

# Exploring the Viability of the Updated World3 Model for Examining the Impact of Computing on Planetary Boundaries

Nara Guliyeva  
nargiz.guliyeva@mail.utoronto.ca  
University of Toronto  
Toronto, Ontario, Canada

Eshta Bhardwaj  
eshta.bhardwaj@mail.utoronto.ca  
University of Toronto  
Toronto, Ontario, Canada

Christoph Becker  
christoph.becker@utoronto.ca  
University of Toronto  
Toronto, Ontario, Canada

## Abstract

The influential *Limits to Growth* report introduced a system dynamics-based model to demonstrate global dynamics of the world's population, industry, natural resources, agriculture, and pollution between 1900-2100. In current times, the rapidly expanding trajectory of data center development, much of it linked to AI, uses increasing amounts of natural resources. The extraordinary amount of resources claimed warrants the question of how computing trajectories contribute to exceeding planetary boundaries. Based on the general robustness of the World3-03 model and its influence in serving as a foundation for current climate frameworks, we explore whether the model is a viable method to quantitatively simulate the impact of data centers on limits to growth. Our paper explores whether the World3-03 model is a feasible method for reflecting on these dynamics by adding new variables to the model in order to simulate a new AI-augmented scenario. We find that through our addition of AI-related variables (such as increasing data center development) impacting pollution in the World3-03 model, we can observe the expected changes to dynamics, demonstrating the viability of the World3-03 model for examining AI's impact on planetary boundaries. We detail future research opportunities for using the World3-03 model to explore the relationships between increasing resource-intensive computing and the resulting impacts to the environment in a quantitative way given its feasibility.

## Keywords

AI impacts, limits to growth, planetary boundaries, system dynamics, World3 model, World3-03 model

## 1 Introduction

The 1972 *Limits to Growth* (LtG) report [46] used a global model, called World3, to show how economic growth would be limited by both increasing population and usage of natural resources and represented this through 12 potential future scenarios. Their World3 model was based on system dynamics, an approach to systems thinking which uses “measurable variables [to] represent structural properties of the world... and organizes these variables with quantitative causal relationships that allow it to predict overall behavior as the emergent property” [8, p. 20]. The model thus provided a method for exploring the dynamics between the “expanding global population and materials economy” and “the [Earth’s] limited carrying capacity” [14]. Given that the model and report were first introduced in 1972 [46], with updates in 1992 [44] and 2004 [43], a consideration of then-emerging concepts of ‘AI’ as a technological

development with global ecological relevance was expectedly not considered.

In current times, the high computational training requirements of so-called ‘AI’ drive substantial energy consumption [12, 40, 66, 68], manufacturing demands [41, 47], and investment needs (e.g., [26, 50]), which could significantly impact resource depletion [2, 3, 41] and pollution levels globally [24]. The trajectory and broader implications of the recent surge in AI are presently widely researched with several studies showcasing that the global proliferation of AI models intensifies environmental pressures [2, 23, 38, 41, 74], thereby necessitating further exploration of the extent of this impact.

Accordingly, this paper explores whether the effect of vastly scaling AI on patterns of world development, particularly environmental outcomes, can be quantified using the World3-03 model. Feedback loops resulting from AI advancements that rely heavily on energy-intensive data centers [13, 15, 23, 54], land use [47], rare and finite raw materials for hardware [41], vast datasets for training [67], and substantial financial investments (e.g., [32, 35, 50, 55, 66, 78]) could accelerate (or mitigate) resource depletion and fundamentally reshape sustainability dynamics.

We examine the following research question:

- (1) Is the *Limits to Growth* World3-03 model a viable method for exploring the impact of computing on planetary boundaries today?

We answer this question by specifically exploring whether and how AI development trajectories can be incorporated into the World3-03 model, which is key to anticipating its long-term consequences on resource availability and stability. We focus here specifically on direct lifecycle impacts resulting from the materials and operations of data centers. We do this by adding new variables to the World3-03 model in order to simulate impacts on pollution as a result of the accelerating development of data centers. Given the rise in data center development causing increased pollution presently and estimated to skyrocket in the future [24], our hypothesis is that by adding variables and setting parameters that increase pollution, we are able to simulate consequences to the relevant stocks in the World3-03 model. Thus we quantify assumptions regarding data center-caused pollution within the World3-03 model to simulate the expected consequences. Our findings confirm our hypothesis and illustrate a method for adjusting the World3-03 model to reflect ecological impact of computing.

We focus on variables that represent data center development to the World3-03 model for a few reasons. Our goal is to explore whether the World3-03 model can quantitatively simulate the impact of computing on limits to growth. Presently, computing’s impacts are discussed and debated most widely pertaining to AI.

‘AI’ is most commonly a marketing term rather than a technical term but currently most often refers to machine learning and large language models, often in the form of chatbots and generative models producing audiovisual and textual materials, including code. While references to AI often encompass things as diverse as the AI industry, AI models, systems, and technologies including generative AI, and the algorithms underlying AI systems, in this paper we focus on how the sector manifests as physical infrastructure, including the accelerating development and operations of hyperscale data centers. Our scope is to explore how data center development may impact pollution (a limit to growth recognized as a planetary boundary today) on a global scale. This is therefore one method by which we can explore how AI’s impact can be added to the World3-03 model.

Through reviving the World3-03 model, the paper encourages further examination of the links between computing and its environmental footprint while identifying new research opportunities. By doing so, it provides the groundwork for future, comprehensive system dynamics perspectives on whether and how accelerating AI and computing broadly push humanity closer to, or further from, planetary boundaries.

## 2 Background

In this section, we first provide an overview of the findings of the *Limits to Growth* report and the development of the World3-03 model. Next, we briefly touch on the critiques surrounding the report’s results, noting that some of these critiques were later revealed to rest on shaky grounds. Finally, we discuss the recent updates to the report and model and its continued general validity recognized throughout the recalibrations.

### 2.1 Limits to Growth and the World3-03 Model

Pioneered by the *Limits to Growth* study in 1972 [46] and followed by two updates in 1992 [44] and 2004 [43], the World3 model is an application of a system dynamics-based approach to explore the interplay of population, industry, resources, agriculture, and pollution over long periods under resource constraints. The model was developed based on Forrester’s exploration of the possibility to build a high-level simulation model of natural resource depletion, pollution crisis, crowding, and food shortage [17].

The equations and technical rationale behind the World3 model are presented in a separate volume, *Dynamics of Growth in a Finite World* [45]. The original (1972) study explores 12 scenarios, which have been described in the Appendix A, Table 4. The debate was coined around 4 key scenarios: Business-as-Usual (BAU) – Standard Run, Business-as-Usual 2 (BAU2) – High Resources, Comprehensive Technology (CT) – Tech Fix, and Stabilized World (SW), each illustrating different pathways for the future [46], explicated further in Table 1. The 1972 LtG uses the term “run” to describe various model simulations while the 2004 study uses the term “scenario” to refer to the same type of model-based exploration. This paper uses the term “scenario” for consistency.

These scenarios emphasize a central message: the planet has a finite carrying capacity (i.e., a maximum population and consumption load that the ecosystem can support indefinitely without irreversible degradation) and thus requires a transition from growth to long-term sustainability.

**Table 1: Assumptions of World3-03 key scenarios, adapted from [25].**

Scenario	Description	Cause
Business-as-Usual (BAU) – Standard Run	No assumptions added to historic averages.	Food production, industrial output, and population expand exponentially until the sharp depletion of resources eventually curtails industrial growth. Collapse is eventually caused by natural resource depletion.
Business-as-Usual 2 (BAU2) – High Resources	Double the natural resources of BAU.	Rapid population growth raises death rates and lowers food production, leaving resources severely depleted despite an initial doubling. Collapse is caused by pollution.
Comprehensive Technology (CT) – Tech Fix	BAU2, plus exceptionally high technological development and adoption rates.	The technology delays the onset of limits significantly. Rising costs for technology eventually cause declines such as in industrial output, but no collapse.
Stabilized World (SW)	CT + changes in societal values and priorities.	The result of sustainability related policies would be population stabilization, as human welfare is on a high level. In this scenario, collapse is not expected.

In their 30-year update, Meadows et al. further reinforced that even by their 1992 update, humanity had already entered overshoot, exceeding the Earth’s carrying capacity, with climate change emerging as a major driver of systemic collapse [43]. Unlike the 1972 edition, their 2004 book explored the results as a choice, highlighting “visioning, networking, truth-telling, learning, and loving” as social tools equal in importance to technological efficiency [43]. Furthermore, the book conveys the need for continual model revision, a message that is directly relevant to our effort to incorporate current computing impacts into World3-03 [43].

The 2004 book further described 10 scenarios that are similar to those outlined in the 1972 and 1992 versions [43]. This version of the study presented the revised model, World3-03, which included new variables such as the human ecological footprint and human welfare. Furthermore, the updated study emphasized the growing importance of technology and its effect on the model. It focused less on exact prediction and more on pattern recognition and underscored that technological advancements alone are not sufficient to prevent collapse [43]. While originally published in 1972, the *Limits to Growth* report remains a relevant and vital influence in

this area of study given its usage as a foundation for now influential frameworks such as the planetary boundaries [63, 65, 69] and doughnut economics [61].

## 2.2 Criticism about the Limits to Growth Report

Several critiques of the LtG surfaced soon after its 1972 release, and those were widely echoed by the general public in popular outlets, such as [58]. The LtG was scrutinized by economists and policy analysts eager to test the robustness of its assumptions and outputs. The critique ranged from methodological concerns to economic arguments that feedback and technological innovation were insufficiently represented. However, subsequent empirical and methodological work has largely overturned those criticisms.

