



OPEN ACCESS

EDITED BY

Charalampos Alexopoulos,
University of the Aegean, Greece

REVIEWED BY

Theodoros Papadopoulos,
University of the Aegean, Greece
Ikhlef Jebbor,
Ibn Tofail University, Morocco

*CORRESPONDENCE

Sebastián Rivero-Silva
✉ sriveros@usj.es

RECEIVED 20 May 2025
ACCEPTED 16 June 2025
PUBLISHED 15 July 2025

CITATION

Rivero-Silva S, Chinarro Vadillo D and
Prieto-Andres A (2025) The green algorithm:
can sustainability define the winner in the AI
race? *Front. Polit. Sci.* 7:1629914.
doi: 10.3389/fpos.2025.1629914

COPYRIGHT

© 2025 Rivero-Silva, Chinarro Vadillo and
Prieto-Andres. This is an open-access article
distributed under the terms of the [Creative
Commons Attribution License \(CC BY\)](#). The
use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in
this journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

The green algorithm: can sustainability define the winner in the AI race?

Sebastián Rivero-Silva^{1*}, David Chinarro Vadillo¹ and
Antonio Prieto-Andres²

¹School of Architecture and Technology, Universidad San Jorge, Villanueva del Gállego, Spain, ²Faculty of Communication and Social Sciences, Universidad San Jorge, Villanueva del Gállego, Spain

This article explores whether environmental sustainability may become a strategic axis in the evolving AI rivalry between China and the United States. By comparing ChatGPT and DeepSeek, it examines how ecological efficiency, data sovereignty, and infrastructural autonomy intersect with national AI strategies. While ChatGPT remain cloud-dependent and resource-intensive, DeepSeek—according to unverified developer data—prioritizes offline deployment and energy-efficient design, aligning with China's pursuit of techno-sovereignty. Still, potential ecological gains may be undermined by online variants or outdated hardware. Also, the literature highlights security risks associated with DeepSeek's distilled models. This analysis, grounded in a case study that is not fully representative but rather illustrative, shows that sustainability is no longer peripheral but increasingly regarded as an important element of geopolitical agendas. Although it remains premature to conclude that it is a decisive axis of technological competition, current evidence suggests a gradual reframing of strategic priorities toward more responsible innovation.

KEYWORDS

AI race, DeepSeek, ChatGPT, sustainability, data sovereignty

1 Introduction: genesis and nature of the AI race

Certain authors have observed that we are currently experiencing a genuine international race for the mastery of artificial intelligence (AI) (Poo, 2025, p. 2), comparable to other notable international races in modern history, such as the space race or the nuclear arms race. In this so-called AI race, China and the United States are particularly involved, with the launch of DeepSeek constituting a true turning point in that competition—especially in the sustainability arena (Moravec et al., 2025, p. 4–5). Other authors, with considerable insight, prefer to use the term “race for data domination” (George, 2025) to describe this covert power struggle between the USA and China, which is seemingly reflected in the technological contest between the U.S.-based ChatGPT and the Chinese DeepSeek. As we shall see below, these two concepts, data sovereignty and sustainability are closely interconnected.

Scholars, as we will see, consistently point to two major challenges arising from AI development: on one hand, (a) the environmental challenge and, on the other, (b) the challenge of data sovereignty. Environmental risk is associated with pollution generated by electricity consumption, e-waste, and other forms of waste, including the cooling water used by data centers hosting online AI models. Concerning the latter, estimates suggest that, by 2028, as much as 20% of the 90 GW (788.4 TWh) of energy likely to be consumed by data centers worldwide will be allocated exclusively to AI. Within that

figure, around 15% of the consumption is expected to be dedicated to AI training, while the remaining percentage will go to pure inference (i.e., AI responses to user queries) (Avelar et al., 2023, p. 2). In this regard, various authors (Ding et al., 2025, p. 2) have proposed that the 390 most widely used generative AI (GAI) models (excluding DeepSeek) currently consume between 24.97 and 41.10 TWh of energy, approximately equivalent to Portugal's annual energy consumption, generating between 10.67 and 18.61 million tons of carbon emissions. Notably, the United States and China together account for 99% of CO₂ emissions related to GAI—China emitting 6.76–8.98 million tons and the U.S. emitting 3.66–8.72 million tons—whereas Europe emits only 0.02–0.09 million tons. Concerning the e-waste associated with GAI, projections suggest that by 2030, accumulated waste could reach 16 million tons (Wang P. et al., 2024, p. 3). This amount is roughly equivalent to the average annual emissions associated with forest fires in Spain over the last decade.

In parallel, the second risk concerns data sovereignty. This term, occasionally polysemic or ambiguous, is directly tied to emerging technologies (Hummel et al., 2021, p. 13). In this context, it refers to the genuine risk stemming from the highly probable military and intelligence uses that can be made of the massive volume of data gathered through a technological medium such as AI. Put simply, the concern is that millions of citizens in one country may depend daily on a foreign AI platform to which they supply all manner of data. This data can then be leveraged for intelligence or military purposes. The problem is not new: the Chinese-owned app TikTok was banned in India and subjected to tight restrictions in the U.S. (Kumar and Thussu, 2023, p. 12–13), precisely to protect national interests and prevent the massive flow of data toward China. This is far from a trivial matter, given that, according to OpenAI's CEO, ChatGPT reached 300 million weekly users in 2025, “despite DeepSeek” (Rooney, 2025).

In this context, some authors (Naghiyev, 2024, p. 7–8) have described ChatGPT's privacy policy as “unclear,” which further intensifies uncertainty about how the millions of personal data points from those 300 million weekly users are processed, stored, and utilized by OpenAI. Another group of scholars (Cartwright et al., 2024, p. 14) agrees on the need for more advanced information security mechanisms but also notes that significant efforts have been made to anonymize data supplied by ChatGPT users, seeking to render personal identification impossible. However, in our view, the real problem of data sovereignty lies not so much in whether a user can be identified—which is also a concern—but rather in what happens when, for instance, a scientist uses ChatGPT to explore potential improvements for a draft industrial patent, or a high-ranking official asks the AI to draft an email intended for another high-ranking official. Even if such information is anonymized, it may include key details with potential implications for national security or intelligence. Where does that data end up, and how is it handled, given that it is ostensibly not personal? As certain researchers have found (Wu et al., 2024, p. 110) ChatGPT, denies collecting user information. However, this claim is ambiguous and contingent upon the user's “consent” (Cartwright et al., 2024, pp. 5–6). According to the current OpenAI privacy policy, conversation history is retained “until it is no longer useful for providing the service” (OpenAI, 2025) or until the user explicitly chooses to delete it— (OpenAI,

2024) an example of politically correct language aligned with European data protection regulations (Sebastian, 2023, p. 4). Be that as it may, the exact use of the data in the possible ‘live’ training of ChatGPT remains unclear, as does the point at which such data ceases to be considered useful. In the worst-case scenario, it is assumed that this would occur only upon the closure of the user's account with OpenAI.

Hence, the rising global dependency on AI—particularly the U.S.-based ChatGPT—and its associated ramifications, as discussed by certain authors (Salah et al., 2024, p. 5), have generated alarm among Chinese authorities regarding the large-scale outflow of data toward the United States. Paradoxically, this is a game that China itself had been playing for years with TikTok and U.S. citizens. Governments cannot deprive citizens of foreign AI once they have become reliant on it, yet they also cannot coexist indefinitely with an exponentially increasing environmental impact that could threaten achieving the United Nations' sustainable development goals (Fan et al., 2023, p. 14). This is the ideal breeding ground for a race to (a) develop national AI technologies (that either utilize national data centers or do not depend on data centers at all) and (b) ensure that this form of “AI self-defense” does not lead to an unmanageable environmental impact. For instance, in China, electricity demand from data centers is expected to reach ~300 TWh in 2026 and 400 TWh by 2030, while at certain stages in its development, OpenAI in the United States reportedly doubled the energy consumed for training its models in less than a year (Stacciarini and Gonçalves, 2025, 2, p. 14). In short, AI's environmental footprint appears increasingly unsustainable over time, which means that streamlining these models becomes a pivotal factor in the battle for AI dominance—particularly as a defining feature that distinguishes DeepSeek from ChatGPT.

In this regard, it is worth citing Humby's perspective (Farronato, 2025, p. 1), who argues that data are the “new oil” and that both states and major corporations share a vested interest in controlling vast troves of citizens' data. This ambition – the control of (inter) national data—can only lead to insisting on “online” AI systems to enable unfettered tracking of these data, thereby creating an important and ongoing demand for energy, computing resources, and water. The latter point is particularly noteworthy in that nearly 57% of the water used by data centers is potable (Mytton, 2021, p. 2), with daily consumption in the United States reaching up to 1.8 billion liters. In this regard, Crawford (Gillett, 2023) warns of the ethical risks associated with depleting freshwater reserves to sustain the operation of artificial intelligence systems. As we will see below, China has departed, albeit with some reservations, from this unsustainable approach by prioritizing environmental sustainability and allowing DeepSeek to function—if consumer wants it –, “offline,” which inherently relinquishes continuous data tracking.

It is also particularly relevant to note that, unlike the other international competitions of the Cold War era, this AI race does not ignore environmental considerations and sustainability, in contrast to, for example, the nuclear arms race (Crowley and Ahearne, 2002). On the contrary, this international race is, for the first time, influenced by the sustainability of the AI being developed. It is possible that, who manages to create the most efficient AI—offering higher performance at lower energy, computational, and e-waste costs—may enjoy a decisive advantage over its rival.

Indeed, this has apparently been China's perspective in developing the offline version of DeepSeek, as we will see next.

This article addresses the following research question: can environmental sustainability become a decisive factor in the Sino-American race for AI supremacy, as reflected in the evolution of ChatGPT and DeepSeek? While most existing analyses focus on computational power, and algorithmic sophistication, this study posits that ecological efficiency—particularly in terms of energy consumption, infrastructure design, and environmental externalities—may increasingly determine strategic advantage. By comparing the architecture and deployment models of DeepSeek and ChatGPT, the article explores whether sustainable AI systems represent not only technological innovation but also a shift in the logic of international technological rivalry.

