Artificial Intelligence and Infrastructural Exceptionality: Data Centers, Water, and Ecological Power

Inteligência Artificial e Excepcionalidade Infraestrutural: Data Centers, Água e Poder Ecológico

Cesar Beck[1]

Abstract: This article examines the environmental materiality of artificial intelligence infrastructures, particularly hyperscale data centers and their dependence on water and energy resources. The study investigates how the expansion of large-scale computational systems reorganizes access to strategic ecological resources while producing new forms of infrastructural power, regulatory asymmetry, and political exceptionality. The central objective is to analyze the ecological and legal implications of AI infrastructures across Environmental Law, Digital Law, Human Rights, and Law and Economics. Methodologically, the research adopts a qualitative and interdisciplinary approach grounded in bibliographic and documentary analysis, drawing on international reports, institutional documents, and contemporary case studies concerning data centers, generative AI, and the political economy of computational scale. The article argues that narratives of innovation, sustainability, technological inevitability, and Silicon Valley accelerationism — exemplified by defense-oriented discourses such as Palantir's *The Technological Republic* — tend to obscure the environmental externalities of AI infrastructures, especially freshwater consumption, energy-intensive cooling, and territorial pressure. The results indicate that hyperscale data centers increasingly operate through infrastructural exceptionality, securing privileged access to environmental resources under weak public scrutiny and flexible regulatory arrangements. The study concludes that artificial intelligence should be understood not merely as a digital phenomenon, but as a material and geopolitical infrastructure capable of reshaping environmental governance, hydropolitical disputes, and democratic control over ecological resources.

Keywords: artificial intelligence infrastructure; data centers; water consumption; Palantir; infrastructural exceptionality.

Resumo: Este artigo examina a materialidade ambiental das infraestruturas de inteligência artificial, especialmente os data centers hiperescaláveis e sua dependência de recursos hídricos e energéticos. O estudo investiga de que maneira a expansão dos sistemas computacionais em larga escala reorganiza o acesso a recursos ecológicos estratégicos, ao mesmo tempo em que produz novas formas de poder infraestrutural, assimetria regulatória e excepcionalidade política. O objetivo central é analisar as implicações ecológicas e jurídicas das infraestruturas de IA nos campos do Direito Ambiental, do Direito Digital, dos Direitos Humanos e do Law and Economics. Metodologicamente, a pesquisa adota

[1] Doctoral Candidate in Law at ATITUS Educação, Brazil, researching Digital Law, Artificial Intelligence governance, Environmental Governance, Constitutional Theory, and Critical Infrastructure Studies. Master of Law in Human Rights from Universidade Regional do Noroeste do Estado do Rio Grande do Sul – UNIJUÍ. Specialist in Contemporary Politics and International Relations (FAAP), Digital Law (UERJ/ITS Rio), Data Science Analytics (USP/ESALQ), Constitutional Law (UCAM), and Entertainment Law (UERJ). CopyrightX Fellow affiliated with Harvard University and ITS Rio. Bachelor of Law from Fundação Getulio Vargas Law School – FGV Direito Rio. Author of *RESET: Novas Fronteiras entre Direito e Novas Tecnologias no Século XXI* (Ed. UNIJUÍ), dedicated to the legal, political, and philosophical implications of emerging technologies. Researcher in Digital Constitutionalism, platform governance, environmental impacts of AI infrastructures, and critical studies of technological power. Lawyer, political commentator, and professor. ORCID: https://orcid.org/0000-0002-8625-6503. E-mail: ocesarbeck@gmail.com.

abordagem qualitativa e interdisciplinar, fundamentada em análise bibliográfica e documental, incorporando relatórios internacionais, documentos institucionais e estudos de caso sobre data centers, IA generativa e economia política da escala computacional. O artigo sustenta que narrativas de inovação, sustentabilidade, inevitabilidade tecnológica e aceleracionismo do Vale do Silício — exemplificadas por discursos de orientação defensiva como *The Technological Republic*, da Palantir — tendem a obscurecer as externalidades ambientais das infraestruturas de IA, especialmente consumo de água doce, resfriamento intensivo em energia e pressão territorial. Os resultados indicam que data centers hiperescaláveis operam progressivamente por meio de excepcionalidade infraestrutural, obtendo acesso privilegiado a recursos ambientais sob reduzido escrutínio público e arranjos regulatórios flexíveis. Conclui-se que a inteligência artificial deve ser compreendida não apenas como fenômeno digital, mas como infraestrutura material e geopolítica capaz de reconfigurar a governança ambiental, as disputas hidropolíticas e o controle democrático sobre recursos ecológicos.

Palavras-chave: infraestrutura de inteligência artificial; data centers; consumo hídrico; Palantir; excepcionalidade infraestrutural.

INTRODUCTION

For more than two decades, digital technologies have been publicly framed through metaphors of immateriality. Expressions such as cloud computing, virtual infrastructure, and digital ecosystems helped consolidate the perception that computational systems exist in a dimension somehow detached from territory, resource extraction, and environmental limitation. Artificial intelligence radicalized this imaginary. Since the popularization of generative AI and large language models, computational systems have increasingly been associated with automation, scalability, and frictionless efficiency, almost as if intelligence itself had escaped material constraints. Yet the opposite may be closer to reality.

Behind the apparent abstraction of AI systems lies an expanding physical architecture composed of hyperscale data centers, cooling systems, electrical substations, semiconductor chains, undersea cables, and mineral extraction networks distributed across strategic territories. What contemporary society calls “the cloud” increasingly resembles something far less ethereal: an industrial infrastructure dependent upon water, electricity, logistics, and continuous thermal regulation (Crawford, 2021, p. 23–27).

There is a peculiar cultural contradiction here. Contemporary digital capitalism frequently presents artificial intelligence through aesthetics of cleanliness and dematerialization, while simultaneously intensifying environmental extraction. In some respects, the imagery resembles the quiet infrastructural dread present in works such as *Blade Runner 2049* and *Death Stranding*, where

technological systems persist through invisible logistical and energetic dependencies hidden beneath the aesthetics of seamless connectivity. Reality, admittedly, is less cinematic — but perhaps not entirely different. AI systems do not float above ecological constraints. They are grounded in them.

Recent empirical studies demonstrate that advanced AI infrastructures require substantial quantities of electricity and freshwater, particularly for cooling operations associated with high-density computational activity **(Li et al., 2023, p. 2–3)**. Simultaneously, the accelerated race for computational scale among technology corporations has intensified territorial disputes surrounding energy allocation, hydric availability, and infrastructural expansion. The environmental issue, therefore, is no longer peripheral to artificial intelligence. It is structural.

Despite growing concern among researchers and environmental institutions, much of the literature surrounding AI sustainability remains fragmented between technical discussions of computational efficiency and isolated carbon-emission metrics. Comparatively less attention has been devoted to the political economy of computational infrastructures themselves: who controls them, which territories sustain them, and under what regulatory conditions they expand. In practice, environmental externalities frequently become secondary within broader geopolitical disputes involving digital sovereignty, technological leadership, and economic competitiveness.

The article investigates precisely this tension — one that increasingly places ecological limits and computational acceleration into direct collision. More specifically, it examines how the global expansion of hyperscale AI infrastructures reorganizes access to strategic ecological resources, especially water and electricity, while simultaneously producing new forms of infrastructural concentration and regulatory asymmetry. The central hypothesis advanced here is that contemporary AI systems depend upon an extractive computational architecture increasingly legitimized through narratives of innovation, sustainability, and technological inevitability. As a consequence, hyperscale infrastructures progressively acquire forms of infrastructural exceptionality through which environmental constraints become flexible whenever they collide with computational acceleration.