A comprehensive critique of the LtG challenges its methodology, data inputs, and the lack of context [9]. Two overarching claims were made: 1) LtG has elected pessimistic parameter values; and 2) the model was itself flawed because small shifts in a few coefficients could swing the results from growth to collapse. Despite some valid points such as that testing models reveal errors in the results, the technical criticism was based on a misunderstanding of feedback-dominated models. For example, the claim that moving the initial year to 1850 would shift model's collapse forward by twenty years ignores the need to recalibrate the entire parameter set to known historical data [1]. It was also argued that the model could not be validated scientifically because its high-level aggregates were abstractions with no direct empirical outputs [9]. However, the empirical results decades later have shown an alignment between several World3 scenarios and observed trends, suggesting that the model's structural insights outweigh its parametric uncertainties [25, 73].

Another influential critique was made by Nordhaus who outlined his critique in three overlapping strands that have since framed much of the cross-disciplinary debate: 1) ad personam/epistemic posture (i.e., long-range simulation is "hybris"), 2) unsubstantiated statements of disbelief (i.e., the values for equations are without any counter evidence); 3) quantifiable/econometric critique (i.e., due to the lack of econometric calibration, results are "numerology") [51]. Nordhaus' claims were rejected by Forrester [18] with the counter arguments that Nordhaus had linearised inherently non-linear relationships and compared incomparable units when appealing to historical data. It was also highlighted that the purpose of LtG was not point forecasting but "structured exploration" of policy-sensitive feedback [18]. After the 1992 update, Nordhaus made another critique that the LtG relies on a rigid system dynamics model that lacks empirical grounding and inadequately incorporates economic feedback and technological innovation [52]. It is worth noting that Nordhaus' arguments rest on economic ideologies that lead to deeply problematic arguments [22], including the suggestion that the 'optimal' temperature increase of the planet occurs at an average of 3.5 degrees Celsius – a level typically described as 'catastrophic' by climate scientists [34, 70].

More recent criticism has stated that claims of systemic collapse are premature and should not be based solely on simulation outcomes and that the World3-03 model is overly abstract and simplified with uncertainties in key variables that are unreliable [6].

Bardi provides an analysis of the key conclusions of the LtG and its updates, arguing that despite criticism, the LtG's findings are valid [1]. Bardi points out that academic criticism of the model stems from a misunderstanding of system dynamics itself and failing to grasp the model's essence. Furthermore, the author states that the public backlash against LtG originates not from methodological flaws but from ideological resistance. Additionally, subsequent updates employing the recent datasets demonstrate the real-world trajectories track remarkably with LtG scenarios, rebutting many early critiques.

## 2.3 Recalibration and Validity

Since the release of the World3 model, several attempts have been made to reassess its validity and compare its trajectories with empirical data. An empirical validation of the World3-03 model compared its results with real-world data from 1970-2000, focusing on three scenarios: BAU, CT, and SW [73]. The analysis demonstrated that observed global trends align closely with the BAU scenario if growth continues unmitigated. However, the effectiveness of technological advancement was questioned in its ability to reverse systemic stress highlighting how solutions such as biofuel, energy efficiency, and others often generate unintended consequences, reinforcing rather than mitigating environmental and economic pressure [73]. These warnings that appear timely given the continued pattern of claims that proven negative consequences of data center development will be offset by unproven promises of positive consequences that are yet to materialize [28].

Another recalibration of the World3-03 model was performed with empirical data from 1995-2012, specifically in the context of Scenario 2 from the original study (i.e., BAU2) [57]. The paper identified key deviations from Scenario 2: 1) industrial-sector productivity has been lower than anticipated; 2) the service sector has significantly outperformed expectations, accelerated by information technology and a growing finance industry; 3) increased food production is attributed to technological innovation with reduced impact on land erosion. These discrepancies suggest that structural changes such as significant shifts towards services and decoupling food outputs from land use are already alleviating stresses highlighted in the original model.

In 2021, a revisit to the LtG was performed using empirical data up to 2020 with projections from the updated World3-03 model [25]. The analysis delved into four scenarios as described in Table 1. Contrary to [73], [25] that argued that the original BAU scenario is closest to the observed trends, it also found that BAU2 and CT emerge as the most closely aligned with empirical data. The CT and BAU2 scenarios differ in how sharply they predict decline: BAU2 presents a severe collapse pattern, while CT exhibits decline under highly optimistic assumptions about the effect of technological advancement. Thus [25] warns that assumptions under CT might be overly optimistic. The paper also highlights that the SW scenario on the sustainable outlook is the least consistent with empirical data.

The most recent recalibration of the World3-03 model was performed by adjusting 35 variables and optimizing the model against eight key empirical datasets using an iterative algorithm. The findings state that the recalibrated trajectories lie between BAU and

BAU2 scenarios [49]. The new trajectories suggest that collapse still occurs due to resource depletion rather than pollution overshoot, as per BAU2, but occurs fifty years later and reaches a higher peak compared to the earlier BAU. The paper [49] thus echoed the message from the LtG on the persistent risk of systemic collapse and aligned with findings from [73].

## 2.4 Computing within Limits

The *Computing within Limits* community established the need to shift focus towards the ecological and planetary boundaries rather than scaling performance 10 years ago, and it continues to be a concern as is seen with the current AI trajectory [7]. Limits to growth as a concept has been used to explore and quantify the “computational limits of deep learning” [71] and more recently AI scaling trends and the resulting consequences [2]. Computing has also been considered as ecocide because of the carbon emissions and energy footprint associated with computation, the combined social and ecological impacts of computing infrastructure that cause resource scarcity, the materiality of computing including severe damages that are possible from improper handling of e-waste, and the tech hype and techno-solutionism that makes computing a facilitator of ecocide [10].

Our paper contributes to the aims of the *Computing within Limits* community to “[explore] ways that new forms of computing [supports] well-being while enabling human civilizations to live within global ecological and material limits” by presenting a method with which to quantify the impacts of AI scaling (by exploring data center development) to limits to growth [48, p. 86]. Thus we revisit the limits to growth concept that grounded the conception of *Computing within Limits*.

## 3 Methods

We first explain our rationale for simulating data center development using the World3-03 model and explain its relevance. We then discuss how we updated the World3-03 model including the reasons for our methods, modelling details, and an illustration of the model parameters and relationships. Finally, we outline a few challenges that we encountered in the process of updating the World3-03 model.

### 3.1 Rationale for AI in the World3-03 Model

It has been found that the ICT sector contributes between 2.1% and 3.9% of global carbon emissions, prompting the need to integrate its growing emissions into broader climate models and policymaking [20]. Similarly, the growing carbon footprint of emerging technologies including AI, blockchain, and IoT is a significant concern as these technologies demand high computational power and infrastructure expansion [20]. AI systems, and specifically GenAI, demand significant computational power leading to high electricity consumption and increased water usage for cooling [31, 36]. Growing climate impacts have been pinpointed along the entire AI supply chain, from hardware manufacturing to data center operations [41].

These factors can significantly affect the global development trajectories thus prompting an assessment of its potential impact. The World3-03 model is a widely recognized system dynamics model

that is ideally suited to capture the reinforcing loops, delays, and potential tipping points that AI growth may trigger. Accordingly, we explore whether explicitly incorporating data center development to the World3-03 model will demonstrate a change in dynamics between causal relationships.

### 3.2 Updates to the World3-03 Model

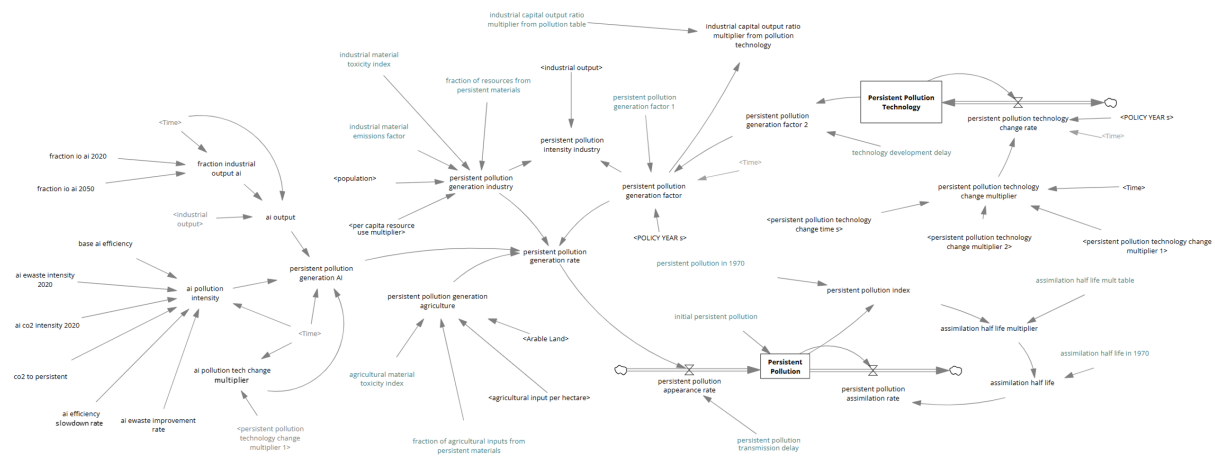
There are several methods that can be used to update the World3-03 model: adding new variables, refining technology parameters, updating model assumptions, or reconstructing AI impacts as a separate sector and connecting it with the existing model.

Prior to explaining our method for updating the World3-03 model, we provide clarifications around the use of key terminology. We define **variables** broadly as dynamic quantities that change over time, and **parameters** as constants or fixed values that define the characteristics of the system and influence how variables interact. **Equation** is defined to be a mathematical expression that defines relationships between different components of the system. **Sector** is defined as a major subsystem that represents a broad functional area, e.g., non-renewable resources, which contains stocks, flows, feedback loops and variables.

