Ecological Economics and Management
https://www.sciltp.com/journals/eem

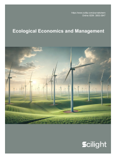

Article

Check for updates

Carbon Intensity Measurement and Green Spillover Effects of Generative Artificial Intelligence

Xinyu Gui [1] and Ritika Chopra [2,*]

[1] Business School, Suzhou University, Suzhou 234000, China
[2] University School of Management and Entrepreneurship, Delhi Technological University (East Campus), Jhilmil Colony, Delhi 110095, India
* Correspondence: ritikachopra67@gmail.com

How To Cite: Gui, X.; Chopra, R. Carbon Intensity Measurement and Green Spillover Effects of Generative Artificial Intelligence. *Ecological Economics and Management* **2026**, *2*(3), 13. https://doi.org/10.53941/eem.2026.100013

Received: 25 May 2026
Revised: 28 June 2026
Accepted: 28 July 2026
Published: 10 August 2026

Abstract: Generative Artificial Intelligence is changing production in an accelerating way. However, given that it has a relatively high demand for computing power and energy consumption, it generates a significant amount of greenhouse gas emissions. Whether it will be a "carbon increaser" or a "carbon reducer" in the process of green transition is still an urgent problem to be solved. This paper will study the carbon emission intensity and the green spillover effects of generative artificial intelligence, determine its environmental costs and green benefits, and provide reference materials for policymakers in their green development plans for the digital economy. From a life-cycle perspective, this paper sets the scope of measurement as hardware production, model training and cloud-based inference, and suggests a system for carbon intensity measurement and analytical models. Select representative large-scale models for empirical estimation. At the same time, a theoretical system explores the transmission paths of green spillover effects and carry out corresponding empirical studies. Based on the above data, the carbon intensity of generative artificial intelligence suggests substantial changes in both scale and structure. Most of the carbon emissions occur during the training stage, and although the proportion from the inference stage has been increasing over time, it is still relatively small. The regional energy structure of the models' deployment appears to be associated with carbon intensity. The results suggest a potential nonlinear pattern between generative artificial intelligence development and regional carbon intensity, although the direction of this relationship remains conditional on energy structure, digital infrastructure, and environmental regulation. The realization of green spillover effects depends on energy structure, digital infrastructure and environmental regulations. Therefore, this paper puts forward some policy recommendations, such as constructing a normalized carbon accounting system, optimizing the spatial arrangement of computing facilities, establishing integrated development mechanisms, enhancing incentive policies, and strengthening international cooperation. The scope of the study is broadened to include the environmental impact of the digital economy, and both theoretical and practical support for the green development of intelligent technologies and low-carbon industrial transformation will be offered.

Keywords: generative artificial intelligence; carbon intensity; green spillover effects; energy efficiency

1. Introduction

The last few years have seen observable advancements in generative artificial intelligence. It is beginning to affect the design of production systems, industrial frameworks, and the conduct of digital economic affairs. Large language models, multimodal generation systems, and intelligent decision- making systems are the latest offerings in the arsenal of tools used in finance, logistics, health, manufacturing, education, and public service. Generative artificial intelligence increases the level of dependence on digital infrastructure and electricity compared with previous generations of digital technologies. There is a growing sense of unease about the electricity required and the carbon emissions generated by the training and ongoing slow deployment of huge models. The increasing integration of AI into production and service systems is part of the broader transformation described in the digital-economics literature [1].

At the time digital transformation began, artificial intelligence was viewed as a "clean" technology because it was typically focused on non-physical, virtual outputs. Growing through a rapid expansion of model parameters, the ongoing race to make further improvements to A.I. is beginning to make artificial intelligence a much less "clean" technology. There is a considerable environmental cost from the training, deployment, and the cooling and infrastructure required to support the use of large models. Research has shown that the training of advanced models is likely to require several million kilowatt-hours of electricity, resulting in a significant increase in carbon emissions.

Due to goals of achieving carbon neutrality and a governance structure focusing on environmental, social, and corporate governance (ESG), the impact of generative artificial intelligence (AI) on the environment has become a hot topic for governments, corporations, and investors alike. Generative AI has the potential to overextend the already burdened computing infrastructure and data centers, ultimately resulting in greater energy consumption and carbon output. Despite this, generative AI may assist in the optimization of industrial systems that will stimulate greater production efficiency and resource consumption. Therefore, the more interesting question may lie in determining whether AI will emit less carbon than it consumes, or vice versa.

It has been mainly determined that a considerable amount of energy and carbon is used and/or produced throughout the various training processes of artificial intelligence. At least some of the research has been devoted to estimating the released carbon from AI systems based on their energy consumption and/or their hardware configuration. Other research has attempted to assess the various impacts, including the impacts of carbon during the complete lifecycle of the hardware and the infrastructures/load that support the hardware. Although these studies have helped answer some of the important questions in this area, many have also fallen short.

First, a unified carbon accounting framework for generative artificial intelligence is yet to be established. Various studies use diverse accounting boundaries, parameters, and methods for estimating emissions, which makes it most difficult to compare empirical results. Second, existing studies focus on the training stage received comparatively less attention. Very large cumulative emissions may be created from large-scale inference services, especially after generative artificial intelligence becomes commercially exploitable. Third, most studies focus on the direct emissions of generative artificial intelligence systems, whereas few studies focus on the effects of the potential green innovations caused by generative artificial intelligence in industrial restructuring, the diffusion of technology, and increase the efficiency of energy.

Generative artificial intelligence may have more multifaceted impacts on the environment compared to conventional digital technologies. The downsides of building more computing infrastructure may be greater emissions in the short run. However, technology and smart infrastructure may improve energy efficiency and reduce carbon emissions in the long run. Therefore, the effects of generative artificial intelligence on the environment may be different from what the linear relationship may suggest.

Due to the points mentioned, the study focuses on the carbon intensity and green spillover effects of generative artificial intelligence using the lens of energy economics and sustainable development. A carbon accounting scheme integrates Model-level energy consumption, data center operational efficiency, and regional electricity carbon intensity. Publicly documented models and experimental AI workloads are used for the empirical estimation.

This study's potential contributions can be found in several aspects. First, by incorporating training energy consumption, inference deployment, and data center efficiency metrics, this study develops a more comprehensive carbon intensity measurement framework for generative artificial intelligence. Second, this study proposes the AI Carbon Intensity Index (AICI) to address the carbon emission performance gap between different models and their respective deployment environments. Third, we broaden the focus in this domain and analyze the potential of generative artificial intelligence to produce green spillover effects from the enhancement of industrial structure and energy efficiency, in addition to serving the purpose of carbon emissions accounting. Lastly, this study endeavors to link the crossroads of technological advancement, energy transition, and the ESG governance framework, to provide further insight into an under-explored area of the digital economy and its environmental ramifications.

The rest of the paper is structured as follows. In Section 2, relevant literature and theories are reviewed. Section 3 presents the carbon accounting framework and the construction of variables. In Section 4, the effects of technological advancement and carbon intensity are analyzed. In Section 5, the potential green spillover effects from generative artificial intelligence are examined. In Section 6, the main results are presented, along with their implications. Section 7 contains the concluding remarks.

2. Literature Review and Theoretical Foundations

2.1. Research on Energy Consumption and Carbon Emissions of Artificial Intelligence

The overarching expansion of artificial intelligence (AI) generates and requires more energy. The effects that these growing energy demands and nonrenewable resource consumption have, and will have, on the digital economy and energy economics are beginning to emerge in the literature. Earlier research concentrated on energy consumption in data centers, communication infrastructure, and cloud computing systems. Now that large-scale models and generative artificial intelligence (AI) have spurred rapid expansion in the field, there are concerns about the carbon footprint that these large-scale models will leave. Researchers are beginning to evaluate the effects of modeling, training, and inference on the environment.

Some believe the rapidly expanding modeling parameters and data significantly increased the demand for electricity in AI systems. AI systems based on the exploded Transformer architecture require massively large GPU and TPU clustered circuitry that may take weeks or even months to train. The systems' demand and scaling will result in increasingly larger effects on the environment [2]. In addition to direct computing loads, cooling systems, power conversion, networking equipment, and hardware production may add to the wider environmental footprint of large AI systems [3–5].