Against this background, the article is guided by the following research question: how does the global expansion of hyperscale artificial intelligence infrastructures reorganize access to strategic ecological resources, particularly water and electricity, while producing new forms of infrastructural power, regulatory asymmetry, and infrastructural exceptionality? The question is not limited to the measurement of ecological impacts, nor to the technical efficiency of computational systems. Rather,

it concerns the juridical and political conditions under which environmental limits are transformed into negotiable variables within the broader race for computational scale, digital sovereignty, and geopolitical advantage.

The relevance of the debate extends beyond conventional sustainability concerns. Artificial intelligence infrastructures now intersect Environmental Law, Digital Constitutionalism, Human Rights, and Law and Economics simultaneously. Conflicts involving water allocation, energy prioritization, and territorial occupation increasingly expose tensions between private computational demands and collective ecological interests, especially in regions marked by infrastructural fragility or resource scarcity. Existing regulatory systems often remain poorly equipped to address these dynamics because they were historically structured around twentieth-century industrial pollution models rather than platform-centered computational infrastructures (Couldry; Mejias, 2019, p. *ix–xx*).

Methodologically, the article adopts a qualitative and interdisciplinary approach grounded in bibliographic and documentary analysis, although the broader ambition is less technical than diagnostic: to understand how computational infrastructures quietly reorganize ecological governance. The investigation combines contributions from Environmental Law, Political Ecology, Critical Infrastructure Studies, and critical political economy, mobilizing institutional reports, environmental disclosures, case studies, and contemporary scholarship concerning artificial intelligence infrastructures. The theoretical framework draws particularly upon Michel Foucault, Ulrich Beck, Karl Polanyi, Ugo Mattei, Shoshana Zuboff, Benjamin Bratton, and Nick Couldry in order to examine the relations between governmentality, risk distribution, commodification, juridical mediation, technological power, environmental governance, and infrastructural concentration.

The paper is divided into two major sections. The first examines the ecological materiality of artificial intelligence, focusing on data centers, cooling systems, energy-intensive computation, and the environmental footprint associated with foundation models. The second analyzes what this article defines as infrastructural exceptionality: the gradual normalization of asymmetrical governance arrangements through which computational infrastructures obtain political priority over ecological restraint.

Ultimately, the argument developed here is relatively simple, although politically uncomfortable: artificial intelligence should not be understood merely as software or informational abstraction. Contemporary AI systems increasingly operate as planetary infrastructures capable of reorganizing

environmental governance, hydropolitical disputes, and the distribution of ecological costs across territories. The cloud, in the end, is not immaterial. It has geography. It has logistics. And increasingly, it has thirst.

2. THE ECOLOGICAL MATERIALITY OF ARTIFICIAL INTELLIGENCE

Artificial intelligence is frequently discussed through the language of abstraction: models, interfaces, automation, cognition, prediction. Yet the contemporary expansion of AI increasingly depends upon intensely physical systems operating through energy extraction, thermal regulation, logistical continuity, and territorial occupation. Beneath the apparent fluidity of digital life lies an infrastructure that is neither virtual nor immaterial. It is industrial. Sometimes brutally so.

There is something almost culturally disorienting about this contradiction. Over the last decade, large technology corporations perfected an aesthetic vocabulary centered on minimalism, seamlessness, and invisible computation. The user sees smooth interfaces, conversational assistants, recommendation systems, synthetic images. What remains hidden are the warehouses, transformers, pipelines, generators, cooling towers, and electrical substations required to sustain computational permanence. In a curious way, contemporary AI infrastructures resemble the landscapes portrayed in *Death Stranding*: hyperconnected systems sustained by fragile logistical networks stretched across environments already approaching exhaustion.

The environmental issue, therefore, cannot be reduced merely to carbon accounting. Artificial intelligence infrastructures increasingly reorganize relations between territory, energy, water, extraction, and political power itself. Computational scale now depends upon ecological scale. And ecological scale, increasingly, becomes subordinated to computational acceleration.

For the purposes of this article, the analysis does not seek to provide an exhaustive comparative survey of all jurisdictions hosting data centers or artificial intelligence infrastructures. Its scope is contextual and geopolitical rather than strictly cartographic: it focuses primarily on hyperscale AI infrastructures associated with the United States–European Union technological axis, while mobilizing comparative references to Latin America and water-stressed regions where they help illuminate broader patterns of ecological externalization, regulatory asymmetry, and infrastructural dependency. This delimitation allows the article to preserve the planetary dimension of computational infrastructure without dissolving the analysis into an indeterminate global inventory.

2.1 Physical infrastructure, data centers, and large-scale computation

The consolidation of artificial intelligence as the dominant technological paradigm of the 2020s reinforced an older illusion inherited from early internet culture: the idea that digital systems operate outside material reality. Expressions such as cloud infrastructure, virtual storage, and digital ecosystems helped normalize the perception that computational activity occurs in an abstract informational layer somehow detached from territory, extraction, and ecological limitation. Yet contemporary AI depends upon an enormous physical apparatus composed of hyperscale data centers, electrical substations, cooling systems, semiconductor supply chains, submarine cables, and mineral-intensive logistical networks distributed unevenly across the planet (Bratton, 2016, p. 3–5).

The contradiction became harder to ignore after the explosion of generative AI systems between 2022 and 2025. The rapid diffusion of large language models dramatically increased global computational demand and accelerated investment in high-density infrastructures capable of sustaining continuous training and inference operations. Microsoft, Amazon, Google, Meta, and OpenAI entered what increasingly resembles an infrastructural arms race centered on computational scale, GPU concentration, and cloud dominance (Khan, 2019, p. 973–977). Unlike earlier internet architectures, frontier AI systems require uninterrupted thermal regulation, massive electrical supply, and sophisticated cooling operations functioning permanently.

The material expansion exposes a curious cultural paradox. Contemporary AI is often marketed aesthetically through images of smooth interfaces, invisible automation, and frictionless intelligence. The infrastructure itself, however, looks closer to industrial gigantism than digital minimalism. Warehouses filled with servers, transformers, generators, fiber-optic systems, and cooling mechanisms sustain what users casually call “the cloud.” Sometimes the contrast feels almost surreal. One opens a chatbot interface that appears cleaner than a minimalist smartphone application while, somewhere else, entire energy grids struggle to stabilize the thermal demands of machine learning clusters.

In this respect, the computational landscapes portrayed in productions such as *Black Mirror*, *Mr. Robot*, or even *Blade Runner 2049* no longer appear entirely speculative. The infrastructure of artificial intelligence increasingly resembles a hidden industrial geography operating behind the illusion of immateriality. Mark Fisher would have recognized the irony immediately: late capitalism no longer hides infrastructure because infrastructure itself has become aestheticized into invisibility.

Benjamin Bratton's concept of The Stack remains particularly useful here because it captures how computational systems reorganize political and territorial relations through layered infrastructures extending far beyond conventional state-centered governance (Bratton, 2016, p. 3–9). Data centers no longer function merely as technical facilities supporting digital communication. They increasingly operate as strategic nodes linking cloud platforms, logistical systems, urban infrastructure, energy networks, and extraction-intensive supply chains. Their operational relevance now exceeds the technological sphere itself.

The transformation also alters the economic logic surrounding environmental resources. Freshwater, electricity, land availability, and cooling capacity become strategic assets for computational competitiveness. As a consequence, regions capable of offering cheap energy, tax incentives, permissive regulation, and hydric availability increasingly attract hyperscale investments. Governments frequently frame such projects as symbols of modernization and digital development while minimizing their ecological implications (Marx, 2024, p. 3–8).

The ideological power of the "cloud" metaphor contributes directly to this invisibility. As Tung-Hui Hu (2015, p. 3–6) argues, cloud computing operates simultaneously as technological infrastructure and ideological abstraction. The metaphor creates symbolic distance between users and the material conditions necessary for digital systems to function. In practice, there is no ethereal cloud. There are industrial facilities consuming electricity, extracting water, dissipating heat, and depending upon continuous logistical stabilization.

