

Article

The Environmental Blind Spot of AI Policy: Energy, Infrastructure, and the Systematic Externalization of Sustainability

Carlos García-Llorente ¹ and Ignacio Olmeda ^{1,2,*}¹ RSI Chair in AI in Banking, University of Alcalá, 28801 Madrid, Spain; cagarc37@ucm.es² Department of Computer Science, University of Alcalá, 28801 Madrid, Spain

* Correspondence: josei.olmeda@uah.es

Abstract

Contemporary artificial intelligence policies systematically externalize environmental costs. Despite divergent governance models, the European Union, the United States, and China converge on the same outcome: none impose binding restrictions on the energy intensity, carbon footprint, or infrastructural expansion of AI systems. This article demonstrates that sustainability is treated as an externality, rather than as a mandatory regulatory constraint, in all major jurisdictions. Focusing on energy consumption, computational infrastructure, and carbon budgets, the analysis shows that current AI policy choices generate predictable patterns of environmental omission and cost externalization. Policy measures aimed at strengthening rights protection and technological autonomy—such as tightening compliance requirements, developing large-scale models, and duplicating infrastructure—are adopted without corresponding limits on energy use or emissions, generating growing tensions with planetary constraints. This article makes three contributions to the literature on AI governance and sustainability. First, it conceptualizes sustainability as a binding material constraint, rather than as a normative objective or efficiency-based goal. Second, through a comparative policy analysis, it shows that despite divergent regulatory styles, the European Union, the United States, and China converge in the absence of enforceable environmental limits applicable to AI systems. Third, it identifies the policy mechanisms—compliance-driven computational expansion, infrastructure duplication, and scale-oriented incentives—that systematically generate environmental externalization across jurisdictions. The article concludes that effective AI policy requires recognizing sustainability as a hard material limit, translated into binding environmental restrictions that condition regulatory design, infrastructure planning, and the permissible scale of computational systems.



Academic Editors: Rohail Hassan and Costinela Fortea

Received: 6 January 2026

Revised: 1 February 2026

Accepted: 2 February 2026

Published: 5 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: artificial intelligence policy; environmental sustainability; energy consumption; computational infrastructure; carbon emissions; AI governance; digital sustainability

1. Introduction

Contemporary approaches to artificial intelligence (AI) policy increasingly recognize environmental sustainability, but tend to treat it as a secondary consideration. This consideration is often accommodated alongside other policy objectives through efficiency improvements, voluntary commitments, or future technological advances. This article challenges that premise by arguing that, under current material and geopolitical conditions, sustainability is not integrated into the design of AI policies, but is systematically

externalized. Energy consumption, carbon budgets, and infrastructure expansion are treated as factors exogenous to AI policy decisions, rather than as hard material constraints that should condition their deployment and scale. This paper addresses the following research question: to what extent do the main AI policy regimes internalize environmental sustainability as a binding constraint on the scale, location, and energy intensity of AI system deployment?

Answering this question requires situating AI policy within the material conditions that make large-scale AI deployment possible. The rapid expansion of AI systems is inseparable from the growth of energy-intensive computational infrastructures, including large data centers, high-performance computing facilities, and global cloud computing networks, which place increasing pressure on electrical systems, water resources, land use, and carbon budgets, especially in contexts where renewable capacity and grid flexibility are limited [1–3]. This material dependence also exposes a broader methodological gap in how the environmental impacts of AI are assessed. As [4] demonstrate, evaluating AI-based solutions requires going beyond energy consumption during the use phase and incorporating the full life cycle of computing equipment, including production, transportation, and end-of-life stages. However, public policy frameworks rarely require such assessments, treating infrastructure expansion and hardware production as elements external to regulatory decisions on the deployment and scale of AI systems.

This pattern is structural rather than accidental. Measures aimed at strengthening rights protection and technological autonomy, such as tightening compliance requirements, developing large-scale models, and duplicating computing infrastructure, are adopted without corresponding limits on energy use, emissions, or absolute resource consumption. As a result, AI policies systematically generate environmental omissions that translate into predictable patterns of cost externalization on energy systems and ecosystems [5,6].

Through a comparative analysis of the European Union, the United States, and China, the article shows that, despite substantial differences in regulatory styles and strategic priorities, the three jurisdictions converge in practice on the same outcome: the absence of binding environmental restrictions applicable to AI systems, including enforceable limits on energy intensity, carbon footprint, or the scale of the computational infrastructure associated with AI deployment. Sustainability is repeatedly invoked in political discourse, but remains weakly operationalized in regulatory instruments [7].

The central contribution of this work is threefold. First, it conceptualizes sustainability as a set of binding material constraints rather than as a normative goal that can be balanced; second, it shows, through a comparative policy analysis, that despite divergent regulatory traditions, the European Union, the United States, and China converge in the absence of specific enforceable environmental limits for AI; and third, it identifies the policy mechanisms that generate this systematic externalization.

The article concludes that addressing the environmental impacts of AI requires a change of perspective in public policy design. Rather than treating sustainability as just another objective to be balanced or integrated alongside others, AI policy must recognize it as a hard material constraint that conditions infrastructure planning, energy allocation, and the permissible scale of computational systems. Without this change, AI policies will continue to externalize environmental costs, compromising the compatibility of digital transformation with transitions toward sustainability. At the same time, it should be recognized that artificial intelligence can contribute to reducing emissions and increasing efficiency in specific contexts. Numerous studies document AI applications aimed at energy optimization, smart management of electrical grids, reducing consumption in industrial processes, and improving logistics and transportation efficiency. There are also regulatory and corporate initiatives aimed at improving the efficiency of data centers, promoting the

use of renewable energy in digital infrastructure, and encouraging more sustainable energy procurement practices, although these are largely voluntary in nature. These developments show that AI is not inherently incompatible with climate goals and that, under certain conditions, it can generate sectoral environmental benefits.

Crucially, the central argument of this article does not deny these potential benefits, but rather points to a structural gap between these specific applications and the overall design of AI policy. Even when positive optimization effects are recognized, current regulatory frameworks do not translate these benefits into binding restrictions on the scale, location, or energy intensity of AI system deployment. In the absence of such limits, efficiency gains tend to be outweighed by the aggregate growth in computational demand, reproducing patterns of environmental externalization. The problem, therefore, does not lie in the lack of environmentally beneficial applications, but in the lack of governance mechanisms that condition the scaling of AI to the material limits of the energy and climate system.