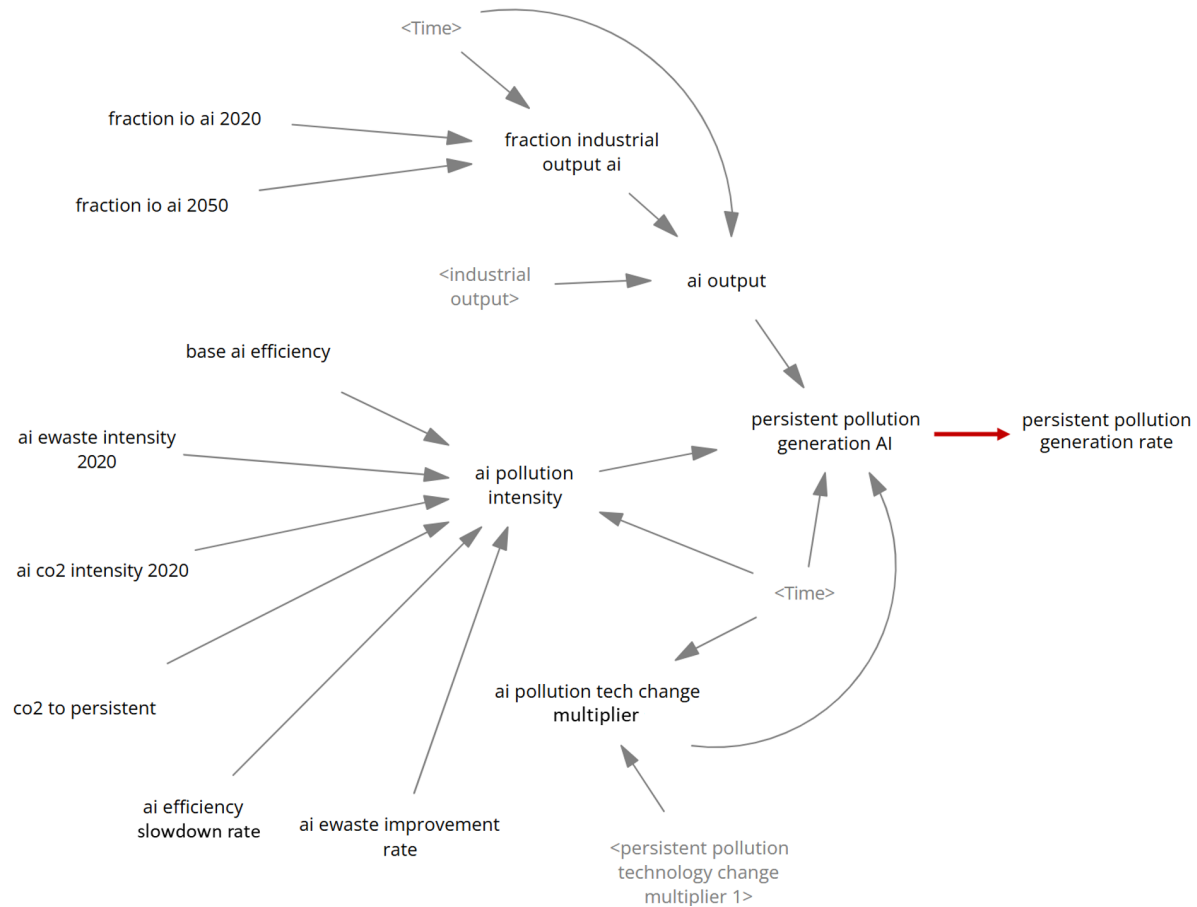
Our chosen method to explore how rapid scaling of AI, through materials and operations of data centers, impacts limits to growth was to update the model by adding variables into the existing persistent pollution sector using the Vensim PLE 10.3.2 software [75]. We specifically make our updates to the 2004 version of the model (i.e., World3-03) and in the paper, we use the term ‘World3-03’ when discussing the 2004 version and ‘World3’ when discussing previous versions. The new variables include both parameters and equations, as well as new connections that were formed between select existing parameters. The entire persistent pollution sector including our additions is pictured in Figure 1a. Figure 1b shows a closer view of the new variables and how they connect to the existing variables in the persistent pollution sector (through a relationship between variables: ‘persistent pollution generation AI’ and ‘persistent pollution generation rate’). We list the additions in Table 5 (description of variables, both parameters and equations), Table 6 (the values of the parameters), and Table 7 (equations) in Appendix A.

We chose to embed AI-relevant variables inside the existing persistent pollution sector of the World3-03 model given the substantive research that currently shows the impact of AI on CO<sub>2</sub> emissions [37, 39]. Furthermore, debates around carbon footprint dominate current regulatory and corporate areas as principal terms of environmental impact [42]. Thus updating the persistent pollution sector was possible given the availability of current data (on how AI impacts emissions). Additionally, this method would introduce the least amount of errors as pollution impacts are already considered within the dynamics simulated by the World3-03 model (as compared to more complex methods discussed next).

In contrast, updating technology-relevant coefficients would be challenging given the difficulty in obtaining how certain values in the model were calculated. Another option would have been to build an AI sector. This is more comprehensive but requires extensive data collection and modeling and thus was considered beyond the scope of this paper.



(a) The persistent pollution sector in the World3-03 model with new variables (left cluster) added to simulate pollution caused by AI scaling.



(b) The added variables to simulate pollution caused by AI scaling are connected with the existing variables in the persistent pollution sector through a relationship between variables: 'persistent pollution generation AI' and 'persistent pollution generation rate'.

Figure 1: A screenshot of the Vensim Software [75] showing the persistent pollution sector along with added AI variables in (a) full view and (b) focused on the relationship connecting new variables to existing variables.

The constants for the parameters were identified based on both recent academic and non-academic sources reporting potential, future AI scaling (see discussion on this in Section 5.1) [4, 11, 19, 29, 53, 59, 60, 72]. Wherever possible we grounded the parameter constants in peer-reviewed literature with others sourced from industry reports; thus the parameters were back solved from the published aggregates. Accordingly, inferred values should be treated as first-pass approximation and not as definitive benchmarks.

### 3.3 Integrating AI Trends into World3-03

The World3-03 model simulates dynamic relationships over a temporal scale and as such computes pollution over time. In order to simulate AI's impacts on pollution, we could not simply create a percentage increase directly but rather had to add dynamic factors that represented AI-caused pollution over time. Adding variables in this manner mirrors the structure of the World3-03 model.

The persistent pollution sector of the BAU scenario of the World3-03 model takes into account two sources of impact (industrial activity and agriculture) through intensity factors, sums the resulting flows, and lets that total accumulate in a single long-lived stock. In the BAU scenario, persistent pollution generation agriculture is calculated by multiplying the total arable land by the agricultural input per hectare, then scaling by the fraction of agricultural inputs from persistent material, and finally adjusting by the agricultural material toxicity index to account for long-term environmental impact. In parallel, persistent pollution generation industry is calculated by multiplying population by per capita resource use multiplier, then scaling by the fraction of resources from persistent materials, and finally applying both the industrial materials emissions factor and industrial toxicity to determine the overall pollution impact. To represent AI impacts without disturbing the integrity of these original feedback loops, we introduced the pathway for AI that follows the logic of the industrial and agricultural activity.

Thus following this logic, persistent pollution generation for AI is calculated as a fraction of industrial output spent on AI, scaled carbon, and e-waste footprint. The empirically grounded share of global industrial output is earmarked for the AI industry. The earmarked economic activity is translated into persistent pollution units by combining 1) an initial operational and embodied carbon coefficient, 2) a parallel coefficient reflecting the lifecycle toxicity of electronic waste, and 3) a conversion constant that maps physical emissions onto the model's dimensionless pollution scale. Both coefficients are allowed to decline at rates calibrated to historical trends in compute efficiency and materials circularity. However, they are bounded below by the same global floor that constrains industrial and agricultural intensities, thereby preserving cross-sectoral comparability. The resulting flow is then added directly to the legacy industrial and agricultural pollution flows before entering the persistent pollution stock. Since all other sectors and parameters remain untouched, any deviation from the canonical trajectories can be attributed unambiguously to the added AI variables. This targeted, transparent augmentation equips World3-03 model to be used to explore whether continued expansion (or policy-driven moderation) of AI development will accelerate, delay, or avert the overshoot dynamics identified in earlier runs, maintaining the model's original coherence.

We executed two scenarios with the World3-03 model over the interval 1900-2100 with a 0.5-year time step. The first scenario was the standard scenario which is the baseline that corresponds to the BAU and the second was the AI-augmented scenario with our added variables. From each scenario, we extracted diagnostic outputs for persistent pollution, i.e., the stock of long-lived pollutants, and human ecological footprint (HEF), i.e., the ratio of humanity's resource throughput to the Earth regenerative capacity. We detail added parameters in Table 5, parameter values in Table 6, and equations in Table 7 in Appendix A.

### 3.4 Challenges

While the initial methods by which the model can be updated appear elementary, there are several nuances involved. We first attempted to update the World3-03 model by adding a new sector on AI development which would begin in 2020. However, considering that building the AI sector in the World3-03 model requires more data inputs and careful consideration of equations, we scaled back to single sector modifications.

Another challenge was in calibrating the constants for the parameters because empirical estimates about data center operations and materials usage, along with other AI impacts, can vary by an order of magnitude depending on the sources, and particularly in industry reports. Thus we selected mid-range baseline values (detailed in Table 6, Appendix A). Additionally, due to the lack of data, some parameter values have been derived from the existing figures, and should be considered as approximations rather than definitive values.