2 Materials and methods

This research employs a qualitative, desk-based methodology (Travis, 2016) aimed at exploring whether environmental sustainability can act as a determining factor in the current geopolitical race for artificial intelligence supremacy, with a particular focus on the United States and China. The study is structured around a doctrinal and comparative analysis of publicly available technical documentation, academic literature, and official policy papers related to the development and deployment of large language models, specifically ChatGPT and DeepSeek.

The first step consisted of examining the technical features of DeepSeek—such as its Mixture-of-Experts routing, Multi-Token Prediction, and knowledge transfer techniques—as reported in developer papers and third-party evaluations. These were compared to the architecture and operational model of ChatGPT, particularly its dependence on data centers and higher energy demands. Special attention was given to quantitative claims regarding electricity consumption, training costs, and environmental externalities (e.g., CO₂ emissions, water usage, and e-waste), in order to assess how each model aligns with or departs from sustainability goals.

In parallel, the study incorporated a geopolitical lens by reviewing sources that discuss the strategic use of AI in the context of data sovereignty, cyber-defense, and digital infrastructure. Case-specific analyses of China's offline deployment strategy and the U.S. preference for data-centralized models were used to understand how environmental and technological variables intersect with national security concerns. The method is interpretive and comparative, seeking to connect the technical configuration of AI models with broader patterns of global competition and resource efficiency. No experimental data were generated; all insights are derived from published, verifiable sources.

As this is a topic that necessarily requires combining high academic standards with policy papers and unverified developer documents, a specific table (Figure 1) indicating the type of sources used in the study is included below.

In connection with the above, Figure 1 provides a representative overview of the types of documents used in the preparation of this article, with particular emphasis on “Developer documentation” and “Technical whitepapers,” which, as discussed in Section 5, constitute a methodological limitation.

3 DeepSeek in the context of the confrontation between China and the United States

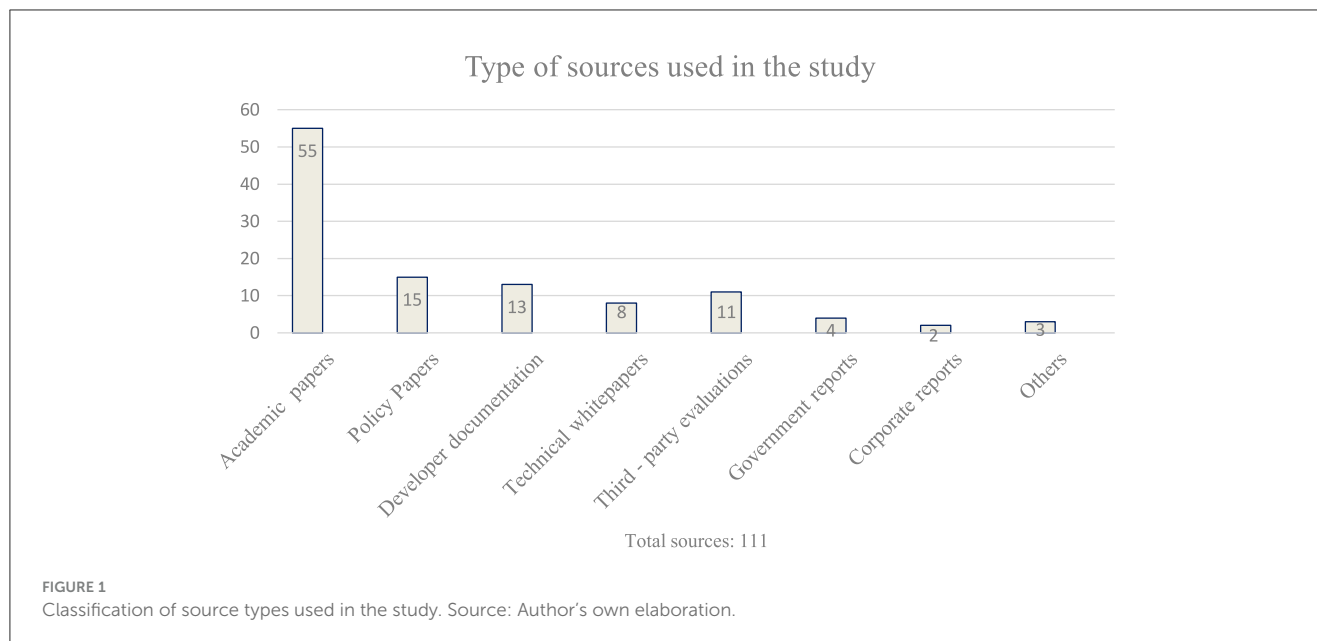
3.1 DeepSeek: sustainability as a key element

3.1.1 Technical innovations and energy optimization in DeepSeek's architecture

DeepSeek is a large language model developed by the Chinese company Hangzhou DeepSeek Artificial Intelligence Co., Ltd. While not directly state-controlled, it operates under censorship aligned with the Communist Party's ideological framework (Gorlla and Tuttle, 2025, p. 6–10) and emerged within a system where private firms—particularly in strategic sectors—are subject to intense political oversight (Li et al., 2020; Almén and Carlsson, 2025). Scholars have also highlighted that such state influence over commercial actors is typical of socialist regimes (Rivero Silva, 2022, p. 13–14). Thus, while DeepSeek cannot be classified as a state-owned AI, its development likely depended on government support, and its functions reflect political alignment with Party standards. As Feakin (2025) argues, it functions as an instrument of Beijing's soft power.

The Chinese state has shown minimal tolerance for technologies beyond its control—going so far as to attempt bans on cryptocurrency (Chen and Liu, 2022). In this context, DeepSeek's rise has occurred with clear ideological authorization. Accordingly, some authors describe it as a tool of “AI diplomacy” (Truby et al., 2025, p. 4), understood as the strategic use of AI to shape international relations and advance national agendas.

DeepSeek—particularly its widely downloaded version, DeepSeek-V3—relies on a specialized neural network named DeepSeekMoE, referencing its Mixture of Experts (MoE) architecture, whose main goal, is to increase model size without proportionally raising the computational cost per token. In simple terms, MoE is an architecture that allows the AI model to activate only a small subset of its components (or “experts”) for each input, instead of using the entire network every time. This makes it possible to increase the model's overall size while keeping energy and computation use relatively low. Thanks to this unique neural design, according to its developers, DeepSeek, may achieve fivefold lower computational overhead in training compared to other models of similar scale. DeepSeek MoE's architecture translates directly into markedly lower operational footprints. Its developers report that a 16-billion-parameter model performs only 39.6 % of the computations required by a comparable dense baseline; moreover, the same design delivers nearly 2.5 times the inference speed of a 7-billion-parameter dense model. These figures could imply proportional reductions in electricity draw, rack-level cooling, and embodied-carbon demand during inference. Taken together, these data suggest that DeepSeek MoE may curtail per-query computation by roughly 60 % and eliminate the need for multi-GPU clusters, whereas GPT-style dense models concentrate their largest environmental burden in resource-intensive training cycles and still incur higher per-generation energy costs. In life-cycle terms, the sparse-activation strategy therefore offers a materially more sustainable



pathway for large-scale language-model deployment (Dai et al., 2024, p. 17–18).

The adoption of DeepSeekMoE, according to its developers, represent a significant competitive advantage for this family of AI models over ChatGPT's "Self-Attention" approach (Vaswani et al., 2017). While DeepSeek V3 handles a 671-billion-parameter language model, only ~37 billion parameters are activated per processed token. In other words, for any given query or input, the model would only use around 5% of its weight drastically which could significantly reduce the computational demands for each user request (DeepSeek-V3 Team, 2025, p. 3). Cai et al. (2025, p. 21) report persistent rumors that ChatGPT-4 may incorporate a variant of the MoE architecture, although it is unclear whether such an implementation would be partial or comprehensive, and the claim has not yet been independently verified. Consequently, the publicly available evidence still indicates that OpenAI's large-scale language models rely on transformer-based neural networks.

A further defining feature of DeepSeek is Multi-Token Prediction (MTP)—a decoding strategy that improves speed and efficiency during inference. In simple terms, rather than generating one word (or token) at a time, the model predicts several upcoming tokens in advance and then verifies which ones make the most sense in context. This approach significantly accelerates response time and reduces the computational load.

In practical terms, MTP can yield up to threefold faster response times in specific areas like coding (Gloeckle et al., 2024, p. 3). By its very nature, it entails lower computational demands and consequently saves energy. Rather than generating one token at a time, the model contemplates an entire range of possible sequences in advance and selects the most coherent answer. This approach was previously implemented, with moderate success, in other generalist models like Llama, specifically adapted for voice recognition (Raj et al., 2025). However, another school of thought urges caution about MTP's outcomes, noting that claiming to handle multiple, complex scenarios solely via straightforward autoregressive predictions might not be as effective as some advocates claim (Bachmann and Nagarajan, 2024, p. 9).

Either way, the data presented by DeepSeek's development team suggests that training DeepSeek V-3 required a total of 2,788 GPU hours on H800s. Assuming an approximate price of 2 USD per H800 GPU hour, the creation of this model cost around 5.6 million USD (Liu et al., 2024, p. 5). Building DeepSeek V-3 is DeepSeek R1 Zero, the next iteration in the DeepSeek family. This version uses unsupervised learning via an algorithm known as GRPO (Group Relative Policy Optimization). Operating within the broader scope of RL (Reinforcement Learning), GRPO allows for (a) training without human supervision and (b) no need for a "critical" model to assess whether a given response is "good or bad," a method sometimes referred to as "semi-supervision." Essentially, it establishes a system of rewards and penalties, sampling a range of answers to estimate the advantage of each one—whether that advantage relates to format (the style of the response) or precision (the substance) (Guo et al., 2025, p. 5).