The environmental implications extend well beyond carbon emissions alone. Data centers increasingly pressure electrical grids, occupy strategic territory, intensify freshwater extraction, and reinforce mineral dependency chains associated with semiconductor production. According to the International Energy Agency (IEA, 2024, p. 32), global electricity demand from data centers may more than double during the present decade due largely to AI expansion. Simultaneously, recent studies indicate that advanced AI systems require significant quantities of water for cooling operations, particularly under high-density computational conditions (Li et al., 2023, p. 2-4).

Another emerging issue involves what recent literature has started describing as the "Data Heat Island Effect," referring to localized thermal pressure generated by hyperscale infrastructures and their interaction with already stressed urban environments (Marinoni et al., 2026, p. 1-3). The expression sounds almost cyberpunk in tone. Unfortunately, it is becoming empirically real.

Existing regulatory systems frequently struggle to address these dynamics adequately because environmental governance remains institutionally fragmented between telecommunications regulation, energy law, environmental licensing, and digital governance frameworks. The result is a structural mismatch between the planetary scale of computational infrastructures and the relatively compartmentalized architecture of legal oversight.

At the same time, computational concentration intensifies asymmetries of infrastructural power. Large technology corporations increasingly control not only platforms and algorithms, but also the physical infrastructures necessary for contemporary digital life itself. Zuboff (2019, p. 8-11) described this tendency as a defining characteristic of surveillance capitalism: the consolidation of asymmetrical relations of informational and economic power through large-scale technological systems. Artificial intelligence infrastructures deepen that logic materially.

From the perspective of Environmental Law, the problem becomes particularly difficult because AI infrastructures rarely resemble classical industrial pollution models. Their externalities are distributed, cumulative, and often politically opaque. Ecological pressure emerges through energy consumption, thermal dissipation, hydric demand, and extraction chains dispersed globally rather than through localized contamination alone.

Artificial intelligence therefore complicates conventional distinctions between digital space and physical territory. Data centers increasingly function as strategic infrastructures capable of reorganizing resource allocation, environmental governance, and territorial priorities. The difference between hyperscale infrastructures and older extractive industries may reside less in material impact than in visibility itself. Mines and pipelines historically occupied explicit positions within environmental conflict. The cloud, by contrast, often remains politically invisible precisely because it appears technologically inevitable.

2.2 Water, energy, and cooling systems

One of the least visible dimensions of artificial intelligence infrastructure concerns something surprisingly mundane: heat. Behind the apparent immateriality of generative AI lies a permanent thermodynamic problem. Large-scale computation generates enormous quantities of thermal energy, and contemporary data centers survive only through continuous cooling operations sustained by electricity, freshwater extraction, and industrial thermal regulation systems.

Public debate surrounding artificial intelligence generally prioritizes algorithmic bias, automation, misinformation, or data protection. Comparatively little attention is devoted to the ecological implications of computational cooling itself. Yet hyperscale infrastructures increasingly depend upon complex thermal stabilization architectures operating uninterruptedly under high-density computational conditions.

The issue is structural rather than secondary. Contemporary generative AI systems require large-scale parallel processing through GPUs and tensor-processing units functioning continuously at extreme computational intensity. Unlike conventional internet infrastructures from the early 2000s, frontier AI models operate through massive energy concentration and constant thermal dissipation (Jones, 2024, p. 698–700). As model scale increases, cooling systems become progressively central to computational viability itself.

Most hyperscale infrastructures currently rely on two dominant cooling models: air cooling and liquid cooling. Traditional air-cooling systems dissipate heat through industrial ventilation and atmospheric circulation mechanisms, while newer liquid-based architectures circulate water or specialized coolants directly through computational components (Masanet et al., 2020, p. 984–985). The latter approach offers superior thermal efficiency under intensive computational loads, but simultaneously deepens freshwater dependency.

Such dependency remains surprisingly opaque. Although major technology corporations increasingly publish ESG disclosures and sustainability reports, methodologies vary considerably and territorial granularity is often absent. Water consumption data frequently appears aggregated globally, making independent environmental assessment difficult. Communities hosting hyperscale infrastructures may therefore experience hydric pressure without possessing meaningful access to localized extraction metrics or operational transparency.

Recent research reinforces the scale of the problem. Li et al. (2023, p. 2-6) estimated that interactions with large language models indirectly consume measurable quantities of freshwater through cooling systems associated with data center operations. Under conditions of accelerated AI expansion, this process may intensify hydric stress globally, especially in territories already affected by drought, desertification, or unstable water distribution systems.

The contradiction is difficult to ignore. Artificial intelligence is frequently marketed through narratives of optimization, sustainability, and smart efficiency while simultaneously depending upon infrastructures characterized by escalating resource consumption. In a strictly metaphorical sense, the problem recalls *Frostpunk*: a society reorganized around heat, scarcity, energy discipline, and the political management of survival. The difference is that AI does not govern a frozen city. It heats the planet while selling itself as cool.

Hydropolitical tensions associated with data center expansion have already emerged in multiple jurisdictions. In parts of Arizona, Nevada, and Utah, local debates surrounding hyperscale facilities increasingly involve freshwater allocation conflicts in regions historically vulnerable to drought conditions. Similar tensions have appeared in Latin America and Europe as governments compete to attract digital infrastructure investments through subsidies, tax incentives, and flexible environmental oversight (IEA, 2024, p. 32-33).

The ecological consequences extend beyond water extraction alone. Large-scale thermal dissipation may also alter localized climatic conditions. Recent studies concerning the so-called “Data Heat Island Effect” indicate that hyperscale AI infrastructures can elevate surrounding land surface temperatures through concentrated thermal emissions generated by computational activity (Marinoni et al., 2026, p. 1-3). Unlike traditional urban heat islands associated with transportation density and metropolitan expansion, these new thermal zones emerge directly from continuous computational concentration.

Electricity consumption follows a similar trajectory. According to projections published by the International Energy Agency, global electricity demand associated with data centers may more than double before 2030, driven largely by AI scaling and cloud expansion (IEA, 2024, p. 32). This creates substantial tension between technological acceleration and climate-transition commitments. Despite corporate sustainability rhetoric, many computational infrastructures continue operating through carbon-intensive electrical matrices dependent upon fossil-fuel-supported grids.

The paradox is politically revealing. International institutions frequently describe artificial intelligence as indispensable for green modernization, sustainable innovation, and efficient governance. The Global Innovation Index 2025, for example, repeatedly associates AI infrastructures with competitiveness and economic modernization (WIPO, 2025, p. 40–43, 54–76). Yet the same

infrastructures increasingly depend upon intensified extraction of water, electricity, territory, and minerals.

From a Law and Economics perspective, this dynamic exposes a familiar pattern of externalization. Environmental costs generated by hyperscale infrastructures become socially distributed while computational concentration and platform revenues remain privatized. Water scarcity, thermal pressure, and energy strain diffuse collectively across territories, whereas economic gains accumulate disproportionately within a limited number of technological conglomerates.

Methodologically, the interaction between Environmental Law and Law and Economics is decisive for this analysis. Environmental Law frames water, energy, territorial occupation, and ecological stability as matters of collective protection, precaution, licensing, and democratic accountability, thereby making visible the normative limits that should constrain infrastructural expansion. Law and Economics, in turn, clarifies how those same infrastructures redistribute costs, incentives, and risks by allowing private actors to internalize computational value while externalizing hydric pressure, thermal burdens, energy strain, and regulatory costs onto territories and communities. The technical analysis of data centers therefore becomes legally relevant precisely because cooling systems, electricity demand, and computational scale are not merely engineering variables; they are institutional mechanisms through which ecological scarcity is converted into an economic condition of artificial intelligence.