2. Environmental Externalization and Material Constraints in AI Policy

This section develops a material approach to contemporary AI policies, focusing on how environmental costs are systematically externalized. Rather than treating sustainability as a normative goal to be balanced against others, the analysis conceptualizes it as a set of binding physical constraints—energy availability, carbon budgets, infrastructure capacity, and material limits—that condition the feasible scale of AI systems. The central argument is that current AI policy options fail to internalize these constraints, generating predictable patterns of environmental omission across jurisdictions.

For analytical clarity, the following key terms are used throughout this paper: ‘AI systems’ refers to both large-scale machine learning models (including foundational models such as GPT, Claude, Gemini, or Llama) and AI applications that require significant computational capacity for training or deployment; ‘computing infrastructure’ refers to the set of data centers, high-performance computing facilities, and cloud computing networks that support the training and operation of AI systems; ‘energy intensity’ measures energy consumption per unit of computation or per task performed; and ‘carbon footprint’ comprises greenhouse gas emissions associated with both the training phase and the inference of models, as well as the production and end-of-life of computing hardware.

2.1. Sustainability as a Material Constraint

The environmental footprint of artificial intelligence is inseparable from the physical infrastructures that make computing possible. The training and deployment of advanced AI systems require large-scale data centers, high-performance computing facilities, and global cloud computing networks, all of which depend on a continuous electricity supply, water-intensive cooling systems, land availability, and mineral-intensive hardware supply chains [1,2,6].

From a sustainability perspective, these dependencies translate into binding constraints. Global carbon budgets compatible with the Paris Agreement impose limits on cumulative emissions, while the pace of renewable energy deployment determines the extent to which additional electricity demand can be decarbonized in the short and medium term [3,8]. Improvements in energy efficiency, while necessary, do not eliminate these constraints when absolute computational demand continues to grow [5,9].

Despite this material reality, AI policy frameworks rarely address energy use, emissions, or infrastructure expansion as conditions that limit the deployment or scale of systems. Instead, sustainability is often addressed through voluntary commitments, efficiency narratives, or ex post reporting mechanisms, leaving the main determinants of environmental impact structurally unregulated.

2.2. Policy Mechanisms That Generate Environmental Externalization

The environmental omission in AI policy is not the result of isolated regulatory gaps, but emerges from identifiable mechanisms that systematically shift environmental costs outside the scope of decision-making. Three mechanisms are particularly relevant.

2.2.1. Enforcement Requirements and Computational Overload

Regulatory measures aimed at strengthening regulatory compliance—such as algorithmic transparency obligations, auditability requirements, traceability of automated decisions, or the implementation of privacy protection techniques—can increase computational complexity and resource usage under certain conditions. Generating ex post explanations, conducting periodic audits, or applying privacy preservation technologies such as federated learning or differential privacy may require additional training, inference, or storage cycles, resulting in increased energy consumption and associated infrastructure load [5,10]. Recent empirical evidence demonstrates that compliance-oriented safety mechanisms in large language models introduce measurable and non-negligible computational overhead. Systematic analyses of RLHF identify computational resource requirements as a fundamental limitation, noting that fine-tuning with human feedback demands powerful hardware, large datasets, and complex pipelines that create barriers to widespread adoption [11,12]. For instance, safety guardrails deployed to moderate harmful content can add hundreds of milliseconds of additional latency per request in common deployment settings, depending on model size and guardrail architecture.

These latency measurements reflect specific implementation architectures reported in [13]. Actual overhead varies substantially depending on model size, guardrail design (e.g., input filtering vs. output validation), deployment infrastructure, and optimization strategies. The key point is not the precise magnitude, which is context-dependent, but the existence of measurable computational overhead associated with safety and compliance mechanisms.

Similarly, reinforcement learning from human feedback (RLHF)—a standard alignment technique used to ensure models refuse harmful requests—requires substantial additional computational resources beyond base model training. Recent work on safety alignment documents a measurable alignment tax, whereby reinforcement-learning-based alignment introduces additional optimization constraints, degrades general capabilities if not carefully controlled, and requires further computational effort to preserve performance [14]. While the magnitude of this overhead varies depending on implementation choices and model architecture, it represents a significant increment to the computational and energy footprint of deployed AI systems. These measurements reflect specific deployment settings and implementation choices, and that the magnitude of latency and energy impacts may vary substantially across models, architectures, and operational contexts. Accordingly, the presence of computational overhead should be understood as a context-dependent effect rather than a universal or fixed property of AI safety and compliance mechanisms.

While these requirements are fully justified from a legal and rights protection perspective, they are rarely evaluated in terms of their aggregate environmental impact. In the absence of explicit limits on energy consumption or emissions, the increased computational load resulting from enforcement tends to be treated as an acceptable collateral cost of regulatory compliance, rather than as a factor that should condition the technical design or deployment of systems. As a result, the environmental costs associated with the application of regulatory safeguards are externalized to energy systems and ecosystems, without being explicitly integrated into AI policy decision-making.

2.2.2. Infrastructure Duplication and Sovereignty-Driven Expansion

A second mechanism of environmental externalization stems from strategies aimed at ensuring technological autonomy through the expansion of domestic computing capacity. Investments in national data centers, semiconductor manufacturing facilities, and locally hosted foundational models generate energy-intensive infrastructure duplication across jurisdictions, especially in contexts of geopolitical competition and technological fragmentation [15,16].

This duplication tends to increase the total environmental footprint when it occurs in conditions characterized by high marginal electricity carbon intensity, low infrastructure utilization factors, or limitations on water and cooling availability. However, the actual environmental impact depends on contextual variables such as the regional energy mix, the operational efficiency of data centers, or the degree of substitution for shared infrastructure. Even so, given that these considerations are rarely incorporated as binding constraints in capacity planning, expansion driven by technological sovereignty objectives often proceeds without a systematic assessment of its aggregate effects on emissions, resource use, and infrastructure pressure. Environmental externalization is not the result of an explicit decision, but rather of the structural omission of material limits in the computational scaling strategy.