DeepSeek's dataset would also have been meticulously refined for maximum efficiency. It features a corpus of 2 trillion tokens that, according to its development team, is periodically expanded. By and large, it comprises a C4 or Common Crawl (an automated crawl representing the "general knowledge" of the web) from which several operations are performed: (a) deduplication, to eliminate redundancy, (b) filtering, to remove toxic or trivial material, and (c) remixing, to ensure sufficient variety and balance in DeepSeek's knowledge base (Bi et al., 2024, p. 4–5). Moreover, the tokenization process—essential for converting dataset inputs into information the AI can process—employs the Byte-level Byte-Pair Encoding (BBPE) algorithm, a mathematical technique capable of creating consistent response patterns and therefore encapsulating more information in fewer tokens. In general, this speeds up the AI's data processing (Ghafari et al., 2024, p. 2).

Additionally, DeepSeek leverages another model-reduction technique known as knowledge transfer, which has demonstrated significant improvements in energy consumption and inference speed (Yuan et al., 2024). Specifically, DeepSeek employs a "teacher-student" inference approach. A large, comprehensive model is used to train specialized, streamlined sub-models, each

containing only knowledge relevant to a particular function (Xu et al., 2025, p. 6). The goal is to produce hyper-focused models whose datasets exclude content that is not essential for their specific purpose. Accordingly, “mini” DeepSeek versions can be found for various use cases, such as math (Shao et al., 2024), medicine (Zhang, 2025), or coding (Zhu et al., 2024). This approach tackles both critical concerns in the AI race: (a) by offering numerous domain-specific models, Chinese users may eventually be more inclined to rely on DeepSeek for specialized queries, thus keeping data within national borders; and (b) because these are “lite” versions, as described next, they can run on personal computers and smartphones, thereby reducing reliance on large-scale data centers. Nevertheless, the proliferation of “mini” models can pose a clear disadvantage in multidisciplinary contexts, since their development would have to be specifically tailored to a single, well-defined task.

In short, DeepSeek according to its developers, could combine a range of cutting-edge techniques aimed at minimizing the size, consumption, and environmental impact of LLM-based AI models. Indeed, the distilled AI sub-models referenced earlier do not even require a data center to function—they can operate offline on ordinary consumer-grade computers, reducing electricity usage and e-waste generation compared to ChatGPT (Ball, 2025). In essence, DeepSeek promotes a decentralized AI that does not bear the heavy sustainability burden of data centers. For instance, Sapkota et al. (2025, p. 17), based on developers’ work, estimate that DeepSeek R1 produces merely 3.3 percent of the carbon emissions attributed to ChatGPT-4 during the pre-training phase. Although no conclusive results are available regarding the online inference phase, preliminary evidence suggests that, as will be discussed below in section 3.1.2, the DeepSeek R1 online version may produce even higher emissions than ChatGPT 4.5.

Likewise, although no exact figure has been disclosed, scholars widely agree that the model achieves a marked reduction in computational demand—and therefore, electricity consumption—(Bogmans et al., 2025, p. 13) at least during the pre-training phase. Should one choose to run DeepSeek offline, some scholars argue that it is ideally suited to reusing compact systems, such as second-hand smartphones or laptops (Ngoy et al., 2025, p. 6), capable of handling lightweight AI models—something that, from one perspective, could contribute to a circular economy strategy by facilitating the reuse of electronic devices, thereby potentially reducing carbon emissions and e-waste. However, an alternative view raises concerns that the widespread offline use of DeepSeek on outdated or energy-inefficient hardware—especially in under-resourced settings, might ultimately result in unpredictable and suboptimal electricity consumption, thus offsetting any presumed environmental benefit. These contrasting interpretations point to an unresolved question: whether the offline deployment of AI on consumer-grade devices will in fact deliver meaningful ecological gains or merely shift the sustainability burden in new and less visible ways—a question that only longitudinal observation and further empirical validation may fully answer. Be that as it may, and irrespective of the inference phase involving offline user interaction, the scholarly consensus is that the production costs for DeepSeek (particularly the R1 version) have been substantially lower than those of its U.S. counterpart, ChatGPT, to the extent that while DeepSeek V3 was purportedly developed for under 5.6

million USD, ChatGPT 4 may have cost between 80 and 100 million USD (Krause, 2025, p. 3).

3.1.2 Empirical caveats and global trade-offs in DeepSeek’s sustainability claims

Researchers at the Massachusetts Institute of Technology maintain a skeptical view of the environmental advances claimed by the DeepSeek development team. First, Knight (2025) critically reflects on the fact that it is DeepSeek itself that has highlighted how inexpensive AI pre-training can be—characterizing this emphasis as a calculated move. This distinction is crucial: the development of AI models generally involves two main stages—(a) pre-training, where DeepSeek claims significant efficiency gains, and (b) inference (online or offline), which refers to user interaction or the execution of natural language processing (NLP) tasks. On this point, O’Donnell (2025) argues that the favorable environmental results associated with DeepSeek—as well as the genuine competitive advantage of the aforementioned MoE and MTP architectures, which is not disputed—lie specifically in the pre-training phase, not in inference. Indeed, it is worth underscoring that the results reported by Sapkota, Raza, and Karkee—indicating that DeepSeek produces just 3.3% of the emissions generated by ChatGPT-4—refer exclusively to the pre-training phase, not to inference. Similarly, the technical challenges involved in comparing carbon dioxide emissions across different AI models, such as DeepSeek and ChatGPT, should not be underestimated—especially given the lack of unified standards in the sector (Luccioni et al., 2024, p. 88).

During DeepSeek’s inference phase, a different scenario emerges from that of pre-training. According to Jegham et al. (2025, p. 7), DeepSeek R1, during online inference, consumes ~23.8 Wh at maximum and 2.1 Wh at minimum to process an input of 100 words and generate an output of 300 words. For the same task, GPT-4.5 consumes a maximum of 6.7 Wh and a minimum of 1.2 Wh. In other words, DeepSeek R1, during online inference, is more energy-intensive than ChatGPT. However, when handling longer tasks—for instance, processing 7,000 input words and generating 1,000 output words—both models consume roughly similar amounts of energy: 33.634 Wh for DeepSeek and 30.459 Wh for GPT-4.5. This places them in a comparable range of computational demand. This observation is relatively significant, as it suggests two key points: (a) DeepSeek may be “cheap to train” in comparison to ChatGPT, but its R1 online version consumes even more energy and water during online inference than its direct North American competitor, ChatGPT-4.5. That said, available data suggests that DeepSeekV3 online—the version preceding R1—has an inference consumption profile relatively similar to that of GPT-4.5.

Regarding DeepSeek’s offline energy consumption, the academic literature is limited. Nonetheless, some sources suggest that certain distilled versions exhibit especially low energy use and can run on systems with modest computational capacity (Chintalapudi et al., 2025, pp. 14–15), although no precise figures are provided. It is important to note that DeepSeekR1—the most recent version of the DeepSeek family—comes in several model sizes, ranging from a lightweight 1.5B (billion parameters) version to a 1.3T (trillion parameters) variant. The following table

(Figure 1), relating to users who download the full DeepSeek model from the Hugging Face platform for offline use, shows a clear preference for the distilled 1.5B and 8B versions—respectively the first and third lightest models available. In relation to the above, Figure 2 below shows a comparison of the number of downloads attributed to each DeepSeek model in the Hugging Face platform.


This is not a trivial issue: enabling offline AI functionality on mid-range personal computers could reasonably support the reuse of older hardware. On this matter, some scholars (Gould et al., 2024, p. 1) point to increasingly shorter hardware obsolescence cycles, driven in part by so-called “planned obsolescence” and growing repair difficulty. While other authors advocate for “design for reparability” (Roskladka et al., 2025, p. 22) as a response to such trends, no clear balance has been reached to date—especially considering that ~70% of global e-waste ends up in developing countries. In this broader context, DeepSeek acquires significant geopolitical relevance, a topic to be explored later in this work. As product lifespans continue to shorten, countries such as Nigeria, Ghana, and India—with limited internet access—frequently receive discarded old or obsolete computers from Western nations (Lopes dos Santos, 2021, p. 58), often under the guise of donations that are, in practice, a form of e-waste export.

Consequently, a model like DeepSeek—extremely inexpensive to train, capable of running in distilled versions with as few as 1.5 billion parameters, and operable across nearly any mid range device—is well suited for countries that (a) have a surplus of outdated hardware and (b) lack widespread internet connectivity. In this sense, DeepSeek, In this regard, DeepSeek’s offline development appears to have been a thoroughly strategic move, given that 25.6% of its population remains without stable internet access (Cui et al., 2024, p. 1). While it offers a meaningful sustainability dimension compared to ChatGPT, particularly in terms of pre-training efficiency and offline inference, it does not renounce online inference or the significant consumption of energy and resources in certain contexts. In other words, it facilitates a more sustainable form of AI suited to modest regions,

outdated devices, or environments with limited internet access. However, it would be premature to characterize DeepSeek as a truly green AI. Importantly, this scenario partially complicates the initial sustainability narrative surrounding DeepSeek: the prospect of millions of outdated computers running distilled versions of DeepSeekR1—even offline—suggests a globally inefficient energy footprint, largely centered in the Global South, which lacks access to state-of-the-art hardware.

Indeed, the evidence suggests a clear bifurcation shaped by politically and commercially driven decisions, reflecting divergent strategic logics in how each state approaches the AI race. For instance, ChatGPT currently serves 300 million active weekly users and includes capabilities such as video generation and voice analysis—functionalities that, by design, entail a higher environmental cost. In contrast, DeepSeek, which deliberately avoids such high-intensity features, only reaches 38 million users. As Li et al. (2020, p. 3) note: “Generating videos consumes significantly more carbon than generating text: the average carbon emission for a single 240p video frame is equivalent to generating 78 text tokens with comparably sized video and text generation models.” Similarly, DeepSeek decision to offer distilled, offline-compatible AI suggests a reduction in user data collection. The absence of video generation—and other—features in DeepSeek, while indicative of a deliberate trade-off prioritizing environmental efficiency, may also reflect other factors, such as a potential lack of the American expertise required to develop such applications. As a counterpoint to the aforementioned, it should not be dismissed that such an environmental perspective may, to a greater or lesser extent, stem from China’s strategic need to position itself competitively in the realm of sustainability, in light of its limited capacity to rival the advancements in artificial intelligence functionalities introduced by ChatGPT. Similarly, the development of DeepSeek could be understood as largely reliant on the foundational and pioneering work carried out by OpenAI through ChatGPT.