Currently, there are two primary approaches to estimating the carbon footprint of AI. The first approach determines the carbon footprint directly through an analysis of the hardware design, power consumption, and the system's runtime. The second approach adopts a broader life-cycle perspective by considering not only model training electricity but also hardware production and the operational infrastructure supporting model deployment [6]. Life cycle assessment has a distinct advantage over direct analysis in determining the expansive effects AI systems have on the environment.

However, there are still multiple limitations to the existing body of research. First, diverse approaches to accounting boundaries and the parameter assumptions used across studies lead to significant segmentation in results for estimating carbon outputs. Second, research on the training phase of models is the most common, and less research has been done on the possible large-scale inferences and their long-term cumulative emissions. Third, most existing studies consider the direct environmental costs and do not consider if Generative Artificial Intelligence (GAI) may indirectly enhance other sectors' energy efficiency or green transformation.

Traditional digital technologies and GAI are both digitally based. GAI, however, has a higher computing requirement and larger energy demands than traditional digital technologies. For this reason, GAI's environmental effects may be more complex than digital technologies.

2.2. Research on Carbon Intensity Measurement

Carbon intensity is used to assess the connection between economic activity and carbon output. Traditionally, research considers carbon intensity from industrial production, regional economic development, and energy consumption structure. Standard metrics comprise carbon output per unit of GDP, industrial output, and energy consumption [7]. The digital economy has transitioned some carbon intensity studies to digital infrastructure such as cloud computing, communication systems, and data centers. Formulated studies suggest that carbon intensity of the digital sectors is influenced by the electricity asymmetry, infrastructural efficiency, and evolving technologies [8]. Considering the carbon intensity of generative artificial intelligence versus traditional manufacturing sectors reveals three distinctive traits. First, from the electricity systems and computing infrastructure, the carbon output may vary widely because of differences in electricity generation within regions. Another thing to note is that the carbon intensity of digital systems isn't steady—its wings up and down with workload changes, like those you see during model training and inference.

Some researchers measure AI-related environmental performance through indicators such as carbon emissions per training task, per unit of computation, or per inference workload [3,9–11]. As helpful as these indicators can be, many current approaches end up with a jumble of hard-to-interpret metrics that shift from one model and deployment scenario to the next. A growing number of studies now make a strong case for integrating data center efficiency metrics—especially Power Usage Effectiveness (PUE)—into AI carbon accounting [12]. In any data center, a big chunk of the total electricity goes to cooling, power conversion, and other auxiliary services. So, if we focus only on the processors and ignore these substantial auxiliary loads, we risk seriously underestimating the environmental footprint of large AI systems.

Taking all the above into consideration, the development of a more integrated carbon intensity measurement system for generative artificial intelligence remains an important and worthwhile research endeavor.

2.3. Research on Green Spillover Effects

The diffusion of generative AI across industries may create external effects beyond the original adopting sector, which is consistent with the broader innovation-diffusion perspective [13]. The term green spillover effect derives from spillover phenomena in the fields of innovation and environmental economics. It generally describes the situation where innovation in technology, upgrading within an industry, or environmental regulation in one area leads to improvements in the environmental performance of adjacent industries or areas [14]. Previous authors have mainly focused on the green spillover effects in manufacturing, energy generation, and innovation subsystems. From contemporary studies, it can be understood that the application of sophisticated technologies can lead to the reduction of energy abuse, enhance the efficiency of resource allocation, and facilitate the transformation of the industry towards low carbon in a more rapid manner through the established linkages and the subsequent diffusion of technology [15].

From the viewpoint of the digital economy, several studies have stated that digital technologies can enhance productivity as well as improve energy systems and management through smart scheduling, data analysis, and industrial integration [9]. An example of this is the intelligent system, which may end up optimizing the efficiency of the logistics networks and of electricity allocation, and as a result, may lead to the reduction of wastage in energy. Furthermore, it is likely that the digital transformation will result in the upgrading of industries and a higher concentration of low carbon industries in the regional economies. The positive effect of digital technologies on the environment may be limited. A number of authors have argued that the construction of digital technology infrastructures requires considerable amounts of electricity and hardware [16]. The creation of data storage centers, systems of communication, and the cloud may lead to an increased demand for electricity in the area where they are constructed. Therefore, the improvement of the digital technologies for the environment may be influenced by the efficiency of the technologies, the energy structure, and the environmental regulations.

Generative A.I. has the potential to create larger spillover effects for a wider array of industries further up the value chain as compared to previous digital innovations. Generative A.I. systems embedded within a value chain can streamline processes, eliminate redundancy, and optimize the entire chain for energy efficiency. Conversely, rapid deployment of large models can create an increased burden on the electricity supply and computing resources. Currently, only a handful of studies focus on the green spillover effects of Generative A.I. The majority of available literature focuses on productivity impacts, effects on the labor market, and impacts on digital transformation. Analysis of environmental spillover effects is still in its early stages

2.4. Theoretical Foundations

2.4.1. Sustainable Development Theory

At the heart of sustainable development theory is the connection among economic advancement, technology, and the preservation of the natural world. From a sustainable-development perspective, technological progress should be assessed in relation to both productivity gains and environmental consequences. Recent evidence on digital technology and green innovation provides an applied basis for this interpretation [17].

There is a possibility that generative artificial intelligence can create both positive and negative implications for the economy and the environment respectively. Positive implications include increased production efficiency and expedited advancement of different industrial sectors facilitated by the progressive nature of artificial intelligence. Negative implications arise from the rapid development of technology and the extensions of computing frameworks leading to increased consumption of electrical energy and greater levels of emitted carbon. Therefore, analyzing generative artificial intelligence from the prong of sustainable development theory in this regard is of great contemporary relevance.

2.4.2. Energy Efficiency Theory

Energy-efficiency analysis focuses on the relationship between useful output and energy input. Evidence from machine-learning training suggests that improvements in hardware, algorithms, and system design can alter this relationship substantially [18]

The use of generative artificial intelligence in the context of intelligent scheduling, predictive systems, and automated optimization can lead the way to enhanced energy efficiency. An example of this is the use of artificial intelligence based systems that can optimize the use of electrical resources which in turn improves the efficient utilization of resources and reduces idle time of equipment. However, the rapid increase in demand for computing resources can offset these efficiency improvements. For this reason, it is unclear whether the use of generative artificial intelligence will ultimately improve energy efficiency.

2.4.3. Technology Spillover Theory

Technology-spillover reasoning suggests that innovation may generate effects beyond the original adopting sector. Empirical research on AI and carbon intensity provides related evidence on heterogeneous and mediating pathways [19]. Through the interrelated industrial method, the mobility of labor, the diffusion of knowledge, and market competition, a series of consequences are produced.

The sharpest tools within generative artificial intelligence systems are those capable of deep penetration of various technologies and systems, as intelligent optimization grows and makes its way into manufacturing, logistics, energy systems, and public service systems. Related evidence also suggests that intelligent development may be associated with energy-conservation and emission-reduction outcomes under particular economic and institutional conditions [20].

2.4.4. ESG Governance Theory

ESG governance theory argues that organizations should integrate ESG aspects into their long-term strategic decisions. More recently, the inclusion of carbon emissions, climate-related risks, and green financing have all emerged as key focus areas within the ESG frameworks. There are concerns about the growing impact of the infrastructure that makes generative artificial intelligence possible, as well as the impact of energy consumption on climate change, and their inclusion in the ESG frameworks is becoming more of a reality. Therefore, analyzing the carbon footprint and environmental impact of generative artificial intelligence may be useful for sustainable investments and may inform the digital governance debate.