2.2.3. Incentives for Scaling and the Absence of Absolute Limits

A third mechanism stems from policy environments that reward performance gains and rapid scaling without imposing absolute limits on resource use. Efficiency improvements in hardware and algorithms are often presented as indicators of environmental progress; however, existing studies consistently show that these gains are offset by increased deployment and demand, leading to an increase in total energy consumption [2,9].

AI policy frameworks reinforce this dynamic by promoting innovation and competitiveness without restricting scale. In the absence of binding limits on energy use, emissions, or infrastructure growth, sustainability remains disconnected from decisions about model size, deployment intensity, and market expansion. The result is a structural pattern in which environmental impacts increase even as relative efficiency improves.

2.3. Implications for Comparative Analysis

These mechanisms provide the analytical basis for the comparative analysis developed in the following section. Although the European Union, the United States, and China differ significantly in their regulatory styles and strategic priorities, they share a common pattern: none of them integrates environmental constraints into AI policy as binding conditions for deployment or scale. Instead, environmental impacts are systematically displaced from the core logic of AI policy, producing convergent outcomes in terms of energy demand, emissions growth, and infrastructure expansion.

3. Comparative Policy Analysis: Environmental Externalization in AI Policy Regimes

This study adopts a qualitative comparative policy analysis approach, focusing on examining how the regulatory and strategic frameworks for artificial intelligence address—or omit—the environmental limits associated with computational scale. The objective is not to quantitatively measure specific environmental impacts. Rather, the analysis identifies structural patterns of governance that condition the integration, or externalization, of material constraints in AI policy design.

The analysis is based on a corpus of public policy and regulatory documents selected for their relevance to AI governance and their interaction with computational infrastructure.

This corpus includes, among others, regulatory texts on artificial intelligence and data protection, industrial and technological sovereignty strategies, policies on data centers and semiconductors, as well as energy and climate planning documents. The selection of sources focused on the period 2019–2024, to capture the recent cycle of AI expansion and associated policy formulation.

The unit of analysis is the treatment that each AI policy regime gives to the material limits of technological deployment, particularly in relation to energy consumption, carbon emissions, and the expansion of computational infrastructure. In this framework, environmental externalization is operationally defined as the absence of binding environmental restrictions that condition the scale, location, or energy intensity of AI systems. Measures that impose enforceable limits, quantitative thresholds, authorization requirements, or effective conditionalities associated with the deployment or scaling of AI systems are considered binding.

To be considered “binding” in the operational sense used in this article, a restriction must meet at least one of the following criteria: (i) establish maximum enforceable thresholds for energy consumption, carbon emissions, or resource use specifically applicable to the training, deployment, or operation of AI systems; (ii) require prior environmental authorizations or assessments whose denial could effectively prevent or condition the deployment of AI computational infrastructure; or (iii) effectively link access to computational infrastructure, public funding, or industrial incentives to compliance with specific quantitative environmental standards for AI. This operational definition allows for the recognition of the existence of general energy regulations, efficiency standards for data centers, or renewable procurement policies, while maintaining the analytical distinction between such regulations and restrictions specifically linked to the scaling of AI.

It is important to clarify that this operational definition does not deny the existence of *de facto* constraints on computational expansion arising from general infrastructure regulation, such as grid connection limits, permitting procedures, water availability, or land-use restrictions applicable to data centers. While such constraints can function as practical limits in specific local contexts, they are not designed or applied as intentional regulatory instruments to condition the scale, location, or energy intensity of artificial intelligence systems as such. In this article, only restrictions that explicitly target AI deployment or computational scaling, or that condition access to infrastructure, funding, or authorization on quantified environmental criteria linked to AI activity, are considered “binding.” This distinction allows the analysis to acknowledge background infrastructure constraints while maintaining analytical focus on the presence—or absence—of environmental limits deliberately integrated into AI policy design.

The comparative analysis is structured around three main criteria:

- (i) the existence or absence of AI-specific regulatory instruments addressing energy use, emissions, or infrastructure footprint;
- (ii) the degree of integration between AI policy and general climate and energy policy frameworks; and
- (iii) institutional and industrial incentives that favor or restrict the expansion of computational capacity. These criteria make it possible to identify common mechanisms of environmental externalization beyond formal differences in regulatory styles.

The comparison is made at the national or supranational level, focusing on the European Union, the United States, and China as the three main AI governance regimes on a global scale. Although there are relevant subnational variations—such as state initiatives in the United States or provincial policies in China—the analysis focuses on central-level frameworks, given that these variations do not currently translate into absolute environmental limits that condition the scale of AI deployment at the systemic level.

This synthetic approach allows us to identify structural convergences in environmental externalization shared between regimes with divergent governance models, without presupposing that an exhaustive analysis of all subnational variations is necessary to demonstrate the central argument of the study. While state, provincial, or sectoral initiatives are relevant to specific implementation contexts, their existence does not alter the systemic pattern identified: the absence of binding environmental restrictions that effectively condition decisions on the scale, location, and deployment of AI systems at the level of general public policy. This methodological approach allows us to identify substantive convergences in the environmental outcomes of formally divergent AI policies, without presupposing normative equivalences or regulatory hierarchies between jurisdictions.

Methodological limitations and scope boundaries. This analysis focuses on policy design and governance structures rather than quantitative environmental impact assessment. While the study identifies patterns of environmental externalization across jurisdictions, it does not quantify specific emissions increases attributable to AI deployment as distinct from general data center growth, nor does it conduct lifecycle assessments of individual AI systems. The analysis examines policy frameworks and their structural omissions, not the granular measurement of energy consumption or carbon footprints associated with particular models or applications. These quantitative dimensions remain important areas for complementary empirical research.

Based on the methodological approach described above, this section examines how the artificial intelligence policies of the European Union, the United States, and China generate differentiated but convergent patterns of environmental externalization. Rather than comparing abstract models of governance, the analysis focuses on how specific regulatory decisions, infrastructure strategies, and industrial policies shape the material footprint of AI systems. The comparison encompasses both explicit regulatory instruments and implicit omissions observable in energy policy, infrastructure planning, and environmental regulation.

3.1. European Union: Regulatory Ambition and Environmental Omission

The European Union has developed one of the most comprehensive regulatory frameworks on artificial intelligence currently in force, with a strong focus on the protection of fundamental rights. The AI Regulation, adopted in 2024, establishes a risk-based system of obligations aimed at preventing discrimination, ensuring transparency, and strengthening procedural safeguards [7,17]. High-risk AI systems are subject to extensive conformity assessments, documentation requirements, human oversight obligations, and enforcement mechanisms backed by significant administrative penalties.