We viewed and modified the model through the Vensim PLE software [75] which generates "warnings" that do not impede being able to run the model but are nevertheless recommended to resolve. We evaluated and resolved several such warnings in the process of updating the model. For example, LOOKUP\_BOUNDS warnings occurred because the model requested a value from the lookup table outside the defined bounds of the table. While running the BAU scenario, the warnings were generated for 12 variables. To improve the validity of the model, LOOKUP\_BOUNDS warnings have been treated by increasing the bounds.

## 4 Findings

We present findings from a comparison between our AI-augmented scenario and the BAU scenario. Based on our added variables to the persistent pollution sector of the World3-03 model, we are able to see updates to the persistent pollution stock. Overall, when compared to the BAU scenario, the AI-augmented scenario results in higher pollution. Comparisons are reported at five benchmark years, i.e., 2020, 2040, 2060, 2080, 2100 and as percentage differences relative to BAU, described further below and listed in full in Table 8 in Appendix A.

We below show how impacts from AI development can be translated into World3-03 stock and flow language. Our objective was not to deliver a point forecast of a future scenario but rather to visualize the kinds of questions that can be asked from the model. We frame our findings below as a proof of concept rather than a definitive prediction.

We first demonstrate that even simulating one minor modification, i.e., that AI affects pollution, the system behaviour was updated (i.e., the persistent pollution stock). To summarize, the persistent pollution stock peaked later and was 4% higher in the AI-augmented scenario as compared to BAU, leading to 45% larger residue in 2100. As a result of this changed behaviour, further changes were triggered in another variable, human ecological footprint (HEF) where HEF remained approximately 11% above the BAU baseline. (While many other variables may have been updated as a result of the persistent pollution stock, we report only the changes to HEF given its direct relevance to pollution, discussed further below). These magnitudes are not decisive predictions, but rather they are signposts that the direction of change matters and the feedbacks remain as relevant today as Meadows et al. emphasized in 2004 [43].

#### 4.1 Variable 1: Persistent Pollution

Figure 2 demonstrates that the persistent pollution stock's output begins to diverge visibly in 2030 for the BAU and AI-augmented scenarios. While a 4% higher peak may appear small, recent advances in attribution science clarify the massive implications of each additional percentage point, even when merely considered in monetary terms: "Each extra percentage point contribution to total 1850–2020 CO<sub>2</sub> and CH<sub>4</sub> emissions generates a further \$834 billion in global economic losses from extreme heat in 1991–2020" [5, p. 897]. World3-03's non-linear feedbacks amplify small stocks which may then magnify into sizable human-welfare impacts. The AI-augmented scenario peaks later and higher, then falls more slowly. In World3-03, the persistent pollution stock feeds into population and food systems through health and land productivity multipliers. A later higher peak is postponed, yet essentially deepens the downstream feedback that triggers the collapse.

#### 4.2 Variable 2: Human Ecological Footprint

We additionally explore the updates to the HEF variable given its direct relevance to pollution and updated output that is triggered from the changes to the persistent pollution stock. HEF is defined as "the sum of three sectors: the arable land used for crop production in agriculture; the urban land used for urban-industrial-transportation infrastructure; and the amount of absorption land required to neutralize the emission of pollutants, assumed to be proportional to the persistent pollution generation rate" [43, p. 293]. The comparison between the BAU and AI-augmented scenarios, in Figure 3, shows that because AI-caused pollution and material throughput intensify feedbacks of BAU, the AI-augmented scenario descends a little faster initially, yet it remains consistently 4–8% above BAU through to 2040. Cumulatively (area between the curves) the world spends approximately 7% more biocapacity than in the baseline over 2020–2070. The AI scenario finishes the century 11% above BAU and still marginally in overshoot. The implication is that, even after collapse, the planetary debt left by an AI-heavy economy is larger and persists longer.

### 5 Discussion

Including AI impacts as modifications in World3-03 opens the door to new questions that the current version cannot answer. In this

section, we first discuss the key takeaways from our findings. Second, we discuss the limitations of our study by highlighting that our results demonstrate the *viability* of using the World3-03 model to quantify relationships between AI impacts and world development but do not yet quantify those relationships itself. Lastly, we outline several research opportunities to encourage the use of the World3-03 model to explore relationships between AI and other variables in a quantitative manner. The research opportunities are ways to deepen the understanding of how AI development interacts with planetary limits by modifying the World3-03 model such as by: updating different variables than we have presented, adding new scenarios, adding new subcomponents, and simulating AI development projections.

#### 5.1 Takeaways


The modifications made to *increase* persistent pollution and the resulting modification to HEF demonstrate that our added assumptions, drawn from evidence, enable us to modify the World3-03 model and obtain expected consequences: AI scaling may deepen and prolong global overshoot in comparison to the Standard Run (BAU). This may further suggest that adding an AI-related layer to a growth-oriented global economy widens the gap between humanity's demands and Earth's regenerative capacity. Our results based on primitive, preliminary explorations of how AI impacts can be incorporated into the World3-03 model demonstrates that this is a promising method to more rigorously quantify the dynamics of the relationships among various systems. We outline several research opportunities to explore these relationships in a later section.

The World3-03 model simulates scenarios based on provided data and assumptions. In current discussions around AI's impact on the environment, there are two asymmetrical approaches that are taken. On the one hand, climate scientists and researchers report on observed and predicted environmental impact of AI using exhaustively researched and empirically validated measures based on conservative assumptions. On the other hand, Big Tech and its supporters make 'climate-positive' assurances with no empirical evidence (e.g., as seen in [33]) and instead create hype [16, 74] around future transformative abilities of AI. In the meanwhile, data center operations have precise measurements of their energy and water impact but such information is purposefully obfuscated [21, 30].

#### 5.2 Limitations

As a proof of concept paper, our objective was to show that the World3-03 model can be re-examined and applied to include AI-related variables, thereby projecting global dynamics in a new light. Future researchers should consider several limitations that delineate the scope of our modifications.


We modified the persistent pollution generation rate, while other variables, such as labour productivity, industrial output, land use, and resource depletion, remained untouched. Hypothetically, these could affect the model as strongly as, or even more strongly than, pollution. Therefore, the resulting output is incomplete by design: it illustrates only one aspect of AI's footprint. These untouched sectors constitute key areas for further exploration, as outlined in the research opportunities section.



**Figure 2: Standard and AI-augmented scenarios for Persistent Pollution Stock.**

Year	% Change from BAU
2020	0.94
2040	3.77
2060	21.69
2080	37.31
2100	45.35

**Table 2: Percent change between BAU and AI-augmented scenarios for Persistent Pollution Stock.**



**Figure 3: Standard (BAU) and AI-augmented scenarios for Human Ecological Footprint**

Year	% Change from BAU
2020	0.01%
2040	8.40%
2060	7.09%
2080	3.83%
2100	10.71%

**Table 3: Percent change between Standard (BAU) and AI-augmented scenarios for Human Ecological Footprint.**

AI-related industries are fast moving and constantly evolving. Empirical estimates have widened substantially over the past five years, and consequently, every coefficient could shift by an order of magnitude as new evidence emerges. The values used for the model, as well as the numerical outputs, should be perceived as illustrative ranges, not calibrated predictions. They may appear pessimistic or optimistic depending on which predictions ultimately prove correct.

On the technical side, the software provided two warnings for current agriculture inputs and food ratio. The warnings suggested a mismatch between the initial conditions set manually and the calculated values based on the model equation. Despite several attempts to treat these warnings, a resolution was not found. However, we do not believe this to impact our findings as the differences were minuscule and therefore unlikely to affect the model dynamic over a long-term horizon.

World3-03 is a globally aggregated model, thus its feedback loops may not capture real-time demands or surface regional pressures on planetary boundaries. While such simplifications keep the equations tractable, they blur regional nuances – precisely the hotspots where limits are likely to be breached first and where targeted policy could be most effective.

### 5.3 Research Opportunities

**Modifying different variables:** Further exploration of updating the World3-03 model to simulate AI impacts on limits to growth can include modifying and adding variables other than pollution, such as non-renewable resources and land development. Data center development demands vast quantities of high-grade silicon and

rare-earth elements [41], thus integrating those flows into the non-renewable-resource stock would reveal potential early-warning breaking points in World3-03. At the same time, the construction of data centers and related logistics infrastructure is converting land at a rapid pace [21, 27]. Embedding a variable representing the impacts of AI development on land use would allow researchers to assess whether terrestrial limits emerge simultaneously with, or even precede, pollution and resource constraints. The updated model could then test scenarios that explore the limits associated with unconstrained AI scaling.

**Establishing scenarios for AI scaling:** Understanding how AI might reshape global trends within planetary boundaries demands more than a single calibrated scenario; it requires a structured set of what-if experiments. Given the diverse range of projections of ecological impacts, potential efficiency gains, and more, it may be useful to establish scenarios of AI scaling in which different ranges of values can be considered for parameters.

We therefore propose a scenario family that can be explored within the current framework: Standard Scenario, Rapid AI Growth, Regulated Growth with a Sustainability Focus, and AI Stagnation. The Standard Scenario (i.e., BAU), unchanged from the classic model, would act as the control against which the remaining three AI scenarios can be compared. The three AI scenarios would be as follows:

- **Rapid AI scaling:** fast resource extraction and high pollution accompany soaring productivity, leading to intense resource use and a potential overshoot-and-collapse dynamic. This case investigates current acceleration narratives and the trade-off between short-term gains and long-term stability.



- Regulated growth with sustainability focus: strict limits on energy use, possibly combined with a shift to renewables, slow deployment but aim for a more balanced, durable trajectory.
- AI stagnation: regulatory caps or raw-material shortages curb adoption, offering a lower-bound estimate of resource use and pollution.
- Efficiency gains and rebound effects: Most positive consequences of AI fall under efficiency gains in other sectors [28]. Incorporating these, and their rebound and other effects, would lend much expressive power to the models at the cost of significant complexity.