Environmental Law struggles to address these dynamics because data centers rarely occupy the symbolic category of extractive industry. They do not resemble mines, pipelines, or petrochemical facilities aesthetically. Their environmental footprint emerges cumulatively through distributed energy consumption, hydric pressure, cooling dependency, and supply-chain extraction dispersed globally.

This ambiguity contributes to what may be described as computational extractivism. Under the discourse of technological necessity, societies progressively normalize the conversion of ecological resources into infrastructural substrates sustaining planetary-scale computation. The cloud, ultimately, does not merely store information. It consumes rivers. It pressures electrical grids. And increasingly, it reorganizes environmental priorities around computational scale.

2.3 The ecological cost of foundation models

The environmental debate surrounding artificial intelligence cannot be limited to data centers viewed as isolated facilities. The ecological issue also resides in the operational logic of foundation models themselves. Large language models, multimodal systems, and generative architectures depend upon permanent computational escalation. Scale is not simply a technical feature of contemporary AI. It is the economic logic sustaining the sector.

The rapid expansion of generative systems after 2022 intensified competition among major technology corporations through model size, GPU concentration, computational capacity, and infrastructural integration. OpenAI, Google, Microsoft, Meta, Anthropic, Amazon, and xAI increasingly compete through access to computation itself. In practical terms, computational scale has become a strategic asset comparable to logistics, energy infrastructure, or industrial capacity.

The environmental consequences of this race often disappear beneath narratives of innovation and digital modernization. Training frontier models may require thousands of GPUs operating continuously for weeks or months, consuming vast quantities of electricity while generating uninterrupted thermal dissipation (Bommasani et al., 2021, p. 10–14). The ecological problem therefore exceeds isolated AI applications. It concerns the normalization of perpetual computational growth.

Benjamin Bratton's concept of planetary-scale computation remains useful because it captures how artificial intelligence increasingly operates through interconnected infrastructures linking cloud systems, semiconductor production, cooling architectures, extraction networks, and logistical chains simultaneously (Bratton, 2016, p. 4–9). Computation ceases to appear as abstraction and reveals itself as material infrastructure.

The environmental footprint associated with foundation models unfolds across several layers at once. Electricity consumption during training and inference operations represents only the most visible dimension. Cooling systems require freshwater extraction; semiconductor production depends upon mineral-intensive supply chains; electronic waste accelerates through constant hardware replacement; territorial concentration intensifies localized ecological pressure. The cloud, despite the metaphor, possesses industrial density.

According to projections published by the International Energy Agency (2024, p. 32), electricity demand associated with AI infrastructures and cloud computation may eventually surpass the annual energy consumption of several industrialized countries if current trajectories persist. At the same time,

hyperscale infrastructures increasingly concentrate in territories already marked by hydric instability and environmental fragility.

The so-called “Data Heat Island Effect” further destabilizes the mythology of digital immateriality. Marinoni et al. (2026, p. 3-5) demonstrate that hyperscale AI infrastructures may elevate surrounding land-surface temperatures through concentrated thermal emissions generated by continuous computational activity. Artificial intelligence infrastructures therefore begin altering localized climatic conditions directly rather than merely contributing abstractly to carbon emissions.

There is something culturally revealing in this contradiction. Contemporary AI is frequently represented aesthetically through images of seamless automation, clean interfaces, and informational fluidity. Yet materially, its infrastructure increasingly resembles industrial gigantism. Less *Black Mirror* here, perhaps, and more *Cyberpunk 2077*: polished interfaces, corporate promises, collapsed ecologies, and an entire political economy hidden beneath the glow of technological inevitability. Digital systems do not eliminate material dependency. They relocate it. Then they aestheticize the relocation.

The environmental burden associated with foundation models also reproduces geopolitical asymmetries. Most advanced AI infrastructures remain concentrated in the Global North, particularly within the United States and selected regions of Europe and East Asia. Yet many of the resources sustaining these systems — lithium, cobalt, rare earth minerals, industrial energy production, and freshwater reserves — depend heavily upon extractive dynamics affecting territories in the Global South.

This geopolitical delimitation is therefore not merely descriptive, but analytical. The United States–European Union axis is treated here as a privileged site of regulatory, corporate, and infrastructural formation, while Latin America and other resource-stressed territories appear as spaces in which the ecological consequences of computational acceleration become particularly visible. The point is not to establish a rigid territorial opposition, but to examine how value, computation, and legal authority tend to concentrate in some regions while environmental pressure, extractive dependency, and infrastructural vulnerability are unevenly distributed across others.

This asymmetry resembles what political ecology literature describes as ecological unequal exchange. Computational infrastructures centralize informational and economic value while

distributing ecological pressure unevenly across territories. Artificial intelligence therefore reproduces older extractive logics through digital infrastructures rather than replacing them.

Opacity intensifies the problem further. Sustainability reports published by major technology corporations frequently emphasize renewable energy investments, carbon-neutrality targets, and efficiency gains while offering comparatively limited transparency regarding localized hydric extraction, thermal emissions, or cumulative ecological impacts associated with model scaling. Artificial intelligence is often represented publicly through narratives of abstraction that conceal the physical systems sustaining computation (Crawford, 2021, p. 23–27).

The discourse surrounding “Green AI” deserves particular scrutiny. Concepts such as efficient computation, carbon-aware infrastructure, and sustainable AI may indeed produce genuine technological improvements. Yet they may also function ideologically by neutralizing criticism through managerial language. Relative efficiency gains frequently coexist with exponential increases in computational scale, adoption, and infrastructural dependency. Efficiency, paradoxically, may stimulate expansion rather than limitation — a dynamic historically associated with the Jevons paradox.

The political implications become increasingly visible in contemporary regulatory disputes. Debates surrounding the European Union AI Act illustrate how environmental concerns and regulatory restrictions frequently encounter resistance whenever they appear capable of slowing computational competitiveness. Under geopolitical pressure, technological infrastructures progressively acquire exceptional political status.

This tendency aligns with Yves-François Le Coadic’s critique of “datafilia”: a technocratic orientation excessively centered on data accumulation and algorithmic mediation detached from broader democratic and human concerns (Villapalos, 2026, p. 2-4). The environmental dimension of that critique is substantial because ecological systems increasingly become subordinated to computational priorities framed as strategically unavoidable.

Ugo Mattei’s diagnosis points in a similar direction. Technological systems, according to Mattei (2025, p. 15-25), progressively substitute traditional forms of institutional mediation through automated informational architectures. The environmental issue surrounding AI therefore concerns more than electricity or water consumption individually. What emerges is a new regime of

infrastructural power in which computational acceleration acquires political legitimacy capable of weakening ecological restraint.

The ecological cost of foundation models ultimately exceeds conventional environmental accounting. At stake is the normalization of a planetary infrastructure whose expansion increasingly appears too profitable, too strategic, and too technologically inevitable to slow down.

3. GREEN GOVERNMENTALITY AND INFRASTRUCTURAL EXCEPTIONALITY

The environmental debate surrounding artificial intelligence no longer concerns only pollution, energy efficiency, or carbon accounting. It also concerns the political management of perception. Major technology corporations do not merely build data centers; they build narratives. They produce sustainability reports, ESG vocabularies, institutional partnerships, and legitimacy strategies capable of softening public scrutiny over the ecological costs of large-scale computation.

This is where Foucault's notion of governmentality becomes useful. Environmental governance does not operate only through prohibition, licensing, or sanctions. It also works through the production of acceptable truths: what counts as innovation, what counts as sustainability, what counts as an unavoidable sacrifice for economic development (Foucault, 2008, p. 15–25). In this sense, AI infrastructure is not merely regulated by law; it also reshapes the language through which regulation becomes politically imaginable.

