AI is discursive, AI is material: DeepSeek and the remaking of the US–China "AI arms race"

Shuxian Liu | 刘 姝贤
School of Information, The University of Texas at Austin

David Gray Widder
School of Information, The University of Texas at Austin

Abstract

This study examines the discourses of DeepSeek R1 through a multi-level comparative approach across two countries and two levels: official narratives and public discussions in the US and China (US congressional hearings and documents, Xinhua coverage, Zhihu posts, and Reddit discussions). We analyze how the key infrastructural components foregrounded by the R1 release—chips, compute, capital investment, export controls, and open-source ecosystems—are articulated across distinct yet entangled discursive fields, and how these articulations reflect broader struggles over sovereignty, scale, and openness in contemporary AI development and geopolitical competition. Drawing on Laclau and Mouffe's (1985) concept of articulation and Carpentier's (2017) concept of the discursive-material knot, we develop the framework of *discursive-material articulation* to examine the recursive mutual shaping of the discursive and the material in AI infrastructures. We argue that DeepSeek functions as a nodal point (Laclau, 1990) around which three interconnected circuits: the sovereignty circuit, the scale circuit, and the openness circuit, organize how AI infrastructure is produced, contested, and governed. Our analysis shows that the same material objects acquire strikingly different identities and meanings depending on the articulatory context, that the uneven political economy of the AI stack shapes the specific mode of articulation available to each formation, and that publics' technically grounded scepticism constitutes a distinct form of political agency operating within but not beneath hegemonic frames. The study offers theoretical insight into the entanglement of discourse and materiality in AI infrastructure, with implications for how AI is studied amid intensifying geopolitical competition.

Keywords: AI infrastructure, discursive-material articulation, DeepSeek, critical discourse analysis, open-source AI, US, China

Introduction

In January 2025, the Chinese AI startup DeepSeek, founded in July 2023 by Liang Wenfeng, launched its R1 model and chatbot. The launch sent shockwaves through the global technology industry. The model's performance rivalled leading Western systems such as OpenAI's GPT-4 at a fraction of the training cost; its weights were released openly on Hugging Face; and it originated in a country subject to strict U.S. export controls on advanced semiconductors—all of which quickly made DeepSeek a focal point for global debates about the future of AI (Allen, 2025).

These discussions are deeply embedded in the prevalent rhetoric of an "AI arms race" between the US and China, which has intensified over the past five years, often imbued with what scholars describe as "yellow techno-peril" (McInerney, 2024) or "techno-orientalism" (Von Blomberg and Liu, 2025). Days after the launch, a US Congressional speech framed it as a "*Sputnik moment*" and "*wake-up call*", raising the alarm that "*China is catching up with the United States on AI…If America falls behind China on AI, we will fall behind everywhere—economically, militarily, scientifically, educationally, everywhere.*" The Sputnik invocation reactivates a specific Cold War discursive formation through which the US fused computational systems, national security, and an imaginary of global containment into a mutually reinforcing political-technical project (Edwards, 1996). In contrast, Xinhua, the state's official news agency in China, invoked "横空出世" (*héng kōng chū shì*, "sudden appearance from nowhere") to mark a rupture that exposed prior American dominance: "*DeepSeek launched a 'surprise attack' on Silicon Valley, overturning the 'stereotype' that in AI's frontier domain, 'America leads, China follows*'." These contrasting framings demonstrate that AI models are not encountered as neutral technical entities but through discursive articulations that confer political meaning and sociotechnical force—articulations whose material consequences can intensify the race for technological superiority and catalyze a military AI arms race with significant ramifications for security and geopolitics (Cave and ÓhÉigeartaigh, 2018).

What the DeepSeek case brings into precise focus is the entanglements of the discursive and material of AI infrastructure–specifically, how discourses foreground, or obscure, the material components that underpin the AI race (e.g., chips, compute, capital, data centers, energy grids, and control regimes), and how such articulations are produced, stabilized, and contested across national contexts and communities. These questions remain underexamined in existing scholarship. DeepSeek's infrastructure is not merely a technical backdrop but a site of dynamic discursive construction and intense geopolitical struggle, exemplified by its direct challenge to assumptions about the effectiveness of US chip export controls and the necessity of massive capital investment in AI infrastructure. The US embargo on Nvidia's high-end chips forced DeepSeek to work with reduced-spec hardware, optimize software to compensate, and engage open-source ecosystems differently from Western competitors (Yuan, 2025). Nvidia's share-price turbulence following the launch (Perrigo and Pillay, 2025) registered in financial

markets a shock that political and institutional arenas were simultaneously absorbing. Debates over DeepSeek's open-source release and distillation practices further complicate the relationship between scale, efficiency, and architectural innovation in AI development (Beer and Jacobsen, 2026). Unpacking how these infrastructural realities are narrated—as Chinese ingenuity, a failure of US policy, a vindication of open-source collaboration, or a threat to be contained—and how these narratives produce material ramifications, is central to understanding power and politics in today's global AI landscape.