However, this regulatory architecture operates in a context of marked material dependence. The European Union relies heavily on non-European cloud computing providers for its computing infrastructure, with the three major US-based providers (Amazon Web Services (AWS), Seattle, WA, USA; Microsoft Azure, Redmond, WA, USA; and Google Cloud, Mountain View, CA, USA) controlling approximately 70% of the European cloud market, while European providers account for only around 15% [18]. The EU also lacks advanced semiconductor manufacturing capacity, with only 6–8% of global production occurring in Europe compared to over 60% in East Asia [15], and depends on external providers for the development of large-scale foundational models. Initiatives aimed at strengthening domestic capacity—such as Gaia-X or the European Chips Act—have so far failed to alter this structural dependence in the short to medium term [16,18,19]. This external dependency reproduces the dynamics of “pollution havens” observed in manufacturing offshoring, where European computing demand is met by infrastructure located in jurisdictions with different energy mixes or regulatory frameworks. The associated climate impacts fall outside European territorial accounting, generating a form of carbon leakage

in the digital sphere, comparable to the debates on consumption-based accounting versus territorial accounting for emissions from manufactured products.

Crucially, the AI Regulation does not address the environmental implications of this configuration. It contains no provisions on energy consumption, carbon footprint reporting, or compatibility with climate goals. Nor does it impose environmental conditions on the location or operation of data centers, model training, or public procurement of AI systems [1,5]. As a result, the expansion of AI deployment in the European Union is proceeding without binding environmental restrictions, despite the Union's own formal climate commitments. Environmental impacts are displaced outside the regulatory scope of AI policy, resulting in systematic omission rather than explicit compensation.

As of early 2025, Gaia-X has faced significant implementation challenges, including limited commercial adoption, difficulties in establishing interoperable infrastructure standards, and continued reliance on non-European hyperscale providers for critical computing capacity. The European Chips Act, while allocating substantial funding for semiconductor manufacturing, has yet to result in operational advanced fabrication facilities that would meaningfully reduce EU dependence on external semiconductor supply in the short to medium term.

The 'pollution haven' hypothesis in environmental economics describes the tendency for polluting industries to relocate to jurisdictions with weaker environmental regulation or enforcement. In the context of digital infrastructure, this dynamic manifests when computing demand in jurisdictions with strict climate commitments is met by data centers and cloud services located in regions with higher carbon intensity electricity grids or less stringent environmental oversight, effectively externalizing emissions outside the demanding jurisdiction's territorial carbon accounting.

3.2. United States: Infrastructural Dominance and Unrestricted Expansion

The United States occupies a unique position in the global AI landscape due to its leading position in critical computational infrastructure. US-based companies control approximately 65–70% of global cloud computing capacity [18]—with AWS, Microsoft Azure, and Google Cloud Platform dominating hyperscale infrastructure—and lead the design of advanced semiconductors, while federal industrial policy actively supports domestic capacity through instruments such as the CHIPS and Science Act [20]. This infrastructural concentration gives the United States substantial control over the development, deployment, and scaling of AI systems.

From a regulatory standpoint, the United States lacks a comprehensive federal framework governing AI or data protection. The protection of rights remains fragmented among sectoral regulations, state-level initiatives, and private litigation, while recent federal actions emphasize voluntary principles and industry self-regulation rather than binding obligations [21,22]. This regulatory stance prioritizes flexibility and innovation, reinforcing the rapid expansion of computing capacity.

Environmental considerations play no restrictive role in this model. There are no federal limits on data center energy consumption, no specific carbon footprint reporting obligations for AI systems, and no environmental conditions attached to public subsidies for semiconductor manufacturing or cloud computing infrastructure [2,9]. Although major cloud providers have announced climate neutrality goals, these commitments coexist with sustained growth in absolute energy demand [3]. The US approach thus generates a pattern of relative efficiency gains combined with absolute increases in resource use, with environmental costs externalized outside the scope of AI policy.

It should be noted that there are relevant initiatives at the state and sectoral levels that indirectly address the environmental impacts of digital infrastructure, such as certain

policies on data centers in California or Virginia, the application of the National Environmental Policy Act (NEPA) to large infrastructure projects, or the recent climate disclosure rules of the Securities and Exchange Commission (SEC). Nevertheless, these instruments operate in a fragmented manner and do not impose specific or binding environmental restrictions on the deployment or scaling of artificial intelligence systems. In particular, they do not directly condition the size of models, the computational intensity of training, or the expansion of AI-related infrastructure, and therefore do not alter the general pattern of environmental externalization identified in the analysis.

3.3. China: Strategic Self-Sufficiency and Growth-Driven Environmental Pressure

China's AI policy is part of a broader strategy of technological self-sufficiency and national security. State-led industrial policy mobilizes large-scale investments to reduce dependence on foreign technologies, especially in cloud infrastructure, semiconductors, and foundational AI models [23,24]. This strategy has given rise to a robust domestic AI ecosystem capable of meeting domestic demand independently of external suppliers.

Regulatory oversight in China prioritizes political control and social stability over individual rights. Mandatory data localization, forced cooperation with security authorities, and content regulation are central elements of the regulatory framework, integrating AI governance into a centralized architecture of state control [25,26].

Environmental sustainability plays a secondary role. Although China has adopted efficiency standards for data centers and encouraged their location in regions with abundant renewable energy, relocating data centers to regions such as Inner Mongolia or Qinghai can reduce carbon intensity when there are high rates of curtailment of renewable energy and limited interconnection with the main power grid, allowing surplus energy that would otherwise be lost to be harnessed [3]. However, the economic efficiency of these policies depends on multiple factors, including utilization rates, quality of grid integration, transmission losses, and cooling requirements in different climate zones. These measures do not impose absolute limits on computational growth.

The rapid expansion toward exaFLOP-scale computing capabilities, with data center capacity growing at estimated double-digit annual rates between 2020–2024 [3], continues to be associated with increases in energy consumption that tend to outpace efficiency improvements [27]. Sustainability functions as a conditional consideration within a growth-oriented strategy, rather than as a binding constraint on AI deployment.