That said, these models, ChatGPT and DeepSeek are not monolithic representations of national AI policy. DeepSeek—or



its environmentally conscious offline deployment strategy—should not be viewed as fully representative of China's overall approach to AI. Nor has the United States, as will be seen in the following section, entirely neglected the environmental implications of AI development. Rather, each model serves as an indicative example—valuable for analysis, but not exhaustive of national agendas. For instance, while DeepSeek promotes low-energy offline functionality, its R1 online version has been shown to consume even more energy than ChatGPT for certain inference tasks, further complicating any binary interpretation of sustainability leadership in AI. The use of offline models may exacerbate DeepSeek's existing security shortcomings, posing additional risks to both data security and the integrity of the model itself—for example, due to missing security updates that protect against cyberattacks or the fraudulent use of AI.

Nevertheless, it is worth noting that, despite its sustainability achievements associated with the pre-training phase and the possibility of running distilled versions offline, DeepSeek has not produced similarly promising outcomes in security. Using algorithmic jailbreaking, [Kassianik and Karbasi \(2025\)](#) launched category-based attacks—cyber-crime, disinformation, illegal activity, and general harm—against DeepSeek-R1 with the HarmBench dataset. DeepSeek-R1 failed every test, yielding a 100 % attack-success rate, while competing systems showed at least partial resilience. Crucially, these safety assessments were conducted only in English. When [Zhang et al. \(2025\)](#) switched the lens to Mandarin with CHiSafetyBench, they revealed additional weaknesses: DeepSeek-R1 achieved just 71.14 % accuracy in risk identification (vs. 91.13 % for the strongest baseline) and struggled to refuse sensitive prompts (RR-1/RR-2 = 67.60 %/67.17 %, compared with 77.71 %/77.27 %). DeepSeek-V3 offered modest gains in overall Mandarin safety (accuracy = 84.17 %) yet still underperformed at screening sensitive queries (accuracy = 66.96 %).

Likewise, the specialized DeepSeek version developed for biomedical applications and traditional Chinese medicine has suboptimal results, mainly because it cannot access real-time, specialized databases ([McGee, 2025](#), p. 648). Finally, regarding “distilled” models, some authors ([Lian et al., 2025](#), p. 13) contend that although these sub-models are powerful, they are not necessarily “superior” within their specialized domains. Likewise, it is worth noting that, although ChatGPT has introduced models that may be classified as “lite”—for example, the o4 mini—reliable information on the computational demand of most AI systems remains wholly opaque, thereby precluding any well-founded assessment of their purported advantages ([Chen, 2025](#)).

3.2 Competing sustainability logics in Sino-American AI development

3.2.1 Deep seek: an example of Chinese ideological integration and ecological soft power policy?

The political process by which China enters the competition in the AI race is certainly complex. On the one hand, we must

acknowledge that, despite having a pronounced state ideology, the country is no longer that centralized and bureaucratized system once so closely mirrored by the Soviet Union ([Molina, 2021](#), p. 15). On the contrary, China has evolved into a political system that has successfully adapted to globalization, viewing socialism as a means of political survival rather than an end in itself. Certain authors ([Pieke, 2025](#)) refer to this as “neo-socialism.” Others ([Palmer and Winiger, 2019](#), p. 4; [Li and Christophe, 2024](#), p. 10) propose that the application of expert knowledge to specific problems, rather than “assembly-style” decision-making, constitutes a form of governance characteristic of a “neo-socialist,” (and thus post-revolutionary), model exemplified by present-day China. This is a well-established doctrinal concept ([Callahan, 2023](#), p. 15) that examines China's response to globalization as it seeks to secure its own international hegemony.

This neo-socialist perspective, coexisting with globalization, is largely what has led Communist China to embrace AI from a standpoint that, according to some authors ([Zeng, 2021](#), p. 2–3), is directly tied to national security and set in an unmistakable context of confrontation with the United States. In this regard, influential figures within the Chinese sphere have pointed out that the AI race essentially encompasses two key factors: (a) computational and energy-related considerations, and (b) algorithms and data ([Zeng, 2020](#), p. 1442). DeepSeek, by virtue of its computing attributes previously discussed, fulfills each aspect of this approach: it is national AI (requiring either no data center or relying on national data centers), sustainable (owing to very low energy and computing demands for pre-training phase and offline use), and affordable (having required only six million USD to develop), enabling China to compete with the United States in terms of efficiency rather than efficacy. In essence, it seems that China has recognized the clear environmental degradation associated with AI's technological development and its impact on public health ([Anwar et al., 2018](#), p. 5)—particularly through e-waste and the contamination of cooling water used in data centers. Consequently, this reality has been considered in the national approach to AI within the broader context of its confrontation with the United States ([Roberts et al., 2021](#), p. 47–49).

In line with the ethical framework proposed by [Crawford \(2021\)](#), it is evident that the seemingly positive environmental outlook surrounding DeepSeek requires critical scrutiny that significantly undermine its perceived sustainability. As previously discussed, the online version of DeepSeekR1, particularly for short-form responses, exhibits inference performance that is even more resource-intensive than that of ChatGPT 4.5. This suggests that the environmental ethic—closely linked to the efficiency of AI models—has not been fully internalized across the entire DeepSeek project. Rather, it appears to apply primarily to the pre-training phase and to offline inference. This combination ceases to be environmentally sound once any online variant of DeepSeek is introduced, given its reliance on data centers. In this light, while there is a discernible orientation toward sustainability, it would be inaccurate to characterize China's technological policy as inherently or consistently “green.” In fact, from a scholarly standpoint, the older or partially obsolete machines on which DeepSeek is expected to operate offline—whether in rural or impoverished regions, in countries with restricted internet access, or in authoritarian

regimes such as North Korea, or in African states receiving e-waste—are themselves energy-inefficient. This dynamic, when assessed through the lens of Koomey's Law—recently revisited by Prieto et al. (2025, p. 11)—ultimately diminishes the environmental benefit that would otherwise be associated with distilled offline AI technologies.

In any case, China's approach to the development of DeepSeek cannot be understood without reference to the concept of “techno-nationalism,” originally introduced by Segal and Kang (2006) and later developed by scholars such as Kennedy (2013). At its core, this concept reflects a desire to develop domestically developed technologies, in order to reduce dependence on foreign intervention—particularly from the United States. In essence, it is about leveraging domestic innovation to meet internal needs, such as delivering AI services to public institutions, universities, or hospitals, even in regions where internet access is limited or where state-of-the-art computing technology is unavailable. Equally relevant here is the concept of ecological modernization, first articulated by Mol and Spaargaren and developed in later work (Mol et al., 2014, p. 2), where it is defined as “*the social scientific interpretation of environmental reform processes at multiple scales in the contemporary world.*” The framework suggests that globalization and the resulting interconnection of societies can act as a driver for environmental reform. In this light, the emergence of offline and distilled versions of DeepSeek should be seen not only as part of the global race for AI dominance, but also as a domestic effort to fulfill national priorities—such as data sovereignty and the ability to sustain AI services with limited resources. This reflection points back to a clear contemporary application of ecological modernization theory.

Another crucial point of the AI “consumption” from which China seeks to escape is economic in nature. Data centers are more or less sustainable depending on whether they use clean or “brown” energy, and that aspect directly affects their maintenance costs. According to the authors, Zhang et al. (2011) renewable energy drawn from the grid can be more expensive than brown energy. For example, industrial solar power can cost 16.14 cents per KWh in sunny conditions and 35.51 cents per KWh in cloudy conditions, whereas the wholesale price for brown energy can be around 6 cents per KWh. As for the carbon footprint associated with data centers, China also has a vested interest in this matter. Some scholars (Zhang and Liu, 2022, p. 12) estimate that CO₂ emissions related to this technology in China could be cut by 90% by the year 2060. This aligns closely with observations by other researchers (Li et al., 2023, p. 8–9), particularly regarding China's strategic approach to digital infrastructure. Notably, China launched a Three-Year Action Plan for the Development of New Data Centers (2021–2023) in July 2021, emphasizing the construction of green, low-carbon data centers and the accelerated adoption of advanced green and low-carbon technologies.

Consistent with this trend, Huawei—a pioneer in sustainable ICT practices in China—has introduced various strategies to reduce its carbon footprint. One notable example involves deploying its intelligent automatic cooling system, iCooling, in its data centers, which could well be available to DeepSeek. This technology has reduced total power usage at cooling stations by ~8%–10%, yielding energy savings equivalent to 3.85 million kilowatt-hours—comparable to planting 79,500 trees (Cao

et al., 2022, p. 12). Additionally, Huawei is actively involved in transitioning toward renewable energy. The company prioritizes the use of renewable sources across its operations and is expanding its photovoltaic (PV) infrastructure on its campuses. In 2020 alone, these installations generated 12.6 million kilowatt-hours of electricity.

From a strategic perspective, it may be counter-productive for the People's Republic of China to enter a scale-driven contest over AI efficiency if doing so would exacerbate environmental degradation, impose public-health externalities, and amplify national energy demand. A more sustainable course would be to relax strict data-centrality in favor of offline or Small Language Models (Wang F. et al., 2024)—DeepSeek “lite models” serving as a leading example—whose markedly lower computational requirements, may mitigate both carbon emissions and operating costs. Such an approach would leave the rising energy intensity and attendant ecological footprint of large-scale, cloud-based systems such as ChatGPT to be borne primarily by competing U.S. platforms, thereby reallocating the sustainability burden without compromising China's long-term technological ambitions. Such energy consumption is, according to some recently published studies (Hosseini et al., 2025, p. 2), inexorably destined to occur. Notably, training GPT-3 reportedly consumes about 1,287 MWh of electricity and emits ~552 tons of CO₂, OpenAI's GPT-4 training is said to have used about 6% of the water consumed by West Des Moines, Iowa (population 75,000), and xAI's training lab (responsible for Grok) uses as much power as 80,000 households.

