT will answer confidently about various ways an ASI could destroy humanity](#).

Transhumanism and TESCREAL Ideology

Many of the technologists, entrepreneurs, and investors in the space hold ["TESCREAL" views](#)¹⁰ and will even admit that they believe the replacement of biological life with digital life is both inevitable and desirable. As such, care for the safety of life as we know it is seen as the ignorant and regressive clinging to the inelegant, limited, and death-bound "wet-ware" biohardware we happen to be currently running on.

- A particularly striking example of this is the publicly stated ultimate ambition of Neuralink, Elon Musk's brain chip company which eventually aims to use these implants to ["merge humans with AI"](#)
- Transhumanism and TESCREAL type views, common in the AI space, are [fundamentally undergirded](#) by [computationalist](#) and [simulationism](#) types of metaphysical beliefs, which view the whole world and reality itself as fundamentally computational
 - The [World Transhumanist Association](#) (H+) and the [transhumanist philosophical project](#) was founded by David Pearce and Nick Bostrom in the late 90s, advocating for technological innovation to strive towards sci-fi type

⁸ See also

⁹ I.e. without precise prompt engineering to "jailbreak" safety programming

¹⁰ Transhumanism, Extropianism, Singularitarianism, Cosmism, Rationalism, Effective Altruism and Longtermism - a mutually reinforcing set of philosophies commonly bundled together in Silicon Valley tech spaces, as described by Timnit Gebru and Emile Torres. For more on the views of TESCREAL ideologies see:

Gebru, T. and Torres, É.P. (2024) 'The TESCREAL bundle: Eugenics and the promise of utopia through artificial general intelligence', First Monday [Preprint]. Available at: <https://doi.org/10.5210/fm.v29i4.13636>.

<https://www.truthdig.com/articles/the-acronym-behind-our-wildest-ai-dreams-and-nightmares/>

advancements for the human condition including cryonics, life extension, space colonization, personality pills, and superintelligent artificial intelligence.

- This has led to the commitment to a “[hedonistic imperative](#)” as part of the [Abolitionist Project](#), which seeks to “phasing out the biology of suffering” for all sentient life via consciousness upload and/or genetically engineering pain receptors
- These kinds of beliefs, combined with techno-optimist visions of human progress also drive a prevalence of disturbingly outlandish views held by many industry leaders and prominent influential thinkers in the AI and tech industry
 - [Compilation of views AI experts have regarding AGI and AI alignment](#)
 - [Compilation of more outlandish AI-related views across the industry](#)

Direct Energy and Material Uses

Within the concentric circles of how AI can and may potentially impact climate and the environment, at the center are direct energy and material demands from hardware itself, particularly server farms.

Key Highlights

- Globally, the digital tech sector consumes about 10% of the world’s electricity consumption, has a carbon footprint about 2x-3x that of the UK or France, accounts for 4% of global GHG emissions and consumes about 5-7% more electricity every year¹¹.
 - [One estimate suggests](#) that global electricity demand from digital technology could grow to 20% of global demand by 2025. The IEA projects that [electricity consumption from AI, data centers and cryptocurrency could double by 2026](#)

¹¹ Stats from Pitron, G. (2023) *The dark cloud: how the digital world is costing the Earth*. Brunswick, Victoria: Scribe Publications.

- Growth of the digital tech sector and particularly AI is currently largely limited by energy demand
- Data center energy demand in 2022 was 17GW from 2,700 data centers, and is projected to reach 32GW by 2029.
 - IEA projections suggest global consumption of electricity for data centers may grow from 460 TWh in 2022 to more than 1000 TWh in 2026, though this is quite possibly an underestimate¹²
 - Data centers in particular account for 2% of global electricity demand.
 - [Google's data centers used 20 percent more water in 2022 than in 2021](#), while Microsoft's water use rose by 34 percent.
 - Power demand projections for power grids [in the US \(second source\)](#) and [internationally](#) have doubled their annual expected growth rate over the five to ten year timeframe as of 2023
- Energy use of ChatGPT
 - "It's been estimated that ChatGPT is responding to something like two hundred million requests per day, and, in so doing, is consuming more than half a million kilowatt-hours of electricity." (For comparison's sake, the average U.S. household consumes twenty-nine kilowatt-hours a day.) (Elizabeth Kolbert, [The Obscene Energy Demands of AI](#))
- Existing machine learning applications have already seen demonstrable Jevons-like dynamics play out.
 - [The use of AI and machine learning assisted translation services](#) did reduce costs in translation services by 40%, but the Bureau of Labor Statistics found an uptick of 19% in translation and interpreter services in 2021.
 - Organizational use of AI for [automating quantified risk policy workflow](#) led to significant increases in revenue per worker (up to 38%), leading the business to hiring more workers

Direct Energy and Material Demands

Booming demand for AI computational capacity has led to a recent surge in demand for necessary infrastructure in the form of data centers, which necessarily have intrinsically high electricity, water and material requirements. These issues haven't escaped the attention of news outlets and analysts, the following are a few select examples:

¹² See in particular the Evaluations of [energy] consumption of data centres in 2018 from the Shift Project [here](#). IEA previous 2018 estimate of ~200TWh global energy consumption from data centres is notably lower than several other independent groups (Shift Project, Borderstep, GreenIT) which estimated a range of 300 - 500TWh.

If this discrepancy holds for the 2022 evaluation, and we take the average of estimates previously as the true value (~380TWh in 2018) then we could expect current data center energy consumption to be closer to ~670TWh for 2022, and the projection for 2026 ~1400TWh.

No consensus exists on either usage figure yet: “Right now, it’s not possible to tell how your AI request for homework help or a picture of an astronaut riding a horse will affect carbon emissions or freshwater stocks [...] the same is true of carbon.”

[As Use of A.I. Soars, So Does the Energy and Water It Requires](#) (E360 Yale, 2024)

Using the training of BERT as a reference experiment, “There is large variation between the least carbon-intensive regions [...] choosing the region in which experiments run can be very impactful (7k grams vs. 26k grams [of CO₂], for the most efficient vs. least efficient regions).”

[Measuring the Carbon Intensity of AI in Cloud Instances](#) (Dodge et al. 2022)

Compute applied to AI has a doubling time of 3.4 months.

[AI and Compute](#) (OpenAI)

“...at present, the emissions from LLMs are relatively insignificant compared to both their popularity and to other everyday activities.”

[An analysis of both hardware production and energy usage](#) (Smith et al. 2023)

Though energy, water and material consumption concerns apply broadly to digital technologies and cloud computing broadly, artificial intelligence in particular poses a distinct concern due to the scale and intensity of computation involved, as well as more particular hardware requirements - notably graphics processing units (GPUs).

Though the majority of attention has thus far been on direct electrical energy consumption and CO₂ emissions, far less has been focused on material footprints of AI data centers. Material demands includes concrete, steel, glass and aluminum for basic infrastructure of data center structures themselves, as well as various metals used in semiconductors and chip manufacturing, notably precious metals including gold, silver and palladium; basic metals such as copper and iron; rare earths such as dysprosium, neodymium and yttrium; and critical raw metals such as indium, titanium and gallium¹³.

From our rough estimate¹⁴, **about 965 thousand metric tons of mining rock waste** was produced **just for the GPUs** (about 10,000 NVIDIA A100s) **used in OpenAI’s custom-built supercomputer to train GPT-3**. This is mostly due to just three elements—indium, gold, dysprosium, and copper—and would plausibly be another magnitude greater if all other components (motherboards, CPUs, network cables, etc.) were included.

¹³ A comprehensive material analysis of chipboards used for servers in data centers was done by the [TEMPRO project](#) from the University of Oldenburg, Germany. See a translated copy of their [2020 report here](#), and [original \(in German\) here](#).

¹⁴ See Appendix entry [Mineral footprint of AI and ML data servers](#) for our calculations

Broad Computation Environmental Footprint

The following are some general notes on the environmental footprint of data centers and large-scale computation. Broadly, the environmental footprint of computation has two sources, operational energy consumption and hardware manufacturing and infrastructure.

- Whilst energy use and carbon emissions from operational demands are decreasing due to algorithmic, software and hardware innovations which boost efficiency, overall carbon footprint of computational systems continues to grow
 - Ultimately, most emissions related to data center equipment comes from hardware manufacturing and infrastructure ([Gupta et al. 2021](#))

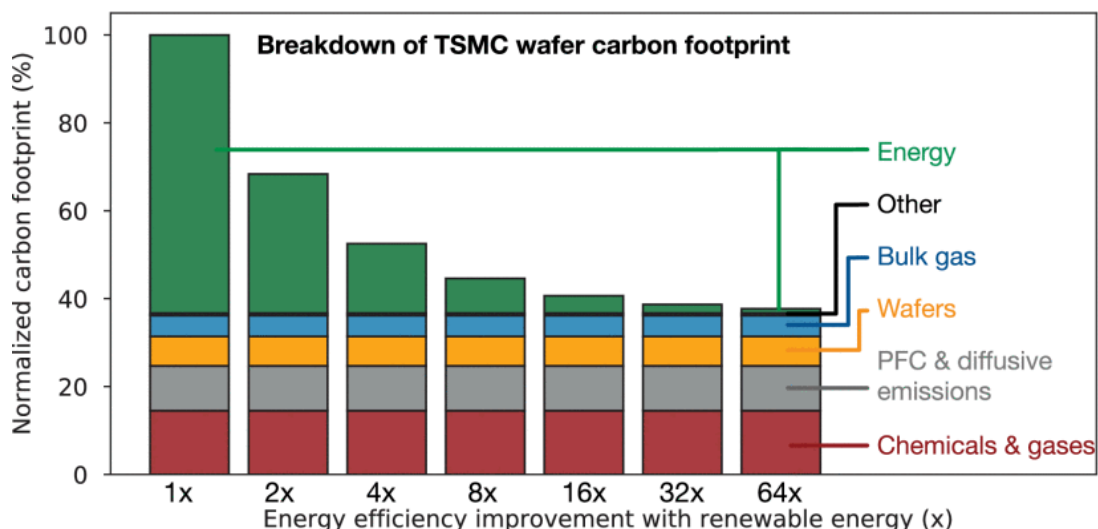


Figure 14 from Gupta et al. shows that a 64× boost in renewable energy reduces the overall carbon output by roughly 2.7× - an ambitious goal. By 2025, TSMC estimates renewable energy will produce 20% of the electricity that drives forthcoming 3nm fabs. Intel already uses renewable energy to meet much of its demand; only 9.7% of the energy consumed by Intel fabs comes from nonrenewable sources. **However note, almost half of wafer production's carbon footprint is insensitive to radical increases in renewable energy, due to the hardware manufacturing process.**

- Data centers and transmission networks account for between [1% - 1.3% of global electricity use](#) (excluding crypto mining), and 0.6% of global carbon emissions.
 - Combined electricity use by Amazon, Microsoft, Google and Meta more than doubled between 2017 and 2021 to about 72 TWh in 2021
 - Amazon, Apple, Google, Meta and Microsoft [together account for over 45GW](#) of corporate renewable purchases globally - over half the global corporate renewables market
- A single average data center consumes the equivalent of heating 50,000 homes yearly.
- Data centers [require large amounts of raw materials](#) including copper, silicon and lithium for server racks, and aluminum, steel and concrete for buildings

- In a [Turner and Townsend Data Centre Cost Index \(2022\) report](#), 95% of respondents noted that material shortages caused delays to data center construction over the last 12 months. The same report also noted a shortage of experienced data center construction teams.
- Data centers also require large amounts of critical and precious materials. RAM, CPUs, drive boards and motherboards require up to 855 g/ton of board for gold, 1802 g/ton silver, 36 g/ton of palladium. Using a rock-to-metal ratio¹⁵ rough calculation, this equates to **about 2.6 tons of mining rock waste per kg of electronic board** used in data servers.
- Major tech companies (at the forefront of AI applications) have previously claimed to produce carbon neutral products (e.g. Apple and Google), due to purchasing of emission credits. Microsoft and Apple have announced they will be carbon neutral by 2030.
 - [Apple has come under fire](#) for not being transparent about supply chain emissions, and no longer requiring suppliers to report GHG inventories.
- Large operators in data centers have regularly claimed power usage effectiveness (PUE) as having fallen, with PUE figures between 1.1 - 1.4 (lower is better), implying a marked improvement in power efficiency over time. An [Uptime Institute analysis](#) tracked PUE stats over 12 years, finding instead that efficiency improvements have stalled since 2013, with a recent deterioration in 2018 to 2019.
 - A number of (speculative) explanations have been raised. E.g. higher extreme temperatures in parts of the world with data centers could account for increased use of cooling and higher PUE. Utilization may have fallen as workloads are moved to public cloud services, so data centers are operating below optimal design efficiency
 - [Even more recent figures from the same agency](#) have shown PUE efficiency figures remained flat, and may have even continued to deteriorate
- At the same time, [there does appear to be more gains possible for energy efficiency of computation](#). Based on peak-output efficiency metric, measured as the efficiency of a processor at its fastest speed by the number of computations performed per kilowatt-hour of electricity, **efficiency has doubled every year and a half for over five decades**.
 - Now, **doubling time for peak-output efficiency is about 2.7 years**, a marked slowdown from historical improvements
 - Other metrics for efficiency indicate further areas for improvement, such as power management when devices are idle or off.
 - The “typical-use efficiency” metric (also in computations per KWh) is calculated over the course of a year by total electricity consumed using a weighted sum of energy use during different modes of processor intensity. By this metric, gains have been steadily improving since 2008, and were predicted to reach a doubling time of 1.5 years by 2020.
- **Electronic waste is the fastest-growing waste stream in the world**, with 57 million tons generated every year ([about the same weight as the Great Wall of China](#))

¹⁵ See [Appendix](#)