We develop the framework of *discursive-material articulation* to examine the recursive mutual shaping of the discursive and the material in AI infrastructures, drawing on Laclau and Mouffe's (1985) concept of articulation and Carpentier's (2017) concept of the *discursive-material knot*. We define discursive-material articulation as "the knotted linkage and recursive interactions of the discursive and the material, unevenly situated and enacted within contingent ideological, political, and economic structures." We elaborate the framework through a multi-level comparative analysis of official narratives and public discussions on DeepSeek in both the US and China, using critical discourse analysis (CDA; Fairclough, 2023). By tracing how official and public actors imagine, legitimize, contest, or (de)politicize AI infrastructure in relation to the same event, we show how AI infrastructure is articulated across distinct yet entangled discursive fields, reflecting broader struggles over sovereignty, scale, and openness. The guiding research questions are:

1. How are the key infrastructural elements of DeepSeek framed across official narratives and public discussions within the US and China?
2. How does the framework of *discursive-material articulation* illuminate the ways discourses articulate, engage with, or obscure the infrastructures of AI?

Our findings demonstrate that DeepSeek functions as a nodal point (Laclau, 1990) around which three interconnected circuits—sovereignty, scale, and openness—organize how AI infrastructure is produced, contested, and governed. Across the four corpora, the same material phenomena (the Nvidia chip, the compute threshold DeepSeek cleared, the R1 weights, distillation, etc.) acquire strikingly different identities and meanings depending on the articulatory context: national security asset or constraint-as-creative-spur, scale imaginary to be defended or strategic miscalculation to be exposed, ungovernable threat or unstoppable commons. Official state discourses on both sides converge on strategic framings aligned with their structural positions in the AI stack, while online public discussions display a more fractured, technically grounded engagement misaligned with either official framing. These findings contribute to emerging scholarship on the situated geopolitics of AI discourse and call for further research into the nuanced dynamics of discursive-material articulation and their ramifications.

Theoretical framework: discursive-material articulation of AI infrastructure

Infrastructure as discursive-material formation

Ever since infrastructure became an object of social scientific study, its discursive dimensions have been seen to have material consequences. Star and Ruhleder (1994) argue that infrastructure is sociotechnical: it consists not only of interdependent technical systems but also of the norms, routines, and forms of knowledge that make those systems function in practice. Larkin (2013) likewise insists that infrastructures must be understood as "semiotic objects": what counts as infrastructure is never self-evident but depends on epistemological and political commitments that determine what is rendered consequential and what is left out. This resonates with growing scholarly attention to the conjunction of discourse and the making of politics and technology (Mager and Katzenbach, 2021) in the field of Science and Technology Studies (STS). Scholars study "expectations and stories about the future" (van Lente, 2016), the role of technological innovations and visionary rhetoric in enterprises (Beckert, 2016), and the discursive struggles around "contested futures" (Brown et al., 2016). Such expectations are both constitutive and performative (Borup et al., 2006). They contribute to the formation of the "social imaginary" of infrastructure—a set of coherent visions by which a group of people collectively imagine their social existence (Taylor, 2004)—and assemble interest groups, shape agendas, and mobilize collective action (Konrad et al., 2016). This dynamic is especially salient for large-scale infrastructures and emerging technologies. Roads and railways, for example, retain an "enchantment" through promises of speed, integration, and economic connection, even when particular projects fail to deliver on those promises (Harvey and Knox, 2012). Regarding AI, in a cross-national study on the narratives and imaginaries of national AI strategies in China, the US, France, and Germany, Bareis and Katzenbach (2022) found similar portrayals of AI as inevitable and massively disrupting technological development by building on rhetorical devices such as a grand legacy and international competition.

Recent scholarship has developed a more nuanced account of how infrastructures and their discourses co-produce one another. Enright (2022) elaborated on the dialectics between *imaginaries of infrastructure*—the contingent representational forms that bring infrastructure into being—and *infrastructural imaginings*, through which senses and meanings are themselves materialized. Rao (2023) theorizes *discursive infrastructure*, which refers not only to the mediated material channels through which knowledge about infrastructure travels, but also to the ways such channels reshape ideological perceptions of what infrastructure is and does. These processes are profoundly affective, constituting subjects through the mobilization of desire, pride, and frustration (Larkin, 2018). Such affective investments have a long historical lineage: Nye's (1994) account of the *technological sublime* shows how large-scale infrastructures, from railroads to electrification to the atomic bomb, have repeatedly been received as objects of awe

and national self-recognition, generating political legitimacy that exceeds their instrumental function. Infrastructure thus does not merely support action or reflect ideology; it actively participates in the articulation of political relations and offers a way of tracing the co-emergence of political and material histories (Collier, 2011; Mukerji, 2009).

Discursive-material articulation of AI infrastructure

Understanding infrastructure as simultaneously material and semiotic, technical and political, sets the conceptual foundation for this study's central goal: how AI's physical components are discursively produced and contested, and the ramifications of this process. As Van Der Vlist et al. (2024) highlight, AI should be understood as an existing technology stack and (emerging) ecosystem *and* a discursive phenomenon—or, as Suchman (2023) puts it, "as a sign invested with social, political and economic capital and with performative effects that serve the interests of those with stakes in the field." Drawing on these arguments and on Laclau and Mouffe's (1985) concept of articulation and Carpentier's (2017) concept of discursive-material knot, we propose the framework of *discursive-material articulation*, which we argue explains current dynamics of AI infrastructure development in their material-political-economic dimensions.

Laclau and Mouffe (1985) theorize articulation as the relational practice through which elements are joined such that their identities are partially produced in the process, from which they insist that discourse is itself material (pp. 107-112). Objects are always within discursive articulations that confer meaning and social force (DeLuca, 1999; Laclau, 1990). Stuart Hall (1986) further emphasizes that articulations are historically contingent and always liable to rearticulation under changed political conditions. Foucault's (1980) concept of *regimes of truth* further specifies the procedures that determine which discourses function as true, who is recognized as capable of producing truth, and how truth and power reinforce one another. Thus, designating certain objects as "AI infrastructure" is not neutral but a discursive operation making some materialities politically legible while rendering others marginal.