3. Carbon Intensity Measurement Framework for Generative Artificial Intelligence

3.1. Carbon Emission Boundaries of Generative Artificial Intelligence

3.1.1. Carbon Emissions Based on a Life-Cycle Perspective

The carbon emissions associated with generative artificial intelligence (AI) go beyond the training phase. The whole operation of large-scale AI systems is a series of process steps with a significant amount of electricity consumption that incurs indirect environmental costs. Current research increasingly suggests that limiting the assessment of the carbon emissions caused by AI systems to the electricity consumption of the training phase is likely to overlook the actual environmental costs of the digital infrastructure [21]. From a life-cycle assessment perspective, the carbon emissions boundaries of generative artificial intelligence encompass hardware production, model training, inference execution, digital infrastructure operation, and equipment disposal. Among these phases, training and inference are the most prominent and quantifiable phases of carbon emissions.

The first and significant phase of indirect emissions is hardware production. Large-scale AI systems require high-performance GPU, TPU, storage, server, and networking hardware. The production of these components, which include the fabrication of advanced semiconductor chips, is one of the most energy- intensive forms of manufacturing [22]. Although emissions from this phase are often difficult to quantify at the model level, they are part of the overall environmental cost of generative AI systems.

The second phase is model training, which existing studies suggest is the most electricity-intensive part of large-scale AI systems [23]. Distributed training requires thousands of high-performance processing units and must operate for long periods of time. Electricity consumption grows as the size of the model increases, the volume of training data grows, and the complexity of the optimization process increases.

The third stage consists of model inference and deployment. Except for training, inference in a single-task model is electrifying, but not to the same extent as training. Yet, inference services are subject to constant interaction with users and large-scale deployment. Unlike the first two forms of AI, generative AI is integrated into most modern applications. The cumulative electrical demand for extended inference use is likely to grow [10]. The operation of the framework is another important factor for indirect emissions. Large AI models require advanced cooling systems, electric transformation, backup power, and communication systems. Auxiliary systems to data centers increase total energy demand from computing resources [23].

The last stage encompasses replacement and recycling, which contribute the smallest share to total emissions, but still add environmental impacts due to electronic waste and hardware recycling.

In light of data availability and empirical feasibility, this paper's quantitative analysis focuses on the training and inference stages. Infrastructure operation will be considered within the extended carbon accounting framework.

3.1.2. Main Sources of Carbon Emissions

Generative AI produces carbon emissions in three major areas.

One major area is the direct electricity consumption of computing hardware. During the training and the actual inference of the model, GPUs, TPUs, CPUs, memory systems, and storage units all use a massive amount of energy. Typically, the demand for electricity grows with the demand for larger and more complex models, training tokens, and computation [24].

The second major area is the energy consumption of supporting systems. Data centers need cooling systems, and stable operations require energy-consuming systems as well as systems for the conversion of electricity. Existing studies have shown that the total consumption of electricity in AI-related infrastructure is often much larger than the IT load due to cooling and supporting systems [4].

Finally, the third major area is the indirect carbon emissions due to the generation of electricity.

There is a large variation regarding electricity and carbon emissions in various areas. Those regions that rely on coal-fired electricity tend to have a much higher carbon intensity of electricity compared to those regions where more renewable sources of electricity are used [5].

Hence, the demand for electricity for generative AI models, the efficiency of the infrastructure, and the energy mix in the region must all be evaluated to determine the carbon intensity.

3.2. Carbon Intensity Measurement Indicators

3.2.1. Total Carbon Emissions

Total carbon emissions are used to measure the overall environmental burden generated during the operation of generative artificial intelligence systems. Existing studies usually estimate AI-related carbon emissions by combining electricity consumption with regional electricity carbon emission factors [11].

In this study, total carbon emissions are calculated as: $CO2 = E_i \times EF_i$, where $CO2_i$ represents the total carbon emissions generated by model i, E_i denotes electricity consumption, and EF_i refers to the regional electricity carbon emission factor.

Plain electricity statistics don't say as much about environmental impact as total carbon emissions do, precisely because emissions capture the effect of a region's electricity structure.

3.2.2. AI Carbon Intensity Index (AICI)

Large models can vary a lot in parameter count and performance, so comparing them only by total carbon emissions doesn't really tell us much about their carbon efficiency. That's why we built the AI Carbon Intensity Index (AICI)it measures carbon emissions per unit of useful model output.

The AICI is computed as follows: $\mathrm{AICI}_i = E_i \times EF_i \,/\, Perf_i$ where $Perf_i$ represents the performance P_i captures the performance of model i, typically measured through task-appropriate metrics like accuracy, BLEU, or inverse perplexity.

A lower AICI basically means the model delivers stronger performance while keeping carbon emissions down. Unlike simple electricity-based metrics, the index does a better job of capturing the trade-off between technological performance and environmental cost.

3.2.3. Unit Parameter Carbon Intensity

To see how model expansion shapes environmental pressure, we built a unit parameter carbon intensity indicator: $PCI_i = CO2_i \,/\, ParamS_i$, where $ParamS_i$ is simply the number of parameters in model i.

We mainly use this indicator to see whether improvements in hardware efficiency and optimization algorithms can bring down per-parameter carbon emissions as models keep scaling up.

Research tells us that while total carbon emissions climb sharply as models get bigger, advances in technology can take some of the edge off per-parameter carbon intensity [25].

3.2.4. Carbon Reduction Efficiency Indicator

We also developed a carbon reduction efficiency metric, so we could compare how different technology choices and optimization strategies stack up.

$CRE_i = \frac{CO2_{baseline} - CO2_i}{CO2_{baseline}}$. Here, $CO2_{baseline}$ is the carbon emissions produced under benchmark conditions.

This metric captures the relative carbon savings you can get from optimization—whether through model compression, efficient scheduling, or switching to low-carbon electricity.

3.3. Data Sources and Variable Construction

3.3.1. Data Sources

We built our empirical analysis largely on public datasets covering AI energy consumption, data center operations, electricity carbon intensity, and regional energy structure. We drew most of our model-level electricity data from public datasets like the LLM Energy

Consumption Dataset and the BUTTER-E Deep Learning Energy Dataset. These sources tell us about training duration, hardware configuration, electricity consumption, and operating conditions for a range of representative AI models. We pulled data center operational indicators—Power Usage Effectiveness (PUE), Carbon Usage Effectiveness (CUE), and renewable electricity usage—from public reports and infrastructure datasets [26]. Data-center operational assumptions were also informed by established energy-use assessments [27].

For regional electricity carbon intensity and renewable energy structure, we leaned heavily on international energy statistics databases. From them, we pulled electricity carbon emission factors and the share of renewables across countries and regions. We also brought in macro-level variables on energy structure, carbon emissions, and digital development to help us later investigate green spillover effects.

The model-level sample is restricted to publicly documented models and experimental workloads for which at least part of the information on energy consumption, runtime, hardware configuration, or operational conditions is available. Proprietary frontier models with undisclosed training and inference configurations are not directly observed. The sample should therefore be regarded as a collection of publicly accessible cases rather than as a statistically representative sample of the entire commercial generative AI industry.

3.3.2. Variable Construction

We use the AI Carbon Intensity Index (AICI) as our dependent variable. Put simply, it captures the carbon emitted per unit of useful model output. We use model parameter scale, training token volume, GPU type, data center PUE, and renewable electricity share as our main explanatory variables. Bigger models and heavier training workloads tend to push electricity demand higher, while gains in infrastructure efficiency and renewable electricity usage can help bring carbon intensity down. We also brought in control variables mainly regional electricity carbon intensity, energy structure indicators, and digital infrastructure variables. We also built a few indicators—AI industry activity level, regional carbon intensity, and energy efficiency—to set the stage for the later green spillover analysis. We mainly use them to see whether the rise of generative AI can indirectly boost environmental performance by spurring industrial upgrading and improving energy efficiency.

4. Empirical Analysis of Carbon Intensity in Generative Artificial Intelligence

4.1. Data Processing and Descriptive Analysis

4.1.1. Data Processing

The first step of the empirical analysis is the standardization and merging of datasets. Since the original data was collected from various public sources, the level of consistency in the recording format and units varied. Electricity indicators were expressed in kilowatt-hours (kWh), whereas carbon emissions were expressed in kilograms of CO_2 equivalents ($kgCO_2e$). To reduce scale differences and the influence of highly skewed observations, selected continuous variables were transformed into logarithmic form. In datasets that capture operational power, electricity consumption is the product of power and operating time. To compute carbon emissions, electricity carbon emission factors for the different regions are imported into the estimation.