- [According to the UN](#), in 2019, the world generated 53.6 metric tons (Mt) of e-waste, only 17% of which is properly collected. The global generation of e-waste grew by 9.2 Mt since 2014 and is projected to grow to 74.7 Mt by 2030, equivalent to doubling in 16 years.
- [A typical North American produces about 20kg of e-waste a year](#), with one-third ending up in developing countries where children, forced labor and destitute “waste pickers” dismantle electronics for gold, silver, copper, and cobalt exposing them to carcinogens and toxins.
- GPU and electronics manufacturing [require large amounts of high quality silica sand for wafer fabrication](#).
 - Sand and gravel are the [second most used natural resources after water](#). Sand is the most used solid material, with about [50 billion tons used per year](#) (largely for infrastructure development).
 - Despite this, there is a large global imbalance as to where high grade metallurgical silica for use in electronics and solar panel manufacturing can be found. High quality deposits are relatively scarce. About 380 Mt of global silica sand was produced in 2022, worth about US\$15 billion ([USGS](#)). China accounts for about 70% of global silicon metal manufacturing.
 - Wafer fabrication involves the refining of silica sand into polycrystalline silicon, made in crucibles which are themselves made from ultra-high purity quartz. The estimated amount of high-purity quartz produced globally is [about 30,000 tons](#).
- After thirty years of a generally falling rate of change in energy demand, power demand projections for power grids [in the US \(second source\)](#) and [internationally](#) have doubled their annual expected growth rate over the five to ten year timeframe as of 2023. These new projections are being driven by electrification, planned data center growth (“being supercharged by the rise of artificial intelligence”) and tech manufacturing, and forecast extreme weather events.
 - “New generative artificial intelligence (GenAI) is a significant driver of BCG’s estimate, with 2 GW of GenAI-related load in the base case and possibly an additional 7 GW of GenAI load online by 2030. At this higher end, BCG estimates that data centers could consume 7.5% of all electricity in the U.S.”
 - In the US, this growth isn’t projected evenly across regions, with highest annual growth in Arizona/Nevada, Michigan/Indiana, and Georgia/the Carolinas showing between 5-10% annual demand growth between 2022-2023 expected to continue projecting forward to at least the end of the decade. A large part of this growth is from data centers and manufacturing in these regions.
 - Emerging markets are consistently showing outsized growth in energy demands: “About 85% of the additional electricity through 2026 is set to come from outside advanced economies, mostly in People’s Republic of China (hereafter, “China”), India and Southeast Asia [...] Particularly in advanced economies and China, electricity demand will be supported by the ongoing electrification of the residential and transport sectors, as well as a notable expansion of the data center sector.” ([source](#))

- Large areas of the U.S. are [at risk of running short of power due to surging energy demand from data centers](#).
 - In 2023 projected energy demand approximately doubled, from 221.5GWh to 563GWh. **Data center energy demand in 2022 was 17GW from 2,700 data centers, and is projected to reach 32GW by 2029.**
 - Nb. Clearly not all data center use is AI. For example, machine learning *training* only [represents 15% of Google's total energy use](#).
 - 2,700 data centers consumed more than 4% of the country's total electricity in 2022 - this will grow to 6% by 2026. Some regions may miss out on economic development opportunities because the grid cannot keep up.
 - The proliferation of data centers is putting more pressure on states to approve new transmission lines, however the amount of transmission installed in the U.S. has dropped sharply since 2013, from 4000 miles per year to now barely 1000 miles per year.

Current AI Footprint

There is currently insufficient information to provide concrete statistics for the environmental footprint of artificial intelligence at present, particularly given that existing data centers may already be used for both AI and non-AI uses (e.g. data storage, general computation, etc.). Nevertheless, some particular points are known:

- The AI industry has [the fastest growing carbon footprint](#).
 - Since 2012, the [most extensive AI training runs](#) have been using exponentially more computing power, doubling every 3.4 months on average.
- Energy use of ChatGPT: "It's been estimated that ChatGPT is responding to something like **two hundred million requests per day**, and, in so doing, is **consuming more than half a million kilowatt-hours of electricity**. (For comparison's sake, the average U.S. household consumes twenty-nine kilowatt-hours a day.) (Elizabeth Kolbert, [The Obscene Energy Demands of AI](#))"
 - Based on Kolbert's numbers above (appear to be from de Vries' quotation of SemiAnalysis stats, see below), this implies **ChatGPT's energy use is ~17,000x the average US household every day**, based just on usage requests.
- AI Data centers also consume extremely large amounts of water for cooling systems.
 - One estimate by Shaolei Ren at UC Riverside suggests that a person engaging in a session with GPT-3 of 10 - 50 responses drives the consumption of half a liter of fresh water.
 - Indications that water use in data centers is accelerating: "Google's data centers used 20 percent more water in 2022 than in 2021, while Microsoft's water use rose by 34 percent." ([As Use of A.I. Soars, So Does the Energy and Water It Requires](#))
- Alex de Vries (2023) published an extensive commentary on [The growing energy footprint of artificial intelligence](#)

- “Research firm [SemiAnalysis](#) suggested that OpenAI required 3,617 of NVIDIA’s HGX A100 servers, with a total of 28,936 graphics processing units (GPUs), to support ChatGPT, **implying an energy demand of 564 MWh per day**. Compared to the **estimated 1,287 MWh used in GPT-3’s training phase**, the inference phase’s energy demand appears considerably higher. Furthermore, Google reported that 60% of AI-related energy consumption from 2019 to 2021 stemmed from inference.”
- “SemiAnalysis estimated that **implementing AI similar to ChatGPT in each Google search would require 512,821 of NVIDIA’s A100 HGX servers**, totaling 4,102,568 GPUs. At a power demand of 6.5 kW per server, this would translate into a **daily electricity consumption of 80 GWh and an annual consumption of 29.2 TWh**”
 - If Google search consumed an annual 29.2 Twh per year due to integrating AI, this would be an equivalent energy consumption as the entire country of Ireland.
- “New Street Research independently arrived at similar estimates, suggesting that Google would need approximately 400,000 servers, which would lead to a daily consumption of 62.4 GWh and an annual consumption of 22.8 TWh.”
- Improving efficiencies in hardware and software are unlikely to offset future energy demands, as this would almost certainly result in increased AI demand triggering a rebound effect in increasing overall resource use due to Jevons Paradox.
- By contrast to de Vries’ analysis, some researchers hold that AI and machine learning’s energy consumption will soon plateau, and then shrink. [Patterson et al. \(2022\)](#) argue that relatively simple “best practices” can reduce energy use by 100x, and CO₂ emissions by 1000x.
 - Patterson et al. also argue that previous published estimates of carbon emissions from training ML models have been off by a factor of 1000x - 100,000x, due to lack of access of internal information

Counterarguments and Cautions

- [Azeem Azhar](#) has argued that we ought to reframe stats such as those noted above on energy and consumption for tech companies like Google in broader contexts, such as economic output efficiency and total global emissions
 - For example, Google is reportedly 8x more carbon-efficient than the global average, and only accounts for 0.025% of global emissions
 - This type of argument runs into several issues, however:
 - Besides the usual kinds of criticisms well-known around [why GDP is a poor metric for human well-being and economic performance](#), these kinds of numbers arise largely as a result of how companies calculate carbon

emissions, and particularly the exclusion of Scope 3 emissions¹⁶ for most tech companies.

- Tech companies generally have [poor transparency across their supply chains](#) and thus struggle to quantify Scope 3 emissions, particularly factors such as [embodied CO₂](#) in materials used (e.g. how much carbon emissions are produced to provide the copper wiring in a server farm). On average, [Scope 3 emissions are 4x higher than from direct operations](#).
 - The comparison between digital tech companies' carbon intensity and that of national economies is itself flawed. Company revenue is not directly comparable with GDP, as the latter accounts for all types of goods and services in a country. National economies necessarily encompass multiple sectors which themselves have various emission intensities, whilst digital tech companies will, by how emissions accounting is done, appear to have less emissions intensity than e.g. agriculture or mining
 - National economies (necessarily) provide far more actual necessary services (food, water, healthcare, education, etc) to citizens than customers of tech companies. This is not captured in the dollar metric used in either GDP or revenue reporting.
- Projections of streaming video carbon impacts as a cautionary example for projections of emission estimates
 - Current attention on emissions intensity of server farms and AI narratively mirrors previous reporting of the carbon impacts from video streaming. In that case, widely reported carbon estimates for watching 30 minutes of streaming video were off [by a factor of 90](#).
 - This core weakness in their argument damaged the ability of organizations like [The Shift Project](#) to impact other effects of streaming video.
 - Note, the [Shift Project published a response](#) to the above linked fact check article from Carbon Brief. They elaborate that though they rightly pointed out an error, it was for a figure that was not from their reports.
 - Note however, AI is not like streaming video in several aspects:
 - Streaming video caps out its value to users at certain resolutions; LLM computations return increasingly useful results at higher environmental costs
 - Potential streaming time is practically limited; value of calls to AI as a service can far exceed a single-threaded human chat use case

¹⁶ GHG emissions accounting standards divide emissions by Scope 1, 2 and 3 emissions.

Scope 1 emissions are from operations directly under operational control such as fuel burned by company vehicles. Scope 2 covers indirect GHG emissions from purchased or acquired electricity. Scope 3 are the rest of the indirect emissions beyond direct operations that are within value/supply chains (e.g. component transportation, distribution, purchased goods and services, etc.)

Acceleration of Extractive Industries

The second concentric circle of environmental and climate impacts after direct energy and material demands is the use of AI to accelerate extractive industries which themselves are directly responsible for significant amounts of environmental damage.

Here, we focus on just the examples of the oil and gas industry along with industrial agriculture, however this also applies to other sectors such as AI in [mining](#), [fisheries](#) and [aquaculture](#), and particularly autonomous robotics for [forestry and logging](#), and [quarrying](#).

The result of these emerging AI “solutions” is that tech corporations (including Accenture, Amazon, Google, IBM, Microsoft and Oracle) are competing to service rising demand for AI technologies from the fossil fuel industry; [what has been described as ‘automating the climate crisis’](#). Putative gains in productivity and efficiency for global supply chains through AI are often described as a “game changer”, or as the “future of” various extractive industries. However [metrics and rhetoric of corporate social responsibility exaggerate benefits and obscure costs](#) of AI in natural resource extraction.

The portrayal of AI as a solution for climate and sustainability in these industries further distances the visibility of pollution and waste in supply chains, utilizes micro-level environmental gains to obviate the need for macro-level changes, and ultimately works to reinforce existing narratives of corporate responsibility, thus legitimizing structural practices which continue to exacerbate the metacrisis.

Oil & Gas Industry

Artificial intelligence is finding wide adoption and use cases across the oil and gas industry. [50% of oil and gas executives have already begun to use AI in their organizations](#). [92% of oil and gas companies](#) are either currently investing in AI, or plan to in the next 2 years. The AI in oil and gas market has been valued at US\$2 billion in 2019, and [expected to reach \\$4.21 billion by 2028](#) (~12% annual growth)

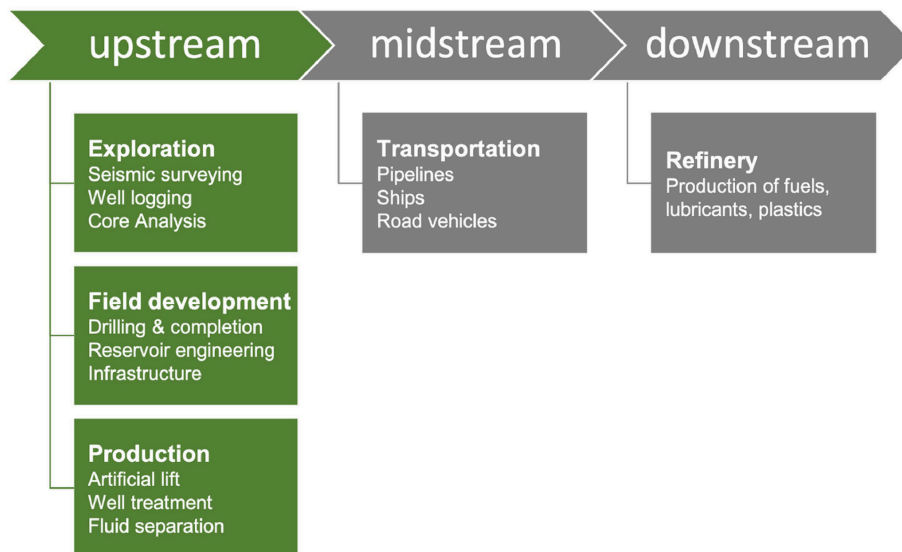
Industry Trends and Collaborations

- [ExxonMobil has started using AI capabilities](#) to support two major areas, predictive maintenance for oil analysis & reduced labor costs, and expedited well development.
- Shell [currently using big data analysis to monitor operations, supply chains and manage its inventory](#), and is planning to utilize AI to improve production outputs and efficiency
- Oil companies have invested in [seismic interpretation methods](#) (using acoustic energy to map geological reservoirs of oil and gas) for decades, however traditional manual trial-and-error approach is arduous and time consuming. The use of AI [has begun to revolutionize the field](#). AI-driven seismic interpretation could **halve** the time needed for well development and planning.
- In various cases, oil and gas companies are also claiming that using AI will bolster their sustainability efforts and reaching net zero faster (e.g. [Shell Energy Podcast](#))
- High-tech firms (including Accenture, Amazon, Google, IBM, Microsoft and Oracle) are competing to service rising demand for AI technologies from the fossil fuel industry – [what has been described as ‘automating the climate crisis’](#).
 - [Sensia](#), established in 2019 is aiming to automate services, drilling, exploration and extend well life to enhance profit margins in oil and gas.
 - [Amazon Web Services offers specialized computing products for oil and gas](#) to identify potential reservoirs faster and cheaper, to optimize production and profitability
- Broadly, one of the [major AI tech trends in oil and gas is the use of digital twin models](#), which are used to represent physical assets (refineries, pipelines, wells) for predictive maintenance, minimizing outages and optimizing performance. These generate cost savings, boosting operations and accelerate growth. The digital twin market by itself is projected to be worth \$154 billion by 2030.
- The GPU manufacturer NVIDIA has multiple direct partnerships for high performance computing (HPC) AI data centers, specialized hardware and software [specifically for oil and gas](#)
 - A recent collaboration [between NVIDIA and Shearwater GeoServices](#) in 2023 [achieved 10x faster performance](#) of seismic processing software key to oil and gas exploration workflows, utilizing the [GH200 Grace Hopper Superchip](#), specifically designed for AI HPC data centers
 - In 2018 NVIDIA announced a [collaboration with Baker Hughes](#) (subsidiary of General Electric) to develop machine learning to “dramatically reduce the cost, finding, extracting, processing and delivering oil”

- Orbital Sidekick, founded as part of NVIDIA's Inception startup program, was the first to use hyperspectral imaging to detect gas pipeline leaks from space, utilizing [NVIDIA's Jetson AGX module as an AI engine](#) which is designed for autonomous machines

AI Applications

Oil and gas industry is segmented into upstream (exploration, subsurface mining and petroleum/gas production), midstream (logistics and transportation), and downstream (processing and conversion to finished product). AI is currently in use, and finding more applications across all these segments of the oil and gas industry¹⁷.