There is also significant variation at the provincial level in China in terms of data center location, energy incentives, and the availability of renewable sources, with policies promoting the relocation of computing infrastructure to less carbon-intensive regions. A systematic analysis of these provincial variations lies beyond the scope of this study, which focuses on central-level policy frameworks. However, these subnational differences do not translate into absolute or enforceable limits on the scale of AI deployment. The expansion of computing capacity continues to be guided by centrally defined objectives of growth and strategic self-sufficiency, with environmental sustainability operating as a criterion of relative efficiency rather than a binding constraint on the overall scale of AI systems.

3.4. Comparative Synthesis: Convergent Outsourcing, Divergent Trajectories

Despite substantial differences in regulatory style, industrial structure, and strategic priorities, the European Union, the United States, and China converge on the same outcome: none of them integrate environmental limits into AI policy as binding conditions on scale, deployment, or infrastructure expansion. In all three cases, sustainability is addressed through indirect, voluntary, or sectorally disconnected measures that leave the main drivers of environmental impact intact.

The trajectories differ. In the European Union, environmental omission results from regulatory compartmentalization and material dependence; in the United States, from infrastructural dominance combined with permissive regulation; and in China, from state-led expansion driven primarily by growth objectives. However, the environmental consequence is similar in all jurisdictions: increased energy demand, expansion of the infrastructure footprint, and a lack of enforceable limits aligned with carbon budgets.

This convergence reinforces the central thesis of the article: contemporary AI policy regimes, regardless of their normative orientation or style of governance, systematically externalize environmental costs by not treating sustainability as a material constraint on computational scale.

Although there are significant subnational and sectoral variations in the three jurisdictions analyzed (Tables 1–3), these do not alter the structural outcome observed: the absence of binding environmental restrictions that effectively condition the computational scale of artificial intelligence.

Table 1. Comparative overview of AI policy regimes and environmental externalization (data and sources updated as of January 2025). Policy orientation and regulatory framework.

Dimension	European Union	United States	China
Main regulatory framework for AI	AI Act (2024) [7], GDPR (2018), Charter of Fundamental Rights	Fragmented framework: BIPA (Illinois), CCPA/CPRA (California), sectoral regulation (HIPAA, FCRA, FERPA)	Cybersecurity Act (2017), Data Security Act (2021), PIPL (2021) [25], Generative AI Regulation (2023) [26]
Type of enforcement	Supranational, legally binding, high administrative penalties	Sectoral and state enforcement; private litigation and agency action (FTC, state attorneys general)	Centralized and administrative; prioritizes state control and national security
Protected subjects	Individuals in EU territory and extraterritorial effects (“Brussels effect”)	Fragmented protection; primacy of corporate freedom of expression	State interests defined as “national security” and “public order”
Dominant regulatory approach	Fundamental rights, algorithmic due process, risk mitigation	Innovation, regulatory flexibility, market leadership	Political control, technological self-sufficiency, social stability

Notes: Main sources [7,17,21,22,25,26,28].

Table 2. Material and industrial capacity for AI deployment.

Dimension	European Union	United States	China
Computing capacity (hyperscale cloud, IaaS) *	Low: high dependence on non-European providers (AWS, Azure, Google Cloud)	High: global leadership in cloud infrastructure (AWS, Azure, Google Cloud)	Medium-high: robust domestic ecosystem (Alibaba, Tencent, Huawei, Baidu)
Semiconductors: ownership and production capacity	Low: <10% global production; absence of advanced manufacturing <5 nm	Medium: leadership in design; external production dependence (TSMC)	Medium: substantial investment; 7 nm capacity with limited yields
Competitive foundational models	Low: relevant projects but no global leadership	High: OpenAI, Google, Anthropic, Meta	High: competitive models in Mandarin; domestic functional autonomy
Effective technological decision-making autonomy	Limited: regulatory sovereignty without full computational sovereignty	High: comprehensive control of the technology chain	High: prioritized strategic self-sufficiency

Notes: Main sources [15,16,18,24,29]. * Cloud market positioning refers to estimated global IaaS and hyperscale cloud market shares, based on industry analyses by Synergy Research Group [18].

Table 3. Integration of environmental limits into AI policy.

Dimension	European Union	United States	China
Integration of energy/emissions into specific AI regulation	No specific binding environmental provisions identified in the AI Act	No specific binding environmental provisions identified in federal AI regulation	No specific binding environmental provisions for AI identified
Carbon budgets applicable to AI	No specific carbon budgets identified for AI	No specific carbon budgets for AI are identified	No specific carbon budgets for AI identified
Specific environmental reporting obligations for AI	No specific enforceable obligations for AI are identified	No specific enforceable obligations for AI are identified	No specific enforceable obligations for AI are identified
Regulatory restrictions on the computational scale of AI	No explicit restrictions linked to environmental limits identified	No explicit restrictions linked to environmental limits identified	No explicit restrictions linked to environmental limits identified
Policies on data centers	EU Code of Conduct (voluntary); no mandatory limits on scale or location	Scattered policies; voluntary corporate commitments	Efficiency standards (PUE) and incentivized location
Observed environmental outcome	Expansion of AI without binding environmental conditions	Absolute growth in energy demand despite efficiency improvements	Rapid growth with efficiency improvements offset by scale

Notes: Main sources [1–3,5,7,9,19,22,26,27].

Methodological Notes

1. The Low/Medium/High categories describe effective material capacity (hyperscale cloud infrastructure, advanced semiconductors, and foundational models), not normative judgments.
2. Cloud market data (Table 2) refers to the IaaS/hyperscale segment. US dominance figure (~65–70%) and EU market distribution (~70% US providers, ~15% European providers) are based on Synergy Research Group Q3 2024 estimates [18]; China figures are derived from the same source and complementary industry analyses.
3. Only measures that impose enforceable thresholds, carbon budgets, prior authorization requirements, or effective conditions that restrict the scale, location, or energy intensity of AI systems are considered “binding environmental limits.”
4. General climate policies, voluntary efficiency standards, or environmental regulations not specific to AI are not considered binding when they do not directly condition decisions on training, deployment, or computational scaling.