The concept of discursive-material knot developed by Carpentier (2017), defined as "non-hierarchical ontology that theorizes the knotted interactions of the discursive and the material as restless and contingent, sometimes incessantly changing shapes and sometimes deeply sedimented" (p.4), further addresses the interconnection between the discursive and the material and their translation in social practice. This means that discursive claims and material referents are mutually constitutive: discourse confers meaning on material objects, while those objects ground the claims that make them politically legible. We retain "articulation" rather than "knot" to foreground the recursive discursive-material formation of infrastructure: discursive functions such as legitimation and contestation produce material implications—policy, investment, deployment—that sediment into new material referents, which then become the ground for contested and reformulated claims.

To name the analytical unit through which these recursions cohere, we adopt the term *circuit*, building on du Gay et al.'s (1997) *circuit of culture*, in which representation, identity, production, consumption, and regulation are theorized as mutually constitutive moments in the life of a technological artefact rather than discrete stages in a linear process. The circuit metaphor allows us to capture how meaning is produced through recursive interconnection rather than unidirectional causation, and how the discursive-material formation of AI infrastructure operates as a *coupled system* whose stability and instability are emergent properties of how its elements articulate together (Hall, 1997). Figure 1 summarizes the framework.

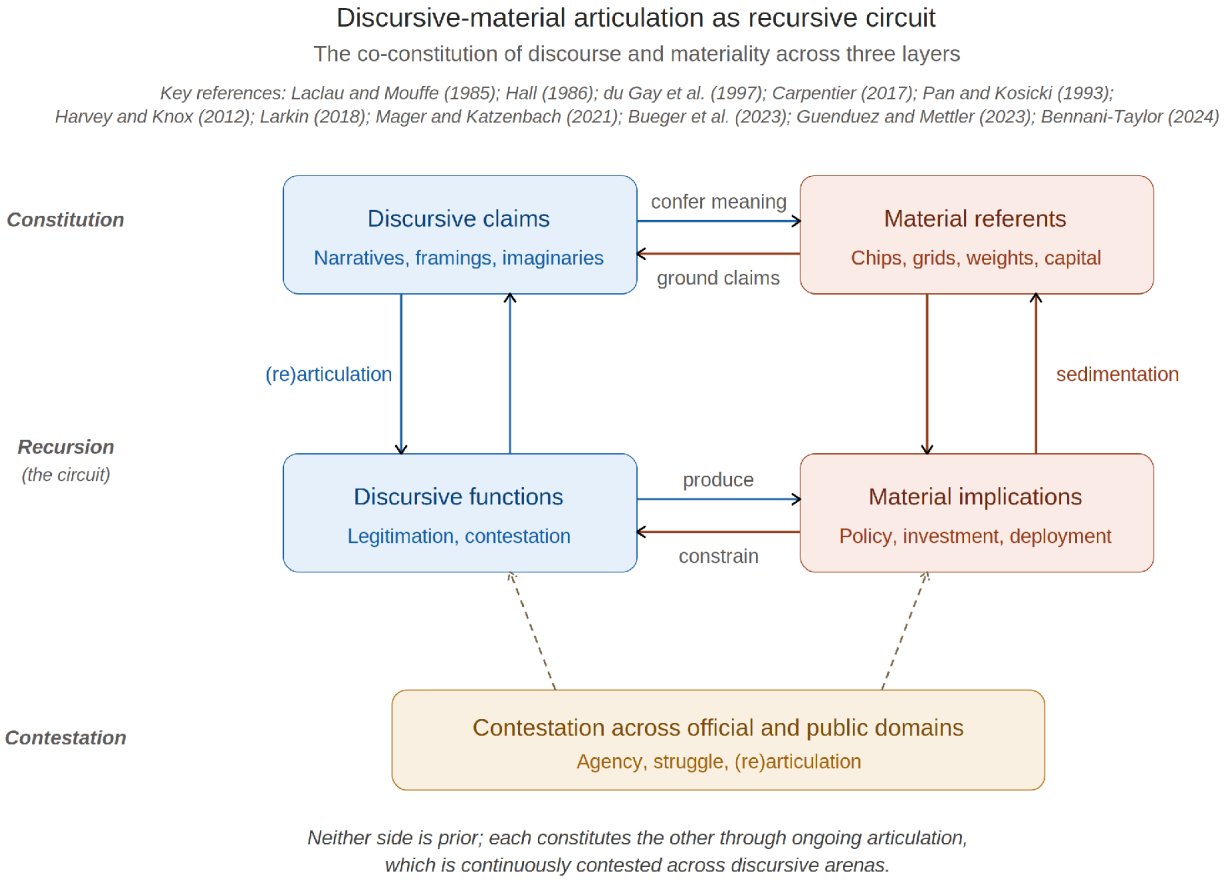


Figure 1. Discourse and materiality mutually constitute one another through recursive articulation and sedimentation, a process continuously contested across official and public domains.