Division of the oil and gas industry into sectors. (Fig 1 - [Koroteev & Tekic, 2021](#))

¹⁷ See the following for comprehensive reviews on AI applications in the oil and gas industry

Koroteev, D. and Tekic, Z. (2021) 'Artificial intelligence in oil and gas upstream: Trends, challenges, and scenarios for the future', *Energy and AI*, 3, p. 100041. <https://doi.org/10.1016/j.egyai.2020.100041>.

Choubey, S. and Karmakar, G.P. (2021) 'Artificial intelligence techniques and their application in oil and gas industry', *Artificial Intelligence Review*, 54(5), pp. 3665–3683. <https://doi.org/10.1007/s10462-020-09935-1>.

Gupta, D. and Shah, M. (2022) 'A comprehensive study on artificial intelligence in oil and gas sector', *Environmental Science and Pollution Research*, 29(34), pp. 50984–50997. <https://doi.org/10.1007/s11356-021-15379-z>.

- Upstream applications of AI in oil and gas include seismic interpretation, petrophysical data analysis and reducing downtime in drilling operations.
 - Subsurface geology can be automatically mapped by machine learning algorithms for rapid acceleration of prospecting new wells, reducing exploration risks and increasing success rates.
 - [In one case](#), machine learning assisted drilling utilizing real-time telemetry reduced “non-profitable time down to 20-50% on average, with an extensive decline of around 90% in failures”.
 - Recent advances in petrophysical interpretation (usually quite time consuming) using automated interpretation algorithms trained on historical wells was 92% accurate, [with up to 1000x increase in speed over manual methods](#).
- Midstream applications primarily focus on optimizing transportation and aiding or automating inspection and repair of pipelines.
 - The global market for pipeline inspection robots (beyond just oil and gas infrastructure) is [expected to exceed \\$2 billion by 2025](#). AI-driven robots are being developed to aid in addressing corrosion, aging, cracks and various kinds of damages
 - AI is also being used to optimize transport flows, such as [maritime route planning](#) and [generative deep learning for decision making in gas networks](#)
- Downstream applications of AI in refineries include “digitalization oil fields”, and optimization of refinery operations via digital twins
 - In 2018, the [Spanish petrochemical company Repsol partnered with Google Cloud](#) to optimize management of its Tarragona refining complex
 - The company Imubit has developed a [“closed loop neural network” platform](#) based on deep reinforcement learning to optimize and manage virtually every process in refineries. Similarly, [AspenTech has developed “smart refinery” technology](#), aiming for completely autonomous and [self-optimizing oil and gas refineries](#).

Industrial Agriculture

The global food system is estimated to contribute between 16.5 - 28.1% of all GHG emissions¹⁸. Industrial agriculture in particular is a leading cause of deforestation¹⁹, biodiversity loss and a major driver of the ongoing Sixth Mass Extinction. Beyond merely exacerbating climate impacts by accelerating growth of large-scale industrial agriculture, AI innovations whilst in some cases may putatively reduce impacts on ecosystems may ultimately result in simply exacerbating all of these biosphere-wide externalities.

Industry Trends

- The autonomous-farming industry is [beginning to boom](#), with about 200 AI-based agriculture startups in the US alone, and [over 1,400 globally](#).
 - E.g. self-driving tractors and combine harvesters, optimized irrigation systems, robot swarms for crop inspection, autonomous sprayers, computer vision for indoor crop analytics, indoor robotic harvesting, machine learning for targeted herbicide applications, AI-monitoring for livestock health, pests, disease and soil quality, yield mapping, automated harvest sorting
 - Much of the above examples fall under the banner of [precision agriculture](#), which is being termed by some as the [Fourth Agricultural Revolution](#). Notably, many of these developments, utilizing blockchain, IoT and AI, are promising to deliver “[climate-smart](#)” [sustainable agriculture production](#).
- The AI in agriculture market is [expected to grow](#) 2.08 billion from 2024 to 5.76 billion by 2029 - a CAGR of 22.5%
- AI and Precision Agriculture could **reduce annual agriculture operating costs by more than 22% globally**, [according to Arkinvest](#).
 - [Combined uses of various](#) AI-enabled autonomous technologies (planters, sprayers etc) have been projected to drop seed, fertilizer and chemical use by 27%. Labor cost could drop by up to 85%, and predictive maintenance + improvements in efficiency could reduce fuel, lubricant, electricity and repairs by 20%
 - ARK estimates that 75-80% of farms will adopt autonomous farming technologies globally.

¹⁸ This [wide range](#) is due to major differences in methodology (e.g. life cycle analysis versus bottom-up or top-down emissions accounting) and system boundaries. Recent well supported high-end estimates for the contribution from animal agriculture in particular is ~24.1% ([Crippa et al. 2021](#))

¹⁹ See:

<https://earthobservatory.nasa.gov/images/148674/sizing-up-how-agriculture-connects-to-deforestation>

[Global annual tree cover loss by driver](#) by Global Forest Watch

Sandker, M. *et al.* (2017) ‘Global deforestation patterns: comparing recent and past forest loss processes through a spatially explicit analysis’, *The International Forestry Review*, 19(3), pp. 350–368
<https://www.jstor.org/stable/44708624>

AI Applications

Virtually all activities involved in industrial agriculture have AI, machine learning and autonomous robotics applications, across cultivation, monitoring and harvesting systems²⁰.

- [Soil management techniques](#) utilizing AI have been developed including prediction of key soil properties like moisture, temperature, erosion and fertilizer leaching
- Particle swarm algorithms, similar to those used developed for military drones, emulate flocks of birds and fish to control autonomous agricultural robots for harvesting and spraying (see e.g. [SwarmFarm robotics](#))
- AI algorithms are critical to [IoT devices used for precision agriculture techniques](#) such as weed identification and targeting, intelligent irrigation systems, pest control and disease prediction, all of which contributes to increased agricultural yields and productivity

System-Wide Jevons Paradox

Summary

AI has the potential to contribute positively to advancements in energy technologies. However, within the context of Jevons Paradox (and the background Garret Relation coupling GDP/energy use), the integration of exponentially advancing AI with our current systems will lead to detrimental effects on the environment on net. The critical issue is not the amount of resources consumed by AI infrastructure directly (e.g. growing computational requirements, server farms, cooling, material requirements etc.), but rather the **exacerbated environmental impacts from economic growth facilitated by AI**. Such accelerated environmental impacts risks pushing us past critical ecological tipping points under the guise of potential benefits.

²⁰ For a comprehensive review of AI applications in agriculture, see

Eli-Chukwu, N.C. (2019) 'Applications of Artificial Intelligence in Agriculture: A Review', *Engineering, Technology & Applied Science Research*, 9(4), pp. 4377–4383.
<https://pdfs.semanticscholar.org/6e08/7108aa8048da8cfc82cdec7071a55bab488.pdf>

Javaid, M. *et al.* (2023) 'Understanding the potential applications of Artificial Intelligence in Agriculture Sector', *Advanced Agrochem*, 2(1), pp. 15–30. <https://doi.org/10.1016/j.aac.2022.10.001>.

Wakchaure, M., Patle, B.K. and Mahindrakar, A.K. (2023) 'Application of AI techniques and robotics in agriculture: A review', *Artificial Intelligence in the Life Sciences*, 3, p. 100057.
<https://doi.org/10.1016/j.aillsci.2023.100057>.

Jevons Paradox dynamics are already playing out with AI in multiple domains

- [AI is reducing the cost](#) of petroleum discovery and refining
- AI-powered translation shows Jevons Paradox dynamics, [increasing rather than decreasing demand](#) for translation services
- There is strong precedence for Jevons Paradox dynamics with agriculture, such as agricultural intensification and yield increases [driving increases in land conversion](#) rather than decreases
 - Precision agriculture may result in more efficient use of herbicides, fungicides (etc.), but Jevons paradox dynamics [may very likely influence greater use of such crop protection products](#), since even if the volume decreases per spraying pass (more efficiency), the total amount of spraying activity may increase as automation makes it easier (greater total consumption)
 - For agriculture in particular, rebound effects are [additionally exacerbated by natural adaptations](#) such as pesticide and herbicide resistance

Rebound and Backfire Effects

Rebound and backfire effects may occur mechanistically at [different economic levels and time frames](#), from decisions made at the household level to global level changes in trade patterns.

The following notes here are not exhaustive of how AI may result in system-wide Jevon's paradox, but should be read as indicative of previous examples with comparable digital technologies, and evidence of rebound effects in industries where AI is expected to boost productivity, efficiency and outputs.

- **Direct and Indirect Effects:** in the standard formulation of Jevons Paradox²¹, increased energy efficiency tends to increase energy consumption by two means.
 - First, increased energy efficiency makes the use of energy relatively cheaper, thus encouraging increased use (the direct rebound effect).
 - Second, increased energy efficiency leads to increased economic growth, which pulls up energy use for the whole economy (the indirect effect). It applies to energy efficiency and resource efficiency.
- There have been extraordinary efficiency improvements for computation over the past fifty years. However, improvements in computing's energy-efficiency do not necessarily, and have not historically, led to reductions in its energy consumption. [If Jevons Paradox occurs, improving computing's energy efficiency may contribute to increasing energy consumption.](#)

²¹ One can also look at Jevons Paradox from a [complexity theory point of view](#); this reveals some inherent problems in the use of quantitative efficiency metrics and can explain systemic drivers of the rebound effect. The linked paper also has strong implications as to why technology innovation can never solve environmental problems.

- Given the exponential growth in computational demand, it's overwhelmingly likely any improvements in energy efficiency will simply further enhance demand through lower costs.
- Previous tech involved in digitisation and dematerialization trends (which are usually claimed to bolster sustainability by “decoupling” economic output from material impacts) have demonstrated remarkable rebound effects
 - [An extensive review of 57 different cases of technologies](#) (e.g. hard disks, photovoltaics, transistors, etc) found that no dematerialization occurs due to demand rebound (Jevon's) impacts which overcompensates technological progress
- Using AI to reduce emissions may backfire in a number of ways
 - Tech advancements can lead to more carbon emissions by driving output growth ([Acemoglu and Restrepo 2020](#))
 - AI advancements could reduce energy costs and energy service prices (see [Electricity and Power Systems](#)), but this in turn may increase consumption of energy in other sectors
- Existing machine learning applications have already seen Jevons-like dynamics play out.
 - For example, [the use of AI and machine learning assisted translation services](#) did reduce costs in translation services by 40%, but the Bureau of Labor Statistics found an uptick of 19% in translation and interpreter services in 2021.
- Transnational corporations have hailed gains in productivity and efficiency of global supply chains through AI as a “game changer”, however [metrics and rhetoric of corporate social responsibility exaggerate benefits and obscure costs](#).
 - There have been at least four major environmental costs from rising power of AI in the global economy: ramped up global production, accelerated consumerism, intensification of mining (particularly areas with human rights issues), and growing volumes of electronic waste in developing countries
 - [McKinsey Global Institute](#) estimates deep learning has the potential to increase annual value of packaged consumer goods by up to US\$500 billion, and the global retail sector by upwards of US\$800 billion
 - AI is also being invested heavily in enhancing advertising revenues. Amazon's recommendation engine generates approximately one-third of annual sales. Shopping malls are increasingly using facial recognition cameras for consumer profiling to optimize impulse purchases.
- **Agriculture:** Per the above points on AI applications in agriculture, it should be expected that autonomous farming technologies will drastically reduce costs and increase agricultural productivity. This would likely lead to an accelerated expansion in agricultural land, worsening impacts on deforestation, species defaunation, extinction, nitrogen runoff, and further n-th order impacts.
 - There is strong precedence for Jevons Paradox dynamics with agriculture, such as agricultural intensification and yield increases [driving increases in land conversion](#) rather than decreases.
 - [Greater yield increases have been shown to lead to higher deforestation rates](#) in Sub-Saharan Africa, Latin America and the Caribbean.