Some datasets were missing hardware configuration and/or operating time, so missing values were estimated by the meaning of the region or mean interpolation methods, depending on the data available for the region. Mean values for highly missing data points were also removed. Most of the different large-scale models of generative AI operate on different hardware architectures and in different electricity systems. So, the merged dataset, which is used in the analysis, is still largely heterogeneous. This is also true for the operating conditions of most generative AI systems.

4.1.2. Descriptive Statistical Results

The descriptive statistical analysis indicates substantial differences in electricity consumption and carbon intensity across different generative AI models. Models with larger parameter scales generally require much higher electricity input during training. But carbon intensity doesn't rise in lockstep with parameter growth. Improvements in optimization algorithms and hardware efficiency help soak up some of the extra electricity demand.

Carbon intensity can also shift a great deal from one deployment region to another. What we found is that when models run in regions with lower electricity carbon intensity, their total carbon emissions tend to be lower under comparable computing workloads, suggesting that the electricity structure shapes the environmental performance of generative AI systems [5,23].

Data center efficiency also varies quite a bit. Facilities with lower PUE values draw less electricity under the same computing load than traditional infrastructure—and earlier studies likewise point out that cooling systems and auxiliary facilities heavily shape total electricity demand in large scale computing setups [4,5,27].

Meanwhile, cumulative electricity demand from the inference stage keeps climbing. A single inference request may not add much carbon on its own, but the sheer volume of user interactions is steadily piling up the environmental burden tied to deployment services.

4.2. Carbon Intensity Measurement Results

4.2.1. Carbon Emissions During the Training Stage

Our results provide indicative evidence that training still dominates carbon emissions in generative AI. As models balloon from billions to trillions of parameters, electricity demand shoots up. That's largely because distributed training needs huge fleets of GPUs running nonstop for long stretches, which makes electricity consumption highly concentrated.

The available observations suggest that model training remains an important source of electricity demand in publicly documented AI workloads. Larger models generally require more computing resources, although the relationship between model scale and carbon emissions is not strictly proportional. Improvements in hardware efficiency, model compression, parameter pruning, and training optimization may reduce emissions per unit of computation, but these efficiency gains do not necessarily offset the increase in aggregate computing demand.

Still, the drop in unit carbon intensity can't fully make up for the surging total electricity demand as models keep getting larger. What we see overall is that total carbon emissions keep rising. Figure 1 provides descriptive background information on the carbon-emission patterns of the enterprise-level dataset used to characterize the broader energy and industrial environment. These patterns should not be interpreted as direct measurements of the training emissions of individual generative AI models.

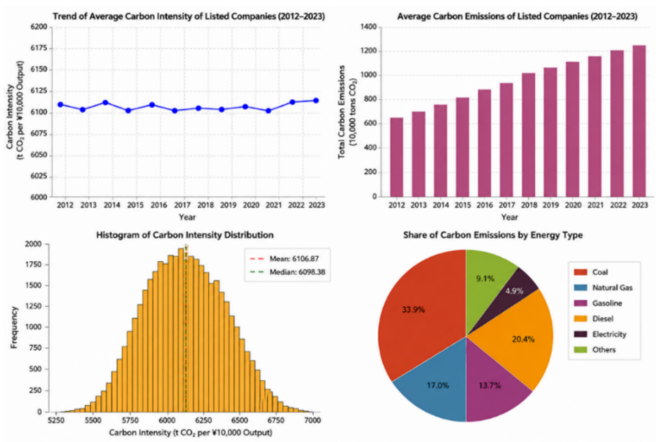

Figure 1. Descriptive carbon-emission patterns in the enterprise-level background dataset, 2012–2023.

4.2.2. Carbon Emissions during the Inference Stage

On a per-task basis, inference generally requires less electricity than model training. Nevertheless, the cumulative environmental burden of inference may become increasingly relevant as generative AI services are deployed across a growing number of applications. Improvements in hardware and algorithmic efficiency can reduce energy use per request, but a simultaneous increase in the number of users, requests, and generated tokens may offset part of these gains.

This pattern is consistent with a possible rebound effect: lower inference costs may support wider adoption, which in turn increases aggregate electricity demand. The available data do not permit a complete measurement of inference activity across proprietary commercial models. The results should therefore be interpreted as indicative of the direction of cumulative demand rather than as a direct estimate of industry-wide inference emissions. Figure 2 presents the descriptive relationships among carbon intensity, energy consumption, and total emissions in the broader background dataset. It illustrates the general association among these variables but does not constitute a direct measurement of inference-related emissions.

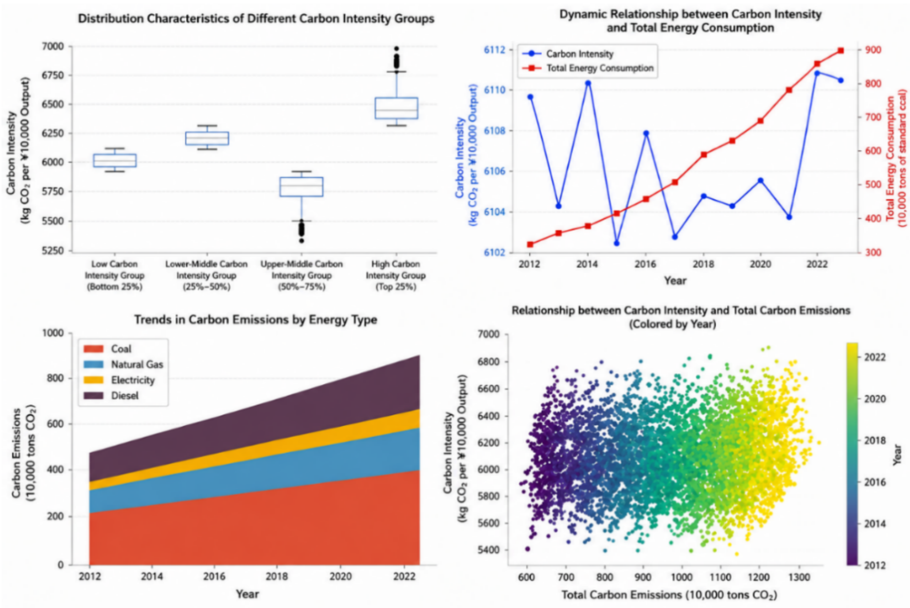

Figure 2. Descriptive relationships among carbon intensity, energy consumption, and total emissions in the background dataset.

4.2.3. Effects of Data Center Efficiency and Energy Structure

How efficiently data centers run really shapes the carbon intensity of generative AI. A lower PUE typically means you draw less electricity for the same computing workload.

What our data indicates is that data centers with advanced cooling and smarter power management draw far less electricity than conventional facilities. Other studies echo this, noting that auxiliary infrastructure often makes up a big chunk of total energy demand in large-scale computing systems.

The regional electricity mix also matters a lot for carbon intensity: models running where renewables make up a larger share tend to emit less carbon than those relying mostly on coal-fired power.

Our scenario simulations point to something clear: even without touching the model architecture or computational workload, shifting from conventional electricity to renewables can cut carbon intensity substantially. So, getting the electricity mix right remains a key way to lower the environmental footprint of generative AI.

4.3. Regression Analysis of Technological Progress Effects

4.3.1. Baseline Regression Results

We explored the link between technological advancement and carbon intensity by treating the AI Carbon Intensity Index (AICI) as the dependent variable in our regression models.

The estimated patterns suggest a positive association between model parameter scale and carbon intensity. It can be inferred that larger-scale models require larger amounts of power, resulting in greater carbon emissions.

Nevertheless, when variables concerning the efficiency of hardware and optimization techniques are included, the coefficients become smaller, illustrating that the burden of expanding models on the environment is alleviated by the progress of technology.