This is why sustainability discourse may become a reputational technology. Corporate reports frequently emphasize renewable energy, offsets, and efficiency indicators while offering limited transparency about localized water extraction, thermal emissions, mineral dependency, or cumulative territorial pressure. The ecological question is displaced from conflict to management. Risk, as Ulrich Beck (1992, p. 19–23) warned, is not eliminated; it is administrated, narrated, and politically distributed.

Polanyi's critique of fictitious commodities also becomes relevant here. Water, energy, territory, and atmospheric stability are progressively subordinated to computational demand, as if ecological

systems were mere inputs for digital acceleration (Polanyi, 2001, p. 71–80). The issue is not simply that AI consumes resources. The deeper problem is that it reorganizes environmental priorities around infrastructures increasingly framed as indispensable.

These theoretical perspectives are not merely juxtaposed; they describe complementary dimensions of the same infrastructural transformation. Foucault helps explain how governance operates through the production of acceptable rationalities — innovation, sustainability, competitiveness, and security — that make infrastructural expansion appear necessary rather than contestable. Beck, in turn, clarifies how the ecological risks generated by such expansion are not eliminated, but socially distributed, administratively managed, and politically normalized. Polanyi reveals the economic violence through which water, energy, territory, and atmospheric stability are treated as quasi-commodified conditions of digital accumulation, while Mattei radicalizes this diagnosis by showing how technological systems may displace juridical mediation itself, reorganizing social life through operational architectures that progressively weaken law's capacity to impose limits (Foucault, 2008, p. 15–25; Beck, 1992, p. 19–23; Polanyi, 2001, p. 71–80; Mattei, 2025, p. 8–12). Taken together, these authors allow infrastructural exceptionality to be understood not as a merely technical phenomenon, but as a juridical-political regime in which ecological scarcity, computational acceleration, and institutional accommodation become mutually reinforcing.

The same political grammar appears, in a more explicit and disturbing register, in the contemporary discourse of defense-oriented technology firms. Palantir's 2026 twenty-two-point summary of The Technological Republic, associated with Alex Karp and Nicholas Zamiska, openly frames software, artificial intelligence, and hard power as central instruments for Western strategic survival (Karp; Zamiska, 2025, p. 19–25; Griffiths, 2026). The document is not about water, energy, or ecological limits. That absence is precisely the point. In this worldview, AI infrastructure appears primarily as civilizational weaponry, geopolitical necessity, and state-corporate destiny. Externalities become background noise. The pipes, grids, aquifers, and sacrificed territories remain offscreen.

Ugo Mattei's diagnosis is useful precisely because he understands technological systems as forces capable of displacing juridical mediation itself (Mattei, 2025, p. 4–8). Applied to environmental governance, the argument becomes sharper: AI infrastructures do not only consume ecological resources; they also reshape the institutional conditions under which ecological decisions are made.

This produces the condition here described as infrastructural exceptionality. The concept refers to the gradual normalization of legal, administrative, economic, and political concessions through which strategic computational infrastructures obtain privileged access to ecological resources, accelerated regulatory treatment, subsidized energy arrangements, and reduced democratic scrutiny without requiring the formal suspension of environmental law. Data centers, in this sense, acquire quasi-strategic status within economic policy: they are not exempted from regulation in an explicit juridical vacuum, but are increasingly governed through flexible arrangements justified by modernization, digital sovereignty, competitiveness, and geopolitical urgency. Environmental limits are therefore not abolished. They are recalibrated around the operational tempo of computational acceleration.

The recent European debate over the postponement of AI Act restrictions illustrates this tendency. Under pressure from industry and member states concerned with global competitiveness, parts of the regulatory framework were delayed or softened. The episode reveals a broader political pattern: whenever digital infrastructures are framed as strategic necessities, regulatory restraint becomes negotiable.

Here a brief detour through popular culture is almost unavoidable. The ideological atmosphere resembles less the clean optimism of innovation policy than the exhausted militarized worlds of *Metal Gear Solid V* or *Cyberpunk 2077*: private contractors, permanent emergency, technological sovereignty, and a political economy in which infrastructure becomes indistinguishable from command. Fisher would probably have recognized the structure immediately. Capitalist realism does not merely say that there is no alternative. It says that even ecological collapse must be administered through the same systems that helped accelerate it.

At this point, environmental governance risks becoming technocratic accommodation. Decisions about water, energy, land, and ecological sacrifice migrate toward hybrid arrangements involving corporations, investment agencies, and executive authorities operating with limited democratic visibility. The cloud ceases to be merely a metaphor. It becomes a geopolitical layer.

The danger, therefore, is not only environmental degradation. It is the normalization of a political grammar in which computational acceleration receives priority over ecological reversibility. Once that happens, sustainability no longer limits expansion. It becomes part of its operating system.

3.1 ESG, Green AI, and discursive neutralization

The environmental legitimacy of artificial intelligence infrastructures increasingly depends not only on technical efficiency, but on narrative management. Hyperscale computation requires electricity, cooling systems, mineral extraction, and territorial concentration at unprecedented scale. Yet the public image surrounding AI expansion is frequently constructed through the language of sustainability, innovation, and responsible modernization. Artificial intelligence is presented simultaneously as environmentally problematic and environmentally indispensable.

This ambiguity is not accidental. It reflects a broader transformation in contemporary governance in which ecological critique is progressively absorbed into managerial vocabularies rather than functioning as a limit to infrastructural expansion. Terms such as “Green AI,” “responsible computation,” “sustainable cloud,” and “carbon-aware infrastructure” now occupy central positions within corporate communication strategies, institutional partnerships, and ESG disclosure frameworks. The result is a discursive environment in which computational acceleration becomes politically easier to justify.

The issue is not that sustainability initiatives are entirely fictitious. Major technology corporations have in fact expanded renewable energy procurement, cooling optimization systems, carbon accounting mechanisms, and efficiency-oriented research programs. Google, Microsoft, and Amazon increasingly publish environmental reports emphasizing decarbonization targets and operational efficiency improvements. However, these measures coexist with continuous expansion of hyperscale infrastructures, rising electricity demand, escalating GPU concentration, and growing freshwater dependency.

A first fracture appears here. The public vocabulary is sustainability; the operational grammar is scale. Computational systems may become individually more efficient while the aggregate environmental footprint of the sector continues increasing. The phenomenon resembles what environmental economics traditionally associates with rebound effects and the Jevons paradox: efficiency gains reduce operational friction, stimulate adoption, and ultimately intensify total consumption rather than limiting it.

The rhetoric surrounding “Green AI” therefore deserves critical scrutiny not because efficiency improvements are irrelevant, but because they may function ideologically. Environmental discourse increasingly operates as a mechanism of reputational stabilization capable of transforming infrastructural expansion into a socially acceptable horizon. Ecological critique is not necessarily

denied outright. More often, it becomes incorporated into optimization frameworks centered on management, metrics, and technical adjustment.

The contemporary AI sector increasingly absorbs criticism by translating structural contradictions into operational problems awaiting technical correction. The environmental crisis surrounding AI infrastructures consequently risks being framed not as a political dispute involving limits, extraction, and distributive conflict, but as a challenge of efficiency engineering. The language changes. The expansion continues.

The visual culture surrounding artificial intelligence reinforces this tendency. Corporate promotional material frequently associates AI infrastructures with images of green landscapes, renewable energy fields, minimalist interfaces, and frictionless automation. Data centers are represented less as industrial facilities than as components of a clean technological future. The aesthetics matter. As documentaries such as *Coded Bias* demonstrated in another context, technological systems frequently derive legitimacy not only from functionality, but from symbolic neutrality and visual abstraction.