We define discursive-material articulation as "the knotted linkage and recursive interactions of the discursive and the material, unevenly situated and enacted within contingent ideological, political, and economic structures." The framework comprises three interlocking dimensions:

At the level of constitution, discursive claims confer meaning on material referents, grounding those claims in narratives, framings, and infrastructural imaginaries. A data center, a semiconductor, or an open-source model is not a fixed object that discourse merely describes; its identity, significance, and social force are constituted through articulation with other elements—technical benchmarks, national security narratives, policy priorities, investment decisions, institutional arrangements, and material allocations (Guenduez and Mettler, 2023). As Bueger et al. (2023) argue, digital sovereignty discourse entangles with semiconductors as a future technology promising safety and innovation, even as the material conditions of its production contradict that sovereignty.

At the level of recursion, discursive claims legitimize, stabilize, or contest sociotechnical assemblages of AI development—actors, ideologies, capital flows, and power relations—while themselves materializing through funding structures, educational priorities, migration regimes, and other arrangements (Bennani-Taylor, 2024). Government resource allocation to imagined futures installs those futures materially, generating a performative lock-in (Mager and Katzenbach, 2021). Discourse thus does not merely reflect competition but actively shapes investment priorities, governance regimes, and public legitimacy, producing material implications that constrain future discursive possibilities, conditioning whether infrastructural promises sustain or collapse into dereliction (Bennett, 2001).

At the level of contestation, the recursive co-constitution of discourse and materiality is a site of conflicting discourses, enabling agency, tension, and struggle (Laclau and Mouffe, 1985). Harvey and Knox (2012) show that beyond top-down institutional narratives, infrastructural promises are reinvigorated through lived, material encounters of stasis, rupture, and blockage; public and bottom-up discursive formations, more heterogeneous and emotionally charged, can reproduce, appropriate, parody, or resist official narratives. Platforms and algorithmic architectures further structure these possibilities, amplifying certain voices, making certain issues visible, and obscuring other material realities (Pan and Kosicki, 1993).

AI discourses in the US–China “AI arms race”

The discursive-material articulation framework is particularly illuminating for the current geopolitics of AI, where infrastructures function simultaneously as a material field of competition and a symbolic mechanism linking technology, sovereignty, and collective identity under new historical and digital conditions. The US–China “AI arms race” is the contemporary formation in which these dynamics are most acute and where official and public discourses are being actively reorganized.

China’s New Infrastructure (*xīn jī jiàn*, 新基建) policies, for instance, materialize through mediated portrayals of “advanced” technologies that benefit everyday life and GDP growth (Rao, 2023), serving as ideological enablers of “infrastructural imaginings” of more productive futures while bearing the impress of socioeconomic realities and an enduring Western-centric agenda of

technological development (Lindtner, 2020). US media and political discourse, by contrast, cast China through techno-orientalist and civilizational-threat frames, turning AI advancement into a national-security imperative articulated through American exceptionalism (McInerney, 2024). Both formations operate within the broader Cold War discursive inheritance through which computational systems, national security, and global containment were fused into a mutually reinforcing political-technical project (Edwards, 1996).

Beyond official discourse, recent empirical work documents the heterogeneity and affective charge of public AI discussions. Zhang (2025) identifies cracks between compliance-focused institutional narratives and ordinary netizens' perceptions in China, as users mobilize satire, outrage, and anxiety to contest issues such as misinformation, plagiarism, and creative authorship—fractures that suggest affective resistance and embodied public reasoning can recalibrate policy attention even under algorithmic suppression. Comparative studies likewise indicate that public engagement with AI varies across national contexts: Liu et al. (2024) find shared optimism among Chinese and US users about ChatGPT's performance and potential, cutting across national discursive formations. Discursive-material articulation thus operates across official and public domains, enacted, contested, and rearticulated in everyday practices filtered through platform affordances (Rao, 2023).

This burgeoning scholarship, however, lacks systematic comparisons of how official and public domains frame the same AI artefact across rival nations, leaving internal discursive dynamics and cross-national relations underexplored. The "AI arms race" narrative has rarely been empirically scrutinized through discursive-material articulation, and the rhetorical mechanisms that reproduce, contest, or complicate it across domains remain unexamined. Further, specific infrastructural components—chips, compute, capital, export controls, open-source ecosystems—have not been theorized as explicit objects of discursive production and contestation across national contexts; their articulation, instrumentalization, or strategic obscuration remains largely untheorized. This study addresses these gaps through a comparative analysis of DeepSeek's framing across US and Chinese official and public discourse, providing the event-specific, cross-national account that AI geopolitics scholarship has yet to produce.

Method

Table 1. Four corpora.

CORPUS	SOURCE	N	DESCRIPTION	KEY GENRES / COMMUNITIES
US official **Federal government**	GPO, Congressional Record, congress.gov, committee websites	48 texts	Eight committee jurisdictions and both chambers of Congress; multiple genres enable analysis of recontextualization across institutional sites	Committee hearings (8 jurisdictions), floor speeches, CRS reports, legislative bills, House Select Committee investigative report, White House press briefings
US public **Reddit**	Reddit (18 subreddits)	36 threads	Range of interpretive communities across English-language public discourse on AI, technology, and geopolitics	Technical practitioners (r/LocalLLaMA, r/LLMDevs); financial investors (r/NVDA_Stock, r/NvidiaStock); futures/acceleration (r/singularity, r/accelerate); China-focused (r/China, r/China_irl); AI product (r/DeepSeek, r/ChatGPT); political (r/asianamerican, r/WayOfTheBern)
Chinese official **Xinhua News Agency**	Xinhua wire service	58 texts	Xinhua has long functioned as a key channel through which party-state narratives are circulated (Repnikova, 2017).	State news agency dispatches; official news reports
Chinese public **Zhihu** 知乎	Zhihu (15 question threads)	134 posts	Question-and-answer format supports extended analytical discussion with a subculture that values technical knowledge (Tan and Liu, 2026)	Capability evaluation; geopolitical significance; organizational critique; infrastructure and cost analysis; open-source debates