4. Public Policy Implications: Internalizing Environmental Limits in AI Policy

The comparative analysis shows a consistent pattern across jurisdictions: despite significant differences in regulatory style, institutional design, and geopolitical strategy, current artificial intelligence policies do not systematically internalize environmental limits. Energy consumption, carbon budgets, and infrastructure expansion remain largely outside the design of AI policy, treated as downstream consequences rather than as constraints that should condition deployment and scale. The central implication for public policy is therefore not the need to balance competing regulatory objectives, but to correct a structural omission: the absence of binding environmental constraints in AI governance.

4.1. Overcoming Techno-Optimism and Voluntary Commitments

The dominant discourse on AI policy continues to rely on narratives of technological optimism, efficiency gains, and voluntary corporate commitments. Sustainability is often

invoked as a long-term goal to be achieved through innovation, greater hardware efficiency, or future deployments of renewable energy. The analysis presented in this article challenges that premise. Efficiency improvements have been consistently outpaced by growth in computational demand, leading to increases in absolute energy consumption across all jurisdictions.

Voluntary sustainability commitments announced by cloud providers and AI developers do not alter this trajectory. In the absence of enforceable limits, competitive pressures and geopolitical incentives favor continued scaling, regardless of environmental impact. As a result, AI policy frameworks based on self-regulation or non-binding guidelines do not address the material determinants of environmental harm. A first implication for policymakers is therefore to abandon reliance on voluntary approaches and recognize that sustainability cannot be achieved through efficiency alone.

4.2. Sustainability as a Binding Condition for AI Deployment and Scale

Treating sustainability as a binding constraint requires a change in public policy design. Environmental limits must be integrated upstream, conditioning decisions on whether, where, and at what scale AI systems are developed and deployed. This implies moving beyond generic climate commitments toward specific regulatory instruments for AI that directly impact energy use, emissions, and infrastructure growth.

Specific options include the introduction of binding carbon budgets for large-scale model training, mandatory carbon footprint reporting for AI systems—with independent verification—and environmental conditionality in public procurement and public financing schemes. For example, carbon budgets could set maximum emission thresholds for foundational model training, differentiated by model size and computational intensity, with ex ante authorization requirements and verified offsetting for executions that exceed predefined limits. Similarly, mandatory reporting frameworks could require standardized disclosure of energy intensity (e.g., kWh per inference operation), hardware lifecycle emissions, and data center Power Usage Effectiveness (PUE), verified through independent audits and aligned with existing climate disclosure regimes. Public procurement could operationalize this conditionality by requiring AI systems deployed in government services to meet sectoral energy efficiency benchmarks, prioritize deployment in low-carbon electricity grids, or demonstrate verified additionality in renewable energy sourcing. Without such instruments, AI policy remains structurally disconnected from climate policy, allowing computational expansion to undermine broader transitions toward sustainability.

Similar regulatory precedents exist in other sectors with high environmental impacts. For instance, emissions performance standards are used in electricity generation, aviation (CORSIA framework), and industrial manufacturing (EU ETS). Vehicle emissions standards differentiate limits by vehicle class and use. Energy efficiency standards for buildings and appliances establish mandatory thresholds with penalties for non-compliance. Such sectoral precedents demonstrate the technical and administrative feasibility of imposing binding environmental limits on specific activities while allowing differentiation based on technical characteristics and intended use.

4.3. Managing Environmental Externalization Under Geopolitical Competition

A central obstacle to integrating environmental limits is geopolitical competition. The analysis shows that no major jurisdiction is willing to impose unilateral restrictions that could reduce its relative technological position. This dynamic leads to a collective outcome in which all actors expand computational capacity while externalizing environmental costs.

However, the cross-border nature of climate impacts means that purely national solutions are insufficient. While comprehensive global coordination may be politically un-

realistic, certain forms of targeted international alignment are feasible and necessary. These include common standards for reporting AI-related emissions, shared methodologies for measuring energy consumption in training and inference, and coordination on minimum environmental requirements for hyperscale data centers. These measures do not require harmonizing AI governance models, but they can reduce the risk of a race toward environmentally unsustainable scaling. In concrete terms, international alignment could take the form of a voluntary multilateral registry of emissions derived from foundational model training, inspired by existing climate finance transparency mechanisms, or the definition of minimum energy efficiency standards for hyperscale facilities involved in cross-border data flows. Although such instruments would not eliminate competitive dynamics, they could introduce minimum expectations for transparency and efficiency that may help mitigate environmentally intensive scaling practices.

4.4. Environmental Limits and Policy Trade-Offs

Integrating sustainability as a binding constraint inevitably affects other dimensions of AI policy. Some computationally intensive practices—such as exhaustive explainability techniques, continuous audits, or redundant infrastructure replication—may be incompatible with strict carbon budgets. This does not imply abandoning fundamental safeguards or strategic capabilities, but it does require prioritizing essential protections and avoiding environmentally costly regulatory excesses.

Similarly, ambitions for complete technological self-sufficiency may need to be reconsidered when they involve large-scale duplication of infrastructure with limited marginal benefits. From an environmental perspective, selective dependence and shared infrastructure may, in some cases, be more sustainable than fragmented expansion driven by sovereignty. Recognizing these limits is not a failure of public policy, but a necessary adjustment to material reality.

4.5. From Aspirational Governance to Constraint-Aware Policy Design

The overall implication of the analysis is that effective AI policy must move from aspirational frameworks to constraint-aware design. Success should not be measured by the number of principles adopted or guidelines published, but by the extent to which AI deployment remains compatible with energy systems, carbon budgets, and planetary boundaries.

Without this shift, AI governance will continue to externalize environmental costs, undermining both climate goals and the long-term viability of digital transformation. Internalizing environmental limits is therefore not a peripheral concern, but a prerequisite for aligning artificial intelligence with transitions toward sustainability.

5. Conclusions

This article has analyzed artificial intelligence policies from the perspective of environmental sustainability understood as a material constraint, rather than a normative aspiration. Based on an analysis of the European Union, the United States, and China, the study shows that, despite their divergent regulatory traditions, differentiated institutional arrangements, and distinct geopolitical strategies, all major AI policy regimes converge on the same outcome: the systematic externalization of the environmental costs associated with computational scale.

In all the jurisdictions analyzed, AI policy frameworks prioritize regulatory compliance, strategic autonomy, or industrial competitiveness, while energy consumption, carbon budgets, and infrastructure expansion remain weakly regulated or completely neglected. Climate goals and AI policy are evolving on parallel but largely disconnected trajectories. As a result, the rapid scaling of AI systems is proceeding without binding limits aligned

with planetary constraints, compromising the compatibility of digital transformation with transitions to sustainability.