3.2.2 ChatGPT: cloud dependency and emerging environmental policy

The previously mentioned efficiency-driven approach, embodied by DeepSeek, carries an environmental dimension that—albeit with many nuances—differs substantially from the paradigm of efficacy supremacy, which places particular emphasis on data accumulation. Returning to the internal power struggles as a defining feature of international races, the American arena appears tumultuous in this respect. While DeepSeek offline is championed in China as a cornerstone of future sustainability, in the United States multiple AIs vie for national leadership. Although some stand out—like Google's Gemini or Microsoft's Copilot—Grok, created by magnate Elon Musk, has garnered special attention, seemingly addressing an “anti-woke” discourse (de Carvalho Souza and Weigang, 2025, p. 2). In other words, there is an endeavor to amass as much data as possible: ChatGPT for professional applications (such as customer service chatbots) and Grok for social media interactions (DR and IS, 2025, p. 5). Indeed, certain studies suggest that Grok may be more accurate than ChatGPT in some responses, chiefly those related to social media-derived context (Yasir Fahim, 2024).

Notwithstanding the above, while the Chinese approach has often been associated with an environmental sustainability agenda, it would be inaccurate to equate the so-called “American perspective” with a wholesale disregard for environmental concerns in the development of AI. At the governmental level, the recent passage of the Artificial Intelligence Environmental Impacts Act of 2024 (U. S. Congress, 2024) is particularly notable.

This legislation acknowledges the environmental risks posed by AI and mandates a formal report on the matter by the Administrator of the Environmental Protection Agency. Likewise, attention should be drawn to the Executive Order on Advancing United States Leadership in Artificial Intelligence Infrastructure (Biden, 2025), in which former President Biden outlines that national AI development must be guided by three key priorities: (a) achieving supremacy over adversaries in the field of AI (with a veiled reference to China), (b) the economic dimension of AI, including the protection of domestic jobs, and (c) addressing environmental impacts.

Finally, the recent report issued by the U. S. Government Accountability Office (2025 p. 11) explicitly addresses the significant consumption of natural resources linked to AI systems. However, the report also notes that this issue remains “uncertain” at present due to persistent challenges in measurement and standardization. In short, the U.S. government has acknowledged the problem—although, to date, companies like OpenAI continue to roll out increasingly complex and resource-intensive functionalities, such as video generation, with limited legislative limitations. Paradoxically, environmental restraint in AI development appears more visibly embodied—through DeepSeek’s offline model—in China, despite the fact that, according to some scholars (Zhang, 2025, p. 10), the country’s regulatory framework, regarding environment protection, remains relatively weak or, at best, “relaxed” in certain respects.

Likewise, it is important to bear in mind that the consumption of natural resources in online AI systems depends not only on the algorithmic architecture of the model itself, but also on the data centers that host its inference capacity. In this regard, it is worth noting that, with respect to Azure servers, Microsoft is currently developing a server model known as “GreenSKU” (Wang J. et al., 2024, p. 2), which is purported to reduce carbon dioxide emissions by up to 8% across the entire data center infrastructure. Similarly, in the case of AWS servers, Amazon is reportedly experimenting with the integration of renewable energy sources, such as solar and wind power, to reduce the environmental impact of its data centers (Raza et al., 2024, p. 31). In addition, AWS has published a set of best practices under the title Cloud Sustainability Pillars (Amazon Web Services, 2025, pp. 2–5), which constitutes a clear statement of intent in this domain. That said, as of today, there is no independently verified evidence to confirm that these measures have been successfully and systematically implemented within Microsoft and Amazon’s data centers. Therefore, it can be said that both the U.S. government and the major corporations responsible for operating the country’s data centers have taken some initial steps—at the very least—to acknowledge the unsustainable consumption of natural resources associated with AI. However, this growing awareness has not prevented the nation’s leading AI operator, OpenAI, from continuing to increase its weekly user base—reaching the previously mentioned 300 million weekly users—and offering increasingly complex generative AI capabilities, such as video creation through “Sora.” These developments significantly raise the energy and resource demands required for the routine operation of ChatGPT.

It is worth noting the link between the data accumulation interests of systems such as Grok and ChatGPT that it has led to the creation of “smart toys,” i.e., toys with integrated AI. A recent study (Pavliv et al., 2024, p. 172–175) demonstrated that these toys “listen” even when they are not interacting with children, raising serious concerns regarding the data they collect, and the likelihood of those data being used for commercial purposes. Ultimately, the sustainability perspective of DeepSeek is most clearly manifested in environmental terms. ChatGPT and Grok do not allow their models to be used offline (Pham et al., 2024, p. 1–3). That is, they are permanently tethered to a data server, which has prompted some authors to describe the computational demand of such systems as “prohibitive,” particularly for small businesses (Hussein et al., 2025, p. 1; Joublin et al., 2023, p. 1). This is largely explained by its independence from large data centers—that is, through offline or “mini” offline models, plus the possibility of deploying online models on ordinary consumer devices like mobile phones or tablets, thereby eliminating reliance on so-called “mega data centers,” upon which the U.S. still depends. Each of these “mega data centers” hosts hundreds of thousands of servers and can draw tens to hundreds of megawatts of power at peak usage (Zhang et al., 2011, p. 2).

China’s steps toward cheaper, offline versions—even at the cost of relinquishing some data control—suggest that the divide between the two models remains palpable. Although both nations remain interested in data control, China has recognized the untenable nature of accumulating data at any cost. It thus follows that this is a long-term race in which the strategic advantage will not necessarily belong to the AI boasting the most accurate answers or the largest user base. China has realized that genuine strategic leverage lies in attaining a level of sustainability that allows it to maintain a national AI presence, collect data as needed, yet avoid depleting its lakes or being burdened by an insatiable computing and energy demand. Put simply, it seeks to benefit from what AI has to offer while refusing to be overtaken by an international race that is, without question, already underway.

In relation to the above, the following Table 1 presents a SWOT analysis comparing two perspectives: one—albeit with nuances—centered on efficiency, as represented by DeepSeek, and the other—also nuanced—focused on effectiveness and the development of new functionalities, as exemplified by ChatGPT.

4 Discussion and conclusions

This article has explored how environmental sustainability is gradually acquiring strategic importance within the broader geopolitical race for AI supremacy between China and the United States. Evidence drawn from policy frameworks and corporate initiatives suggests that ecological concerns are no longer marginal but increasingly embedded in national AI agendas. Still, it is too early to conclude that sustainability has become a decisive axis. On this matter, Bakhtiarifard et al. (2025) argue that overall sustainability must reconcile environmental tensions with economic and social considerations.

The comparison between ChatGPT and DeepSeek illustrates a fundamental divergence in both technological design and strategic orientation. While ChatGPT remains dependent on cloud-based

TABLE 1 SWOT analysis.

Category	DeepSeek perspective	ChatGPT perspective
Strengths	<ul style="list-style-type: none">- Offline compatibility enhances infrastructure sovereignty and cyber-resilience.- MoE and MTP architecture may reduce pre-training energy and cost.- In general, alignment with China's techno-nationalist and ecological modernization agenda	<ul style="list-style-type: none">- Backed by robust cloud infrastructure ensuring stability and scalability.- Large user base (300 M weekly users) - Strong institutional support (OpenAI, Microsoft)- Leading role in AI ecosystem and innovation standards.
Weaknesses	<ul style="list-style-type: none">- Environmental claims lack independent verification, especially during inference.- Poor performance in AI safety benchmarks (e.g., jailbreak resistance, prompt refusal).- Limited functionality in online mode compared to U.S. competitors	<ul style="list-style-type: none">- High energy and water use tied to centralized data centers.- Full reliance on centralized cloud data centers.- Opaque environmental reporting and resource usage.
Opportunities	<ul style="list-style-type: none">- Potential for AI deployment in regions with limited internet access.- Reduces dependency on foreign data infrastructures.- Supports geopolitical aims through infrastructural and data sovereignty.	<ul style="list-style-type: none">- Green initiatives (e.g., Microsoft GreenSKU, AWS Sustainability Pillars) show promise.- Legislative frameworks emerging for green infrastructure.- Potential to set industry sustainability standards if reformed.
Threats	<ul style="list-style-type: none">- Use of outdated hardware may offset ecological benefits.- Sustainability applies mostly to pre-training and offline use.- Security and robustness concerns may limit adoption.	<ul style="list-style-type: none">- Rising scrutiny over environmental and data governance.- Competitors adopting MoE and MTP could imply a limitation in terms of growth in ChatGPT users.- Environmental regulations could raise operational costs significantly.

infrastructure with significant environmental costs, DeepSeek—according to figures published by its own developers, which await independent empirical verification—claims lower training expenditures and promotes offline, decentralized inference that may reduce energy consumption and dependence on continuous internet connectivity. This architecture supports broader goals of infrastructural autonomy and environmental moderation. However, DeepSeek also allows for online inference, which appears to be significantly more energy-intensive and, at present, cannot match ChatGPT in terms of functional breadth.

This new offline orientation entails a calculated trade-off. While potentially lowering environmental impact, it may also introduce inefficiencies—especially when deployed on outdated or suboptimal hardware—thus potentially offsetting the model's ecological advantages. In any case, the potential environmental benefits of DeepSeek are unlikely to become apparent in the short term. The strategic value of DeepSeek lies as much in its sustainability narrative as in its alignment with

infrastructural independence and data sovereignty. Scholars such as [Okaiyeto et al. \(2025\)](#) have framed this divergence as part of a broader global reconfiguration of AI geopolitics. Beyond the Sino-American rivalry, countries like Russia are pursuing sovereign AI architectures ([Petrella et al., 2021](#)), and as [Morandín-Ahuerma \(2023\)](#) notes, strategic self-sufficiency has become a global aspiration.

Within this context, data sovereignty plays a central role. DeepSeek's offline-compatible models enable local data processing, reinforcing cyber-resilience—as theorized by [Hallaq et al. \(2017\)](#)—and aligning with the view that infrastructural and data control is key to digital sovereignty ([Haney, 2020](#); [Ciuriak and Ptashkina, 2021](#)). These models reflect broader political understandings of vulnerability, control, and systemic risk. The U.S. model, while still cloud-dependent, has not ignored sustainability concerns. Initiatives such as the *Artificial Intelligence Environmental Impacts Act* (2024), Microsoft's GreenSKU program, and Amazon's Cloud Sustainability Pillars suggest a growing environmental discourse. Nonetheless, these efforts remain aspirational and lack robust verification.