Research design

We employ a comparative qualitative design using Critical Discourse Analysis (CDA) (Fairclough, 2023). For data sources, we constructed four corpora spanning two national contexts (the US and China) and two institutional registers (official and public). The multi-sited design is theoretically warranted by the shared material infrastructure that both national AI programs respond to—making their discourses structurally conditioned—and by the official/public distinction, which captures how institutional constraints and public affordances differently shape discursive constructions. The primary analytical window is January 20 to April 20, 2025, the three months immediately following DeepSeek R1's release, when dominant framings and nodal points were being established across the two national formations. However, several US official texts from May 2025 are included, notably: three US congressional hearings: "Winning the AI Race," "Protecting Our Edge," and "DeepSeek: A Deep Dive". All texts collected and analyzed are publicly available documents (Franzke et al., 2020). Table 1 shows the four corpora.

The four corpora differ in language, institutional status, and moderation regime. We treat this heterogeneity as a methodological asset rather than a liability: the aim is not to compare equivalent texts but to examine how the same event was constructed differently across institutional positions, national contexts, and community norms.

Each text was subjected to three-level CDA following Fairclough's model (2023). We first examined vocabulary choices in their original language, attending to politically specific formulations that carry ideological content. Secondly, we analyzed how claims travel across genres, acquiring or losing institutional force in the process. Finally, we interpreted how texts constitute AI infrastructure as a political stake and enable specific governance interventions.

CDA requires explicit reflexivity, as researchers are not neutral observers but part of the discursive field under analysis. The first author grew up in China, is a native Chinese speaker fluent in English, and has been academically trained in China, the UK, and the US, providing familiarity with the affect and cultural metaphors that translation can flatten while requiring reflexive attention to the blind spots that proximity to one of the discursive fields entails. The last author is a US citizen who grew up in Asia and is an early-career professor at a public university in a US state whose leaders increasingly cast Chinese AI technology, including DeepSeek, as suspect. We discussed these institutional pressures throughout the research process.

Findings: the discursive-material articulation of DeepSeek

The following sections analyze how the discursive-material articulation of DeepSeek's infrastructure operates across our findings through three interconnected circuits: the sovereignty circuit, the scale circuit, and the openness circuit.

The sovereignty circuit: chip as chokepoint, constraint, creative spur, and affective infrastructure

Produced in a highly concentrated and relatively inelastic global supply chain, semiconductor chips are essential for shaping the speed, cost, and feasibility of AI model training and inference (Khan and Mann 2020). They are also sociomaterial assemblages loaded with values inscribed through the language of digital sovereignty that legitimizes trade, economic, and research policies (Monsees, 2025).

Semiconductor chips have become strategic assets within a broader turn toward economic nationalism and US–China technological rivalry. US export controls were designed to limit China's access to advanced compute to address national security concerns (Hauge et al., 2025). Washington subsequently requested that Taiwan Semiconductor Manufacturing Company (TSMC) withhold advanced chips from Huawei, whose ties to the Chinese government were framed as a security concern (Chiang, 2023). China, in turn, has pursued domestic substitution: amassing chipmaking inputs and expanding domestic manufacturing and AI capabilities. Debates over whether US chip controls can be effectively implemented or whether they may accelerate Chinese efforts toward supply-chain substitution have intensified (Sastry et al., 2024). These dynamics demonstrate how semiconductors are becoming part of a constellation of digital sovereignty power, translated into various political and economic projects to establish autonomous digital infrastructures in a hyperconnected world (Musiani, 2022). *Sovereignty circuit* thus describes how semiconductors are linked to national security, chokepoint power, domestic substitution, and strategic autonomy through both discourse and material practice.

Competing constitutions of the same chip

The Nvidia H100, the core infrastructure layer enabling the training, scaling, and deployment of large models, is constituted as a national asset, dependency marker, engineering constraint, and investor instrument across the four corpora. In the US official narrative, the H100 chip is framed as a national security asset whose movement must be controlled at the border. The H800—a reduced-spec chip designed for the Chinese market after earlier US controls—appears as evidence that export control is vulnerable to technical workarounds. This produces an illegitimate-derivative narrative in which DeepSeek is constructed not as autonomous Chinese innovation, but as a derivative achievement allegedly enabled by two forms of unauthorized access to US model capabilities through distillation and to US hardware through restricted Nvidia chips. In this framing, DeepSeek's success is politically troubling because it appears to convert American technical assets into a Chinese geopolitical advantage.

What makes the US official chip discourse analytically distinctive is that the security argument extends both outward to foreign adversaries and inward to domestic industry. The *DeepSeek Unmasked* report's Nvidia Singapore revenue chart functions as a material inscription that makes this double exposure visible—a commercial financial document recontextualized as

evidence that Nvidia itself was "*deliberately designing chips to skirt U.S. export controls*" and "*undermining the intent of U.S. policy*." Singapore's existence as a transshipment hub extends the geography of enforcement outward, while H800 extends enforcement inward to the company's own engineering decisions. The threat is constituted as operating through globally distributed material networks that neither the Chinese border nor the US Entity List (which aims to control chip export) can adequately contain. This requires a governance response that reaches corporate design choices as well as foreign acquisition channels.