- Similar effects have been found due to [improvements in efficient irrigation](#) for water management, increasing groundwater extraction and driving switches to higher-revenue but more water intensive crops.
- **Counterargument:** Jevons Paradox for agriculture gets more nuanced when taking into account cross-regional Total Factor Productivity (TFP) growth. Whilst improvements in regional TFP are associated with regional land expansion, [Villoria \(2019\)](#) argues this expansion is more than compensated by improvements in TFP in the rest of the world.
 - They claim that in the absence of TFP growth, an additional 125 Mha would have been needed to satisfy food demand from 2001 - 2010, half of which would have been in the most biodiverse biomes of the world.
 - They also find that recent rates of TFP growth seem insufficient to counteract short term increases in demand from 2020 - 2025. From this point of view, AI improvements to agriculture could hold promise to maintain net reduction in land conversion from indirect effects of extra-regional TFP growth.
- **Counterargument:** The theoretical foundations of sustainable agriculture development is supported by the [Borlaug Hypothesis](#) (also referred to as [land-sparing hypothesis](#)), which roughly states that if new agricultural technologies boost per hectare crop productivity, more can be produced on the same or even a smaller amount of land. Higher crop yields increase the supply on agricultural output markets, where prices drop and thus reduce the incentives for cropland expansion. This effectively runs directly counter to the Jevons Paradox.
 - **Response:** In reality, empirical research on the land-sparing effects of improvements in agricultural total factor productivity is quite mixed. [Recent work seems to contradict the Borlaug Hypothesis](#), showing that in the long-run agricultural TFP has a U-shaped relationship with CO₂ emissions; TFP growth improves environmental quality up to a point, after which further growth degrades the environment. [Similarly, Garcia et al. \(2020\)](#) showed that low- and middle-income countries with export-oriented agriculture showed strong rebound effects with intensification driving cropland expansion.
- AI-enabled developing country growth
 - Hypothesis: LLMs as personal assistants and translators, especially of less represented languages, will drive additional market capture in emerging markets
 - Emerging markets are [positioned to drive the largest economic growth and growth in power demand and emissions](#)
 - LLM-powered mobile operating systems will expose more market participants to software as services through translation and voice interfaces
 - Expanding demand for energy intensive services should be expected in developing economies

- Multi-level [moral-psychological rebound effects](#)²²:
 - **Economy-wide techno-optimism**: Perceived improvements through technological innovations - and particularly in the case of AI where strong marketing pushes are distinctly utilizing “sustainability”, “green”, “nature positive”, and “climate friendly” frames - make create a general sense that environmental problems can be solved by technical AI innovations, thus not requiring system-level nor individual-level behavioral changes
 - **Indirect consumer perception effect**²³: Efficiency (or other) improvements - whether real or not - may improve the perception from consumers that a particular product or service has a lower consumption footprint (carbon, water, energy, material, etc.), and thus reduce their restraint in consumption of other products.
 - **Direct consumer perception effect**: A more environmentally positive/sustainable image for a given product or service may reduce restraint in greater consumption of the same product or service

Positive Views on Jevons Paradox and AI

- Techno-optimist view that Jevons paradox actually supports the economic case for AI
 - “Broadly, we believe that a drop in marginal value of creation will massively drive demand. Historically, in fact, the Jevons paradox consistently proves true: When the marginal cost of a good with elastic demand (e.g., compute or distribution) goes down, the demand more than increases to compensate. **The result is more jobs, more economic expansion, and better goods for consumers.** This was the case with the microchip and the Internet, and it’ll happen with generative AI, too.” ([A16Z blog](#))

²² For more comprehensive analysis of socio-psychological mechanisms for rebound effects, see:

Reimers, H. *et al.* (2021) ‘Indirect rebound effects on the consumer level: A state-of-the-art literature review’, *Cleaner and Responsible Consumption*, 3, p. 100032. <https://doi.org/10.1016/j.clrc.2021.100032>.

Reimers, H., Lasarov, W. and Hoffmann, S. (2022) ‘Moral-psychological mechanisms of rebound effects from a consumer-centered perspective: A conceptualization and research directions’, *Frontiers in Psychology*, 13. <https://doi.org/10.3389/fpsyg.2022.886384>.

²³ This is also known as “moral licensing” in psychological literature, and has been empirically observed in relatively simple cases such as people [using more office and hand towel paper when a recycling bin is available](#).

In the indirect case, moral licensing can “spillover” across consumption categories such as in the case of eco-labelled goods, where [research demonstrated that participants who preferred more “green products” displayed less environmental behavior](#).

Garrett Relation and AI

The Garrett Relation postulates a fixed relationship between energy growth and global GDP accumulated over history ([Garrett, 2011](#)). It is supported by both observational evidence and a thermodynamic justification. Broadly speaking, insofar as AI may accelerate economic growth as such, if the Garrett Relation holds then it also necessarily implies unavoidable increases in primary energy consumption, regardless of purported efficiency gains.

- **Observational evidence:** Measuring indexed changes to rate of global energy consumption (E , exajoules) and cumulative production²⁴ since 1970 (W , trillion US\$ 2019), one finds that the ratio $W / E = \lambda$ is roughly constant; approximately **\$5.50 ± 0.21 trillion USD (2019) per exajoule of energy consumed each year** ([Garrett et al. 2022](#))²⁵. This is equivalent to about **182 kilojoules per cumulative dollar of economic production**.
 - This ratio has been remarkably stable over the 50 year period from which it's been calculated. The standard deviation of λ has been measured at ~3%, over a time period where GDP increased by 238% and cumulative production by 111%. The stability of this relation over time suggests this is not merely a correlation.
- **Thermodynamic argument:** Civilization is an open thermodynamic system, where external energy sources and raw materials are used and dissipated as waste. Historical cumulative GDP resides in the capacity to use energy to sustain flows of communication and materials along networks connecting people, capital and the physical environment. An individual capital element has no value without associated consumption of primary energy to actively connect it to other elements of civilization.
 - See also [Tim Garrett's discussion of the thermodynamics of civilization and collapse here](#)
- AI represents an enormous possible accelerant of economic growth for the global economy. [PwC projects](#) that AI could potentially contribute upwards of \$15.7 trillion by 2030.
 - Using a 2016 adjusted value²⁶ for λ (5.33 trillion USD per exajoule), and a differential relation between economic production to the rate of increase in world primary energy consumption, we estimate that **additional energy required by the added economic growth from AI would be 2.66 EJ/yr, in 2030**²⁷. This is

²⁴ It's important to stress this relation holds for *cumulative production*, not GDP. Cumulative production is the time integral of economic production over the entirety of human history. To analogize with the metabolism of the human body, this is like the relation between body mass and caloric requirement. Annual economic production, by contrast, would be like how much weight is gained in a year.

²⁵ For a least-squares fit to a logarithm of W and E , the relation is $W = 5.47 * E^{1.00}$. As a linear fit, the relation is $W = 5.67 * E - 66$. For the decadal mean, $w \sim 5.5 \pm 0.21$, which applies better to the inferred relation to the rate of increase in world primary energy consumption, $Y = dW/dt = w dE/dt$

²⁶ [Inflation calculator](#) gives \$1 (2019) = \$0.94 (2016), so $\lambda_{2016} = \lambda_{2019} * 0.94 = 5.33$ (using $\lambda_{2019} = 5.67$ from the linear fit above)

²⁷ Using the inflation adjusted decadal mean scaling constant, this gives $dE/dt = 15.7/5.33 \sim 2.66$ EJ/yr

equivalent to adding about 739TWh in one year, which is approximately the entire nuclear energy generation of the US²⁸.

- The degree to which additional consumed power from AI driven GDP growth results in CO₂ emissions will depend on the future energy mix in question.
 - Assuming current policy projection in the IEA STEPS scenario with a rough average carbon intensity by 2030 of 400 gCO₂/kWh, then this becomes about 2.96x10¹⁴g CO₂, or about 29.6 MtCO₂, per year
- Note the above estimate only accounts for the necessary energy production from *added economic activity from AI*, not the energy requirements of AI itself.
 - The PwC projection does not provide an estimate for how much added computational load this would require, but given IEA projections suggested an increase of 540 TWh just from 2022 to 2026, **it would not be surprising if the real figure were 2x - 5x the estimate above, possibly >3000 TWh, which is larger than current global nuclear power generation (about 2552 TWh as of 2023)**.
 - This estimate should be taken with a significant degree of caution however, as the Garret relation is estimated based on highly smoothed data over the history of the global economy, and does not resolve short-term and fine-scale behaviors (such as year-to-year variability due to recessions or pandemics)²⁹.

²⁸ 779.19 TWh as of 2023 ([IAEA](#))

²⁹ Nevertheless, Garret et al. (2022) do calculate a running decadal mean of 5.9 ± 2.2 USD trillion per EJ (note the larger variability), which is quite close to the calculated constant of 5.50 ± 0.21. Based on the uncertainty, our estimated range should be between 1.94 – 4.24 EJ/yr.

AI Applications for Climate, Limitations and Risk

The following section considers some of the “high-impact” areas described by [Rolnick et al. \(2022\)](#) in the paper *Tackling Climate Change with Machine Learning*, a comprehensive review of the topic. Their work is also presented in the website [climatechange.ai](#), which hosts an [interactive list of summaries](#) for each application area.

Here, we discuss to what degree a number of “high leverage” examples³⁰ actually do constitute real solutions, where these may be quite marginal improvements, strong limitations for advancing a given ML application, and potential for exacerbating risks to climate or other symptoms of the metacrisis.

It should be noted that Rolnick et al. restrict their review to primarily *machine learning*, so only a few examples address *generative* AI in a limited few cases (primarily material science), and so should not be conflated with AI in general nor large language models. Moreover, they recognise that ML isn’t a silver bullet as well as various shortcomings, such as being used to accelerate fossil fuel exploration and energy demand as described above.

Summary of ML Climate Solution Space

- There are various ML-driven solutions **that may usefully contribute to climate change mitigation or adaptation**:
 - Smart cities, reducing food waste, and better concrete and steel are some notable, plausible high-impact areas.
 - Optimizing supply chains to reduce food waste and ML-designed concrete formulas alone could, optimistically, reduce 4% of global GHG emissions by themselves
 - Generative algorithms for battery material discovery are demonstrably useful, and ML approaches are broadly applicable to battery R&D as a whole due to the combinatorial nature of the problem space.
 - Battery tech is a necessary area for more sustainable decarbonization pathways to for less material resource intensity
 - ML emulation of cloud physics for climate models could significantly improve estimates of equilibrium climate sensitivity, which is the primary source of uncertainty for long-range global mean temperature change projections

³⁰ The examples selected here aim to be broad enough to address the most substantial application areas and to get a sense of how useful AI and machine learning may be for climate change generally, though we cannot claim to be comprehensive.

There are several more “high leverage” examples described by Rolnick that have not been included here for the sake of time but are worth investigating further. These include: Tracking and reducing illegal deforestation, ML for disaster response, Facilitating behavior change, Policy data and evaluating effectiveness and Financial climate analytics (e.g. carbon markets, investment risks, etc.)

- ML approaches for forecasting short- and long-term food security risks is obviously useful, but runs into limits for cases of multi-breadbasket failures - which will be more likely in the future
- Most purported solutions are useful at first order, but have **major caveats, challenges, risks and/or are contingent on other factors** coming into play
 - Areas where AI/ML can improve efficiency such as smart buildings, transport use & demand forecasting are liable to have Jevons rebound effects, which reduce overall efficacy of interventions
 - Nuclear fusion R&D can be assisted by ML/AI techniques but probably won't be relevant to near-term decarbonization pathways due to broader socio-economic and scaling considerations
 - Optimizing life-cycle fossil fuel emissions via supply chain optimization and reducing methane leakage, whilst worthwhile at first order, are extremely liable to boost the greenwashing efforts of fossil fuel companies
 - Better ML-driven modeling of ice sheet loss and sea level rise would be useful but is limited by current probabilistic coastal adaptation planning decision-making frameworks
 - Better ML-assisted ecosystem and biodiversity monitoring and big data approaches would be great for addressing major data gaps, but demonstrative efforts for biodiversity climate adaptation management are broadly lacking. The prevalent underlying regime of monitoring with limited or no action implies the need for systemic changes in biodiversity conservation
- **Some have specific, major systemic risk implications:**
 - ML-driven power system forecasting may not perform well with high renewable penetration and introduce greater systemic risks from data system fragility to cybersecurity vulnerabilities
 - ML + IoT for smart buildings and cities introduces systemic cybersecurity risks for domestic and urban areas. Moreover, the same tech which is useful for smart city infrastructure has dual use for authoritarian governance.
- **Various AI/ML ideas would simply accelerate or support currently deleterious efforts**
 - For example, better remote sensing of carbon stocks in forests sounds good on paper, but ties into fundamentally broken and ultimately worthless carbon offsets in the REDD+ sector³¹

³¹ [REDD+](#) stands for Reducing emissions from deforestation and forest degradation in developing countries, and was established as a framework under the Paris Agreement, and has been the center of numerous controversies:

[REDD+ projects falling far short of claimed carbon cuts, study finds](#) (Mongabay, 2023)

[SBTi policy and lessons from the 2023 redd+ controversy](#) (Illuminem, 2023)

[Revealed: more than 90% of rainforest carbon offsets by biggest certifier are worthless, analysis shows](#) (Guardian, 2023)

See [REDD-Monitor](#) for comprehensive up to date reporting on issues with the REDD+ sector

- **Some have unclear or as yet undetermined overall climate and systemic implications**
 - AI/ML-driven precision agriculture shows promise but is plausibly caught up in a hype cycle at present. The compatibility between precision ag tech and more truly regenerative agriculture and permaculture is unclear
 - Better estimates of carbon emissions and stocks from forestry, land use, agriculture (etc.) is probably useful but presents relatively few high impact leverage points

Emission Reductions

Firstly, there are some slightly more general arguments on how AI could accelerate emissions reductions not predicated on specific applications.

[Zhong et al. \(2023\)](#) notes a number of relatively straightforward ways AI may help reduce or inadvertently increase emissions:

- AI can enhance the ability of enterprises to limit emissions (e.g. Digital Twin-driven carbon control for manufacturing processes; [Zhang et al. 2019](#))
- AI can promote technological innovation which drives industrial upgrading; increasing productivity, reducing energy consumption and carbon intensity in industrial processes
- Enhanced efficiency of resource allocation, in turn reducing carbon emissions

One can also make a kind of market-driven argument for a faster renewable energy transition via AI + economy of scale argument (steelman):

- Destabilizing demand for energy by AI, crypto high energy requirements will drive a more efficient rollout of carbon neutral energy generation

- Market efficiencies occur when increases in scale & demand, specifically of infrastructure builders in countries where demand increases the most
- Recent large cost reductions in solar and wind energy ([estimated](#) 55% and 85% for solar and wind respectively, and 85% for batteries from 2010 – 2020), along with [globally competitive levelized cost of electricity \(LCOE\) against oil and gas](#), would suggest that the build out for required electrical generation for AI data servers would mostly be renewables
- Moreover, as rapid changes in demand for energy in the grid occur, power companies need to scale up and improve expansion efficiency
 - Increased efficiency of renewable energy infrastructure building lowers the cost and increases viability of replacing aging coal and natural gas plants, speeding up a clean energy transition

Rebuttal for AI-Driven Emissions Reductions

- The lists above of innovation areas where AI may accelerate emission reductions across various sectors lean heavily on areas of efficiency improvements and optimization; standard areas where Jevons Paradox often arises.
 - [Various studies](#) across macroeconomic to micro-level scales find persistent empirical evidence of rebound and backfire effects where improvements in energy efficiency drives increases in energy consumption..
 - [At the national scale](#) this has been found to be a backfire effect upwards of 208%³² for China. [Meso-scale studies on energy-specific technology gains across the U.S. economy](#) found substantial backfire effects across 30 sectors; 120% long-term rebound in energy use in the utility sectors, services (179%), agriculture (381%), communication (104%), and government enterprises (182%). The overall rebound level was 172%.
 - If the majority of AI applications for enhancing mitigation are aimed at efficiency gains, then it's likely given the historical trend these would merely backfire
- Focusing purely on emission reduction improvements can run afoul of "[carbon tunnel vision](#)". Technological solutions in mainstream mitigation policies such as electric vehicles, solar, wind and batteries have extraordinary material costs
 - Between 2022 - 2050, the energy transition could require the production of [>6.5 billion tonnes of end-use materials](#); primarily steel, copper and aluminum.
- Historically, whilst renewables are composing a growing share of overall energy production, **these are not replacing fossil fuels** but [simply expanding the overall amount of energy that is produced](#)

³² See [Appendix](#) for an explanation of interpreting rebound and backfire effect percentages.