Opacity intensifies the problem. Sustainability reports published by major technology corporations frequently emphasize carbon neutrality targets and renewable investments while providing comparatively limited transparency regarding localized water extraction, thermal emissions, mineral dependency, or cumulative ecological pressure associated with model scaling. Carbon accounting itself may become selectively narrow. Metrics centered exclusively on emissions often obscure broader environmental dimensions involving territory, hydric stress, energy prioritization, and ecological asymmetry.

Recent empirical research reinforces this concern. Studies published by the International Energy Agency (2024, p. 32) indicate that electricity demand associated with AI infrastructures may expand dramatically before 2030 due to accelerated computational scaling. Simultaneously, research conducted by Li et al. (2023, p. 2-6) demonstrates that large language models indirectly consume substantial quantities of freshwater through cooling operations associated with hyperscale data centers. Yet these dimensions frequently occupy secondary positions within public sustainability narratives.

The political consequence is subtle but significant. Once artificial intelligence becomes discursively associated with green modernization, climate optimization, and efficient governance,

opposition to infrastructural expansion risks appearing economically irrational or technologically regressive. Environmental critique becomes politically displaced. The debate gradually shifts from whether expansion should occur toward how expansion may be optimized.

This process aligns partially with Foucault's understanding of governmentality. Governance increasingly operates through the production of acceptable rationalities rather than through direct coercion alone (Foucault, 2008, p. 15–25). Infrastructural expansion becomes naturalized through discourses of inevitability, competitiveness, and innovation urgency. Artificial intelligence consequently acquires a peculiar political condition: it appears simultaneously as technological infrastructure and civilizational necessity.

The discourse of inevitability becomes especially visible in geopolitical competition. Governments increasingly justify regulatory flexibilization, subsidized electricity arrangements, and accelerated licensing procedures through arguments involving digital sovereignty and strategic competitiveness. Under such conditions, ecological restraint risks being reframed as an obstacle to modernization itself.

The recent public discourse surrounding Palantir makes this grammar unusually explicit. In April 2026, the company circulated a twenty-two-point summary of *The Technological Republic*, associated with Alexander Karp and Nicholas Zamiska, in which software, artificial intelligence, hard power, and national security were presented as central to Western strategic survival. The document's emphasis on AI weapons, military capacity, and civilizational competition reveals a wider Silicon Valley accelerationist posture: technological expansion appears not as a political choice, but as a historical obligation (Karp; Zamiska, 2025, p. 11-15; Griffiths, 2026). Its environmental silence is revealing. Water does not appear as a limit. Energy does not appear as a political conflict. Territory does not appear as sacrifice. The infrastructure is assumed.

Palantir is analytically relevant, therefore, not because it represents the entire AI industry, but because it condenses with unusual clarity a broader political grammar of infrastructural acceleration. Its discourse converts software from a commercial product into a civilizational instrument, artificial intelligence from a technological system into a strategic weapon, and computational infrastructure from an economic facility into a condition of geopolitical survival. This rhetorical conversion is central to infrastructural exceptionality: once AI infrastructure is narrated as defense, sovereignty, and historical necessity, the material substrates that sustain it — water, electricity, land, cooling systems,

mineral chains, and environmental licensing — are displaced from the field of democratic contestation into the background of strategic administration. The result is not the disappearance of ecological costs, but their demotion within a hierarchy of urgency organized around security, competitiveness, and civilizational endurance.

That assumption matters. Once AI is framed as geopolitical destiny, ecological costs become administratively secondary. The question is no longer whether societies should reorganize water, electricity, land, and minerals around computational growth. The question becomes how fast that reorganization can be made operational. The externalities remain, but they are pushed outside the dramatic frame — like the background machinery of a video game city that keeps glowing because somewhere else, unseen, the grid keeps burning.

The European debate surrounding the AI Act illustrates a softer version of the same tendency. Political pressure from industrial sectors and member states concerned with competitiveness contributed to the postponement or flexibilization of several regulatory provisions connected to AI deployment. Officially, the argument centered on innovation-friendly governance and economic necessity. Implicitly, however, the episode revealed the growing difficulty of imposing substantial limitations upon infrastructures framed as strategically indispensable.

This is why ESG discourse surrounding AI infrastructures cannot be interpreted merely as corporate communication. It increasingly functions as a governance technology. Environmental concern becomes translated into managerial indicators compatible with continuous expansion. Sustainability remains present rhetorically, but progressively detached from ecological limitation as a political principle.

The risk, therefore, is not simply greenwashing in its traditional sense. The deeper issue concerns discursive neutralization itself: the gradual transformation of ecological conflict into operational administration. Artificial intelligence infrastructures continue expanding materially while sustainability language simultaneously reduces the political visibility of that expansion.

Under these conditions, the environmental debate surrounding AI risks becoming trapped within a paradoxical logic in which the infrastructures contributing to ecological pressure are continuously represented as the infrastructures destined to solve it.

3.2 Data centers as infrastructures of exception

The consolidation of artificial intelligence infrastructure increasingly reveals a phenomenon that exceeds ordinary disputes between economic growth and environmental protection. Data centers no longer operate merely as private technological facilities. They increasingly function as infrastructures treated politically as too strategic to slow down.

Under these conditions, ecological limits begin to appear negotiable. Environmental licensing procedures are accelerated, tax exemptions multiplied, electricity contracts subsidized, and water allocation priorities flexibilized in the name of innovation, competitiveness, and digital sovereignty. What emerges is not the formal suspension of environmental law, but its gradual adaptation to the operational tempo of computational expansion.

The process resembles what Giorgio Agamben (2005, p. 1–4) described as the normalization of exceptionality. The exception no longer appears extraordinary; it becomes administrative routine. In the context of artificial intelligence, however, the mechanism is infrastructural rather than security-oriented. Ecological safeguards remain formally intact while their practical rigidity weakens whenever computational acceleration is framed as strategically indispensable.

This helps explain why hyperscale infrastructures increasingly occupy an ambiguous legal position. Governments continue recognizing environmental concerns publicly, yet regulatory systems simultaneously tolerate exceptional concessions to maintain competitiveness within the global AI race. The discourse of necessity reorganizes the meaning of limitation itself.

The European debate surrounding the AI Act illustrates this tendency clearly. The postponement and flexibilization of several regulatory provisions after pressure from industrial sectors revealed how rapidly technological competitiveness may condition legal accommodation. Officially, the objective was innovation-friendly governance. Politically, the message was broader: computational acceleration had become difficult to constrain without appearing economically vulnerable.

This transformation produces a subtle inversion within environmental governance. Historically, ecological law emerged to impose limits where market incentives proved insufficient to protect collective goods. Under infrastructural exceptionality, however, governance becomes increasingly

managerial rather than restrictive. The objective shifts from limiting expansion toward optimizing its continuity.

ESG discourse contributes significantly to this dynamic. Sustainability language often operates less as ecological restraint than as reputational stabilization. Expressions such as “Green AI,” “responsible cloud infrastructure,” and “carbon-neutral computation” help neutralize political resistance while large-scale expansion proceeds materially. The contradiction is not hidden exactly. It is absorbed.

Foucault’s notion of governmentality remains useful here because it illuminates how power frequently operates through the production of acceptable rationalities rather than prohibition alone (Foucault, 2008, p. 15–25). Artificial intelligence infrastructures become associated with modernization, efficiency, climate innovation, and strategic necessity. Once framed under these categories, environmental criticism risks appearing anti-developmental or technologically reactionary.

The invisibility of digital infrastructures intensifies the problem further. Unlike pipelines, petrochemical complexes, or mining operations, data centers preserve an aesthetic of cleanliness and abstraction. Their environmental footprint remains partially displaced from public perception. Yet the material infrastructure persists stubbornly beneath the interface: cooling systems consume water, GPUs demand electricity, semiconductors require extractive supply chains, and logistical networks occupy territory physically.