The volume of tokens also imposes a similar effect on carbon intensity, and the scale of the effect seems to be more modest. In this case, it can be concluded that the amount of electricity necessary for conducting the tasks is inversely proportional to the efficiency of the computations and the arrangement of the infrastructure.

Moreover, lower PUE values are associated with lower AICI estimates, which is consistent with the expected role of infrastructure efficiency in reducing operational electricity demand.

4.3.2. Robustness and Heterogeneity Analysis

Several alternative specifications were examined to assess whether the main patterns were sensitive to indicator transformations, extreme observations, or regression assumptions. First, logarithmic transformations of carbon intensity indicators were introduced to reduce the influence of extreme observations. Second, alternative carbon intensity indicators were used for comparison. Third, robust regression methods were applied to address potential heteroscedasticity problems. The main conclusions remain relatively stable under different model specifications.

Heterogeneity analysis further indicates that the estimated association between technological progress and carbon intensity appears to vary across electricity structures. In regions with higher renewable electricity penetration, improvements in computational efficiency produce stronger carbon reduction effects. But in regions that still lean heavily on fossil fuel power, the environmental gains from technological progress tend to be noticeably weaker.

Differences across infrastructure conditions are also relatively obvious. Regions with more advanced digital infrastructure and higher-quality data center systems generally exhibit lower carbon intensity under similar computing demand conditions.

These findings suggest that technological progress alone may not fully solve the carbon emission problem associated with generative artificial intelligence. The realization of low-carbon transformation still depends heavily on electricity structure optimization and green infrastructure development.

5. Green Spillover Effects of Generative Artificial Intelligence

5.1. Mechanisms of Green Spillover Effects

5.1.1. Technology Diffusion and Intelligent Optimization

There are more factors than the consumption of electricity used by AI technologies when evaluating the effect of generative AI on the environment. When large-scale AI technologies begin to penetrate more industries, their impacts begin to expand into production systems and resource allocation within industries. Although more evaluation is still needed, emerging theories posit digital technologies balance more efficient energy use, and better technology integration and distribution.

Generative AI is more adaptable and applicable across more industries than many digital technologies. These systems may optimize all aspects of production and resource utilization, including scheduling, demand, logistics, and even energy consumption, at the same time. Intelligent scheduling employed in many manufacturing industries may optimize production to reduce idle time and unnecessary consumption of energy. AI-based optimizations in logistics and transportation may achieve better utilization of fuel.

Thus, even though generative AI consumes a lot of electricity, the diffusion of the technology may help reduce energy consumption and improve the efficiency of resource utilization across many industries.

5.1.2. Industrial Upgrading and Structural Transformation

Generative artificial intelligence can alter environmental impact as a result of restructuring within an industry. Most of the literature reviewed posits that a digital service economy requires fewer resources or lower levels of traditional high-energy activities. Traditionally high-energy industries are transformed or phased out entirely.

Intelligent systems, computing, and AI production have expedited changes in the nature of economic activities. Under these conditions, some branches of production become more automated and digitalized, and more information-intensive production continues to grow.

This process inevitably has dual environmental impacts. Shorter-term focused production and increased demand for computing can elevate carbon emissions and lead to increased power consumption. However, in the longer term, the structural upgrading of industry can facilitate improvements to production that reduce environmental impacts sufficiently.

Longer-term structural and production changes within the industry can lead to waste reductions that outpace the environmental impacts from additional supporting structures. The impacts of generative AI on production and the environment are not mono-directional and can have complex and rapid changes.

5.1.3. Energy Efficiency Improvement

Energy efficiency gains offer another key pathway. What research tells us is that intelligent systems can make energy use more efficient—by monitoring in real time, scheduling adaptively, and optimizing based on predictions.

Generative AI can help make electricity allocation and energy system operations more efficient. Take AI-assisted smart grids, for instance: they can help integrate more renewables and cut the electricity wasted when supply and demand doesn't match intelligent building systems can tweak cooling, heating, and lighting on the fly, based on how they're actually being used. In industrial systems, AI-assisted optimization can help ramp up production efficiency while cutting down on energy that would otherwise go to waste.

Meanwhile, AI-driven optimization can help renewable electricity systems coordinate better, giving an indirect boost to the low-carbon energy transition.

5.2. Econometric Models of Green Spillover Effects

5.2.1. Baseline Model

To see whether generative AI has broader environmental effects, we set up a panel regression model with regional carbon intensity as the explained variable.

$$CFI_{it} = \alpha + \beta_1 AI_{it} + \beta_2 AI_{it}^2 + \beta_3 X_{it} + \mu_i + \lambda_t + \varepsilon_{it}$$

where CFI_{it} is regional carbon intensity, AI_{it} captures the development level of generative AI, AI_{it} picks up possible nonlinear effects, and X_{it} is a vector of control variables.

The quadratic term is introduced to explore whether the relationship between generative AI development and regional carbon intensity may be nonlinear. Given the limitations of the available regional indicators, the specification is interpreted as an exploratory test rather than definitive evidence of an inverted-U relationship. Existing research on digital technologies and environmental performance often finds that infrastructure expansion initially increases energy demand, while technological optimization may gradually improve environmental efficiency over time.

5.2.2. Mediation Effect Model

To identify the transmission paths of green spillover effects, mediation effect models are further introduced.

$$M_{it} = y_0 + y_1 AI_{it} + y_2 X_{it} + \mu_i + \lambda_t + V_{it}$$

$$CFI_{it} = \delta_0 + \delta_1 AI_{it} + \delta_2 M_{it} + \delta_3 X_{it} + \mu_i + \lambda_t + \eta_{it}$$

where M_{it} captures the mediating variables we focus on— mainly industrial upgrading and energy efficiency gains. We use these models to see whether generative *AI* indirectly shapes regional environmental performance by spurring structural transformation and efficiency improvements.

5.2.3. Heterogeneity Analysis

Generative AI's environmental impact isn't uniform—it can look quite different across industries, enterprises, and regions. So we approached heterogeneity from a few angles. We started with enterprise scale. Large firms tend to have stronger digital infrastructure and deeper pockets for tech investment, which can amplify green spillover effects. Then we looked at industry characteristics. Technology-intensive industries usually get more of an environmental boost from AI-assisted optimization than traditional energy-intensive ones. Finally, we examined differences in regional electricity structure.

Regions where renewables make up a larger share of electricity tend to see stronger environmental gains, simply because lower carbon intensity per unit of electricity eases the environmental burden tied to expanding computing.

5.3. Empirical Results of Green Spillover Effects

5.3.1. Indicative Nonlinear Patterns

The estimated results suggest a potential nonlinear pattern between generative AI development and regional carbon intensity. During the initial expansion of computing infrastructure, additional electricity demand from data centers, computing systems, and model deployment may place upward pressure on regional carbon intensity. At more advanced stages of digital and technological development, intelligent optimization, industrial upgrading, and improvements in energy management may begin to offset part of this pressure.

This pattern is consistent with an inverted-U interpretation, but it should not be treated as conclusive evidence of a universal turning point. The direction and magnitude of the relationship are likely to depend on regional electricity structures, infrastructure efficiency, the measurement of AI development, and other institutional conditions. The estimates therefore provide indicative evidence of possible nonlinearity rather than establishing a deterministic transition from carbon increase to carbon reduction.

5.3.2. Mediation Effect Results

The mechanism analysis points to industrial upgrading and energy-efficiency improvement as two possible transmission channels. Generative AI may support industrial restructuring by facilitating digital transformation and the expansion of technology-intensive activities. AI-assisted optimization may also improve electricity allocation and reduce avoidable energy use in some production settings.

These mechanisms provide a plausible explanation for the potential green spillover effects discussed in this study. However, the available evidence does not permit a precise decomposition of the indirect effect or a definitive causal interpretation. Industrial upgrading and energy efficiency should therefore be understood as theoretically and empirically plausible channels rather than as conclusively established mediators.

5.3.3. Heterogeneity Results

The heterogeneity analysis identifies pronounced cross-regional and cross-industry variation.

The more pronounced green spillover effects in large enterprises stem from their advanced digital infrastructure and greater capacity for technological investment. Technology-intensive industries similarly exhibit larger environmental optimization effects than their traditional energy-intensive counterparts.