- ³³As such, even if the AI market-driven acceleration argument above held, this would not guarantee that renewable energy growth would actually displace existing (and also growing) fossil fuel power generation

Electricity and Power Systems

Various niche and broad applications for electricity and power systems from energy forecasting and ML-based battery material discovery to emissions modeling, optimal installation of rooftop solar and managing rural microgrids.

Forecasting supply & demand

- **Key idea:** more advanced ML methods will be used to forecast electricity supply and demand, allowing for better real-time electricity scheduling and long-term system planning.
 - Various ML-based electricity forecasting methods already exist and some relatively basic ones using support vector machine regression (SVMs) have already been implemented in cities like New York and Guangdong
 - More advanced methods are being developed which could integrate innovations in weather forecasting, and optimize for other targets such as minimizing overall costs & GHG emissions
- **Risks/Limitations:** One major risk is that the incorrect use of the trained models can lead to a critical deterioration in forecast accuracy and the appearance of errors unacceptable for the models' operation
 - Training data corruption could occur from any number of reasons: data providers, failed collection systems (meters, sensors, data transmission), read-write errors, implementation error due to human factors. This makes ML-based forecasting models which integrate more data equivalently more fragile.
- Improved prediction accuracy does not necessarily provide reliability and resilience, particularly when real-world operating conditions change
- It's unclear how well even ML based forecasting is going to operate in high-penetration scenarios for renewables due to low grid inertia. Higher percentages of renewable energy sources mean less rotational inertia from conventional generators, lowering inertial support³⁴ for the entire grid.
 - This consequently makes the entire grid more susceptible to sudden generation loss, short circuits and other grid-scale faults. Overall risk of grid frequency instability and cascading failures increase.

³³ Note, this is a contentious point however. See for example:

Karlilar Pata, S. and Balcilar, M. (2024) 'Decarbonizing energy: Evaluating fossil fuel displacement by renewables in OECD countries', Environmental Science and Pollution Research, 31(21), pp. 31304–31313. <https://doi.org/10.1007/s11356-024-33324-8>, which argues that displacing 1% of fossil fuels is achieved by increasing about 1.15% of renewable generation capacity

³⁴ Grid inertia is important since it resists drop in grid frequency, a measure of the balance between electricity supply and demand.

- In extremely low inertia systems, small frequency variations can cause cascading tripping of generation units under load shedding, with blackouts under worst-case.
- There's plausible solutions to this (e.g. synthetic inertial support from batteries, fast frequency response services for extreme events, etc), but [the challenges are significant](#)

Accelerating materials science

- **Key ideas:** ML models can be used to design new materials for areas such as [photoabsorber materials for solar fuels](#) (synthetic fuels produced from sunlight or solar heat), battery storage materials and electrocatalysts; and non-electrical system applications such as alternative cements, better CO₂ sorbents, etc.
- **Useful for batteries:** ML for battery materials research probably isn't hype. An extensive, authoritative review on the topic by [Lombardo et al. \(2022\)](#) shows and argues that ML algorithms across the board assist in overcoming the main limitations in battery optimization - namely combinatorial explosion in electrode/cell manufacturing. Though there are existing challenges and limitations like overfitting on limited data sets, this is one of the few areas where ML-assisted research has fairly high potential.
- **Marginal benefit:**
 - For *solar fuels*, regardless of the fuel in question (hydrogen or carbon-based) the use-case for climate is [highly debatable](#). In either case, [relying on combustible fuels risks furthering current fossil fuel dependency](#), particularly in cases where direct electrification is already possible. Carbon-based fuels would theoretically close the loop on carbon emissions, but not lower them. The reliability of reported experimental results in photocatalytic reduction [have been systemically questioned](#).
- **Risks:**
 - **Omni-use potential:** As has been previously explored in the context of synthetic biology, the generative models that could be putatively used to discover new battery materials and CO₂ sorbents can just as well be used for accelerating the pharmaceutical industry via novel drug discovery, and chemical/biological weapons manufacturing. The [latest paper cited by](#) Rolnick et al. on this point demonstrates their molecule generation method for functional drug-like molecules.

Nuclear Fission and AI

- **Key ideas:** AI could be used to boost the efficiency of nuclear power production, particularly with the use of digital twins. The International Atomic Energy Agency has established a [coordinated research project](#) to explore how AI can help expedite deployment of small modular reactors.

- Machine learning has already been applied in the nuclear industry for some time for areas like predictive maintenance
- Further efficiency gains could be found in automation using robotic AI systems to handle routine tasks, optimizing fuel consumption and maximizing reactor output
- [Other applications](#) include fault management & diagnosis and identification of accident scenarios
- Overall it doesn't appear that AI and machine learning is a "game changer" in the nuclear fission space, but certainly helpful in niche applications.
 - Efficiency gains in fission are unlikely to run into Jevons paradox. These gains are largely to do with the efficiency of construction, operation, and maintenance – not energy output, and even in these cases appear to be marginal and/or hindered by black box issues, data availability and explainability³⁵
 - Capital costs and construction time for nuclear power are extremely high, so unless AI were to discover a radically cheaper and more efficient way of doing small modular reactors (speculative, but apparently [part of the IAEA project](#)), it's difficult to see either substantial improvements in fission reactor rollout or cascading rebound effects.

Nuclear Fusion and AI

- **Key ideas:** ML can and has been previously used to accelerate R&D in nuclear fusion efforts, such as experimental design, parameter exploration, reactor monitoring, predicting plasma instabilities and speculatively steering reactor control
- **True but overblown:** This is all broadly accurate, however most of these applications are largely time/cost saving on marginal elements.
 - **Genuine results:** The most impressive example thus far is [ML predicting and avoiding tearing mode instabilities](#), which can otherwise terminate plasma operation and cause reactor damage. Controlling tearing mode instabilities is a necessary but not sufficient part of making fusion power happen – in the case of magnetic confinement using tokamak designs, at least³⁶.
 - In the private fusion space there's likely to be greater adoption of AI/ML-enabled advances. [Back in 2017](#) one of the major Fusion companies Tri-Alpha Energy partnered with Google to implement an algorithm to vastly improve the efficiency in experiment selection and design, which led to the discovery of a confinement regime with >50%

³⁵ Huang, Q. et al. (2023) 'A review of the application of artificial intelligence to nuclear reactors: Where we are and what's next', Heliyon, 9(3), p. e13883. <https://doi.org/10.1016/j.heliyon.2023.e13883>

³⁶ Notably, most private fusion startups separate from mainstream academic efforts are [pursuing non-tokamak designs](#)

reduction in energy loss rate. ([see associated paper](#)). The algorithm used by Tri-Alpha Energy for their experiment design [compressed two weeks of work into an afternoon](#) – *useful* but not *groundbreaking*.

- Notably, TAE Technologies has pushed back its projected commercialization date for a nuclear fusion power plant [several times over its existence](#); from 2015 - 2020, to 2020 - 2025, and now [“early 2030s”](#) for the expected first prototype power plant.
- **Ongoing AI - Fusion R&D programs:** The International Atomic Energy Agency established a [5-year Coordinated Research Project to accelerate Fusion R&D with AI/ML](#), so the crossover between these areas is taken quite seriously. See also this review paper on [Deep Learning applications to Nuclear Fusion research](#)
- **Techno-optimist hype bubbles:** Most discussion and media on AI accelerating progress on nuclear fusion reads like techno-optimist hype (cf. the kinds of posts on X on [“nuclear fusion and machine learning”](#)). At the same time, there are legitimate developments in the field enabled by AI/ML.
 - At the same time, reporting generally overplays the hype. E.g. This CNN article [Nuclear fusion: Scientists say they can use AI to solve a key problem in the quest for clean energy](#), reports on the tearing mode instability control via AI noted above.
- **Success is contingent on other non-AI factors:** How well fusion energy is fit to play a role in decarbonization scenarios this century is also [debatable](#). Current reactor designs are not economically viable, and large-scale adoption is going to be highly dependent on national government stances.
 - [Various issues](#), such as intermediate-level radioactive waste, non-existent supply chains, complex necessary technical support infrastructure, lithium and beryllium material constraints (etc.), make fusion adoption timelines, *even with* R&D breakthroughs, probably quite long and of less impactful relevance for climate mitigation over the next ~50 years.
 - Reactor designs lack economic competitiveness in future energy market models, and construction timescales are extremely long ([Griffiths et al. 2022](#)). Although fusion will probably play a role in the energy mix by the late century, this is [predicated on there being successful commercialization by 2030](#).
- **Small possibility of major AI-driven fusion breakthrough with superconductors:** Though not noted in the Rolnick et al. review, one commonly speculated use of generative AI (material discovery algorithms) is for room-temperature superconductors. [ML-driven search for superconducting materials](#) is already happening. [Historically, superconductor breakthroughs have spurred fusion development](#).
 - One of the last breakthroughs in fusion was the discovery and use of [REBCO high-temperature superconductors](#). These allowed for a new iteration of the MIT Alcator C-Mod Tokamak, which ultimately led to the [Commonwealth Fusion System's SPARC design](#) and [other much smaller](#) (and hence more economical) tokamaks. As such, there's precedent for even better superconductors accelerating fusion reactor design.

- There are several major caveats with superconductor breakthroughs for fusion:
 - Firstly, room-temperature superconductors may not necessarily work in high magnetic field settings. The magnetic field strength in magnetic confinement fusion reactors ranges from 5 - 24 Tesla (T). [To put these in context](#), a 5T magnetic field is powerful enough to lift 32,869 refrigerators, 10T is powerful enough to lift 102 Statues of Liberty, 15T could lift 31 Eiffel Towers, and a 20T field could lift over 700 SpaceX Falcon 9 rockets. It's unknown whether a room-temperature superconductor would even exist within the operating paradigms of conventional fusion reactors.
 - Secondly, even after discovery, HTS superconductors take decades of work to become operational in applications - REBCO superconductivity was discovered in 1987 but only became manufactured into tapes in 2010. Thirdly, even if you can get room temperature superconductors working in high-field settings and solve manufacturing hurdles, the engineering challenges from magnetically induced forces and torques are non-trivial and major limiting factors.
- **Systemic Jevon's accelerant:** Conversely, even in the most optimistic scenario of fast fusion R&D breakthroughs, high technology diffusion and wide adoption, this would plausibly act as a step-change in energy surplus, or at least a continuation of fossil fuel-level EROIs (preliminary estimates of fusion EROI are between 27 - 170). Given previous [energy system transitions have also resulted in a system-wide acceleration of economic production, extraction, and consumption](#), a comparable fusion-driven civilization-scale energy surplus would presumably have the same effect.
 - Without other systemic changes to underlying drivers of ecological overshoot and the metacrisis, further maximization of economic metabolism and energy flux via fusion, with concomitant greater material consumption and ecological damage, would simply accelerate.
 - In a way, this could be a *worst-case* scenario in disguise, as the massive boon of fusion in the near term would plausibly validate the memetic appeal of the naive progress narrative, further entrenching Molochian generative dynamics.
- **Geopolitical risks of fusion as a prestige technology:** As with AI and quantum computing, major breakthroughs in fusion R&D would very likely immediately make it yet another technological stream of geopolitical tension and great power competition. The first country to crack net gain fusion will hold a [significant first-mover advantage and leverage](#) in nuclear energy.
 - [Militarized fusion energy](#) would unlock various possibilities over fission for large-scale use of directed energy weapons, electric weapons, and railguns. It only took seven years after the first fission-generated commercial electricity for the first fission-powered aircraft carrier (*USS Enterprise*) to come into service in 1961

Reducing life-cycle fossil fuel emissions

- **Key idea:** One suggestion is to use ML along with sensor and/or satellite data for natural gas pipeline monitoring and prevent/detect methane leaks. Others include reducing emissions from freight transportation of fuels and optimizing power plant parameters to reduce CO₂ emissions.
- **Obvious greenwashing potential:** Though Rolnick et al. mention this one also has “uncertain impact” and that such projects “should be pursued with great care so as not to impede or prolong the transition to a low-carbon electricity system”, this misses the obvious near unavoidable systemic use of new tech by the fossil fuel industry to greenwash themselves to appear as having lower emissions.
 - Using AI/ML to avoid fugitive methane emissions from pipelines would in itself be a good thing - a [2022 analysis by EDF estimated](#) between 1.25 - 2.6 million metric tons of methane from pipelines a year in the US alone, between 3.7 - 8 times greater than EPA estimates, far higher than reported leaks, and [making it as bad as coal](#). At the same time however natural gas has [long been played as a “climate-friendly” alternative to coal](#). Lobby groups have [previously tried to get natural gas redefined as “green energy”](#). ML pipeline optimization to avoid leakages would no doubt add yet another AI-tech spin to this narrative warfare.
 - New leak detection tech [has already been played up by industry](#) for meeting climate goals. See also this [Australian Petroleum Production and Exploration Association report](#) which foregrounds leak detection in reducing industry methane emissions
 - A similar logic applies for other areas like optimized freight transport and power plant emissions. [This Boston Consulting Group report](#) is a good example of the kind of AI/ML strategy oil and gas is likely to proceed with for optimizing their emissions.

Transportation

Various applications include ML analysis of transportation data, modeling transport demand, enabling shared mobility in rideshare services, optimized freight routing, autonomous vehicles, grid integration and R&D for EVs, etc.