**Adding new AI sector:** Another promising research direction is to build a stand-alone AI sector with its own stocks (e.g., AI capital, compute capacity, critical material inventories, etc.) and flows (hardware turnover, e-waste generation, electricity demand, etc.) and then connect its key variables to the existing World3-03 loops. This modular approach would place each AI variable in a single, transparent location, making its influence on the wider system easy to trace. This method would allow researchers to tune the AI sector without disturbing unrelated sectors, turning the model into a living toolkit for exploring how AI scaling may accelerate or alleviate pressure on planetary boundaries.

**Simulating current AI projections:** A next stage research agenda could be a scenario that links the AI capital stock to industrial output and, in turn, captures the feedback between investments and economic growth. In essence, this new scenario, with explicit AI capital stock from venture capital and corporate investments, would feed directly into World3-03 loops. This scenario will visualize how current financial projections and investments in AI might push or relax the planet's biophysical limits. Feeding "bullish" capital investment trajectories for the AI share into the model could reveal how rapidly the global system tips into collapse.

**Moving towards planetary boundaries and pathways:** The World3-03 model offers none or very limited considerations of geographic specifics, the dynamics of planetary boundaries and tipping points, and other planetary ecosystem dynamics [64, 65, 69, 77]. Future work should explore planetary boundaries and Shared Socioeconomic Pathways (SSPs) as modelling and scenario baseline for the questions explored here, moving beyond limitations of the World3-03 model [56, 62].

## 6 Conclusion

AI systems impact all lives on this planet in economic, social, and ecological ways. A system dynamics approach through a planetary-scale World3 model offers a practical method for exploring macro-level consequences. This includes quantitative modelling of impacts on planet-scale limits.

If AI scaling ever were to reach the absurd 'growth' conjured up by Big Tech figures – as for example Eric Schmidt's recent claim that AI will eventually demand '99 percent' of the world's electricity consumption [76] – it would most certainly shatter planetary boundaries. This makes it an urgent challenge to *Computing within Limits*.

In this paper, we have demonstrated how the World3-03 model still remains a feasible method for simulating system-level dynamics

of limits to growth, including how the accelerating development and operations of hyperscale data centers at a global scale impacts these limits. Using an exploratory approach, we added variables to the World3-03 model that would increase pollution based on recent evidence and prediction of its acceleration due to AI scaling. Accordingly, we were able to simulate an increase in pollution figures in the World3-03 model and thus established a proof of concept for the World3-03 model's continuing relevance. Future applications of the model can provide a quantitative approach to exploring 'what-if' questions regarding future AI and computing impacts.

Currently, predictions and metrics surrounding AI often boast of its innovative abilities while lacking critical distance and empirical evidence, and industry players obscure the details surrounding its harms. Here we take the opposite approach: we begin by modelling what is widely documented, i.e., the lifecycle impacts. Once beneficial impacts can be proven, they should also be integrated.

## References

- [1] Ugo Bardi. 2011. *The Limits to Growth Revisited*. Springer, New York, NY. doi:10.1007/978-1-4419-9416-5
- [2] Eshta Bhardwaj, Rohan Alexander, and Christoph Becker. 2025. Limits to AI Growth: The Ecological and Social Consequences of Scaling. doi:10.48550/arXiv.2501.17980 arXiv:2501.17980 [cs].
- [3] Benedetta Brevini. 2021. *Is AI Good for the Planet?* John Wiley & Sons.
- [4] Jacques Bughin, Jeongmin Seong, James Manyika, and Michael Chui. 2018. Modeling the global economic impact of AI [. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy>
- [5] Christopher W. Callahan and Justin S. Mankin. 2025. Carbon majors and the scientific case for climate liability. *Nature* 640, 8060 (April 2025), 893–901. doi:10.1038/s41586-025-08751-3 Publisher: Nature Publishing Group.
- [6] Rodrigo Castro. 2012. Arguments on the Imminence of Global Collapse Are Premature when Based on Simulation Models. *GAIA - Ecological Perspectives for Science and Society* 21, 4 (2012), 271–273. doi:10.14512/gaia.21.4.9
- [7] Jay Chen. 2015. Computing within limits and ICTD. *First Monday* (2015). doi:10.5210/fm.v20i8.6124
- [8] Christoph Becker. 2023. *Insolvent: How to Reorient Computing for Just Sustainability*. The MIT Press. <https://direct.mit.edu/books/book/5594/InsolventHow-to-Reorient-Computing-for-Just>
- [9] H. S. D. Cole. 1973. *Models of doom; a critique of The limits to growth*. New York, Universe Books.
- [10] Rob Comber and Elina Eriksson. 2023. Computing as Ecocide. In *Computing within Limits*. LIMITS. doi:10.21428/bf6fb269.9fcd0c0
- [11] Nordic Computer. [n. d.]. What Is the Effect of Datacenter Hardware on the Climate? [https://nordiccomputer.com/resources/what-is-the-effect-of-datacenter-hardware-on-the-climate?utm\\_source=chatgpt.com](https://nordiccomputer.com/resources/what-is-the-effect-of-datacenter-hardware-on-the-climate?utm_source=chatgpt.com)
- [12] Kate Crawford. 2024. Generative AI's environmental costs are soaring – and mostly secret. *Nature* 626, 8000 (2024), 693–693. doi:10.1038/d41586-024-00478-x
- [13] Miyuru Dayarathna, Yonggang Wen, and Rui Fan. 2015. Data center energy consumption modeling: A survey. *IEEE Communications surveys & tutorials* 18, 1 (2015), 732–794. Publisher: IEEE.
- [14] Donella Meadows, Jorgen Randers, and Dennis Meadows. 2004. A Synopsis: Limits to Growth: The 30-Year Update. <https://donellameadows.org/archives/a-synopsis-limits-to-growth-the-30-year-update/>
- [15] Eric Masanet and Nuoa Lei. 2020. How Much Energy Do Data Centers Really Use? <https://energyinnovation.org/2020/03/17/how-much-energy-do-data-centers-really-use/>
- [16] Luciano Floridi. 2024. Why the AI Hype is Another Tech Bubble. *Philosophy & Technology* 37, 4 (2024), 128. doi:10.1007/s13347-024-00817-w
- [17] Jay Wright Forrester. 1971. *World Dynamics* (2nd edition ed.). Pegasus Communications.
- [18] Jay W. Forrester, Gilbert W. Low, and Nathaniel J. Mass. 1974. The debate on World Dynamics: A response to Nordhaus. *Policy Sciences* 5, 2 (1974), 169–190. doi:10.1007/BF00148039
- [19] Vanessa Forti, Cornelis P. Balde, Ruediger Kuehr, and Garam Bel. 2020. *The Global E-waste Monitor 2020: Quantities, flows and the circular economy potential*. United Nations University/United Nations Institute for Training and Research, International Telecommunication Union, and International Solid Waste Association. <https://collections.unu.edu/view/UNU:7737>