The carbon reduction effects associated with generative AI are significantly moderated by regional electricity structures. In regions where renewable electricity accounts for a larger share, these effects are more pronounced, whereas in fossil-fuel dependent regions, the gains from computational efficiency improvements are substantially offset by high carbon electricity intensity.

The environmental performance of generative artificial intelligence is significantly shaped by infrastructure quality. Regions equipped with more advanced data center systems and digital infrastructure tend to register lower carbon intensity under comparable computing demand.

The environmental impacts of generative artificial intelligence are shaped by electricity systems, infrastructure conditions, and industrial structures rather than being determined solely by the AI technologies themselves. As shown in Figure 3, both AI-related coefficients and grid emission factors vary considerably across regions.

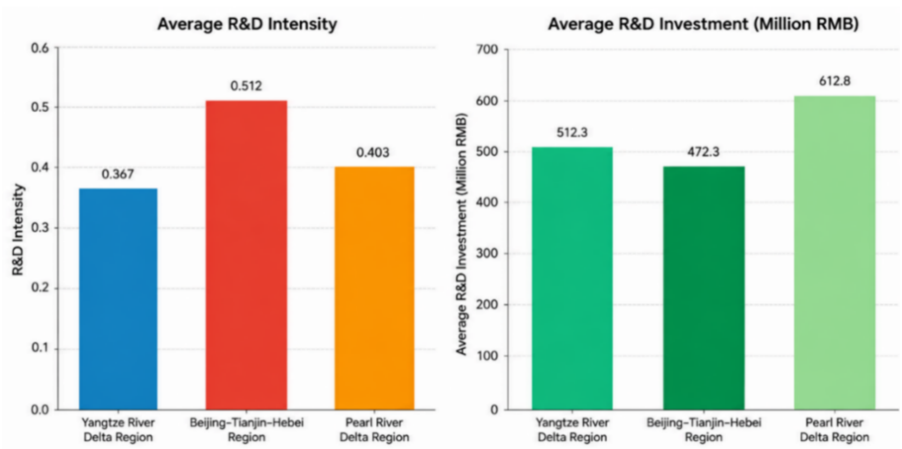

Figure 3. Differences in AI coefficients and grid emission factors across regions.

6. Discussion

6.1. Carbon Intensity Characteristics of Generative Artificial Intelligence

Empirical evidence suggests that the carbon intensity of generative AI models is jointly determined by the efficiency of the underlying hardware and algorithms and the carbon intensity of the available electricity supply. In the context of large-scale models, training computation drives up electricity consumption, and although technological progress can reduce carbon intensity per unit of computation, the total carbon emissions remain substantial. This pattern aligns with the literature on digital economy and energy consumption. Early research portrayed AI as a relatively low-carbon digital technology, given that its outputs were largely confined to the digital sphere. The emergence of large models, however, has fundamentally altered this assessment, as the environmental consequences of massive, distributed computing infrastructures have become unmistakable.

Concurrently, the evidence indicates sustained gains in computational efficiency stemming from technological advancement. Newer hardware architectures and optimization algorithms lower electricity consumption per computation task, while model compression techniques partially alleviate the environmental burden imposed by parameter expansion. The relationship between model scale and carbon intensity is, therefore, not strictly linear.

The inference stage presents a further critical characteristic: per-request electricity consumption remains relatively low, yet cumulative electricity demand escalates with large-scale user interaction and the proliferation of deployment scenarios. Prior research identifies a potential rebound effect in which efficiency improvements stimulate broader application demand, thereby sustaining upward pressure on total consumption.

Therefore, evaluating the environmental impacts of generative artificial intelligence only from the perspective of training emissions may underestimate the long-term environmental pressure associated with deployment and inference systems.

6.2. Green Spillover Effects and Environmental Transformation

The analysis suggests that generative AI may generate broader green spillover effects under certain conditions. The estimated pattern is consistent with a potentially nonlinear relationship between AI development and regional carbon intensity. In the early stages of infrastructure expansion, additional electricity demand may increase environmental pressure. As AI-enabled optimization and industrial upgrading deepen, part of this pressure may be offset by efficiency gains. This interpretation remains conditional and should not be understood as evidence that AI development automatically produces a decline in regional carbon intensity.

What this indicates is that the negative externalities of generative AI are not static. The majority of current literature focuses on the indirect effects of the AI systems on electricity use. However, the negative impact of computing infrastructure may be offset by the broader economic effects of smart optimization.

Industrial upgrading and energy-efficiency improvement emerge as two plausible mechanisms through which green spillover effects may occur. AI systems can optimize production scheduling, the allocation of electricity, and the coordination of operations, reducing waste of resources across the entire industrial system. Simultaneously, the digital transformation of the economy speeds the development of technology- intensive industries, and reduces the reliance on traditional, high-energy-consumption production systems.

That said, the positive green spillover effects of generative AI do not occur by design. Their occurrence is also very much dependent on the electricity structure, the quality of infrastructure, and the degree of environmental regulations in place. In locations where renewable electricity is present to a higher degree, the use of AI may demonstrate a greater ability to decrease carbon emissions, as low- carbon electricity can alleviate the negative environmental effects of the increased use of computing.

Conversely, in areas that continue to depend significantly on electricity generated from fossil fuels, the environmental impact of the rapidly growing demand for electricity may not be balanced out by progress in computational efficiency.

6.3. Implications for ESG Governance and Green Investment

The rapid growth of generative AI creates problems for ESG governance and sustainable investment. Current ESG evaluation frameworks focus on manufacturing, transport, and energy. Recently, the environmental costs of the computing and digital platform industries are also beginning to attract assessments. AI’s environmental performance is shown to be beyond the evaluation of the capability of technology and the size of the market. The structure of the electricity, the efficiency of the data centers, and the operational efficiency of the AI infrastructures are all critical to the generation systems of AI.

Therefore, metrics for AI infrastructure-related emissions will begin to be included in ESG assessments. Compared to overall emissions, the AI Carbon Intensity Index (AICI) proposed in this study strikes a more appropriate balance between the environmental costs and the technological outcome.

Given the green spillovers, investments in a sustainable low-carbon AI infrastructure may yield more environmental benefits beyond the immediate financial returns. Data centers with low carbon emissions and coupled with renewable electricity, as well as low-carbon computing, will be avenues for impactful green investments.

The market for AI systems remains volatile, with uncertainty regarding their environmental costs. Some firms may focus on "green AI" while still being highly carbon intensive. For ESG governance to evolve, stricter carbon disclosure policies and stable assessments of AI systems are needed.

6.4. Limitations and Future Research Directions

We've put together a fairly complete framework for analyzing the carbon intensity and green spillover effects of generative AI, but it still has its limits.

For starters, we still must estimate some model-level electricity consumption data indirectly, simply because the complete operational records for frontier models aren't fully public. As a result, our carbon accounting may not be as precise as we'd like, simply because the necessary data aren't always publicly available.

Another limitation is that we mainly focus on the training and inference stages, while hardware manufacturing and recycling are still discussed in a more qualitative way. Bringing in more granular data on semiconductor manufacturing and infrastructure operations would help make lifecycle carbon accounting more accurate down the line.

The measurement of AI development levels and—green spillover effects—remains partially dependent on proxy variables, as the rapid evolution of the generative AI industry introduces temporal instability in the underlying indicators.

The present analysis is limited to macro-level environmental impacts. Enterprise-level behavioral mechanisms, regional policy-heterogeneity, and cross-national comparisons of the environmental governance of generative AI systems represent important directions for future research.

7. Policy Implications

7.1. Toward a Standardized Carbon Accounting Framework for Generative Artificial Intelligence

Empirical evidence demonstrates that the carbon emissions of generative AI systems are strongly correlated with model scale, infrastructure efficiency, and the electricity structure. China's national computing-hub policy also emphasizes the coordination of computing infrastructure, energy resources, and regional deployment [28]. However, substantial variation in accounting boundaries and estimation methods across existing studies impedes direct comparisons of empirical results across regions and enterprises.