Xinhua offers an almost inverted account where the same chip restriction is framed as a constraint turned into creative pressure. One Xinhua account states DeepSeek's achievement

> lies in the US ban on exporting high-end computing cards [ie, H100] to China, which forced companies like DeepSeek to deeply mine the computational potential of "reduced-spec" cards like the H800. These are challenges that Western enterprises like OpenAI have never encountered.

In this narrative, export control becomes the condition of innovation. Xinhua further frames US controls as short-sighted and hegemonic, and the restriction is transformed into proof of resilience. DeepSeek also becomes a moral-technological achievement: evidence not only of engineering capability but also of Chinese start-ups' capacity to contribute to the global AI commons despite hegemonic constraint.

The poetics of the chip: political aesthetics and affective infrastructure

The poetic rhetoric of the chip in Xinhua's narratives warrants particular attention. The phrase "中国芯" (*Zhōngguó xīn*) exemplifies the principle: "芯" (*xīn*, "chip/core") is homophonic with "心" (*xīn*, "heart/mind/civilizational soul"), fusing the material object and the affective-national subject at the level of the signifier itself. Adjacent formulations include "卡脖子" (*kǎ bózi*, "strangling at the chokepoint") which constitutes the chip restriction as physical violence against the national body; "国之重器/大国重器" (*guó zhī zhòng qì*/*dà guó zhòng qì*, "important implements of the nation/great power") draws on a longer political vocabulary in which "重器" (*zhòng qì*, "important vessel/implement") referred to ritually and politically charged objects, most notably bronze ritual vessels, bound up with authority, hierarchy, and dynastic memory. In this sense, the semiconductor is invested with desire, ambition, restoration, and civilizational significance.

In contrast, the US chip discourse is instrumental and more managerial (asset, chokepoint, vulnerability, restriction, Entity List). The vocabulary is bureaucratic-strategic, designed to manage and control. The chip is a tool, a lever, a risk. Two formations are not merely proposing different policies toward the same material object but constituting different material-affective objects through incompatible political aesthetics (Larkin, 2013). Consequently, the domestication of the technical stack and self-determination operate not merely as a policy objective or an engineering roadmap but as an affective mandate, part of a broader project of

digital sovereignty oriented toward protecting economic prosperity, security, and culture (Musiani, 2022).

Skepticism, financial logic, and the limits of official discourses

Public discourses exhibit a more fractured and technically grounded relationship to the semiconductor substrate, operating in productive tension with, rather than fully aligning to their respective official formations, and from different subject positions and with different political stakes.

Many Zhihu posts defend DeepSeek against the US illegitimate-derivative narrative while simultaneously acknowledging dependency on Nvidia's CUDA (Compute Unified Device Architecture). Celebration of DeepSeek's success despite export controls is partial, qualified, and coexists with technical analysis. The chip is articulated as an engineering specification, with varying degrees of skepticism; posts identify FP8, PTX-level optimization, and CUDA binding with a precision that both official corpora elide. One post states it plainly:

> DeepSeek-V3, from training through to inference deployment optimization, is completely built on NVIDIA CUDA… DeepSeek-R1/V3, in many respects, is deeply bound to NVIDIA GPUs.

The post specifies why: Huawei's Ascend 910B "*doesn't have FP8, and inter-chip efficiency is even lower than the H800*," which is a reference to the mixed-precision training technique central to DeepSeek's efficiency gains.

Infrastructure capacity calculation constitutes the most visible technical counter to the optimism of Xinhua's narratives:

> DeepSeek is likely to be trained and inferred on old A800s and H800s; 80% of the software is still built on Nvidia's CUDA ecosystem; completely switching the overall architecture to the Huawei system is not realistic in the short term.

Another post reframes Nvidia-specificity as a practical problem to navigate:

> Nvidia chips are restricted for sale to us; the OpenAI route is blocked—DeepSeek's appearance means we can deploy on AMD, Huawei, Moore Threads chips, even on phones. This is a path to survival.

Chip dependency is constituted here as a navigation problem: the question is which alternative deployment paths open up, not which sovereignty claims become possible. Contrasting with Xinhua's promise of "*100% domestic supply chain throughout the technical pathway*," most Zhihu discussions are sympathetic to the wider achievement while technically pragmatic and calculatedly skeptical of such claims, showing critical engagement from within the techno-nationalist framework rather than direct opposition. Chip dependency is treated as an engineering problem with a timeline, not as a refutation of Chinese AI capability. Xinhua's re-articulation of constraints as the necessary process of developing domestic infrastructural

independence is itself a response to the kind of technical specificity that would complicate the sovereignty performance.

On Reddit, the semiconductor is constituted primarily through financial logic and as infrastructure that can be circumvented through local deployment:

> China pushed DeepSeek to use Huawei chips for political reasons, it failed, and they had to go back to Nvidia. That's the most coherent summary of China's chip strategy I've seen.

Another added:

> Losing the China market and letting Huawei take all their market share there is an existential threat to Nvidia. Eventually, perhaps 3-5 years down the road, Huawei will flood the market with cheaper AI chips that work good enough and they will have the software ecosystem solved.