Comparative analysis shows that this outcome is not accidental. It arises from the interaction between geopolitical competition, industrial incentives, and the absence of specific environmental regulation for AI. No relevant jurisdiction is willing to impose unilateral restrictions on computational expansion that could reduce its relative position, even when aggregate outcomes threaten shared climate goals. Efficiency improvements, voluntary commitments, and general climate statements have proven insufficient to offset the growth in absolute energy demand driven by AI deployment. Recent comparative evidence shows that the adoption of artificial intelligence is associated with higher carbon emissions even after controlling for renewable energy deployment and institutional quality, indicating that governance reforms and efficiency gains alone do not neutralize the environmental impact of AI-driven computational expansion [30]. Complementary analyses at the urban level further confirm that the spread of AI is linked to a deterioration in carbon emissions performance, despite efforts at energy transition and industrial transformation, highlighting the persistence of scale effects [31]. Recent empirical research on artificial intelligence, cloud computing, and digital infrastructures further shows that digital expansion generates net increases in energy demand and emissions even under ESG-oriented governance frameworks, reinforcing the relevance of cumulative scale effects in digital sustainability [32]. Even within the literature on sustainability-oriented AI, this dynamic is beginning to be recognized, although it is usually framed in terms of mitigation strategies rather than binding constraints on computational scale [33].

Looking ahead, current regulatory and policy trajectories point towards an increasing emphasis on AI safety, accountability, and compliance across the AI lifecycle, particularly for large-scale and high-impact systems. While these developments are typically framed in legal or ethical terms, they also carry important technical implications. As shown by recent empirical work discussed above, compliance-oriented mechanisms are not computationally neutral and tend to increase evaluation, validation, and control requirements. Strengthening compliance is therefore likely to entail higher computational intensity at the system level, with corresponding implications for energy use and infrastructure demand.

The central contribution of this article is to reframe the challenge of AI sustainability as a problem of public policy design rather than technological optimization. Environmental impacts are not an unintended side effect of AI development, but the predictable result of governance choices that treat sustainability as external to decisions about scale, infrastructure, and deployment [4]. Without explicit mechanisms to internalize environmental limits, AI policy will continue to reproduce patterns of omission and externalization.

From a public policy perspective, aligning AI with sustainability requires integrating binding environmental constraints into the governance of computational systems. This includes enforceable carbon budgets for large-scale model training, verifiable emissions reporting obligations, environmental conditionality in public procurement and financing, and explicit consideration of energy systems and infrastructure location in AI deployment decisions. These measures do not eliminate political or economic tensions, but they define the conditions under which AI development can remain physically viable.

This analysis has several limitations. First, although the article demonstrates the existence of systematic environmental externalization between jurisdictions, it does not quantify the magnitude of emissions growth attributable to AI deployment compared to other digital services. Second, the analysis focuses on policy omissions, rather than specific cases where AI has demonstrably reduced emissions through optimization applications. These limitations point to clear avenues for future empirical research.

Ultimately, the sustainability of artificial intelligence is not solely a question of ethical alignment or future innovation. It is a question of whether societies are willing to govern digital expansion within the material limits of the planetary system. Without this shift, AI policy risks accelerating environmental pressures at precisely the moment when global emissions must decline. Treating sustainability as a binding constraint is therefore not an optional refinement of AI governance, but a prerequisite for ensuring that artificial intelligence contributes to—rather than undermines—long-term sustainability goals.

Author Contributions: C.G.-L. and I.O. contributed equally to all aspects of this research, including conceptualization, methodology, formal analysis, investigation, writing—original draft preparation, writing—review and editing, visualization, and supervision. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Jones, N. How to Stop Data Centers from Gobbling Up the World's Electricity. *Nature* **2018**, *561*, 163–166. [CrossRef] [PubMed]
2. Masanet, E.; Shehabi, A.; Lei, N.; Smith, S.; Koomey, J. Recalibrating Global Data Center Energy-Use Estimate. *Science* **2020**, *367*, 984–986. [CrossRef] [PubMed]
3. International Energy Agency. *Electricity 2024: Analysis and Forecast to 2026*; IEA: Paris, France, 2024. Available online: <https://www.iea.org/reports/electricity-2024> (accessed on 5 January 2026).
4. Ligozat, A.-L.; Lefevre, J.; Bugeau, A.; Combaz, J. Unraveling the Hidden Environmental Impacts of AI Solutions for Environment Life Cycle Assessment of AI Solutions. *Sustainability* **2022**, *14*, 5172. [CrossRef]
5. Strubell, E.; Ganesh, A.; McCallum, A. Energy and Policy Considerations for Deep Learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, 28 July–2 August 2019; Association for Computational Linguistics: San Diego, CA, USA, 2019; pp. 3645–3650. [CrossRef]
6. Crawford, K. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*; Yale University Press: New Haven, CT, USA, 2021.
7. European Union. *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*; L 1689; Official Journal of the European Union: Luxembourg, 2024. Available online: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689> (accessed on 5 January 2026).
8. United Nations Framework Convention on Climate Change (UNFCCC). *Paris Agreement*; United Nations: Paris, France, 2015. Available online: https://unfccc.int/sites/default/files/english_paris_agreement.pdf (accessed on 5 January 2026).
9. Patterson, D.; Gonzalez, J.; Le, Q.; Liang, C.; Munguia, L.M.; Rothchild, D.; So, D.; Texier, M.; Dean, J. Carbon Emissions and Large Neural Network Training. *arXiv* **2021**, arXiv:2104.10350. [CrossRef]
10. Bender, E.M.; Gebru, T.; McMillan-Major, A.; Shmitchell, S. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21); ACM: New York, NY, USA, 2021; pp. 610–623. [CrossRef]
11. Casper, S.; Davies, X.; Shi, C.; Krendl Gilbert, T.; Scheurer, J.; Rando, J.; Freedman, R.; Korbak, T.; Lindner, D.; Freire, P.; et al. Open Problems and Fundamental Limitations of Reinforcement Learning from Human Feedback. *arXiv* **2023**, arXiv:2307.15217. [CrossRef]
12. Dahlgren Lindström, A.; Methnani, L.; Krause, L.; Ericson, P.; Martínez de Rituerto de Troya, Í.; Coelho Mollo, D.; Dobbe, R. Helpful, Harmless, Honest? Sociotechnical Limits of AI Alignment and Safety through Reinforcement Learning from Human Feedback. *Ethics Inf. Technol.* **2025**, *27*, 28. [CrossRef] [PubMed]
13. Lee, S.; Lee, D.B.; Wagner, D.; Kang, M.; Seong, H.; Bocklet, T.; Lee, J.; Hwang, S.J. SafeRoute: Adaptive Model Selection for Efficient and Accurate Safety Guardrails in Large Language Models. *arXiv* **2025**, arXiv:2502.12464. [CrossRef]
14. Niu, Y.; Chen, Z.; Liu, X.; Wang, Z.; Zhang, D.; Ding, N. Mitigating the Safety Alignment Tax with Null-Space Constrained Policy Optimization. *arXiv* **2025**, arXiv:2512.11391. [CrossRef]