The divergence in accounting practices underscores the urgency of developing a standardized carbon accounting framework tailored to generative AI. Existing carbon accounting systems remain predominantly oriented toward traditional industrial sectors, while the environmental impacts of computing infrastructure and AI deployment have received comparatively limited scrutiny.

A comprehensive carbon accounting framework for generative AI should integrate model training, inference deployment, data center operations, and electricity structure into a single analytical architecture. Integrated carbon intensity indicators provide a more accurate representation of the actual environmental burden of AI systems than electricity consumption data alone.

Strengthening carbon disclosure mechanisms can mitigate uncertainty in ESG evaluations and green investment decisions. Large-scale AI enterprises and cloud computing platforms are expected to incorporate indicators of electricity consumption, renewable electricity utilization, and infrastructure efficiency into future sustainability reports.

7.2. Optimizing the Spatial Distribution of Computing Infrastructure

Empirical evidence identifies electricity structure as an observable determinant of carbon intensity in generative AI systems. Under equivalent computing demand, a higher renewable electricity share is associated with lower carbon intensity.

The environmental performance of an AI system is determined not only by its computational efficiency, but also substantially by the electricity supply structure of the region in which it operates.

Under this condition, the spatial arrangement of computing infrastructure becomes increasingly important. In regions where renewable electricity resources are relatively abundant, deploying large-scale computing facilities may reduce carbon intensity more effectively. By contrast, if AI infrastructure expansion mainly relies on fossil-fuel-based electricity systems, improvements in computational efficiency may not fully offset the environmental pressure generated by increasing electricity demand.

Therefore, future computing infrastructure planning may need to coordinate more closely with renewable electricity development and regional energy transition strategies.

At the same time, improving data center operational efficiency remains an important direction for reducing carbon intensity. Existing studies generally suggest that cooling systems and auxiliary facilities substantially affect total electricity demand in large-scale computing systems. Therefore, reducing PUE levels and improving infrastructure efficiency may further reduce associated with AI expansion.

7.3. Encouraging Low-Carbon Technological Optimization

The empirical results further indicate that technological progress partially reduces the carbon intensity associated with model expansion. Improvements in hardware systems, optimization algorithms, and model compression technologies all contribute to reducing electricity consumption precomputation task.

Efficiency gains alone, however, prove insufficient to fully address the carbon emission challenges posed by generative artificial intelligence. Declining computational costs accelerate the proliferation of application scenarios, generating a rebound effect in electricity demand.

Consequently, future technological optimization must extend beyond model performance to prioritize low-carbon operational efficiency. Lightweight models, parameter compression, adaptive scheduling systems, and energy-efficient hardware architectures thus represent essential trajectories for sustainable AI development.

Moreover, intelligent electricity management systems enhance the coordination efficiency between AI infrastructure and renewable electricity systems. So, AI-assisted optimization could end up being part of a bigger push toward low-carbon energy.

7.4. Strengthening ESG Governance and Green Investment Guidance

The rapid expansion of generative AI is also starting to reshape the environmental governance of digital industries. In the ESG world, the focus has long been on manufacturing, transportation, and conventional energy systems, while the environmental footprint of digital infrastructure has drawn far less attention.

What our findings tell us is that an AI system’s environmental performance can’t be assessed just by looking at market scale or technological capability. Infrastructure efficiency, electricity structure, and carbon management systems also go a long way in shaping actual carbon intensity.

Given this, carbon intensity indicators for generative AI could well find their way into future ESG governance systems. Companies running large-scale AI infrastructure are likely to face growing pressure to disclose carbon data and undergo sustainability evaluations

Meanwhile, green investment is likely to flow more toward low-carbon computing infrastructure, better integration of renewables, and efficient data-center technologies. Compared with traditional high-energy-consumption industries, low-carbon AI infrastructure may gradually become an emerging direction in sustainable investment markets.

Market-based mechanisms continue to face substantial uncertainty in the environmental assessment of AI technologies. Some enterprises may employ “green AI” narratives without achieving corresponding reductions in infrastructure-related carbon emissions. Advancing environmental disclosure standards and strengthening carbon supervision therefore represent important priorities for the next stage of digital governance. The possibility that environmental regulation and efficiency-oriented innovation may reinforce one another is broadly consistent with the environment–competitiveness perspective [29].

7.5. Strengthening International Cooperation on Low-Carbon AI Governance

The environmental toll of generative AI doesn’t stop at national or regional borders. Large-scale AI systems usually run on global supply chains, cross-border cloud services, and international electricity- intensive infrastructure networks.

This means that keeping the carbon footprint of generative AI in check will likely call for broader international coordination. Carbon accounting standards and ESG disclosure rules still vary a lot from country to country, and that only adds to the uncertainty in global environmental governance.

Going forward, international cooperation will likely center on a few core areas: unified carbon accounting standards, renewable electricity coordination mechanisms, low-carbon data center technologies, and sustainable computing infrastructure.

On top of that, international cooperation on low-carbon electricity systems and green digital infrastructure could help digital transformation and carbon neutrality goals move forward together.

8. Conclusions

As generative AI expands at breakneck speed, it's beginning to reshape the relationship between digital technology, electricity use, and carbon emissions. Unlike earlier digital systems, large-scale AI models need far more computing power and place a much heavier load on infrastructure. These dynamics render the environmental footprint of generative artificial intelligence increasingly difficult to overlook.

By integrating model-level electricity consumption data, data center operational indicators, and regional electricity mix data, we construct a carbon intensity measurement framework for generative artificial intelligence. We then examine its green spillover effects through the lens of energy economics and sustainable development.

What our findings show is that the carbon intensity of generative AI hinges heavily on three things: model scale, infrastructure efficiency, and the electricity mix. Training still accounts for most of the carbon emissions, simply because large-scale distributed computing is so electricity hungry. Although optimization algorithms and hardware efficiency improvements reduce unit carbon intensity to some extent, total carbon emissions still continue to rise together with increasing computational demand.

At the same time, the inference stage is becoming increasingly important in the overall carbon emission structure of generative artificial intelligence. Single-task inference generally consumes relatively little electricity, but rapidly expanding deployment scenarios and user interaction frequency continuously increase cumulative electricity demand. Under this condition, evaluating only training- related emissions may underestimate the long-term environmental burden associated with AI deployment systems

The analysis further indicates that the environmental impacts of generative artificial intelligence are not completely negative. The regression results show a relationship between generative artificial intelligence and regional carbon intensity is consistent with a possible transition from initial carbon pressure to conditional efficiency gains. In the early stage of infrastructure expansion, increasing electricity demand tends to increase carbon emissions. However, as technological penetration deepens, intelligent optimization, industrial upgrading, and energy efficiency improvement gradually weaken the environmental pressure associated with computing expansion.

The mechanism analysis suggests that industrial upgrading and energy-efficiency improvement may contribute to the indirect environmental effects of generative AI. AI-assisted systems enhance production scheduling, electricity allocation, and operational coordination across broader industrial systems, thereby reducing energy consumption that would otherwise be wasted.

However, the environmental benefits derived from generative AI remain highly contingent on the electricity structure and infrastructure conditions. In regions where renewables account for a larger proportion of electricity, carbon reduction effects are more pronounced, because a low-carbon electricity mix mitigates the environmental burden imposed by computing demand. But in regions that still lean heavily on fossil-fuel electricity, just making computing more efficient may not be enough to offset the rise in electricity use.

This tells us that the environmental performance of generative AI can't be assessed just by looking at model capability or computational efficiency. The actual carbon intensity of AI systems is driven to a large extent by electricity systems, data center operations, renewable electricity integration, and infrastructure governance.

What our findings point to is that future low-carbon governance of generative AI will likely need to center on a few key areas: better carbon accounting systems, smarter spatial distribution of computing infrastructure, tighter integration of renewables, more low-carbon technology optimization, and stronger ESG disclosure mechanisms.

Meanwhile, international cooperation is likely to matter more, simply because large-scale AI systems lean so heavily on global infrastructure networks and cross-border digital services. So, going forward, governance frameworks are likely to call for more coordinated standards on carbon accounting, green infrastructure, and low-carbon computing.