- [20] Charlotte Freitag, Mike Berners-Lee, Kelly Widdicks, Bran Knowles, Gordon S. Blair, and Adrian Friday. 2021. The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. *Patterns* 2, 9 (2021), 100340. doi:10.1016/j.patter.2021.100340
- [21] Miranda Gabbott. 2024. Why We Don't Know AI's True Water Footprint. <https://techpolicy.press/why-we-dont-know-ais-true-water-footprint>
- [22] Stephen M Gardiner. 2014. *A Perfect Moral Storm: The Ethical Tragedy of Climate Change*. Oxford University Press, Cary. OCLC: 960165657.
- [23] Gianluca Guidi, Francesca Dominici, Jonathan Gilmour, Kevin Butler, Eric Bell, Scott Delaney, and Falco J. Bargagli-Stoffi. 2024. Environmental Burden of United States Data Centers in the Artificial Intelligence Era. doi:10.48550/arXiv.2411.09786 arXiv:2411.09786 [cs].
- [24] Yuelin Han, Zhifeng Wu, Pengfei Li, Adam Wierman, and Shaolei Ren. 2024. The Unpaid Toll: Quantifying the Public Health Impact of AI. doi:10.48550/arXiv.2412.06288 arXiv:2412.06288 [cs].
- [25] Gaya Herrington. 2021. Update to limits to growth: Comparing the World3 model with empirical data. *Journal of Industrial Ecology* 25, 3 (2021), 614–626. doi:10.1111/jiec.13084 \_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jiec.13084>
- [26] John Herrman. 2024. AI Investors Are Starting to Wonder: Is This Just a Bubble? <https://nymag.com/intelligencer/article/ai-investors-are-starting-to-wonder-is-this-just-a-bubble.html>
- [27] Mél Hogan. 2015. Data flows and water woes: The Utah Data Center. *Big Data & Society* 2, 2 (2015). doi:10.1177/2053951715592429
- [28] IEA. 2025. *Energy and AI*. International Energy Authority, Paris. [www.iea.org/reports/energy-and-ai](http://www.iea.org/reports/energy-and-ai)
- [29] IEEE. 2020. *International Roadmap for Devices and Systems*. <https://irds.ieee.org/editions/2020>
- [30] David Jeans. 2021. Data In The Dark: How Big Tech Secretly Secured \$800 Million In Tax Breaks For Data Centers. <https://www.forbes.com/sites/davidjeans/2021/08/19/data-in-the-dark-how-big-tech-secretly-secured-800-million-in-tax-breaks-for-data-centers/>
- [31] Lynn H. Kaack, Priya L. Donti, Emma Strubell, George Kamiya, Felix Creutzig, and David Rolnick. 2022. Aligning artificial intelligence with climate change mitigation. *Nature Climate Change* 12, 6 (2022), 518–527. doi:10.1038/s41558-022-01377-7 Publisher: Nature Publishing Group.
- [32] Keach Hagey. 2024. Sam Altman Seeks Trillions of Dollars to Reshape Business of Chips and AI. <https://www.wsj.com/tech/ai/sam-altman-seeks-trillions-of-dollars-to-reshape-business-of-chips-and-ai-89ab3db0>
- [33] Ketan Joshi. 2025. Is AI really going to meet 20% of our 2030 climate goals? <https://web.archive.org/web/20250417183135/https://www.linkedin.com/pulse/ai-really-going-meet-20-our-2030-climate-goals-ketan-joshi-xhtw/>
- [34] Christopher Ketcham. 2023. When Idiot Savants Do Climate Economics. <https://theintercept.com/2023/10/29/william-nordhaus-climate-economics/>
- [35] Spencer Kimball. 2024. Microsoft signs deal to invest more than \$10 billion on renewable energy capacity to power data centers. *CNBC* (2024). <https://www.cnbc.com/2024/05/01/microsoft-brookfield-to-develop-more-than-10point5-gigawatts-of-renewable-energy.html> Section: Energy.
- [36] Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. 2023. Making AI Less “Thirsty”: Uncovering and Addressing the Secret Water Footprint of AI Models. arXiv:2304.03271 (2023). doi:10.48550/arXiv.2304.03271 arXiv:2304.03271 [cs].
- [37] Alexandra Sasha Luccioni and Alex Hernandez-Garcia. 2023. Counting Carbon: A Survey of Factors Influencing the Emissions of Machine Learning. doi:10.48550/arXiv:2302.08476
- [38] Alexandra Sasha Luccioni, Emma Strubell, and Kate Crawford. 2025. From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarized Environmental Debate. arXiv:2501.16548 (2025). doi:10.48550/arXiv.2501.16548 arXiv:2501.16548 [cs].
- [39] Alexandra Sasha Luccioni, Sylvain Viguier, and Anne-Laure Ligozat. 2023. Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model. *Journal of Machine Learning Research* 24, 253 (2023), 1–15. <http://jmlr.org/papers/v24/23-0069.html>
- [40] Sasha Luccioni, Yacine Jernite, and Emma Strubell. 2024. Power Hungry Processing: Watts Driving the Cost of AI Deployment?. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*. Association for Computing Machinery, New York, NY, USA, 85–99. doi:10.1145/3630106.3658542
- [41] Sasha Luccioni, Bruna Trevelin, and Margaret Mitchell. 2024. The Environmental Impacts of AI - Primer. <https://huggingface.co/blog/sasha/ai-environment-primer>
- [42] Santosh K. Mahapatra, Tobias Schoenherr, and Jayanth Jayaram. 2021. An assessment of factors contributing to firms' carbon footprint reduction efforts. *International Journal of Production Economics* 235 (2021), 108073. doi:10.1016/j.ijpe.2021.108073
- [43] Donella Meadows, Jorgen Randers, and Dennis Meadows. 2004. *Limits to Growth: The 30-Year Update*. Chelsea Green Publishing. Google-Books-ID: QRyQ-INGW6oC.
- [44] Donella H. Meadows, Dennis L. Meadows, and Jørgen Randers. 1992. *Beyond the limits: confronting global collapse envisioning a sustainable future*. Chelsea Green Publ. Company, Vermont.
- [45] Donella H. Meadows, Dennis L. Meadows, and Jørgen Randers. 1993. Beyond the Limits: An Executive Summary. *Bulletin of Science, Technology & Society* 13, 1 (1993), 3–14. doi:10.1177/027046769301300102 Publisher: SAGE Publications Inc.
- [46] Donella H. Meadows, Dennis L. Meadows, Jørgen Randers, and William W. Behrens Iii. 1972. *The Limits to Growth - Club of Rome*. <https://policycommons.net/artifacts/1529440/the-limits-to-growth/2219251/> Publisher: Club of Rome.
- [47] James Muldoon, Callum Cant, Mark Graham, and Funda Ustek Spilda. 2023. The poverty of ethical AI: impact sourcing and AI supply chains. *AI & SOCIETY* (2023). doi:10.1007/s00146-023-01824-9
- [48] Bonnie Nardi, Bill Tomlinson, Donald J. Patterson, Jay Chen, Daniel Pargman, Barath Raghavan, and Birgit Penzenstadler. 2018. Computing within limits. *Commun. ACM* 61, 10 (2018), 86–93. Publisher: ACM New York, NY, USA.
- [49] Arjuna Nebel, Alexander Kling, Ruben Willamowski, and Tim Schell. 2024. Recalibration of limits to growth: An update of the World3 model. *Journal of Industrial Ecology* 28, 1 (2024), 87–99. doi:10.1111/jiec.13442 \_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jiec.13442>
- [50] Nestor Maslej, Loredana Fattorini, Raymond Perrault, Vanessa Parli, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Juan Carlos Nieves, Yoav Shoham, Russell Wald, and Jack Clark. 2024. *The AI Index 2024 Annual Report*. Technical Report. Institute for Human-Centered AI, Stanford University, Stanford, CA. <https://aiindex.stanford.edu/report/>
- [51] William D. Nordhaus. 1973. World Dynamics: Measurement Without Data. *The Economic Journal* 83, 332 (1973), 1156–1183. doi:10.2307/2230846
- [52] William D. Nordhaus, Robert N. Stavins, and Martin L. Weitzman. 1992. Lethal Model 2: The Limits to Growth Revisited. *Brookings Papers on Economic Activity* 1992, 2 (1992), 1–59. doi:10.2307/2534581
- [53] OECD. 2019. *Artificial Intelligence in Society*. Paris. <https://doi.org/10.1787/eedfee77-en>
- [54] Antonio Olivo. 2024. Internet data centers are fueling drive to old power source: Coal. *Washington Post* (2024). <https://www.washingtonpost.com/business/interactive/2024/data-centers-internet-power-source-coal/>
- [55] Ondrej Burckacký, Julia Dragon, and Nikolaus Lehmann. 2022. The semiconductor decade: A trillion-dollar industry. <https://www.mckinsey.com/industries/semiconductors/our-insights/the-semiconductor-decade-a-trillion-dollar-industry>
- [56] Brian C. O'Neill, Elmar Kriegler, Kristie L. Ebi, Eric Kemp-Benedict, Keywan Riahi, Dale S. Rothman, Bas J. van Ruijven, Detlef P. van Vuuren, Joern Birkmann, Kasper Kok, Marc Levy, and William Solecki. 2017. The roads ahead: Narratives for shared socioeconomic pathways describing world futures in the 21st century. *Global Environmental Change* 42 (Jan. 2017), 169–180. doi:10.1016/j.gloenvcha.2015.01.004
- [57] Roberto Pasqualino, Aled W. Jones, Irene Monasterolo, and Alexander Phillips. 2015. Understanding Global Systems Today—A Calibration of the World3-03 Model between 1995 and 2012. *Sustainability* 7, 8 (2015), 9864–9889. doi:10.3390/su7089864 Number: 8 Publisher: Multidisciplinary Digital Publishing Institute.
- [58] ByPeter Passell and Marc Roberts and Leonard Ross. 1972. The Limits to Growth. *The New York Times* (1972). <https://www.nytimes.com/1972/04/02/archives/the-limits-to-growth-a-report-for-the-club-of-romes-project-on-the.html>
- [59] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. 2021. Carbon Emissions and Large Neural Network Training. arXiv:2104.10350 (2021). doi:10.48550/arXiv.2104.10350 arXiv:2104.10350 [cs].
- [60] Anand S. Rao and Gerard Verweij. 2017. *Sizing the prize What's the real value of AI for your business and how can you capitalise?*
- [61] Kate Raworth. 2017. *Doughnut economics: seven ways to think like a 21st century economist*. Chelsea Green Publishing, White River Junction, Vermont.
- [62] Keywan Riahi, Detlef P. van Vuuren, Elmar Kriegler, Jae Edmonds, Brian C. O'Neill, Shinichiro Fujimori, Nico Bauer, Katherine Calvin, Rob Dellink, Oliver Fricko, Wolfgang Lutz, Alexander Popp, Jesus Crespo Cuaresma, Samir Kc, Marian Leimbach, Leïwen Jiang, Tom Kram, Shilpa Rao, Johannes Emmerling, Kristie Ebi, Tomoko Hasegawa, Petr Havlik, Florian Humpenöder, Lara Aleluia Da Silva, Steve Smith, Elke Stehfest, Valentina Bosetti, Jiyong Eom, David Gernaat, Toshihiko Masui, Joeri Rogelj, Jessica Streffer, Laurent Drouet, Volker Krey, Gunnar Luderer, Mathijs Harmsen, Kiyoshi Takahashi, Lavinia Baumstark, Jonathan C. Doelman, Mikiko Kainuma, Zbigniew Klimont, Giacomo Marangoni, Hermann Lotze-Campen, Michael Obersteiner, Andrzej Tabeau, and Massimo Tavoni. 2017. The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global Environmental Change* 42 (Jan. 2017), 153–168. doi:10.1016/j.gloenvcha.2016.05.009
- [63] Johan Rockström, Louis Kotzé, Svetlana Milutinović, Frank Biermann, Victor Brovkin, Jonathan Donges, Jonas Ebesson, Duncan French, Joyeeta Gupta, Rakhyn Kim, Timothy Lenton, Dominic Lenzi, Nebojsa Nakicenovic, Barbara Neumann, Fabian Schuppert, Ricarda Winkelmann, Klaus Bosselmann, Carl Folke, Wolfgang Lucht, David Schlosberg, Katherine Richardson, and Will Steffen. 2024. The planetary commons: A new paradigm for safeguarding Earth-regulating