There is an irony here that dystopian fiction often perceived earlier than public governance debates. Films such as *Elysium* and games such as *Cyberpunk 2077* imagined technological orders in which the clean interface of privilege depends upon exhausted territories elsewhere. Contemporary AI infrastructures increasingly resemble that condition: seamless systems in the foreground, massive territorial metabolism behind the screen.

The political risk emerges when invisibility combines with indispensability. Once computational infrastructures become framed as essential civilizational systems, democratic capacity to impose ecological limits weakens considerably. Environmental governance enters a structurally fragile position because restriction itself becomes interpreted as geopolitical disadvantage.

At the level of infrastructural governance, the accelerationist discourse emerging from parts of Silicon Valley gives this risk a more concrete institutional form. Palantir's 2026 manifesto-style summary of The Technological Republic framed software, artificial intelligence, national defense, and hard power as inseparable elements of Western strategic survival, while critics described its tone as militarized, civilizational, and openly hostile to softer models of technological governance (Karp; Zamiska, 2025, p. 19–25; Griffiths, 2026; Karpf, 2026). The point is not that every data center is a military facility. It is subtler than that. Once AI infrastructure is absorbed into the vocabulary of war, sovereignty, and survival, ordinary environmental restraint begins to look almost unserious.

This is the decisive mutation. A data center may be privately owned, territorially located, and environmentally licensed, but politically it becomes narrated as part of a civilizational front line. Water contracts, energy prioritization, tax benefits, and accelerated licensing no longer appear as concessions to corporate power. They appear as the price of remaining relevant in the computational order. In such a framework, ecological externalities do not vanish. They are demoted.

This tendency is reinforced by infrastructural concentration. OpenAI, Microsoft, Google, Amazon, Meta, Nvidia, and defense-oriented software firms increasingly control not only software ecosystems, but also computational architectures, cloud infrastructures, semiconductor chains, and AI deployment capacity. Nick Srnicek's diagnosis of platform capitalism becomes even more concentrated in the AI era because frontier model development now requires enormous computational scale inaccessible to most states, universities, or smaller firms (Srnicek, 2017, p. 40–52).

The consequence is a new form of infrastructural dependency. States pursue technological autonomy while simultaneously depending materially upon infrastructures controlled by transnational conglomerates. This partially explains why governments frequently tolerate exceptional concessions to hyperscale corporations. Exclusion from the computational frontier increasingly appears politically intolerable.

Benjamin Bratton anticipated aspects of this transformation when describing planetary computation as an infrastructural layer capable of reorganizing sovereignty operationally (Bratton, 2016, p. 5–10). AI infrastructures do not merely function inside political systems; they increasingly condition the possibilities of those systems. Cloud dependence, computational concentration, and algorithmic mediation become structural dimensions of governance itself.

Environmental conflicts surrounding data centers therefore concern more than administrative licensing disputes. They increasingly involve infrastructural sovereignty: who controls computational capacity, who absorbs ecological pressure, and who possesses authority to define acceptable environmental sacrifice.

The issue intensifies because frontier AI models require cumulative escalation. More computational scale demands more electricity, more cooling, more hardware turnover, more territorial concentration, and more infrastructural integration. Environmental externalities and economic monopolization consequently advance together.

This dynamic destabilizes early digital imaginaries associated with decentralization and democratized participation. The contemporary AI ecosystem increasingly reveals the opposite trajectory: concentrated infrastructure, opaque governance, and asymmetrical access to computational power. The internet once promised horizontal participation. The AI economy increasingly resembles vertical integration.

Ugo Mattei's argument becomes particularly relevant in this context. Technological systems progressively displace traditional juridical mediation through informational architectures capable of organizing behavior directly (Mattei, 2025, p. 4–17). The issue is therefore not simply technological mediation of law, but the emergence of infrastructures whose operational logic begins performing quasi-regulatory functions historically associated with democratic institutions.

The environmental implications remain insufficiently theorized. Water systems, electrical grids, and ecological capacities increasingly sustain infrastructures governed beyond ordinary territorial accountability mechanisms. Ecological limits continue existing physically, but politically they become flexible whenever computational acceleration acquires strategic status.

The ecological critique of artificial intelligence therefore cannot remain restricted to carbon accounting or efficiency metrics alone. The deeper issue concerns the emergence of a political regime in which certain infrastructures become too central, too profitable, and too geopolitically important to constrain effectively.

Artificial intelligence does not simply introduce new technologies into existing legal systems. It pressures democratic institutions to normalize exceptional infrastructural arrangements under the

rhetoric of inevitability. The long-term risk may not be ecological degradation alone, but the gradual erosion of the political capacity to impose limits upon technologically mediated accumulation itself.

CONCLUSION

Artificial intelligence infrastructures challenge one of the most persistent myths of digital society: the belief that computation floats above material constraints. It does not. Large-scale AI systems depend on electricity, freshwater, minerals, land, cooling operations, logistics, and regulatory tolerance. The cloud has never been immaterial. It was only made to look that way.

This article argued that hyperscale AI infrastructures cannot be analyzed through sustainability indicators or carbon-efficiency metrics alone. Their environmental relevance lies in a broader reorganization of ecological resources. Water systems, electrical grids, territorial planning, and mineral supply chains increasingly become aligned with computational imperatives shaped by platform capitalism and frontier AI development.

That movement exposes a mismatch within traditional Environmental Law. Classical regulatory frameworks were largely designed for visible pollution, localized damage, and identifiable causal chains. AI infrastructures operate differently. Their impacts are cumulative, distributed, technically mediated, and often hidden behind contractual opacity, corporate ESG language, and the aesthetic neutrality of digital systems.

The central contribution of this article is the concept of infrastructural exceptionality. As argued throughout the analysis, data centers increasingly receive preferential regulatory treatment, accelerated licensing, privileged access to energy and water, subsidized infrastructural arrangements, and reduced democratic scrutiny under narratives of innovation, competitiveness, digital sovereignty, and geopolitical urgency. The exception does not arrive as a dramatic suspension of law. It arrives quietly, as administrative accommodation: environmental limits are not abolished, but recalibrated around the operational tempo of computational acceleration.

This is where the political problem becomes sharper. Environmental externalities are socially distributed, while computational capacity remains concentrated in a small number of corporations controlling cloud architectures, semiconductor chains, proprietary models, and deployment ecosystems. The result is not merely an ecological imbalance, but a redistribution of power.

The territorial dimension reinforces the diagnosis. Regions marked by hydric stress, political dependency, or weak bargaining capacity may become attractive precisely because they offer the conditions demanded by hyperscale expansion. In this sense, AI infrastructures reproduce patterns of ecological unequal exchange: some territories provide water, energy, minerals, and climatic absorption; others capture data, computation, and value.

The research question can therefore be answered in these terms: the global expansion of hyperscale artificial intelligence infrastructures reorganizes access to ecological resources not merely through consumption, but through legal, economic, and geopolitical mechanisms that transform water, energy, land, and cooling capacity into strategic conditions of computational power. Within the contextual scope adopted here — centered on the United States–European Union technological axis, with comparative attention to Latin America and resource-stressed territories — infrastructural expansion reveals how ecological scarcity becomes administratively negotiable whenever it is subordinated to digital sovereignty, platform concentration, and geopolitical competition.

Sustainability discourse also plays an ambivalent role. “Green AI,” “carbon neutrality,” and “sustainable cloud infrastructure” may indicate relevant technical improvements, but they may also neutralize critique by translating ecological conflict into managerial optimization. Like the clean interfaces of dystopian cinema or the polished corporate cities of contemporary videogames, the surface promises control while the machinery underneath keeps extracting.