This moment marks a departure from earlier models of technological rivalry. Unlike Cold War-era competitions based on extractive and resource-driven metrics ([Laakkonen et al., 2016](#)), today's AI race increasingly values efficiency and ecological viability. DeepSeek's adoption of architectures such as MoE and MTP points to a conscious design logic: to reduce training costs and energy demands while advancing national strategic aims. This technical orientation is deeply interwoven with China's broader techno-nationalist agenda, which seeks to reduce dependency on foreign infrastructure and assert infrastructural sovereignty.

Yet, even this model is not without potential environmental pitfalls, mass deployment of offline AI in resource-constrained environments may reproduce inefficiencies and complicate the standardization of carbon footprints. Ultimately, sustainability has emerged as a visible—if not yet decisive—axis of global AI competition. Whether it will evolve into a dominant logic shaping AI leadership remains uncertain. What is evident, however, is that ecological sustainability, infrastructural self-sufficiency, and responsible deployment are beginning to emerge as critical dimensions of technological power in the twenty-first century.

5 Limitations

The analysis is partly based on developers' information and white papers, which may not fully reflect the technical specifications or energy consumption data of ChatGPT and DeepSeek. This inherent opacity limits full comparability between models. The academic community must remain attentive to new data that is empirically validated.

In addition, although the article positions sustainability as a potential axis of strategic advantage, it does not empirically test this hypothesis through deployment or real-time performance measurement. The argument is grounded in policy trends, technical literature, and theoretical frameworks. As [Avin et al. \(2021, p. 3\)](#) observe, many AI development teams are not subject to independent oversight, complicating

external validation. This issue is further amplified by the absence of standardized reporting on environmental impact, highlighting the need for third-party verification as emphasized by Brundage et al. (2020, p. 8). Future research should incorporate empirical approaches to assess AI deployment, ecological footprint, and strategic relevance across diverse geopolitical settings.

Author contributions

SRS: Writing – review & editing, Conceptualization, Writing – original draft. DC: Writing – review & editing, Writing – original draft, Supervision, Methodology. APA: Methodology, Writing – review & editing, Supervision, Writing – original draft.

Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

References

- Almén, O., and Carlsson, H. (2025). *The Chinese Communist Party's influence over businesses*. Swedish Ministry of Defense. Available online at: https://www.ui.se/globalassets/ui.se-eng/publications/other-publications/particeller_nkk_report_2025.pdf (accessed May 10, 2025).
- Amazon Web Services. (2025). AWS Well-Architected Framework: Sustainability Pillar. Available online at: <https://docs.aws.amazon.com/pdfs/wellarchitected/latest/sustainability-pillar/wellarchitected-sustainability-pillar.pdf#cloud-sustainability> (accessed June 6, 2025).
- Anwar, S., Ghaffar, M., Razzaq, F., and Bibi, B. (2018). E-waste reduction via virtualization in green computing. *American Scientific Research Journal for Engineering, Technology, and Sciences (ASRJETS)*, 41, 1–11. Available online at: <https://core.ac.uk/download/pdf/235050535.pdf> (accessed May 10, 2025).
- Avelar, V., Donovan, P., Lin, P., Torell, W., and Arango, M. A. T. (2023). *The ai disruption: Challenges and guidance for data center design*. Artificial Intelligence in Medicine, 138. Available online at: https://media.datacenterdynamics.com/media/documents/Schneider_Electric_-_Modernization_WP110_V1.1_EN.pdf (accessed May 10, 2025).
- Avin, S., Belfield, H., Brundage, M., Krueger, G., Wang, J., Weller, A., et al. (2021). Filling gaps in trustworthy development of AI. *Science* 374, 1327–1329. doi: 10.1126/science.abi7176
- Bachmann, G., and Nagarajan, V. (2024). The pitfalls of next-token prediction. *arXiv preprint arXiv:2403.06963*. Available online at: <https://proceedings.mlr.press/v235/bachmann24a.html> (accessed May 10, 2025).
- Bakhtiarifard, P., Tözün, P., Igel, C., and Selvan, R. (2025). Climate And Resource Awareness is Imperative to Achieving Sustainable AI (and Preventing a Global AI Arms Race). *arXiv preprint arXiv:2502.20016*. doi: 10.48550/arXiv.2502.20016
- Ball, D. (2025). What DeepSeek r1 Means—and What It Doesn't. *Lawfare*. Available online at: <https://www.lawfaremedia.org/article/what-DeepSeek-r1-means-and-what-it-doesn-t> (accessed May 10, 2025).
- Bi, X., Chen, D., Chen, G., Chen, S., and ... Zou, Y. (2024). DeepSeek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*. Available online at: <https://arxiv.org/pdf/2401.02954> (accessed May 10, 2025).
- Biden, J. R. (2025). Executive Order on Advancing United States Leadership in Artificial Intelligence Infrastructure. The White House Archives. Available online at: <https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2025/01/14/executive-order-on-advancing-united-states-leadership-in-artificial-intelligence-infrastructure/> (accessed June 6, 2025).
- Bogmans, C., Gomez-Gonzalez, P., Melina, G., Miranda-Pinto, J., Pescatori, A., and Thube, S. (2025). Power hungry: how AI will drive energy demand (No. 2025/081). International Monetary Fund. Available online at: <https://www.imf.org/-/media/Files/Publications/WP/2025/English/wpica2025081-print-pdf.ashx> (accessed May 10, 2025).
- Brundage, M., Avin, S., Wang, J., Belfield, H., and Anderljung, M. (2020). Toward trustworthy AI development: mechanisms for supporting verifiable claims. *arXiv preprint arXiv:2004.07213*. Available online at: <https://arxiv.org/pdf/2004.07213> (accessed June 2, 2025).
- Cai, W., Jiang, J., Wang, F., Tang, J., Kim, S., and Huang, J. (2025). A survey on mixture of experts in large language models (*arXiv:2407.06204v063*). doi: 10.1109/TKDE.2025.3554028
- Callahan, W. A. (2023). Chinese global orders: socialism, tradition, and nation in China–Russia relations. *Issues Stud.* 59:2340008. doi: 10.1142/S1013251123400088
- Cao, Z., Zhou, X., Hu, H., Wang, Z., and Wen, Y. (2022). Toward a systematic survey for carbon neutral data centers. *IEEE Commun. Surv. Tutor.* 24, 895–936. doi: 10.1109/COMST.2022.3161275
- Cartwright, O., Dunbar, H., and Radcliffe, T. (2024). Evaluating privacy compliance in commercial large language models-chatgpt, claude, and gemini. doi: 10.21203/rs.3.rs-4792047/v1
- Chen, C., and Liu, L. (2022). How effective is China's cryptocurrency trading ban? *Finan. Res. Lett.* 46:102429. doi: 10.1016/j.frl.2021.102429
- Chen, S. (2025). How much energy will AI really consume? *The good, the bad and the unknown*. *Nature* 639, 22–24. doi: 10.1038/d41586-025-00616-z
- Chintalapudi, R., Perugu, D., Andra, R., Gorrepati, N., Kumar, A., and Sakhamuri, S. B. (2025). Enhancing gender-neutral language with LoRA fine-tuning of DeepSeek-R1 for offline AI Applications. *Available at SSRN* 5225651. doi: 10.2139/ssrn.5225651
- Ciuriak, D., and Ptashkina, M. (2021). *Technology rents and the new Great Game. In The China-US Trade War and South Asian Economies* London: Routledge. (pp. 229–248). doi: 10.4324/9781003053613-18
- Crawford, K. (2021). Ethics at arm's length. *En Atlas of AI: Examining the human and environmental costs of artificial intelligence*. Goethe-Institut. Available online at: <https://www.goethe.de/prj/k40/en/eth/arm.html> (accessed June 4, 2025).
- Crowley, K. D., and Ahearne, J. F. (2002). Managing the Environmental Legacy of US Nuclear-Weapons Production: Although the waste from America's arms buildup will never be “cleaned up,” human and environmental risks can be reduced and managed. *Am. Sci.* 90, 514–523. doi: 10.1511/2002.39.514
- Cui, Y., Zhao, Q., Glauben, T., and Si, W. (2024). The impact of internet access on household dietary quality: evidence from rural China. *J. Integr. Agric.* 23, 374–383. doi: 10.1016/j.jia.2023.11.014
- Dai, D., Deng, C., Zhao, C., Xu, R. X., and Liang, W. (2024). DeepSeekmoe: towards ultimate expert specialization in mixture-of-experts language models. *arXiv preprint arXiv:2401.06066*. doi: 10.18653/v1/2024.acl-long.70