ML for Electric Vehicles

- **Key ideas:** There are a number of roles ML algorithms can and likely will play in EV tech; battery management, hybrid EV optimization, improving vehicle-to-grid technology, etc...
- **EVs are not really a climate solution:** [Nate’s roundtable with Arthur Berman, Pedro Prieto & Simon Michaux](#) covers this really well. Going over all the reasons why this is the case is out of the scope here, though some notable highlights:

- Michaux: [His modeling](#) found that a 30% market share of EVs in Europe by 2030³⁷ would require an *additional* 892.1 TWh of non-fossil fuel electrical generation (~14,921 avg power stations), and more copper, nickel, lithium, cobalt and vanadium than their respective global metal production in 2019 - equivalent to 280 years of global mining production for lithium, 175 years of global mining for vanadium, etc...
- Prieto: Passenger transport only accounts for about a quarter of oil use (in the case of Spain). The rest is going to areas like kerosene for aviation, LPG and diesel for trucks, machinery, agriculture, boilers, lubricants, etc. EVs are by themselves insufficient to decarbonize transport or oil use more broadly.
- [Berman](#): Passenger transport is only about 8% of global emissions, the amount of money³⁸ and effort towards EVs does not make sense in this context. The structure of fractionating towers and the oil refining process means we're not going to stop producing or using gasoline just because of EVs; depending on geology gasoline accounts for 30 - 45% of every barrel of oil - if we're still using oil for other applications (lubricants, asphalt, diesel & jet oil, etc.) EVs don't suddenly remove this gasoline from oil processing
- **Other socio-technical implications:** Importing advanced tech into the fundamental operation of passenger vehicles isn't limited to batteries and optimization; one of the current side-effects of how the automotive industry is going about its EV strategy is [stuffing them with increasing amounts of infotainment and digital capabilities](#); everything from Zoom meetings and streaming services to TikTok integration. On average, Americans [spend seven weeks a year in cars and 4.5 years in the average lifespan](#). This would likely extend further as infotainment in EVs would be designed to dopamine hijack drivers while waiting for it to charge. Consequently, this extends social media reach to a non-negligible fraction of human lives.
- **Legitimate benefits from ML:** At the same time, insofar as EVs will continue to grow in market penetration globally, ML algorithms particularly for grid integration makes logical sense. In the most optimistic case, EV's [could provide significant amounts of short-term grid storage](#) (possibly all in 2030 for very optimistic assumptions and 32 - 62 TWh by 2050) - the achievement of which almost [certainly requires extensive use of ML for electrical grid integration](#), though this is an extremely complex endeavor.

Transport use and demand modeling

- **Key ideas:** ML can (and is) being used to provide information for mobility patterns in transport planning tools, short-term forecasting of public transit ridership, reduce congestion on airport runways, and even using mobile phone sensors for urban topology and walking route choices. Broadly, these can be used to better inform efficient infrastructure decisions, and ultimately lower GHGs emissions.

³⁷ This refers to the [EV30@30 Campaign](#) and the IEA projection after the Paris Agreement in 2016 which predicted 30% EV penetration by 2030.

³⁸ [\\$90 billion in 2018](#) from the automotive industry

- **Unclear climate benefit:** Though noted as “high leverage”, many of the grab bag ideas under this banner described by Rolnick et al. seem fairly marginal with unclear direct effects on GHG emissions. How one goes from modeling walking routes in urban topology to lower emissions is somewhat speculative (presumably designing more walkable cities and suburbs would lower motor transport use?). Similarly, it’s unlikely minor reductions in emissions from optimizing airport runway congestion is comparable to the intrinsic impact of air travel as a whole³⁹.
- **Inherent optimization tradeoff for public transport:** Public transport network planning can be formulated as a [multi-objective optimization problem](#), for example between cost, emissions and travel time. It has been shown there is a negative relationship between carbon emissions and travel time & network resource use on the Pareto front⁴⁰. Regardless of whether this is solved with traditional nonlinear optimization solvers or machine-learning based methods, the unavoidable tradeoff means a more optimized transport network for GHG emissions is less convenient for passengers.
 - In the worst case, a public transport network optimized solely on GHG emissions would be significantly less convenient (e.g. less frequent services), leading to less users of the network, hence greater use of more GHG intensive road and taxi transport which could reverse any emission gains entirely.

Buildings and Cities

Smart Buildings

- **Key idea:** ML can be employed to optimize intelligent control systems for HVAC, lights, refrigeration (etc) to reduce energy usage by adapting to use patterns and electricity grid signals.
 - About 21% of global final energy consumption and 17% of global emissions have been estimated as attributable to residential energy use ([UNEP, 2022](#)); hence the logic for ML-driven smart buildings being a major lever for reducing emissions.
- **Jevon’s rebound effect:** Rolnick et al. themselves note for smart buildings that one of the dangers is a (Jevon’s) rebound effect, [which is well documented for consumer end-use efficiency gains](#). Though they note “in certain cases...additional building energy consumption typically ranging between 10% and 20%”, real rebound effects are likely much higher. Meta analyses of consumer energy services in developed countries find direct rebound effects ranging from 10% - 40% for water heating to 2 - 60% for space heating, for example.
 - **Economy-wide rebound effects** are usually much larger than assumed too. Across over 30 studies of energy efficiency rebound effects, [the economy-wide](#)

³⁹ [Apparently](#) about 5 - 20% of emissions from aircraft ground movement, depending on the mode of taxi. “Holding” is only 10% - so any efficiency improvements would only be a fraction of these.

⁴⁰ “The main reason is that taxi (online car-hailing) travel can save a lot of time, but its unit distance emissions are greater than other public transportation modes. While bus and rail transit with their large capacity and low emission advantages often make passengers travel longer due to route planning and transfer.”

[rebound usually exceeds 50%](#). In other words, economy-wide n-th order effects altogether erode more than half the energy savings from any improvements to energy efficiency

- **Rebound effects for smart homes:** [a review of work on smart homes](#) found that rebound effects are something of a blind spot in the field (<1% of research in the field considers rebound effects when studying energy efficiency of smart homes). What little work there is does indicate rebound effects (obviously) occur for smart homes, but tentatively suggest a full backfire effect does not occur (rebound doesn't cancel out efficiency gains entirely)
- **Global footprint of Internet of Things (IoT):** Smart homes are fundamentally predicated on a massive growth in IoT technologies; smart meters, fridges, AC, toasters, etc. **In 2023 there were 16.7 billion connected IoT devices** - an 18% growth from 14.3 billion in 2022. All of these require components such as semiconductors, actuators and sensor devices with relatively high embedded emissions from material & energy⁴¹ used in production and manufacturing – which could well outweigh the environmental savings of using those devices⁴².
 - [Primary energy demand just from IoT semiconductor production is projected to increase from 2 EJ in 2016 to 35 EJ in 2025](#) (see graph in Appendix - sharp exponential)
 - A [life cycle analysis of embodied carbon footprint from IoT edge devices](#) estimated 22 to 562 MtCO₂-eq/year in 2027, depending on number of devices and whether they were simple or complex (~150x difference in embodied emissions). **Worst case scenario was ~1000 Mt CO₂-eq/yr, equivalent to almost 20% of US 2022 emissions.**
 - Presumably, machine learning-driven smart homes would require more complex devices - making higher end estimates above more plausible.
 - Additionally, exponentially larger volumes of data are necessary for cyber-physical IoT systems, entailing greater energy use from cloud computing servers.

Smart Cities

- **Key ideas:** Extending ML-driven IoT to city-wide applications opens up a [wide array of avenues](#): coordination of [smart energy systems](#) for city district heating and cooling networks, power generations, EV charging stations, public lighting etc. Cyber-physical IoT sensor systems can optimize waste management, water quality monitoring. Big data analytics for urban computing which integrates credit card and social media data can [help support data-driven sustainability policies](#).
 - All of this putatively opens up the possibility of “eco-friendly” smart cities, with minimal environmental and climate impacts
 - As of 2022, about [250 smart city projects have been launched in 178 cities worldwide](#)

⁴¹ See [Energy footprint of IoT Devices](#) for a graph of energy demand in the Appendix.

⁴² See quote from [Ipsen et al. \(2019\) below](#) in the case of Smart Cities,

- **Benefits:**
 - **Largely sensible climate-ML application:** Given that cities account for 75% of natural resource consumption and 60 - 80% of GHGs despite only occupying 3% of Earth's land area, optimizing resource efficiency at city-scale makes abundant sense, particularly in the context of rapidly urbanizing populations ([~7 billion by 2050](#)). Even relatively small efficiency gains would be significant and worthwhile.
 - **Existing promising examples:** A variety of pilot projects are already beginning to demonstrate environmental gains from smart cities. For example, Smart city projects have been initiated since the [mid-2000s in South Korea](#). The [100 climate neutral and smart cities initiative in the EU](#) appears to be making [promising progress](#).
 - **Insensitive to Jevons effects?** - It's possible given the scale of application that efficiency gains for city-wide infrastructure would not necessarily run afoul of Jevons paradox. Since efficiency savings are distributed across the city and not an individual agent, there aren't any obvious causal pathways for rebound effects, at least at first order.
 - **Compounding sublinear scaling:** [Cities tend to have various scaling relations](#) between size/volume input & output variables; these tend to be (theoretically) sublinear for resource use & emissions, and superlinear for GDP and innovation output. Though there are some mixed findings for emissions broadly (some previous work finding only [linear scaling for US cities and emissions](#), though [prevalent sublinear scaling for Chinese cities](#)), in the cases where emissions (and other waste like MSW and air pollution) *is* sublinear, increased efficiency gains from smart cities would therefore compound with growing city size.
 - **Major benefits across other dimensions:** In best case implementation, smart cities could optimize across a number of environmental factors simultaneously such as water quality, waste management and air quality. One major co-benefit is the use of smart city data and IoT to better enable [urban resilience to climate extremes](#).
- **Issues:**
 - **Negative overall environmental impacts:** Critically, the idea of eco-friendly smart cities does not pan out in reality when taking into account full material and energy requirements for urban metabolism according to life cycle assessment.
 - [Ipsen et al. \(2019\)](#) found that *"implementation of Smart City Solutions generally has a negative influence on the environmental sustainability performance of an urban system. The limited positive influence from the Smart City Solutions is due to burden shifting from the direct impacts of the urban system to embedded impacts which are out of sight for most policy makers"*
 - **Accelerated impacts from economic growth:** Previous work critical of smart cities has pointed out that "smart" urban development [tends to reinforce](#)

- [neoliberal economic growth and consumerist culture](#) due to its focus on services for more affluent populations⁴³.
- **Systemic risks:** At the same time, as is usual with complexification, [smart multi-vector energy system integration](#) also increases risks of cascading failures. The failure of one energy network if connected to a broader smart city management system could trip the rest. Similarly, interconnected energy and infrastructure networks also open up greater vulnerabilities for cyber-attacks, with much more severe consequences as more public infrastructure may be dependent on it (e.g. municipal water supply, traffic systems, public transport, etc).
 - **Boosting dystopia attractor:** One of the obvious concerns with smart cities are [privacy issues](#) and [surveillance capitalism](#). The [development of smart cities in China has been central to broader authoritarian capabilities](#) such as social credit systems and draconic applications of facial recognition tech. The integration of IoT in city infrastructure goes hand in hand with authoritarian surveillance and control, under the guise of public safety and security. Smart cities pose additional risks in [fostering and consolidating digital authoritarianism in developing African countries](#). It's unclear whether and how this ultimately plays out for smart cities in Western democratic countries.

Industry

Industrial production, logistics and construction are some of the main hard-to-abate sectors of the economy. ML-approaches have been proposed for a variety of industrial GHG reduction areas by optimizing supply chains, reducing overproduction of goods⁴⁴, etc...

Reducing Food Waste

- **Key idea:** One high-leverage industry ML idea is to reduce food waste by optimizing delivery routes ([freight routing & consolidation](#)), improving demand forecasting, better refrigeration system management, and food quality sensors to avoid spoilage.
- **Putatively high impact:** A [2011 FAO estimate](#) suggests about a third of all food produced (~1.3 billion tons) for human consumption is wasted, with a footprint of about 4.4 Gt CO₂-eq (8% of global emissions) and 250 km³ of water. [WWF \(2021\)](#) estimates this is closer to 40% of all food. Of this, FAO estimates ~40% of food waste in

⁴³ I cannot find any quantitative estimates of this sociopsychological effect – it's also unclear how one would measure this robustly. Nevertheless, the concept that a more tech-infused urban environment would lead to greater implicit acceptance and pervasiveness of consumerism appears conceptually sensible.

⁴⁴ Supply chain optimization is an obvious and plausibly useful area for ML applications, but I couldn't find any solid quantitative estimates of high level emission/energy savings, or otherwise non-dry technical details, so did not include it here. Note also this type of efficiency optimization would still be vulnerable to rebound effects, though hypothetically reducing overproduction of goods by predicting demand could potentially limit these.

industrialized nations occurs in transit⁴⁵ between harvest and processing/retail, and a further 40% at the end of the supply chain. Depending on how much ML can reduce food waste across supply chains, this could be a significant lever for reducing unnecessary waste, emissions, water use, fertilizer etc.

- One quantitative estimate of ML-based [demand forecasting for catering services](#) has been previously modeled to reduce food waste between 14 - 52%. Taking the upper end and generously assuming this generalizes across food demand waste (~40%), this would avoid about 0.26 Gt CO₂-eq per year, about 1% of global annual emissions. If this applied to transport losses to (also ~40%), could optimistically add up to 2% of global emissions.
- **Caveat - High data uncertainty:** Research and analysis of food loss and waste [has extremely high uncertainty](#) and major dependencies on secondary data, so the noted above estimates should be treated with caution.