- systems in the Anthropocene. *Proceedings of the National Academy of Sciences* 121, 5 (2024), e2301531121. doi:10.1073/pnas.2301531121
- [64] Johan Rockström, Louis Kotzé, Svetlana Milutinović, Frank Biermann, Victor Brovkin, Jonathan Donges, Jonas Ebbesson, Duncan French, Joyeeta Gupta, Rakhyun Kim, Timothy Lenton, Dominic Lenzi, Nebojsa Nakicenovic, Barbara Neumann, Fabian Schuppert, Ricarda Winkelmann, Klaus Bosselmann, Carl Folke, Wolfgang Lucht, David Schlosberg, Katherine Richardson, and Will Steffen. 2024. The planetary commons: A new paradigm for safeguarding Earth-regulating systems in the Anthropocene. *Proceedings of the National Academy of Sciences* 121, 5 (Jan. 2024), e2301531121. doi:10.1073/pnas.2301531121 Publisher: Proceedings of the National Academy of Sciences.
- [65] Johan Rockström, Will Steffen, Kevin Noone, Åsa Persson, F. Stuart Chapin, Eric Lambin, Timothy M. Lenton, Marten Scheffer, Carl Folke, Hans Joachim Schellnhuber, Björn Nykvist, Cynthia A. de Wit, Terry Hughes, Sander van der Leeuw, Henning Rodhe, Sverker Sörlin, Peter K. Snyder, Robert Costanza, Uno Svedin, Malin Falkenmark, Louise Karlberg, Robert W. Corell, Victoria J. Fabry, James Hansen, Brian Walker, Diana Liverman, Katherine Richardson, Paul Crutzen, and Jonathan Foley. 2009. Planetary Boundaries: Exploring the Safe Operating Space for Humanity. *Ecology and Society* 14, 2 (2009). <https://www.jstor.org/stable/26268316> Publisher: Resilience Alliance Inc..
- [66] Emil Sayegh. 2024. The Billion-Dollar AI Gamble: Data Centers As The New High-Stakes Game. <https://www.forbes.com/sites/emilsayegh/2024/09/30/the-billion-dollar-ai-gamble-data-centers-as-the-new-high-stakes-game/>
- [67] Jaime Sevilla, Lennart Heim, Anson Ho, Tamay Besiroglu, Marius Hobbahn, and Pablo Villalobos. 2022. Compute Trends Across Three Eras of Machine Learning. doi:10.48550/arXiv.2202.05924
- [68] Jaime Sevilla, Tamay Besiroglu, Ben Cottier, Josh You, Edu Roldán, Pablo Villalobos, and Ege Erdil. 2024. *Can AI Scaling Continue Through 2030?* Technical Report. epoch.ai. <https://epochai.org/blog/can-ai-scaling-continue-through-2030>
- [69] Will Steffen, Katherine Richardson, Johan Rockström, Sarah E. Cornell, Ingo Fetzer, Elena M. Bennett, Reinette Biggs, Stephen R. Carpenter, Wim de Vries, Cynthia A. de Wit, Carl Folke, Dieter Gerten, Jens Heinke, Georgina M. Mace, Linn M. Persson, Veerabhadran Ramanathan, Belinda Reyers, and Sverker Sörlin. 2015. Planetary boundaries: Guiding human development on a changing planet. *Science* 347, 6223 (2015), 1259855. doi:10.1126/science.1259855
- [70] Nicholas Stern, Stiglitz Charlotte Taylor, Joseph, , and Charlotte Taylor. 2022. The economics of immense risk, urgent action and radical change: towards new approaches to the economics of climate change. *Journal of Economic Methodology* 29, 3 (July 2022), 181–216. doi:10.1080/1350178X.2022.2040740 Publisher: Routledge \_eprint: <https://doi.org/10.1080/1350178X.2022.2040740>.
- [71] Neil Thompson, Kristjan Greenewald, Keeheon Lee, and Gabriel F. Manso. 2023. The Computational Limits of Deep Learning. In *Computing within Limits*. LIMITS. doi:10.21428/bf6fb269.1f033948
- [72] Neil C. Thompson and Svenja Spanuth. 2021. The decline of computers as a general purpose technology. *Commun. ACM* 64, 3 (2021), 64–72. doi:10.1145/3430936
- [73] Graham M. Turner. 2008. A comparison of *The Limits to Growth* with 30 years of reality. *Global Environmental Change* 18, 3 (2008), 397–411. doi:10.1016/j.gloenvcha.2008.05.001
- [74] Gaël Varoquaux, Alexandra Sasha Luccioni, and Meredith Whittaker. 2025. Hype, Sustainability, and the Price of the Bigger-is-Better Paradigm in AI. doi:10.48550/arXiv.2409.14160 arXiv:2409.14160 [cs].
- [75] Ventana Systems Inc. 2025. Vensim. [https://vensim.com/documentation/vensim-10\\_3.html](https://vensim.com/documentation/vensim-10_3.html)
- [76] Joe Wilkins. 2025. Former Google CEO Tells Congress That 99 Percent of All Electricity Will Be Used to Power Superintelligent AI. <https://futurism.com/google-ceo-congress-electricity-ai-superintelligence>
- [77] Amanda Williams, Paolo Perego, and Gail Whiteman. 2024. Boundary Conditions for Organizations in the Anthropocene: A Review of the Planetary Boundaries Framework 10 Years On. *Journal of Management Studies* n/a, n/a (2024). doi:10.1111/joms.13150 \_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/joms.13150>.
- [78] Matteo Wong. 2024. Silicon Valley's Trillion-Dollar Leap of Faith. <https://www.theatlantic.com/technology/archive/2024/07/ai-companies-unprofitable/679278/>

# Appendix

Table 4: Scenarios adapted from the LtG of 1972 [6].

Scenario	Name	Description
1	Standard Run	The “standard” world model assumes no major shifts in the systems shaping global development. Food, industry, and population grow fast until a declining resource base halts industrial growth. Delays cause pollution and population to rise briefly after the peak. Growth stops as death rates climb due to less food and healthcare.
2	Natural Resource Reserved Doubled	To test the model’s resource assumption, 1900 reserves were doubled, with all other conditions unchanged from the Standard Run. Industrialization reaches a higher peak as resources last longer, but the larger industrial base generates pollution faster than the environment can absorb. Pollution spikes, driving up death rates and cutting food outputs. By the end, resources are nearly exhausted, despite the initial doubling.
3	“Unlimited” Resources	Resource depletion is avoided by assuming unlimited nuclear energy and double resource reserves, enabling widespread recycling and substitution. But if these are the only changes, growth still ends due to increasing pollution.
4	“Unlimited” Resources and Pollution Controls	A new technological change was added to reduce pollution per unit of output to one-fourth of its 1970 level, aiming to solve earlier resource and pollution issues. This allows growth until arable land becomes the limit. Food per person drops, and industry slows as capital shifts to agriculture.
5	“Unlimited” Resources, Pollution Controls, and Increased Agricultural Productivity	To prevent the earlier food crisis, land yield was doubled in 1975, alongside prior pollution and resource policies. These combined changes remove key growth limits, leading to high population and industrial levels. But despite lower pollution per unit, total output grows so much that pollution still triggers collapse.
6	“Unlimited” Resources, Pollution Controls, and “Perfect” Birth Control	Instead of boosting food output, the model tests more effective voluntary birth control to ease the food issue. Population still grows, just more

		slowly, and the crisis is only delayed by 10–20 years.
7	“Unlimited” Resources, Pollution Controls, Increased Agricultural Productivity, and “Perfect” Birth Control	Four tech policies are added to avoid collapse: full resource use with 75% recycling, pollution cut to one-fourth, land yield doubled, and widespread access to birth control. Still, industrial growth stops and death rates rise as resources run low, pollution builds, and food drops.
8	Stabilized Population	This scenario matches the standard model, except population is stabilized after 1975 by balancing birth and death rates. The unchecked industrial feedback loop drives continued exponential growth in output, food, and services. Eventually, resource depletion causes industrial collapse.
9	Stabilized Population and Capital	Capital growth is limited by making investment equal to depreciation, added to the population stabilization policy. This halts exponential growth and creates a temporary steady state. But without resource-saving technology, high levels of population and capital still drain resources quickly. As resources fall, industrial output drops, and capital becomes less efficient, with more used for extraction than production.
10	Stabilized I	Technological and growth-limiting policies are combined to create a long-term sustainable equilibrium. Tech measures include recycling, pollution control, longer capital life, and soil restoration. Value shifts emphasize food and services over industry. Births equal deaths, and industrial investment matches depreciation. Industrial output per person reaches three times the 1970 average.
11	Stabilized II	If strict growth limits are lifted and population and capital are regulated within system delays, equilibrium shifts to higher population and lower industrial output per person. With perfect birth control and a desired family size of two by 1975, the birth rate slowly nears the death rate due to age-structure delays.
12	Stabilizing Policies Introduced in the Year 2000	When 1975 policies are delayed until 2000, the equilibrium becomes unsustainable. Population and industry grow too high, leading to food and resource shortages before 2100.

Table 5: New and existing variables for the persistent pollution sector of World3-03.