15. Boston Consulting Group (BCG); Semiconductor Industry Association (SIA). *Strengthening the Global Semiconductor Supply Chain in an Uncertain Era*; BCG: Boston, MA, USA; SIA: Washington, DC, USA, 2021. Available online: https://www.semiconductors.org/wp-content/uploads/2021/05/BCG-x-SIA-Strengthening-the-Global-Semiconductor-Value-Chain-April-2021_1.pdf (accessed on 5 January 2026).
16. Semiconductor Industry Association (SIA). *2024 State of the U.S. Semiconductor Industry*; SIA: Washington, DC, USA, 2024. Available online: <https://www.semiconductors.org/resources/2024-state-of-the-u-s-semiconductor-industry/> (accessed on 5 January 2026).
17. Veale, M.; Borgesius, F.Z. Demystifying the Draft EU Artificial Intelligence Act. *Comput. Law Rev. Int.* **2021**, *22*, 97–112. [CrossRef]
18. Synergy Research Group. *Cloud Infrastructure Services Market Share; Q3 2024 Market Report*; Synergy Research Group: Reno, NV, USA, 2024. Available online: <https://www.srgresearch.com/> (accessed on 5 January 2026).
19. European Commission. *European Chips Act*; COM(2022) 45 final; European Commission: Brussels, Belgium, 2022. Available online: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52022DC0045> (accessed on 5 January 2026).
20. United States Congress. *CHIPS and Science Act of 2022*; Public Law 117–167; United States Congress: Washington, DC, USA, 2022. Available online: <https://www.congress.gov/bill/117th-congress/house-bill/4346> (accessed on 5 January 2026).
21. Cath, C.; Wachter, S.; Mittelstadt, B.; Taddeo, M.; Floridi, L. Artificial Intelligence and the ‘Good Society’: The US, EU, and UK Approach. *Sci. Eng. Ethics* **2018**, *24*, 505–528. [CrossRef] [PubMed]
22. White House. *Executive Order 14110 on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*; Federal Register: Washington, DC, USA, 2023; Volume 88. Available online: <https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence> (accessed on 5 January 2026).
23. State Council of the People’s Republic of China. *14th Five-Year Plan for National Economic and Social Development and Long-Range Objectives Through the Year 2035*; State Council: Beijing, China, 2021. Available online: <https://cset.georgetown.edu/publication/china-14th-five-year-plan/> (accessed on 5 January 2026).
24. Allen, G.C. *Understanding China’s AI Strategy: Clues to Chinese Strategic Thinking on Artificial Intelligence and National Security*; Center for a New American Security: Washington, DC, USA, 2019. Available online: <https://www.cnas.org/publications/reports/understanding-chinas-ai-strategy> (accessed on 5 January 2026).
25. People’s Republic of China. *Personal Information Protection Law of the People’s Republic of China*; Order No. 91 of the President; National People’s Congress: Beijing, China, 2021. Available online: <http://www.npc.gov.cn/npc/c30834/202108/a8c4e3672c74491a80b53a172bb753fe.shtml> (accessed on 5 January 2026).
26. Cyberspace Administration of China. *Measures for the Management of Generative Artificial Intelligence Services*; (effective 15 August 2023); Cyberspace Administration of China: Beijing, China, 2023. Available online: http://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm (accessed on 5 January 2026).
27. Wu, C.J.; Raghavan, P.; Gupta, U.; Pentland, A.; Jordan, M.I.; Weller, A.; Green, B.; Hecht, B.; Narayanan, A.; Daumé, H., III; et al. Sustainable AI: Environmental Implications, Challenges and Opportunities. Available online: https://proceedings.mlsys.org/paper_files/paper/2022/hash/462211f67c7d858f663355eff93b745e-Abstract.html (accessed on 5 January 2026).
28. State of Illinois. *Biometric Information Privacy Act*; 740 ILCS 14; State of Illinois: Springfield, IL, USA, 2008. Available online: <https://www.ilga.gov/legislation/ilcs/ilcs3.asp?ActID=3004&ChapterID=57> (accessed on 5 January 2026).
29. Council on Foreign Relations Independent Task Force. *Innovation and National Security: Keeping Our Edge*; Task Force Report No. 77; Council on Foreign Relations: New York, NY, USA, 2019. Available online: <https://www.cfr.org/report/keeping-our-edge/> (accessed on 5 January 2026).
30. Iqbal, A.; Zhang, W.; Jahangir, S. Building a Sustainable Future: The Nexus between Artificial Intelligence, Renewable Energy, Green Human Capital, Geopolitical Risk, and Carbon Emissions through the Moderating Role of Institutional Quality. *Sustainability* **2025**, *17*, 990. [CrossRef]
31. Qu, F.; She, W. Artificial Intelligence Technology and Regional Carbon Emission Performance: Does Energy Transition or Industrial Transformation Matter? *Sustainability* **2025**, *17*, 1844. [CrossRef]
32. Wang, K.-H.; Jiang, X.-Y.; Tang, Y. Artificial intelligence, cloud computing, blockchain, and the energy market in the era of energy transition. *Energy Econ.* **2025**, *151*, 108975. [CrossRef]
33. Toderas, M. Artificial Intelligence for Sustainability: A Systematic Review and Critical Analysis of AI Applications, Challenges, and Future Directions. *Sustainability* **2025**, *17*, 8049. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.