Better Construction Materials

- **Key idea:** Cement and steel production [account for ~13.5% of global GHG emissions](#). ML could presumably reduce emissions in these sectors by optimizing and redesigning industrial processes, using generative design to develop structural products which require less raw material, optimizing 3D printing and additive manufacturing, designing [new low emission concrete formulas](#), and speculatively, [discovering new “climate friendly” materials](#).
- **Sensible application for concrete formulas:** The [2022 Meta project](#) found a replacement concrete formula using fly ash and slag with 40% lower carbon impact that could replace up to 70% of cement use. Assuming their findings are true and applicable in real world construction, given concrete alone accounts for 5 - 8% of global carbon emissions, on the higher end this alone could eliminate ~2% of global emissions (~0.75 GtCO₂).
- **Possible but marginal impacts for steel:** Current applications of AI/ML for steel production appear limited to optimizing processes (e.g. [minimizing raw materials required for alloy production](#)). Reportedly, Fero Labs has helped steel manufacturers reduce mined ingredients by 34% and avoided 450,000 tons of CO₂ per year, which if scaled to all US steel production would avoid 11.9 million tons of CO₂ per year, roughly equivalent to New York’s annual emissions. This isn’t insignificant, but marginal (order of magnitude) compared to avoided emissions in the concrete example.
- **Possible Jevons Rebound:** Depending on the replacement, it’s possible “climate-friendly” concrete (and optimized steel production), by virtue of being cheaper and/or less energy intensive, could fall prey to Jevon’s rebound effects.
 - Though the Meta AI project for concrete didn’t describe energy/cost gains, [similar results from other groups](#) have found ~40% lower energy consumption and ~25% lower cost.

⁴⁵ Though [more recent estimates](#) suggest a much wider range depending on food type from <1% - 15%. [This looks different again](#) for the US according to ReFED’s analysis.

- Given construction materials have seen major price increases since COVID ([22% for steel and 15% for concrete in 2023](#)), it would not be surprising to see a fairly sizable rebound effect from green concrete alternatives and cheaper steel - especially when considering this in the context of [current global housing shortages](#).

Farms and Forests

Remote Sensing Emissions

- **Key idea:** ML + hyperspectral remote sensing could improve temporal and spatial data gaps for more precise emissions inventories from agriculture and forestry
- **Logical use of ML:** This is already being actively developed and explored in the [remote sensing field for a variety of applications](#), including emissions/[carbon storage estimates](#). Although there are some challenges, the field is rapidly advancing and broadly makes sense to make the most use of hyperspectral remote sensing.
 - **Few leverage points:** The issue is there are relatively few leverage points for making use of this information for actually reducing emissions. Though emissions accounting is broadly important for research and policy (e.g. national emission inventories), the only leverage points are probably for carbon dioxide removal accreditation, such as [verifying soil carbon changes](#). Though accurate remote sensing techniques for soil carbon sequestration could help reduce verification costs, there are [numerous other challenges preventing such policies being implemented](#).

Estimating forest carbon offsets

- **Key idea:** The need to model and price carbon stored in forests requires accurate assessment of carbon stocks in above-ground biomass. This can be estimated with LiDAR and UAVs, though isn't scalable to all forests. [ML can extrapolate LiDAR measurements via satellite imagery](#)
- **Forest carbon offsets are deeply flawed:** Forest carbon offsets are a very popular market-driven approach falling under the REDD+ framework from the Paris Agreement and one of the main sources of carbon offsets for private companies and governments, but [are deeply flawed](#). About [90% of rainforest carbon offsets are worthless](#). Forests themselves aren't always the most stable carbon stores (e.g. fires), direct and market leakage effects often shift deforestation and GHG emissions from offset projects to other locations, and [counterfactual baseline scenarios are often gamed](#) by developers ([and entire governments](#)) to artificially inflate their carbon credits.
 - In the worst case, ML-enhanced forest carbon measurements are going to signal boost and perpetuate offset schemes and add further artificial

vener to their integrity. Overall, [REDD+ is currently exacerbating climate problems](#), and ML would only further accelerate these

Precision Agriculture⁴⁶

- **Key idea:** ML can be applied to precision agriculture systems in numerous areas; robotic platform algorithms, optimized irrigation systems, disease detection, weed detection, soil sensing (see above), demand forecast for guiding seed planting. Broadly, all categories of precision agriculture (guidance, recording and reacting tech) [assist in reducing GHG emissions](#). Broadly increasing farm output efficiency could result in agricultural emission reductions.
- **Jevons Rebound:** As [previously noted in the AI - climate & Jevons paradox brief](#), there is very strong evidence for rebound and backfire effects for agriculture - advancements in precision & autonomous agriculture boosted by AI & ML are liable to accelerate this, along with associated climate & broader environmental externalities.
- **Limited demonstration:** Precision ag tech has seen quite slow adoption rates, particularly due to high upfront costs and maintenance requirements. As such, [there is \(reportedly\) still a scarcity of data](#) that comprehensively demonstrates precision ag tech efficacy. It is plausible much of precision agriculture is still [caught in a hype cycle](#)
- **Emission reductions require changing paradigms not tech:** The review authors themselves note that actually significantly reducing agricultural emissions may require completely different agricultural paradigms such as regenerative agriculture, agroecology, silvopasture and tree intercropping. It's unclear to what degree the slew of tech under the banner of precision agriculture are compatible with such approaches.
 - Some note that precision ag tech can be compatible or promote regenerative practices; e.g. ultra-light tractors to avoid soil compaction, fine-accuracy robotics for multi-cropping, etc. A number of precision tech companies are pursuing this line for marketing purposes already (e.g. [Cognizant](#), [Syngenta](#)). See also [O'Donohue et al. \(2024\) Digital Regenerative Agriculture](#)
 - At the same time, others argue that so far, precision agriculture has been more associated with traditional farming practices like monocropping. Broadly [the debate on agriculture has already been polarized](#) between conventional and organic/low-tech approaches. It's plausible that many of the most effective regenerative ag and permaculture practices aren't easily compatible with robotic automation or reducible to precise metrics (e.g. from remote sensing) for optimization
- **Enabling conventional industrial practices to further outcompete alternative farming approaches:** market dynamics make it far more likely that AI and precision agriculture tech will only be accessible to large industrial farms due to high capital costs,

⁴⁶ This topic is extremely complex and nuanced, and the limited points made here barely scratch the surface. Though only briefly noted, the possibility for digital (and AI) enabled regenerative agriculture and permaculture should not be discounted.

At the same time, market dynamic risks of industrial outcompetition of smallholders, Jevons rebounds plus material and energy intensity of precision ag tech itself all needs to be weighed against putative benefits for crop yields (including food insecure regions) and reducing biodiversity impacts.

rather than smallholder farms. Given that smallholders [produce one-third of the world's food](#), if AI precision agriculture enables industrial farms to further outcompete smallholders (and assuming industrial-scale operations are significantly less likely to take up regenerative practices), this could result in alternative farming practices becoming less likely in the future.

- Conversely, there are falling costs of some precision ag-relevant tech, including UAVs and IoT sensors, which could help precision agriculture (and possibly precision ag that is compatible with regenerative practices) become [more accessible to smallholders](#)

Carbon Dioxide Removal

Direct Air Capture CO₂ Sequestration

- **Key ideas:** Captured CO₂ (e.g. from direct air capture - DAC) has to be sequestered and stored, usually by direct injection into geological formations like aquifers or old oil and gas reservoirs. ML could help in a few ways; identifying potential storage locations via subsurface imaging from seismographs, maintaining active sequestration sites from noisy sensor measurements, and [monitoring for CO₂ leaks](#).
- **Extraordinarily deep flaws with direct air capture:** Direct air capture and sequestration tech face extraordinary scaling challenges. The [mass of carbon in the atmosphere](#) has gone from 590 GtC to 886 GtC in 2022. Removing carbon at the [gigaton-scale](#) (up to 38 Gt/yr) would require more infrastructure (pipelines, wells, energy sources, etc) [than the entire current fossil fuel industry combined, and is arguably absurd](#).
 - Global levels of CDR via various methods (DACs, enhanced weathering, afforestation, etc) have been estimated to be required at 10 Gt/yr by 2050, and 20 Gt/yr by 2100 for 1.5C pathways. In 2022 total DAC sequestration was about ~11 Mt/yr. Even with market-driven exponential growth⁴⁷, DAC is unlikely to reach the required levels. [Globally coordinated “emergency DAC” would require an investment of up to 2% of global GDP annually](#) to only remove 2 GtCO₂/yr, 10% of what would be necessary. DAC and geological sequestration is also [the most expensive form of carbon dioxide removal](#), with an absolute floor of ~\$100/ton, but more often \$500/ton.
 - [Further problems include:](#) business case for it (currently) is enhanced oil recovery and/or subsidies, the high energy requirements either detract from renewable build out or just emit more CO₂ if it's fossil fuel-powered, and solvents used for DAC are extremely toxic.

⁴⁷ Technically, there is some exponent which does meet this target, but to put this in context even in the globally coordinated scenario described by [Hanna et al. \(2021\)](#) for emergency DAC deployment, they consider a central estimate of 20% - up to 50% *per year*. This would be truly extraordinary for a technology with relatively limited market applications beyond synthetic fuels, plastics and carbonating beverages.

Solar Geoengineering

Some commentators have drawn a link between the prospect of [solar geoengineering](#) (including techniques such as stratospheric aerosol injection or marine cloud brightening) and AI which could assist research, development and deployment of such technologies

- **Key ideas:**
 - **Aerosol design space:** Stratospheric aerosols for solar geoengineering are generally modeled to be sulfates (e.g. SO₂, H₂SO₄), however a [number of other possible aerosols](#) could be used such as titanium dioxide, calcite and synthetic diamond. ML could be used to explore the aerosol design space for specific [chemical, material and optical properties](#) to avoid experimentation and brute-force simulation.
 - **Climate simulation:** ML could be used to quantify and constrain uncertainty in climate modeling of solar geoengineering, and reduce computational load. [AI-based climate emulators](#) could be used to significantly speed up scenario exploration for solar geoengineering design.
 - The design space for marine cloud brightening is extremely large, as ships (or drones) which spray sea salt into clouds to make them brighter could conceivably be done anywhere in the ocean; with varying effects on the climate due to teleconnections. ML has [begun to be used](#) to map out what the regional impacts and teleconnections are from MCB in order to constrain the [physical climate response](#) in models.
 - Practical deployment of stratospheric aerosols or marine cloud brightening would require some [feedback-control algorithm](#) to determine how much aerosols (or sea spray) to inject, along with [other parameters such as latitude and seasonality](#). These all necessarily run into trade-offs. AI could conceivably be used to “fine-tune” geoengineering interventions to optimize for certain outcomes, such as multiple objectives (temperature and precipitation targets, reduced agricultural impacts, etc). [Early work in the field](#) has already begun to explore this area.
- **Solar geoengineering of any form is not a cure-all “solution” to climate change**
 - Whilst current modeling does indicate that techniques such as stratospheric aerosol injection or marine cloud brightening could restore global average temperatures, these have [significant side effects and regional climate trade-offs](#).
 - Aerosols in the stratosphere are not “anti-CO₂”, and a geoengineered climate will not have the same [precipitation patterns or atmospheric circulations](#) as one which is not geoengineered
 - While AI control-based algorithms and ML-facilitated climate modeling may provide an additional veneer of confidence and control over complexity, no degree of sophistication will cover all unknowns in the climate system.
 - Moreover, larger social system concerns like [geopolitical risks](#) (securitization, unilateral deployment) and deterrence of efforts on mitigation (“[moral hazard](#)”) cannot be addressed by AI

Climate Modeling

Clouds and aerosols

- **Key idea:** Cloud physics and radiative effects are some of the hardest and most uncertain areas in climate modeling. Sub-grid scale (<4km) physical modeling is too computationally expensive, so usually global circulation models (GCMs) deal with this by convective parameterization which simplifies and aggregates physical processes, but these have systemic biases and result in large error propagation. [Most of the spread in equilibrium climate sensitivity \(ECS\) estimates](#) can be traced back to cloud representations. Deep neural networks can be trained to [emulate high resolution cloud simulations](#)
- **Promising development with policy relevance for ECS:** This is a genuinely high leverage area for ML to advance an extremely important area of climate physics and modeling. ECS is one of the single most important unknown variables for climate projections - it constrains what the most likely global average temperature rise is for a given concentration of CO₂.
- **Non-stationarity error and ML limitations:** Climate change is inherently non-stationary - the underlying statistics of mean state and variability are changing over time. This means [ML-based models trained on data in current climate may not perform well in warmer climates](#), as climate dynamics stray further out of the scope of training sets.
 - There are some emerging methods to deal with non-stationarity issues like [“climate-invariant” machine learning](#) by incorporating physical climate processes directly into ML algorithms.

Ice sheets and sea level rise

- **Key idea:** Ice sheet dynamics and sea level rise projections are another major uncertain area in climate modeling, [particularly for Antarctic ice sheets where models are frequently inconsistent with observations](#). Theoretically, ML could be used to leverage new data driven models using satellite and video data to predict changes in sea ice extent, ocean heat mixing, etc.
- **Makes sense, though major data and non-stationarity limitations:** This sort of work is already being pursued using deep neural networks (e.g. [MELTNET](#)), however data limitations means these ML approaches have only really been trained on synthetic data from existing ocean and cryosphere models, or otherwise very limited training sets. Implausible behavior outside of training conditions is a particular concern for ice sheet modeling since [non-stationarity, interconnectivity via atmosphere-ocean coupling and tipping point behavior](#) substantially increase long-term risks and may not be tractable via ML approaches.
- **Limited actual leverage without better decision frameworks:** The degree to which better prediction of ice sheet dynamics can guide climate policy and coastal planning for sea level rise is [primarily constrained by current decision procedures](#). Sea level rise projections are obviously important for coastal infrastructure, adaptation planning,

insurance modeling, etc., most approaches use static, probabilistic ranges which fail to grapple with deep uncertainty (e.g. emission pathways).

- Prediction-first cost-benefit analysis type approaches simply fail in a context where, for example, extreme sea level distribution uncertainty spans three orders of magnitude (e.g. Extreme sea levels for New York spanning from ~2m - >4.5m and annual exceedance estimates ranging from 0.01 - ~10 events per year). While [some places like the Netherlands have pioneered adaptation planning](#), which takes into account decision-making under deep uncertainty, better ML-driven sea level rise projections aren't going to be of much greater help until these sorts of decision frameworks become more prevalent.