The chip here is neither a national asset nor a dependency vulnerability but a market position to be defended against an approaching competitive timeline that Chinese alternatives are progressively closing. Some Reddit practitioners further demonstrate that the Nvidia dependency can be partially severed at the user level. Running DeepSeek on consumer-grade gaming GPUs (a GeForce RTX 3090 or 4090) lets individual users deploy the model locally rather than through cloud services, addressing both official concerns simultaneously: the US worry about data flowing to Chinese servers dissolves because no connection to Chinese infrastructure exists, while the Chinese concern about Nvidia dependency compromising domestication is rendered politically inert at the level of individual inference, where no sovereignty performance is being staged. Local deployment thereby constitutes a chip articulation that acknowledges Nvidia dependency as materially real while dissolving its political stakes by relocating the site of use from the national-strategic register to the domestic-personal register.

In the sovereignty circuit, the chip is enrolled in competing material programs: US official discourse deploys a logistics-of-exclusion grammar in which export controls produce reduced-spec chips that produce evidence of leakage, authorizing further controls. Xinhua instead deploys an aesthetics-of-restoration grammar, investing semiconductor domestication with affective mandate, desire, and civilizational weight, while public discourse renders these articulations more technically explicit and politically unsettled.

The scale circuit: power, compute investment, and disruptions of scale imaginary

The development of frontier AI systems has become increasingly synonymous with large compute budgets (Besiroglu et al., 2024). Investment decisions, hardware procurement, policy frameworks, and geopolitical positions were organized around it (Bennani-Taylor, 2024). By 2024, this logic had materialized on an extraordinary scale: Goldman Sachs estimated that

generative AI was driving roughly US$1 trillion in prospective capital expenditure (Goldman Sachs, 2024).

We refer to this formation as the *scale circuit*: a mode of discursive-material articulation that links claims about AI capability to compute budgets, energy systems, investment narratives, and geopolitical positioning. Against this backdrop, DeepSeek's efficiency is this study's most consequential material-discursive event. The V3 and R1 training runs required a combined 2.94 million H800 GPU-hours, approximately 2,500–3,000 MWh, roughly 5% of GPT-4's estimated energy consumption and 20 times less, using 2,048 chips compared to GPT-4's 25,000 (DeepSeek-AI et al., 2025; Guo et al., 2025). Frontier performance was demonstrably achievable at a fraction of the compute previously assumed necessary, and the challenge this posed was potentially delegitimizing, technically, financially, and geopolitically. What is analytically significant is not the technical efficiency per se but how each discourse community managed the threat to an imaginary in which it held fundamentally different material stakes. Before January 20, 2025, the assumption that frontier AI required massive compute was institutionally settled; after that, it became a political stake requiring active defense or active refutation. The same material fact was articulated into radically different discursive-material chains depending on each formation's structural position and investment in the imaginary under challenge.

Absorption versus vindication: two official strategies

The US official corpus and the Chinese official corpus both engage the efficiency challenge directly but perform opposite discursive operations on it.

In US official discourse, the challenge is absorbed rather than allowed to unsettle the underlying development model. The House Energy and Commerce Committee hearing title, *Converting Energy into Intelligence: The Future of AI Technology, Human Discovery, and American Global Competitiveness* (April 9, 2025), already frames the problem as one of sustaining scale, not questioning it. Rep. Fedorchak's floor speech (February 13, 2025) makes this logic explicit:

> On January 20, China unveiled DeepSeek-R1… This proves that China is rapidly closing the gap, and we can't afford to fall behind.

The acknowledgement is immediately converted into a warrant for further energy infrastructure:

> Energy is now the limiting factor for building the data centers that AI uses, which is why, to win the race for AI, we need to unleash American energy.

In this formulation, efficiency does not undermine the scale model but intensifies the need to scale faster. The sovereignty argument extends further down the material stack of the supply chain. Rep. Cline's linkage of AI competitiveness to rare earth minerals binds electricity, geology, compute, and national power together:

> We are dependent on China for rare earth minerals that are so important to power our grid, powering the technologies that support AI… talking to countries like Greenland and Ukraine.

Scale becomes a national policy horizon that absorbs its own material preconditions, resonating with broader work showing that AI infrastructures are increasingly articulated through energy systems, logistics, and extractive supply chains rather than through software alone (Muldoon et al., 2026).

Xinhua performs the inverse operation. Efficiency is not a challenge to manage but the vindication of a constrained-hardware development path that export controls had inadvertently forced. Xinhua explicitly names and discredits the cultural logic of scale:

> Rather than relying on major technological breakthroughs, more of a "brute force scaling" strategy is used — building larger models, using larger datasets, deploying larger computing power. DeepSeek-R1's disruptiveness lies in questioning this approach, demonstrating that the best generative AI models can be achieved with less computing power and less financial investment.

The US path is thereby constituted as strategically mistaken: practitioners who did not face hardware constraints did not develop the algorithmic innovations that such constraints can produce. The restriction intended to limit Chinese AI capability is reframed, in this articulation, as having produced a more disciplined approach to AI development.

Yet Xinhua does not abandon compute investment but rearticulates it:

> The algorithmic efficiency improvement brought by DeepSeek has not suppressed computing power demand; on the contrary, more users and scenarios have joined, driving large model popularization and driving data center, edge, and endpoint computing construction.

The material consequence of the efficiency claim—less energy per model—is rearticulated as the material precondition of a new scaling claim: more users require more total infrastructure. This is the discursive operation of *Jevons's paradox* (Jevons, 1865), in which efficiency gains in a resource generate rather than suppress aggregate demand for it. Applied to compute, the rearticulation resolves the tension between celebrating efficiency and justifying ongoing compute investment by constituting them as a virtuous circuit rather than a contradiction. Projects such as the Zhengzhou Compute Towers (Chen and Zhang, 2026) materialize this logic primarily as regional growth and industrial upgrading rather than as security infrastructure, anchoring compute investment in developmental competition and ecosystem building.