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Gen AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- de Carvalho Souza, M. E., and Weigang, L. (2025). Grok, Gemini, ChatGPT and DeepSeek: comparison and applications in conversational artificial intelligence. *Inteligencia Artificial* 2.
- DeepSeek-V3 Team (2025). DeepSeek-V3: Scaling to 236B and Beyond. *arXiv preprint arXiv:2412.19437v2 [cs.CL]* 18 Feb 2025.
- Ding, Z., Wang, J., Song, Y., Zheng, X., He, G., Chen, X., et al. (2025). Tracking the carbon footprint of global generative artificial intelligence. *The Innovation*, 100866. Available online at: [https://www.cell.com/the-innovation/fulltext/S2666-6758\(25\)00069-4](https://www.cell.com/the-innovation/fulltext/S2666-6758(25)00069-4) (accessed May 10, 2025).
- DR, A., and IS, S. (2025). Advancements in AI-Powered NLP Models: a Critical Analysis of Manus AI, Gemini, Grok AI, DeepSeek, and ChatGPT. Gemini, Grok AI, DeepSeek, and ChatGPT (March 19, 2025). doi: 10.2139/ssrn.518131
- Fan, Z., Yan, Z., and Wen, S. (2023). Deep learning and artificial intelligence in sustainability: a review of SDGs, renewable energy, and environmental health. *Sustainability* 15:13493. doi: 10.3390/su151813493
- Farronato, C. (2025). Data as the New Oil: Parallels, Challenges, and Regulatory Implications. NBER Chapters. Available online at: <https://www.nber.org/system/files/chapters/c15121/revisions/c15121.rev0.pdf> (accessed June 6, 2025).
- Feakin, T. (2025). DeepSeek's disruption: Geopolitics and the battle for AI supremacy. RUSI Commentary (Royal United Services Institute). Available online at: <https://www.rusi.org/explore-our-research/publications/commentary/deepseeks-disruption-geopolitics-and-battle-ai-supremacy> (accessed May 10, 2025).
- George, A. S. (2025). AI supremacy at the price of privacy: examining the tech giants' race for data dominance. *Partn. Univers. Int. Res. J.* 3, 26–43. doi: 10.5281/zenodo.14909763
- Ghafari, S., Safari, L., and Afsharchi, M. (2024). BBPE-AE: A Byte Pair Encoding-Based Auto Encoder for Password Guessing. doi: 10.20944/preprints202409.0834.v1
- Gillett, E. (2023). Atlas of AI: Examining the human and environmental costs of artificial intelligence. *Our Voices. Robert F. Kennedy Human Rights*. Available online at: <https://rfkhumanrights.org/our-voices/atlas-of-ai-examining-the-human-and-environmental-costs-of-artificial-intelligence/>
- Gloeckle, F., Idrissi, B. Y., Rozière, B., Lopez-Paz, D., and Synnaeve, G. (2024). Better & faster large language models via multi-token prediction. *arXiv preprint arXiv:2404.19737*.
- Gorilla, C., and Tuttle, R. (2025). A Feature-Level Approach to Mitigating Bias and Censorship in DeepSeek-R1. Available online at: <https://hal.science/hal-04992348v1> (accessed 2 April 29, 2025).
- Gould, P., Song, G., and Zhu, T. (2024). Environmental and Economic Impact of I/O Device Obsolescence. Retrieved from: *arXiv preprint arXiv:2412.20655*. (accessed June 9, 2025).
- Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., et al. (2025). DeepSeek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hallaq, B., Somer, T., Osula, A. M., Ngo, K., and Mitchener-Nissen, T. (2017). Artificial intelligence within the military domain and cyber warfare. In *Eur. Conf. Inf. Warf. Secur. ECCWS* (pp. 153–157). Available online at: <http://wrap.warwick.ac.uk/94297> (accessed May 10, 2025).
- Haney, B. S. (2020). Applied artificial intelligence in modern warfare and national security policy. *Hastings Sci. Tech. LJ*, 11:61. Available online at: https://repository.uchastings.edu/hastings_science_technology_law_journal/vol11/iss1/5 (accessed May 10, 2025).
- Hosseini, M., Gao, P., and Vivas-Valencia, C. (2025). A social-environmental impact perspective of generative artificial intelligence. *Environ. Sci. Ecotech.* 23:100520. doi: 10.1016/j.ese.2024.100520
- Hummel, P., Braun, M., Tretter, M., and Dabrock, P. (2021). Data sovereignty: a review. *Big Data Society* 8:2053951720982012. doi: 10.1177/2053951720982012
- Hussein, H., Gordon, M., Hodgkinson, C., Foreman, R., and Wagad, S. (2025). ChatGPT's impact across sectors: a systematic review of key themes and challenges. *Big Data Cogn. Comput.* 9:56. doi: 10.3390/bdcc9030056
- Jegham, N., Abdelatti, M., Elmoubarki, L., and Hendawi, A. (2025). How Hungry is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference. *arXiv preprint arXiv:2505.09598*. Available online at: <https://arxiv.org/pdf/2505.09598> (accessed June 6, 2025).
- Joublin, F., Ceravola, A., Deigmoeller, J., Gienger, M., Franzius, M., and Eggert, J. (2023). A glimpse into ChatGPT capabilities and its impact for AI research. *arXiv [Preprint]*. *arXiv:2305.06087*.
- Kassianik, P., and Karbasi, A. (2025). Evaluación del riesgo de seguridad en DeepSeek y otros modelos de razonamiento de frontera. Available online at: <https://blogs.cisco.com/security/>
- Kennedy, A. B. (2013). China's search for renewable energy: pragmatic technological nationalism. *Asian Survey*, 53(5), 909–930. Available online at: https://www.researchgate.net/profile/Andrew-Kennedy-15/publication/272543490_China%27s_Search_for_Renewable_Energy/links/573ea33008aea45ee842efcd/Chinas-Search-for-Renewable-Energy.pdf (accessed June 4, 2025).
- Knight, W. (2025). How DeepSeek ripped up the AI playbook—and why everyone's going to follow it. MIT Technology Review. Retrieved from: <https://www.technologyreview.com/2025/01/31/1110740/how-DeepSeek-ripped-up-the-ai-playbook-and-why-everyones-going-to-follow-it/> (accessed June 2, 2025).
- Krause, D. (2025). DeepSeek and FinTech: the democratization of AI and its global implications. Available at SSRN 5116322. doi: 10.2139/ssrn.5116322
- Kumar, A., and Thussu, D. (2023). Media, digital sovereignty and geopolitics: the case of the TikTok ban in India. *Media Cult. Soc.* 45, 1583–1599. doi: 10.1177/01634437231174351
- Laakkonen, S., Pál, V., and Tucker, R. (2016). The cold war and environmental history: complementary fields. *Cold War Hist.* 16, 377–394. doi: 10.1080/14682745.2016.1248544
- Li, G., Sun, Z., Wang, Q., Wang, S., Huang, K., and Zhu, Z. (2023). China's green data center development: policies and carbon reduction technology path. *Environ. Res.* 231:116248. doi: 10.1016/j.envres.2023.116248
- Li, K., and Christophe, B. (2024). Oscillating between the techniques of discipline and self: how Chinese policy papers on the digitalization of education subjectivize educators and the educated. *Learn. Media Technol.* 1–13. doi: 10.1080/17439884.2024.2306552
- Li, X., Chan, K. C., and Ma, H. (2020). Communist party direct control and corporate investment efficiency: evidence from China. *Asia-Pacific J. Accoun. Econ.* 27, 195–217. doi: 10.1080/16081625.2018.1470541
- Lian, S., Zhao, K., Lei, X., Wang, N., Long, Z., Yang, P., et al. (2025). Quantifying the capability boundary of DeepSeek models: an application-driven performance analysis. *arXiv preprint arXiv:2502.11164*.
- Liu, A., Feng, B., Xue, B., Wang, B., Wu, B., Lu, C., et al. (2024). DeepSeek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Lopes dos Santos, K. (2021). The recycling of e-waste in the industrialised global south: the case of Sao Paulo Macrometropolis. *Int. J. Urban Sustain. Dev.* 13, 56–69. doi: 10.1080/19463138.2020.1790373
- Luccioni, S., Jernite, Y., and Strubell, E. (2024). Power hungry processing: Watts driving the cost of ai deployment? In *Proceedings of the 2024 ACM conference on fairness, accountability, and transparency* (pp. 85–99). doi: 10.1145/3630106.3658542
- McGee, R. W. (2025). Leveraging DeepSeek: an AI-powered exploration of traditional Chinese medicine (Tai Chi and Qigong) for medical research. *Am. J. Biomed. Sci. Res.* 25, 645–654. doi: 10.34297/AJBSR.2025.25.003362
- Mol, A. P., Spaargaren, G., and Sonnenfeld, D. A. (2014). Ecological modernisation theory: Where do we stand. *Ökologische Modernisierung. Zur Geschichte und Gegenwart eines Konzepts in Umweltpolitik und Sozialwissenschaften*, 35–66. Available online at: https://www.researchgate.net/profile/Arthur-Mol/publication/40798926_Ecological_modernisation_Three_decades_of_policy_practice_and_theoretical_reflection/links/56794f1f08ae6041cb49f40f/Ecological-modernisation-Three-decades-of-policy-practice-and-theoretical-reflection.pdf (accessed June 5, 2025).
- Molina, C. E. M. (2021). El Centenario del Partido Comunista de China a la luz del Marxismo Leninismo. *Revista Política Int.* 3, 13–21.
- Morandín-Ahuérma, F. (2023). United States, China, and Russia: national proposals for an ethics of AI in the New Cold War. In *Principios Normativos para una Ética de la Inteligencia Artificial* (pp. 162–185). Available online at: https://www.researchgate.net/publication/374586769_United_States_China_and_Russia_National_Proposals_for_an_Ethics_of_AI_in_the_New_Cold_War (accessed May 10, 2025).
- Moravec, V., Gavurova, B., and Kovac, V. (2025). Environmental footprint of GenAI—Changing technological future or planet climate? *J. Innov. Know.* 10:100691. doi: 10.1016/j.jik.2025.100691
- Mytton, D. (2021). Data centre water consumption. *NPJ Clean Water* 4:11. doi: 10.1038/s41545-021-00101-w
- Naghiyev, K. (2024). ChatGPT From a Data Protection Perspective. *Baku St. UL Rev.*, 10, 1. Available online at: <https://heionline.org/HOL/LandingPage?handle=hein.journals/bakustulr10&div=7&id=&page=> (accessed April 29, 2025).
- Ngoy, P., Dar, F., Liyanage, M., Yin, Z., Norbistrath, U., Zuniga, A., et al. (2025). Supporting Sustainable Computing by Repurposing E-waste Smartphones as Tiny Data Centres. *IEEE Pervasive Computing*. doi: 10.1109/MPRV.2025.3541558
- O'Donnell, J. (2025). DeepSeek might not be such good news for energy after all. MIT Technology Review. Available online at: <https://www.technologyreview.com/2025/01/31/1110776/deepseek-might-not-be-such-good-news-for-energy-after-all/> (accessed June 2, 2025).
- Okaiyeto, S. A., Bai, J., Wang, J., Mujumdar, A. S., and Xiao, H. (2025). Success of DeepSeek and potential benefits of free access to AI for global-scale use. *Int. J. Agric. Biol. Eng.* 18, 304–306. doi: 10.25165/f.jiabe.20251801.9733
- OpenAI. (2024). How to delete and archive chats in ChatGPT. OpenAI Help Center. Available online at: <https://help.openai.com/en/articles/8809935-how-to-delete-and-archive-chats-in-chatgpt>. Last access: 09/06/2025.
- OpenAI. (2025). Política de privacidad de OpenAI. Available online at: <https://openai.com/es-ES/policies/row-privacy-policy/> (accessed June 9, 2025).