Variables	Type	Description	Importance
<i>co2_to_persist</i>	New Parameter	An empirical conversion constant used in World3 to translate a mass of greenhouse gases into the model's abstract persistent-pollution units (pu).	It is essential because without the conversion all new AI terms would carry incompatible units.
<i>ai_co2_intensity_2020</i>	New Parameter	Direct operational, plus embodied CO <sub>2</sub> released per dollar of AI infrastructure and model-training expenditure in 2020.	Captures the energy/CO <sub>2</sub> footprint of hyperscale data-centres and other build-outs. This variable is a key leverage point for climate-aware computers.
<i>ai_ewaste_intensity_2020</i>	New Parameter	Life-cycle CO <sub>2</sub> -equivalent to toxic metals, Per-and polyfluoroalkyl substances (PFAS), and others in AI hardware discarded in 2020.	This variable has been added because AI's footprint is not only energy but also material toxicity.
<i>frac_io_ai_2020</i> , <i>frac_io_ai_2050</i>	New Parameter	Fraction of industrial output of AI in 2020 and 2050. These variables are the empirically anchored share of global industrial output devoted to AI hardware & services in 2020 and its saturation level in 2050.	These levers, not intensity, drive how big the AI sector becomes inside the macroeconomy.
<i>industrial output</i>	Existing variable	This stock is the global dollar flow generated by the industrial capital sector.	When a recession, capital shortage, or overshoot-and-collapse episode drives industrial output down, the amount of compute that can be bought also falls. Thus, AI- related pollution rises and falls consistently with the rest of the economy.
<i>base_ai_efficiency</i>	New Parameter	This variable is an annual fractional improvement in compute efficiency, i.e. how quickly AI hardware is getting more energy-efficient over time.	This is a negative exponent, which means that as AI gets more efficient, it uses less electricity for the same economic output which reduces CO <sub>2</sub> emissions per dollar of AI services.
<i>ai_efficiency_slodown_rate</i>	New Parameter	Decay of the improvement rate.	Without a slowdown, AI could appear to become infinitely clean, which is unrealistic.

<i>ai_ewaste_improvement_rate</i>	New Parameter	This parameter lowers the AI_Ewaste_intensity part of AI_pollution_intensity each year to reflect progress in circular design, reuse and recycling.	It is essential because, even if computers get cleaner, material toxicity can dominate the long-lived pollution stock.
<i>Fraction_industrial_output_ai</i>	New Equation	Time-path of the AI share.	This variable helps avoid a discontinuous jump in 2020 and keeps the capital-allocation loop numerically stable.
<i>ai_output</i>	New Equation	This variable converts a share of the global economy into an absolute activity level.	It links the AI branch to the existing capital feedback: if the world economy collapses, AI activity collapses with it.
<i>ai_pollution_intensity</i>	New Equation	This variable scales today's footprint to any future year before the tech multiplier is applied.	It allows experimenters to tune learning-by-doing, i.e. make the 3 % and 5 % slopes steeper or shallower without touching the conversion logic.
<i>persistent_pollution_tech_change_multiplier</i>	Existing variable	Global S-curve lookup that halves effective intensity every ~25 yr down to a floor of 0.5.	Re-used so AI inherits the same “planetary clean-tech progress” assumption.
<i>ai_pollution_tech_change_multiplier</i>	New Equation	This variable helps prevent AI from falling below the global floor.	This lets the model decide whether AI is held to the same floor or allowed to fall lower/higher in scenario studies.
<i>persistent_pollution_generation_ai</i>	New Equation	The new inflow that feeds the long-lived pollution stock.	It activates exactly when generative-AI deployment begins. Automatically fades after 2100, so late-century behaviour is not distorted.
<i>Time</i>	Existing variable	This variable reflects the year.	It helps to set the time bounds within other variables.

Table 6: Values for the New Parameters, derived from [1–5, 8–11].

Parameters	Suggested Value	Units	Explanation
<i>co2_to_persist</i>	2.3 e-4	Dmnl	Fraction of a one-off CO <sub>2</sub> pulse that remains in the atmosphere/ocean “long enough” to show up in the LtG persistent pollution stock. Derived from [7].
<i>ai_CO2_intensity_2020</i>	1.5 e-1	Dmnl	Direct operational CO <sub>2</sub> emitted per 2020-\$ of AI output. Derived from [10].
<i>ai_ewaste_intensity_2020</i>	3.5 e-4	Dmnl	Embodied-carbon, end-of-life & recycling impacts per 2020-\$ of AI hardware spend. Derived from [8].
<i>frac_io_ai_2020</i>	1.3 e-2	Dmnl	Share of total industrial output that is AI/compute in 2020. Derived from [9].
<i>frac_io_ai_2050</i>	6.0 e-2	Dmnl	Target share of industrial output devoted to AI by 2050. Derived from [1, 5].
<i>base_ai_efficiency_improvement</i>	2.5 e-1	Dmnl	Annual fractional improvement in compute efficiency. Derived from [3, 4].
<i>ai_efficiency_slowdown_rate</i>	4.0 e-2	1/year	How fast the above improvement rate decays. Derived from [11].
<i>ai_ewaste_improvement_rate</i>	3.0 e-2	1/year	Annual reduction in e-waste CO <sub>2</sub> -eq per \$ due to circularity & reuse. Derived from [2] .



Table 7: New and updated equations.

Variable	Equations	Unit
<i>fraction_industrial_output_AI</i>	= fraction io ai 2020 + (fraction io ai 2050 - fraction io ai 2020) / (1 + EXP(-(Time - 2035) / 5))	Dmnl
<i>ai_output</i>	= IF THEN ELSE(Time >= 2020, fraction industrial output ai * industrial output, 0)	\$/year
<i>ai_pollution_intensity</i>	= ( ai co2 intensity 2020 * (1 + EXP(-base ai efficiency * (1 - MAX(0, Time - 2020) * ai efficiency slowdown rate))) * MAX(0, Time - 2020) + IF THEN ELSE( Time < 2070, ai ewaste intensity 2020 * (1 - ai ewaste improvement rate * MAX(0, Time - 2020)), 2e-05 ) ) * co2 to persistent	Pollution units/\$
<i>ai_pollution_tech_change_multiplier</i>	= MIN( 5 , MAX( 0.7 * persistent pollution technology change multiplier 1, 0.3 * (1 + 0.1 * (Time- 2020)) ) )	Dmnl
<i>persistent_pollution_generation_ai</i>	= IF THEN ELSE( Time < 2020 :OR: Time > 2100 , 0 , ai output * ai pollution intensity / ai pollution tech change multiplier )	Pollution units/year
<i>persistent pollution generation rate (existing)</i>	= ( persistent pollution generation industry + persistent pollution generation agriculture + persistent pollution generation AI) * ( persistent pollution generation factor )	Pollution units/year

Table 8: Comparison between the standard scenario (BAU) and AI-augmented scenario.

Time (year)	Persistent Pollution: AI-Augmented	Persistent Pollution: BAU	% Change from BAU
2020	976172000	967120000	0.94
2025	1.176E+09	1165000000	0.99
2030	1.369E+09	1352590000	1.21
2035	1.495E+09	1464870000	2.03
2040	1.506E+09	1451000000	3.77
2045	1.413E+09	1323730000	6.77
2050	1.26E+09	1134610000	11.08
2055	1.085E+09	932903000	16.29
2060	910573000	748269000	21.69
2065	749359000	591696000	26.65
2070	607107000	463889000	30.87
2075	486134000	361763000	34.38
2080	386071000	281165000	37.31
2085	304886000	218014000	39.85
2090	239870000	168755000	42.14
2095	188261000	130473000	44.29
2100	147537000	101504000	45.35

# References

- [1] Anand S. Rao and Gerard Verweij. 2017. *Sizing the prize What's the real value of AI for your business and how can you capitalise?* PWC.
- [2] Vanessa Forti, Cornelis P. Balde, Ruediger Kuehr, and Garam Bel. 2020. *The Global E-waste Monitor 2020: Quantities, flows and the circular economy potential*. United Nations University/United Nations Institute for Training and Research, International Telecommunication Union, and International Solid Waste Association. Retrieved April 29, 2025 from <https://collections.unu.edu/view/UNU:7737>
- [3] John L. Hennessy and David A. Patterson. 2019. A new golden age for computer architecture. *Commun ACM* 62, 2 (2019), 48–60. <https://doi.org/10.1145/3282307>
- [4] IEEE. 2020. *International Roadmap for Devices and Systems*. Retrieved April 30, 2025 from <https://irds.ieee.org/editions/2020>
- [5] Jacques Bughin, Jeongmin Seong, James Manyika, and Michael Chui. 2018. Modeling the global economic impact of AI |. *McKinsey*. Retrieved April 30, 2025 from <https://www.mckinsey.com/capabilities/quantumblack/our-insights/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy>
- [6] Donella H. Meadows, Dennis L. Meadows, Jørgen Randers, and William W. Behrens Iii. 1972. *The Limits to Growth - Club of Rome*. Club of Rome. Retrieved March 31, 2023 from <https://policycommons.net/artifacts/1529440/the-limits-to-growth/2219251/>
- [7] Donella Meadows, Jorgen Randers, and Dennis Meadows. 2004. *Limits to Growth: The 30-Year Update*. Chelsea Green Publishing.
- [8] Nordic Computer. What Is the Effect of Datacenter Hardware on the Climate? Retrieved April 30, 2025 from [https://nordiccomputer.com/resources/what-is-the-effect-of-datacenter-hardware-on-the-climate?utm\\_source=chatgpt.com](https://nordiccomputer.com/resources/what-is-the-effect-of-datacenter-hardware-on-the-climate?utm_source=chatgpt.com)
- [9] OECD. 2019. *Artificial Intelligence in Society*. OECD, Paris. Retrieved April 30, 2025 from <https://doi.org/10.1787/eedfee77-en>
- [10] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. 2021. Carbon Emissions and Large Neural Network Training. <https://doi.org/10.48550/arXiv.2104.10350>
- [11] Neil C. Thompson and Svenja Spanuth. 2021. The decline of computers as a general purpose technology. *Commun ACM* 64, 3 (2021), 64–72. <https://doi.org/10.1145/3430936>