Broader Societal Impacts

Ecosystem and biodiversity monitoring

- **Key ideas:** Ecosystem monitoring has traditionally been done and limited by on-the-ground observations. Remote sensing, environmental sensor networks and marine robots could all be useful and accelerated by ML via data collection, transfer learning and missing data imputation. Similarly, species population estimates could be done far more efficiently with computer vision approaches and leverage citizen science for large datasets (e.g. birdwatching)
- **Some current limitations, but useful for addressing major data gaps:** The lack of robust, accurate and comprehensive data for ecosystem monitoring is one of the major limitations for biodiversity conservation. Though there are challenges in [effectively using big data from citizen science](#) for biodiversity tracking, and [many ML studies of ecosystem services have methodological issues](#), there are still [significant opportunities across the board](#), particularly in areas such as geospatial data.
- **Limited climate and adaptation relevance?** Whilst a lot of the above examples are directly useful for biodiversity conservation planning and monitoring, it's not entirely clear how high leverage this would be for climate efforts as such. Even for areas we are currently monitoring and ostensibly managing such as marine protected areas, [one review found only a minute fraction \(~4%\) of management actions](#) actually engaging in any adaptation measures to enhance biodiversity conservation against climate change. Half of those were just monitoring measures. As such, simply enhancing monitoring via ML does not guarantee we are going to do any better for actually protecting species and ecosystems without a change to underlying conservation approaches.

Food Security

- **Key Ideas:** ML big data approaches could be useful for monitoring food insecurity risk in real time to forecast near-term shortages and risk areas in the long term. Data can incorporate disparate sources from consumer purchases, manual crop surveys, aerial imagery and meteorological data. Crop diseases can be identified from plant photos with computer vision. Large scale ML simulations could emulate long-run risks in crop yield via biological and ecological models.
- **Extremely useful, though limited in extreme cases:** There's [significant work being done](#) on applying machine learning for food security currently, with attention from [World Bank](#) and [World Food Program USA](#). The [use of non-traditional data](#) (mobile phone data, satellite imagery, money transactions, tweets etc.) appear to be effective in ML approaches to explain significant amounts (~73%) of variation in insufficient food consumption in crisis situations.
 - At the same time, increasing climate variability and hazards along with other major drivers of food insecurity such as socio-political instability and conflict means we're increasingly entering an era of [simultaneous multi-breadbasket failure](#). Large, global scale shocks to food supply may not be buffered by existing trade connections. It's unclear to what degree better ML-driven predictions will aid national resilience strategies in this regard without fundamental changes to food supplies.

Appendix

Energy forecast of ICT and Data Centers

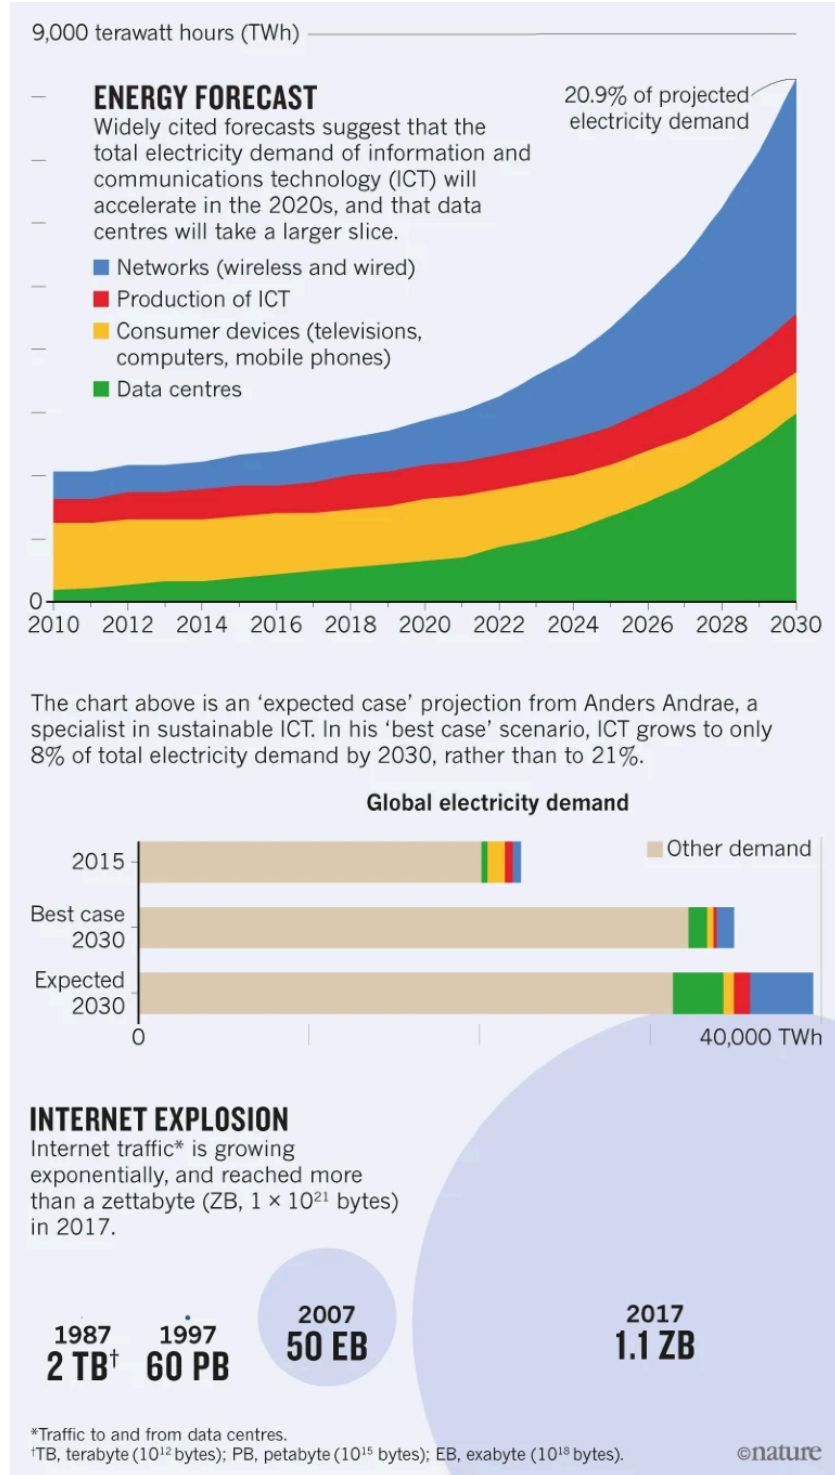



Figure from [Jones \(2018\)](#)

Energy footprint of IoT Devices



[Das & Mao \(2020\)](#) Figure 10. Global total primary energy footprint of electronics in IoT devices.

Mineral footprint of AI and ML data servers

Using [rock-to-metal ratio](#) for mine waste of the above metals, we can estimate how much rock waste has to be moved per ton of board.

Critical materials in data center boards from [Tempro project](#) ([see main report - in German](#))

855 g/ton of board for gold (RAM & CPUs), 1802 g/ton silver, 36 g/ton of palladium (mother boards)

Drive boards also contain rare earths (dysprosium, neodymium, yttrium)

Rock to metal ratio in supplementary material of [Nassar et al. study](#); [es1c07875_si_001.pdf](#)

- Gold rock-to-metal ratio: ~3,000,000:1 (eq to 3 metric tons (3000kg) rock per gram of gold)
- Silver rock-to-metal ratio: ~22,000:1 (eq to 22 kg rock per gram of silver)
- Palladium rock-to-metal ratio: ~688,000:1 (eq to 668 kg rock per gram of palladium)

Translating to kg of rock per kg board (/1000 ratios above, multiply by rock-to-metal ratio)

- 2565 kg rock per kg of board (gold)
- 39.64 kg rock per kg of board (silver)
- 24.1 kg rock per kg of board (palladium)
- Total: 2628.74 kg (2.6 tons) rock per kg of board

How many kg of computer board per average data center?

—

2nd attempt (05/31):

Decompose to:

Kg per board * board per server * servers in data center

Number of servers in an average data center?

Can do basic average, constraint per area or constraint by power capacity

- Average full-scale data center runs [about 100,000 servers](#) (for comparison, Amazon Web Services uses at least 454,400 servers in seven data centers globally)
- [Area estimate](#): 1 million ft² data center with 65% of site used for servers using 52U racks can have 2,688,638 1U servers, 1,344,318 2U servers, or 896,212 3U servers
- Power-based estimate: Typical server rack uses 7 KW, higher powered racks use 25 KW. Largest data center in the world, [Switch Tahoe is capped at 55 KW per cabinet](#).
 - Average U height is 52 units per rack

- Taking the Switch Takoe capacity of 850 MW, this works out to the following range:
 - $850,000 \text{ kW} / 7 \text{ kW} = 121,428 \text{ racks} * 52(\text{U}) = 6,314,256 \text{ 1U servers}$
 - $850,000 \text{ kW} / 25 \text{ kW} = 34,000 \text{ racks} * 52(\text{U}) = 1,768,000 \text{ 1U servers}$
 - $850,000 \text{ kW} / 55\text{kW} = 15,454 \text{ racks} * 52(\text{U}) = 803,608 \text{ 1U servers}$

- Boards per server?
 - Note: cannot figure this out. Underdetermined information

See attached translated PDF of TEMPRO project report

■ [Tempro_Endbericht_final_2020_05_14-gooqtranslate.pdf](#) ([original](#))

Raw material	RAM/CPU's	Rock-to-metal ratio (X:1)	Per 10,000 GPUs (kg)	Total rock displaced (ton)
Precious metals				
Au (g/t)	855	3000000	30.951	9.29E+04
Ag (g/t)	1240	22000	44.888	9.88E+02
Pd (g/t)	10	688000	0.362	2.49E+02
Basic metals				
Al (kg/t)	16	7	253.4	1.77E+00
Cu (kg/t)	490	513	18570.6	9.53E+03
Fe (kg/t)	133	9	325.8	2.93E+00
Rare Earths				
Dy (g/t)	22	1.70E+04	796.4	1.35E+04
Nd (g/t)	72	1.90E+03	2606.4	4.95E+03
Y (g/t)	21	3.50E+02	760.2	2.66E+02
Critical raw materials				
In (g/t)	4258	152589.7271	5523.748119	8.43E+05
Ti (g/t)	1652.9	99	3.5838	3.55E-01
Ga (g/t)	256.1	513	18.5706	9.53E+00
			Total	965,255.15

Synthesis and simplification of Table 9 in TEMPRO report plus rock-to-metal ratios from Nassar et al. Rare earth rock-to-metal ratios from [Nassar et al. \(2023\)](#)

AI and ML Servers?

- Deep learning models generally [have much higher processing requirements for data centers](#), demanding advanced hardware including GPUs and [TPUs](#) (tensor processing units) - the latter being specialized circuits designed to accelerate AI/ML workloads
 - E.g. [more dense GPU clustering](#): *As these racks get denser — have more GPUs stuffed into them and more cabling run to them — they become incredibly heavy. As the amount of power they require increases, more energy needs to be brought into the data center. And as they generate incredible amounts of heat and are grouped closer together, liquid cooling alternatives to traditional forced-air cooling need to be provided.*
 - [Higher rack densities](#): existing data centers average 10KW/rack - AI catered server racks around 5 - 10x this (50kW - 100kW)
 - [Other notes on Nvidia GPU materials and lifecycle](#)

OpenAI's data servers (nb. Not full specs from Microsoft's custom AI supercomputer infrastructure in Azure cloud for OpenAI)

- >285,000 CPU cores
- 10,000 GPUs
 - Assume all are NVIDIA [Tesla A100](#): 3.62kg per GPU
 - $3.62\text{kg} * 10,000 \text{ GPUs} = 36,200 \text{ kg}$
- 400 GB/s network connectivity for each GPU server

Based on metals found in other standard data servers (see table above):

About 965 thousand metric tons of mining rock waste just for the GPUs in OpenAI's custom built supercomputer

- Mostly due to indium, gold, dysprosium and copper
- Plausibly another order of magnitude greater if including all other components

Would be good to do a specific material/energy case study on the [“Stargate” Microsoft - AI project](#). \$100 billion supercomputer data center to support OpenAI's models, set to launch in 2028. Dwarfs current largest supercomputer centers.

A more comprehensive analysis should take into account all components in data centers. See e.g. the lifecycle analysis done by [Penaherrera et al. \(2019\)](#) (particularly Fig 1).

Jevon's Paradox, Rebound and Backfire Effect

In studies of Jevons paradox, the effect is usually decomposed into direct, indirect and economy-wide rebound effects.

Direct rebound is when improvement in efficiency of some process/device (e.g more fuel efficient car, more energy efficient light bulbs) leads to greater use of that same process/device; (e.g. more miles driven, more light bulbs bought and used).

Indirect rebound is when the savings from energy efficiency are used to increase consumption in other areas (e.g. savings on a power bill used for more air travel, food, discretionary spending, etc)

Economy-wide rebound is the total effect, usually modeled in general equilibrium models by the change in supply-demand balance

Finally, **backfire** is when the result of energy efficiency investment is the consumption of energy/emissions increases above what it was prior to the change (>100% total effect). Though Jevons postulated the rebound effect is always greater than 100%, [there isn't enough evidence to assert this in all cases](#).

From York et al. (2022) [The rebound effect and the challenge of moving beyond fossil fuels: A review of empirical and theoretical research](#):

*Subsequent to Jevons's work, economists have focused on the rebound effect, a broader class of phenomena than the paradox. **A rebound occurs when efficiency increases by X%, but resource consumption decreases by less than this percentage.** For example, if the efficiency of lightbulbs used in a house improved by 10%, but energy consumption associated with lighting decreased by only 7% (perhaps due to leaving lights on for a longer period of time), then a rebound has occurred since some of the anticipated energy conservation was not realized. This can be defined mathematically as the ratio of the "lost benefit" to the expected benefit. In the hypothetical scenario presented here, there is a 30% $[(10-7)/10 = 0.30, \text{ or } 30\%]$ rebound, based on the difference between the "expected" decline of 10% and the actual decline of only 7%. **When the rebound is greater than 100%, indicating that there was growth in resource consumption corresponding to rising efficiency, then we observe the scenario characterized by the Jevons paradox, sometimes described as a "backfire."** For example, if the 10% improvement in the efficiency of lightbulbs led to a 1% increase in energy consumption, there is a "lost benefit" of 11% and, therefore, a rebound of 110% $[(10-(-1))/10 = 11/10 = 1.10, \text{ or } 110\%]$, since instead of the "expected" decrease of 10% there is an increase of 1% (i.e., a - 1% decrease).*