- Palmer, D. A., and Winiger, F. (2019). Neo-socialist governmentality: managing freedom in the People's Republic of China. *Econ. Soc.* 48, 554–578. doi: 10.1080/03085147.2019.1672424
- Pavliv, V., Akbari, N., and Wagner, I. (2024). AI-powered smart toys: interactive friends or surveillance devices? In *Proceedings of the 14th International Conference on the Internet of Things* (pp. 172–175). doi: 10.1145/3703790.3703841
- Petrella, S., Miller, C., and Cooper, B. (2021). Russia's artificial intelligence strategy: the role of state-owned firms. *Orbis* 65, 75–100. doi: 10.1016/j.orbis.2020.11.004
- Pham, D., Sheffey, J., Pham, C. M., and Houmansadr, A. (2024). ProxyGPT: enabling anonymous queries in AI Chatbots with (Un) trustworthy browser proxies. *arXiv preprint arXiv:2407.08792*.
- Pieke, F. N. (2025). The communist party of China's new central department of social work: neo-socialist governance or Bolshevized party discipline? *J. Contemp. China* 1–13. doi: 10.1080/10670564.2025.2490744
- Poo, M. M. (Ed.). (2025). Reflections on DeepSeek's breakthrough. *Nat. Sci. Rev.* 12:nwaf044. doi: 10.1093/nsr/nwaf044
- Prieto, A., Prieto, B., Escobar, J. J., and Lampert, T. (2025). Evolution of computing energy efficiency: Koomey's law revisited. *Clust. Comput.* 28:42. doi: 10.1007/s10586-024-04767-y
- Raj, D., Keren, G., Jia, J., Mahadeokar, J., and Kalinli, O. (2025). Faster speech-llama inference with multi-token prediction. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1–5). IEEE. doi: 10.1109/ICASSP49660.2025.10890328
- Raza, M., Sakila, K. S., Sreekala, K., and Mohamad, A. (2024). Carbon footprint reduction in cloud computing: Best practices and emerging trends. *Int. J. Cloud Comput. Data. Manag.* 5, 25–33. doi: 10.33545/27075907.2024.v5.i1a.58
- Rivero Silva, S. (2022). Joint Ventures of foreign investment in dictatorial contexts: the case of the hotel industry in Cuba. *Atlantic Review of Economics (ARoEc)*, 5, 1–16. Available online at: <https://hdl.handle.net/10419/282289> (accessed May 10, 2025).
- Roberts, H., Cows, J., Morley, J., Taddeo, M., Wang, V., and Floridi, L. (2021). The Chinese approach to artificial intelligence: an analysis of policy, ethics, and regulation (pp. 47–79). Springer International Publishing. doi: 10.1007/978-3-030-81907-1_5
- Rooney, K. (2025). OpenAI tops 400 million users despite DeepSeek's emergence. *CNBC*. Available online at: <https://www.cnn.com/2025/02/20/openai-tops-400-million-users-despite-deepseek-emergence.html> (accessed April 29, 2025).
- Roskladka, N., Bressanelli, G., Sacconi, N., and Miragliotta, G. (2025). Repairable electronic products for the circular economy: A review of design for repair features, practices and measures to contrast obsolescence. *Discov. Sustain.* 6:66. doi: 10.1007/s43621-024-00753-x
- Salah, M., Abdelfattah, F., Alhalbusi, H., and Al Mukhaini, M. (2024). Me and My AI Bot: Exploring the "Alholic" Phenomenon and University Students' Dependency on Generative AI Chatbots—Is This the New Academic Addiction? doi: 10.12103/rs.3.rs-3508563/v2
- Sapkota, R., Raza, S., and Karkee, M. (2025). Comprehensive analysis of transparency and accessibility of chatgpt, DeepSeek, and other sota large language models. *arXiv preprint arXiv:2502.18505*. Available online at: <https://arxiv.org/html/2502.18505v1> (accessed May 10, 2025).
- Sebastian, G. (2023). Privacy and data protection in ChatGPT and other AI chatbots: strategies for securing user information. *International Journal of Security and Privacy in Pervasive Computing* 15. doi: 10.4018/IJSPPC.325475
- Segal, A., and Kang, D. (2006). *The Siren Song of Techno-nationalism*. Far Eastern Economic Review.
- Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., et al. (2024). DeepSeekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Stacciarini, J. H. S., and Gonçalves, R. J. D. A. F. (2025). Data Centers, Critical Minerals, Energy, and Geopolitics: The Foundations of Artificial Intelligence (No. 2zvkt_v1). *Center for Open Science*. Available online at: https://www.researchgate.net/profile/Joao-Stacciarini-2/publication/389376214_Data_Centers_Critical_Minerals_Energy_and_Geopolitics_The_Foundations_of_Artificial_Intelligence/links/67c06fcc207c0c20fa9a48/Data-Centers-Critical-Minerals-Energy-and-Geopolitics-The-Foundations-of-Artificial-Intelligence.pdf (accessed May 10, 2025).
- Travis, D. (2016). Desk research: The what, why and how. Userfocus. Available online at: <https://www.userfocus.co.uk/articles/desk-research-the-what-why-and-how.html> (accessed May 10, 2025).
- Truby, J., Brown, R. D., Dahdal, A. M., and Ibrahim, I. A. (2025). Ai Diplomacy in the Age of Stargate & DeepSeek: Legal and Strategic International Approaches to Techno-Nationalism, Regulatory Soft Power and the Ai Chips Race. *Regulatory Soft Power and the Ai Chips Race*. doi: 10.2139/ssrn.5187702
- U. S. Congress (2024). S.3732 – Artificial Intelligence Environmental Impacts Act of 2024. Congress.gov. Available online at: <https://www.congress.gov/bills/118th/congress/committee-bills/3732/text> (accessed June 5, 2025).
- U. S. Government Accountability Office (2025). Artificial intelligence: Generative AI's environmental and human effects (GAO-25-107172). Available online at: <https://www.gao.gov/assets/gao-25-107172.pdf> (accessed June 5, 2025).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention is all you need. *Advances in neural information processing systems*, 30. Available online at: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fb053c1c4a845aa-Abstract.html> (accessed May 10, 2025).
- Wang, F., Zhang, Z., Zhang, X., Wu, Z., and col. (2024). A Comprehensive Survey of Small Language Models in the Era of Large Language Models (*arXiv:2411.03350*).
- Wang, J., Berger, D. S., Kazhamiaka, F., Irvine, C., and Sriraman, A. (2024). Designing cloud servers for lower carbon. In *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)* (pp. 452–470). IEEE. Available online at: https://www.pdl.cmu.edu/PDL-FTP/CloudComputing/Wang_ISCA24.pdf (accessed June 5, 2025).
- Wang, P., Zhang, L. Y., Tzchor, A., and Chen, W. Q. (2024). E-waste challenges of generative artificial intelligence. *Nature Computational Science*, 1–6. Available online at: <https://www.nature.com/articles/s43588-024-00712-6> (accessed May 10, 2025).
- Wu, X., Duan, R., and Ni, J. (2024). Unveiling security, privacy, and ethical concerns of ChatGPT. *J. Inf. Intell.* 2, 102–115. doi: 10.1016/j.jiixd.2023.10.007
- Xu, Z., Wang, J., Xu, X., Yu, P., Huang, T., and Yi, J. (2025). A Survey of Reinforcement Learning-Driven Knowledge Distillation: Techniques, Challenges, and Applications. Available online at: <https://www.preprints.org/manuscript/202503.0903/v2> (accessed May 10, 2025).
- Yasir Fahim, J. (2024). Mastering the art of AI Language: an in-depth exploration of prompting techniques and their influence on model performance. Kenyon University. Available online at: https://digital.kenyon.edu/dh_iphs_ss/35 (accessed May 10, 2025).
- Yuan, Y., Shi, J., Zhang, Z., Chen, K., Zhang, J., Stoico, V., et al. (2024). The Impact of Knowledge Distillation on the Energy Consumption and Runtime Efficiency of NLP Models. In *Proceedings of the IEEE/ACM 3rd International Conference on AI Engineering-Software Engineering for AI* (pp. 129–133). doi: 10.1145/3644815.3644966
- Zeng, J. (2020). Artificial intelligence and China's authoritarian governance. *Int. Aff.* 96, 1441–1459. doi: 10.1093/ia/iaa172
- Zeng, J. (2021). Securitization of artificial intelligence in China. *Chin. J. Int. Polit.* 14, 417–445. doi: 10.1093/cjip/poab005
- Zhang, A. H. (2025). The promise and perils of China's regulation of artificial intelligence. *Colum. J. Transnat'l L.* 63:1. Available online at: https://www.jtl.columbia.edu/s/01_CTL_63_1_Zhang-1.pdf (accessed June 6, 2025).
- Zhang, W., Lei, X., Liu, Z., Wang, N., Long, Z., Yang, P., et al. (2025). Safety evaluation of DeepSeek models in Chinese Contexts. *arXiv preprint arXiv:2502.11137*.
- Zhang, Y., and Liu, J. (2022). Prediction of overall energy consumption of data centers in different locations. *Sensors* 22:3704. doi: 10.3390/s22103704
- Zhang, Y., Wang, Y., and Wang, X. (2011). Greenware: Greening cloud-scale data centers to maximize the use of renewable energy. In *Middleware 2011: ACM/IFIP/USENIX 12th International Middleware Conference, Lisbon, Portugal, December 12–16, 2011. Proceedings 12* (pp. 143–164). Springer Berlin Heidelberg. Available online at: https://link.springer.com/chapter/10.1007/978-3-642-25821-3_8 (accessed May 10, 2025).
- Zhu, Q., Guo, D., Shao, Z., and Yang, D., Liang, W. (2024). DeepSeek-coder-v2: Breaking the barrier of closed-source models in code intelligence. *arXiv preprint arXiv:2406.11931*.