The environmental question surrounding AI is therefore inseparable from democratic governance. The issue is not only whether data centers consume too much water or electricity. The deeper question is who decides how much ecological sacrifice is acceptable, which territories must absorb it, and whether democratic institutions still possess the authority to impose limits on infrastructures presented as inevitable.

The recent accelerationist rhetoric of defense-oriented Silicon Valley makes this problem more visible. When AI is framed as hard power, civilizational duty, military infrastructure, or strategic destiny, ecological restraint begins to look politically secondary. Palantir’s manifesto-style discourse around *The Technological Republic* is useful here not because it caused the problem, but because it states the grammar bluntly: software becomes sovereignty; AI becomes deterrence; infrastructure becomes destiny. Water, energy, and territory remain outside the heroic frame. They appear only as background conditions of technological survival.

For Digital Constitutionalism, the implication is direct. If computational infrastructures increasingly organize communication, economic coordination, scientific research, military capacity, and environmental allocation, they begin to approximate quasi-sovereign systems. They do not merely operate within legal orders; they condition the practical capacity of law to regulate technological accumulation.

The article therefore concludes that artificial intelligence should not be understood merely as software, automation, or informational abstraction. It is a planetary infrastructure with ecological weight, territorial consequences, and political priority. The challenge for Environmental Law, Digital Law, Political Ecology, and Human Rights is no longer only to make AI cleaner. It is to make it governable.

The cloud, in the end, has territory. It has infrastructure. It has thirst. And unless democratic law recovers the power to say no, it will also have priority over ecological reversibility.

REFERENCES

AGAMBEN, Giorgio. *State of exception*. Chicago: University of Chicago Press, 2005.

BECK, Ulrich. *Risk society: towards a new modernity*. London: Sage Publications, 1992.

BLACK MIRROR. Created by Charlie Brooker. United Kingdom: Channel 4 / Netflix, 2011–present.

BLADE RUNNER 2049. Directed by Denis Villeneuve. United States: Warner Bros. Pictures, 2017.

BOMMASANI, Rishi et al. *On the opportunities and risks of foundation models*. Stanford: Stanford University, 2021. Available at: https://crfm.stanford.edu/report.html. Accessed on: 14 Apr. 2026.

BRATTON, Benjamin H. *The Stack: on software and sovereignty*. Cambridge: MIT Press, 2016.

CODED BIAS. Directed by Shalini Kantayya. United States: 7th Empire Media, 2020.

COULDRY, Nick; MEJIAS, Ulises A. *The costs of connection: how data is colonizing human life and appropriating it for capitalism*. Stanford: Stanford University Press, 2019.

CRAWFORD, Kate. *Atlas of AI: power, politics, and the planetary costs of artificial intelligence*. New Haven: Yale University Press, 2021.

CYBERPUNK 2077. Developed by CD Projekt Red. Poland: CD Projekt, 2020.

DEATH STRANDING. Directed by Hideo Kojima. Japan: Kojima Productions, 2019.

ELYSIUM. Directed by Neill Blomkamp. United States: TriStar Pictures, 2013.

FISHER, Mark. *Capitalist realism: is there no alternative?* Winchester: Zero Books, 2009.

FOUCAULT, Michel. *Nascimento da biopolítica*. São Paulo: Martins Fontes, 2008.

FROSTPUNK. Developed by 11 bit studios. Poland: 11 bit studios, 2018.

GRIFFITHS, Brent D. Palantir's summary of CEO Alexander Karp's manifesto is generating buzz: read the 22 bullet points. *Business Insider*, New York, 20 Apr. 2026. Available at: https://www.businessinsider.com/palantir-manifesto-alex-karp-technological-republic-summary-2026-4. Accessed on: 25 May 2026.

HU, Tung-Hui. *A prehistory of the cloud*. Cambridge: MIT Press, 2015.

IEA – INTERNATIONAL ENERGY AGENCY. *Electricity 2024: analysis and forecast to 2026*. Paris: IEA, 2024. Available at: https://www.iea.org/reports/electricity-2024. Accessed on: 03 May 2026.

JONES, Nicola. The AI boom could use a shocking amount of electricity. *Nature*, London, vol. 628, no. 8006, p. 698–700, 2024. Available at: https://www.nature.com/articles/d41586-024-00478-x. Accessed on: 11 May 2026.

KARP, Alexander C.; ZAMISKA, Nicholas W. *The technological republic: hard power, soft belief, and the future of the West*. New York: Crown Currency, 2025.

KARPF, Dave. Palantir's Manifesto Is as Subtle as a MAGA Hat. *Tech Policy Press*, Washington, 21 Apr. 2026. Available at: https://techpolicy.press/palantirs-manifesto-is-as-subtle-as-a-maga-hat/. Accessed on: 25 May 2026.

KHAN, Lina M. The separation of platforms and commerce. *Columbia Law Review*, New York, v. 119, n. 4, p. 973–1098, 2019. Available at: https://columbialawreview.org/content/the-separation-of-platforms-and-commerce/. Accessed on: 18 Apr. 2026.

LI, Pengfei et al. Making AI less "thirsty": uncovering and addressing the secret water footprint of AI models. *arXiv preprint*, 2023. Available at: https://arxiv.org/abs/2304.03271. Accessed on: 07 May 2026.

MARINONI, Andrea et al. *The Data Heat Island Effect: quantifying the impact of AI data centers in a warming world*. *arXiv preprint*, 2026. Available at: https://arxiv.org/abs/2603.20897. Accessed on: 16 May 2026.

MARX, Paris. *Road to nowhere: what Silicon Valley gets wrong about the future of transportation*. London: Verso, 2024.

MASANET, Eric et al. Recalibrating global data center energy-use estimates. *Science*, Washington, vol. 367, no. 6481, p. 984–986, 2020. Available at: https://www.science.org/doi/10.1126/science.aba3758. Accessed on: 22 Apr. 2026.

MATTEI, Ugo. *O fim do direito: a grande substituição tecnológica na nova era*. Milão: Feltrinelli, 2025.

METAL GEAR SOLID V: THE PHANTOM PAIN. Directed by Hideo Kojima. Japan: Kojima Productions; Konami, 2015.

MR. ROBOT. Created by Sam Esmail. United States: USA Network, 2015–2019.

POLANYI, Karl. *A grande transformação: as origens da nossa época*. 2. ed. Rio de Janeiro: Elsevier, 2001.

POLITICO. EU clinches deal to roll back AI restrictions. Brussels, 7 May 2026. Available at: https://www.politico.eu/article/eu-clinches-deal-to-roll-back-ai-restrictions/. Accessed on: 21 May 2026.

SRNICEK, Nick. *Platform capitalism*. Cambridge: Polity Press, 2017.

UNIVERSITY OF OTTAWA. Technofascism: Palantir's Manifesto. Ottawa, 23 Apr. 2026. Available at: https://www.uottawa.ca/en/news-all/academic-expertise-technofascism-palantirs-manifesto. Accessed on: 25 May 2026.

VILLAPALOS, Víctor. "Devemos defender algoritmos democráticos e evitar sucumbir a uma abordagem centrada apenas em dados, uma 'datafilia'": interview with Professor Yves-François Le Coadic. *Archivoz Magazine*, 11 May 2026. Available at: https://www.archivozmagazine.org/en/we-must-defend-democratic-algorithms-and-avoid-succumbing-to-a-data-centered-approach-a-dataphilia-interview-with-professor-yves-francois-le-coadic/. Accessed on: 24 May 2026.

WIPO – WORLD INTELLECTUAL PROPERTY ORGANIZATION. *Global Innovation Index 2025: innovation at a crossroads*. Geneva: WIPO, 2025. Available at: https://www.wipo.int/publications/en/details.jsp?id=4807. Accessed on: 09 Apr. 2026.

ZUBOFF, Shoshana. *The age of surveillance capitalism*. New York: PublicAffairs, 2019.