Our study still has its limits, of course. For one thing, some model-level electricity data still come from indirect estimates, simply because full operational records aren't publicly available. For another, we couldn't look closely at hardware manufacturing and equipment recycling—the data just isn't available. Richer data on infrastructure and semiconductor manufacturing could make life-style carbon accounting more accurate down the line. In short, the environmental footprint of generative AI doesn't fit neatly into a simple "high-carbon" or "low-carbon" box. Electricity demand significantly, intelligent optimization and industrial upgrading may gradually generate broader green spillover effects under appropriate infrastructure and electricity conditions.

Author Contributions

X.G. designed the study framework, ran the empirical analysis, processed the data, and wrote the first draft of the manuscript. R.C. helped with manuscript review, polished the language, and contributed to interpreting the results. All authors have read and agreed to the published version of the manuscript.

Funding

This research was funded by the 2025 “Four New” Research and Reform Practice Program of Suzhou University (project: “Research on the Practical Path of Integrating Labor Education into ‘College Students’ Career Planning and Development’ from the Perspective of New Quality Productive Forces,” grant no. szxy2025xwjy02) and by the 2025 National College Students Innovation and Entrepreneurship Training Program (grant nos. 202510379028 and 202510379032).

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

This study draws on multiple publicly available datasets covering AI-related energy consumption, data-center infrastructure, electricity systems, and carbon emissions. The model-level observations are limited to publicly documented models and experimental workloads. Proprietary commercial models for which training configurations, inference volumes, or energy-consumption data are not disclosed are not directly included. The resulting dataset therefore supports comparative and scenario-based analysis but should not be interpreted as a complete representation of the global commercial generative AI industry. The processed data and analysis code may be provided by the corresponding author upon reasonable request.

Conflicts of Interest

The authors declare no conflict of interest.

Use of AI and AI-Assisted Technologies

No AI tools were utilized for this paper.

References

1. Goldfarb, A.; Tucker, C. Digital Economics. *J. Econ. Lit.* **2019**, *57*, 3–43.
2. Strubell, E.; Ganesh, A.; McCallum, A. Energy and Policy Considerations for Deep Learning in NLP. *arXiv* **2019**, arXiv:1906.02243.
3. Henderson, P.; Hu, J.; Romoff, J.; et al. Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning. *J. Mach. Learn. Res.* **2020**, *21*, 1–43.
4. Jones, N. How to Stop Data Centres from Gobbling Up the World’s Electricity. *Nature* **2018**, *561*, 163–166.
5. Masanet, E.; Shehabi, A.; Lei, N.; et al. Recalibrating Global Data Center Energy-Use Estimates. *Science* **2020**, *367*, 984–986.
6. Luccioni, A.S.; Viguier, S.; Ligozat, A.-L. Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model. *J. Mach. Learn. Res.* **2023**, *24*, 1–15.
7. Raupach, M.R.; Marland, G.; Ciais, P.; et al. Global and Regional Drivers of Accelerating CO_2 Emissions. *Proc. Natl. Acad. Sci. USA* **2007**, *104*, 10288–10293. https://doi.org/10.1073/pnas.0700609104.
8. Li, J.G.; Gong, F.M.; Wang, Y.B. How Does Artificial Intelligence Technology Innovation Affect Carbon Emission Reduction in Manufacturing? *China Popul. Resour. Environ.* **2025**, *35*, 14–24. (In Chinese)
9. Lacoste, A.; Luccioni, A.; Schmidt, V.; et al. Quantifying the Carbon Footprint of Machine Learning. *arXiv* **2019**, arXiv:1910.09700.
10. Budennyy, S.; Lazarev, V.; Zakharenko, N.; et al. Eco2AI: Carbon Emissions Tracking of Machine Learning Models as the First Step towards Sustainable AI. *arXiv* **2022**, arXiv:2208.00406.
11. Luccioni, A.S.; Jernite, Y.; Strubell, E. Power Hungry Processing: Watts Driving the Cost of AI Deployment. *arXiv* **2023**, arXiv:2311.16863.

12. Sun, Z.Q.; Yang, R. Research on the Impact of AI Technology Innovation on Regional Carbon Emissions: Mechanism Identification and Rebound Effect. *Sci. Technol. Manag. Res.* **2024**, *44*, 168–177. (In Chinese)
13. Rogers, E.M. *Diffusion of Innovations*, 5th ed.; Free Press: New York, NY, USA, 2003.
14. Wang, X.W.; Gong, Y.L.; Xie, T.P. Mechanism and Spatial Effects of Artificial Intelligence on Provincial Carbon Emissions. *Enterp. Econ.* **2025**, *44*, 140–150. (In Chinese)
15. Li, L. The *Environmental* Spillovers of Buyers' Digital Transformation: Evidence from China. *Technol. Forecast. Soc. Chang*. **2024**, *209*, 123828. https://doi.org/10.1016/j.techfore.2024.123828.
16. Liu, S.; Cao, Y.J. Digital Carbon Neutralization: An Empirical Study on the Dual Effects of Artificial #br# Intelligence in Promoting Industrial Carbon Reduction#br. *J. Jiangsu Univ.* **2025**, *27*, 26–39. (In Chinese)
17. Mensah, J. Sustainable Development: Meaning, History, Principles, Pillars, and Implications for Human Action—Literature Review. *Cogent Soc. Sci*. **2019**, *5*, 1653531. https://doi.org/10.1080/23311886.2019.1653531.
18. Patterson, D.; Gonzalez, J.; Van Der Merwe, J. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. *Nat. Mach. Intell.* **2021**, *3*, 875–876.
19. Zhao, P.; Gao, Y.; Wu, M.; et al. How Artificial Intelligence Affects Carbon Intensity: Heterogeneous and Mediating Analyses. *Environ. Dev. Sustain.* **2024**, *28*, 2301–2325.
20. Cao, Z.; Ding, Z.J. Energy Conservation and Emission Reduction Effects of Intelligent Development: Theoretical and Empirical Tests. *Stat. Decis.* **2023**, *39*, 174–178. (In Chinese)
21. Kaack, L.H.; Donti, P.L.; Strubell, E.; et al. Aligning Artificial Intelligence with Climate Change Mitigation. *Nat. Clim. Chang*. **2022**, *12*, 518–527. https://doi.org/10.1038/s41558-022-01377-7.
22. Schwartz, R.; Dodge, J.; Smith, N.A.; et al. Green AI. *Commun. ACM* **2020**, *63*, 54–63.
23. de Vries, A. The Growing Energy Footprint of Artificial Intelligence. *Joule* **2023**, *7*, 2191–2194.
24. de Vries-Gao, A. The Carbon and Water Footprints of Data Centers and What This Could Mean for Artificial Intelligence. *Patterns* **2026**, *7*, 101430. https://doi.org/10.1016/j.patter.2025.101430.
25. Vercellino, R.; Willard, J.; Campos, G.; et al. Measurement of Generative AI Workload Power Profiles for Whole-Facility Data Center Infrastructure Planning. *arXiv* **2026**, arXiv:2604.07345.
26. Elsworth, C.; Huang, K.; Patterson, D.; et al. Measuring the Environmental Impact of Delivering AI at Scale. *arXiv* **2025**, arXiv:2508.15734.
27. Shehabi, A.; Smith, S.; Sartor, D.; et al. *United States Data Center Energy Usage Report*; Lawrence Berkeley National Laboratory: Berkeley, CA, USA, 2016.
28. National Development and Reform Commission; Cyberspace Administration of China; Ministry of Industry and Information Technology; National Energy Administration. Implementation Plan for the Computing Hubs of the National Integrated Big Data Center Collaborative Innovation System; FGGT [2021] No. 709; Beijing, China, 2021. Available online: https://www.ndrc.gov.cn/xxgk/zcfb/tz/202105/t20210526_1280838.html (accessed on 25 June 2026).
29. Porter, M.E.; van der Linde, C. Toward a New Conception of the Environment-Competitiveness Relationship. *J. Econ. Perspect.* **1995**, *9*, 97–118.