Public discourse and the capital critique

The public discourse produces a third position: the scale imaginary appears not as a technical necessity, but as a political-economic project tied to speculative valuation, infrastructural rents, and the concentration of AI power (Widder and Hicks, 2024).

Zhihu posts repeatedly frame DeepSeek's significance less as a breakthrough in raw capability than as an event that punctures the financial narrative surrounding US firms. One post argues

that DeepSeek's real achievement is to "burst the compute bubble" that had inflated US tech stocks for years:

> The new generation of US tech people reinforced a stereotyped narrative formula: advanced AI = high compute + high investment + elite talent; only America has the strength… But DeepSeek's appearance proved this formula does not hold… After training, it can run independently without relying on cloud models, which turned many investors' dreams of wealth into foam.

Others describe the prevailing US model as a "compute capital game," in which sovereign AI ambitions in other countries are forced into dependence on US firms' infrastructure, APIs, or chip ecosystems. In this reading, the scale imaginary serves not simply technical optimization, but revenue capture and geopolitical leverage. This critique is also turned inward toward China. Posts noting that DeepSeek emerged from a *"side-project company"* rather than other Chinese firms *"that received tech subsidies and policy support, occupying traditional internet advantages"* suggest that the problem criticized here is not uniquely American but a similar assumption that compute concentration and institutional backing would naturally produce frontier results. DeepSeek's emergence becomes evidence that both US hyperscaler narratives and Chinese state-backed incumbency can misrecognize where innovation actually occurs, with implications for how resources are allocated under the scale imaginary in both contexts.

On Reddit, the energy sovereignty argument that US official discourse deploys as a warrant for domestic buildout is inverted:

> China kept advancing its electrical grid and its capabilities while the USA rested on its laurels, and now China is eclipsing us in all elements of energy except perhaps extraction**.** China has destroyed the USA in power generation.

While this does not amount to a coherent counter-imaginary, it does reveal the instability of the dominant one: the scale imaginary deployed by Fedorchak as a warrant for US energy policy is turned against that discourse by the recognition that China has been building at scale while the US debated its regulatory framework.

DeepSeek's demonstration of frontier performance at drastically reduced cost supplied material evidence contradicting the dominant scale imaginary. The US absorbed the shock by intensifying scale; China inverted its hierarchy while preserving aggregate demand through rearticulation; public communities named the circuit as a circuit rather than inhabiting it. What the case thus shows is that a compute-scale imaginary long treated as neutral technical consensus is, under pressure, revealed as a deeply contested political stake, organizing capital, state support, and infrastructure investment in ways that exceed technical necessity.

The openness circuit: open source as contested articulation site

When DeepSeek released its R1 model weights on Hugging Face, a single downloadable artefact carried markedly different political meanings across the four corpora: a national security threat,

a contribution to the global AI commons, a geopolitical instrument, a movement victory, and a contested legal object—differences stemming from how the weights were situated within distinct discursive and infrastructural arrangements. The model weight thus functions as a boundary object (Star and Griesemer, 1989): stable enough to remain recognizable across sites, yet flexible enough to be taken up within different institutional and political logics. Model weights are enacted, not merely interpreted, through infrastructures, deployment practices, legal classifications, and governance interventions; their political significance lies not in representation alone but in how they circulate, where they can run, what access they enable, and the control they appear to require. Openness in AI is not binary but layered across weights, training data, code, documentation, licenses, and use restrictions (Widder et al., 2024). We refer to this formation as the *openness circuit*: a recursive discursive-material articulation linking the release of model weights to competing narratives of commons, threat, market disruption, and legitimacy, while reorganizing the material conditions of circulation and use.

A central site of contestation within this circuit is *distillation*: the family of practices through which capabilities are transferred from one model to another, including synthetic-instruction generation, rationale-augmented supervision, and the use of stronger proprietary models' outputs to train weaker or open ones (Wang et al., 2022). The method is technically routine and increasingly central to the cost structure of frontier AI development, yet its governance is deeply unsettled. A growing legal literature argues that attempts to prohibit downstream reuse through terms of use rest on uncertain foundations, leaving the boundary between legitimate distillation, competitive imitation, and unlawful appropriation politically as well as legally contested (Henderson and Lemley, 2024). Others characterize distillation as a form of attack that renders AI "less safe, less original, and less open" and propose technical countermeasures to disincentivize large-scale distillation from frontier proprietary systems (Trockman and Savani, 2026). The release of DeepSeek R1 has complicated the concept of distillation, significantly influencing the entire imaginary of AI, with material consequences for how the relationship between scale, efficiency, and architectural innovation is understood (Beer and Jacobsen, 2026). Distillation, as our analysis shows, operates as both a material practice and a discursive nodal point around which competing claims about innovation, legitimacy, and theft are organized.

Openness as strategy: competing geopolitical readings

Across all four corpora, every formation, including those that evaluate the open-source release positively, constitutes it as a strategic act rather than an altruistic contribution to shared knowledge. What differs is not whether strategy is present but how it is evaluated and what geopolitical or economic work it is seen to perform.

In US official discourse, "*the PRC's strategy to capture AI market share with low-cost, open-source models*" is targeted as the primary threat requiring federal procurement prohibition. Open source is the weapon; low cost is the delivery mechanism; market capture is the objective. Xinhua names the same strategic reading from the opposite side:

> DeepSeek's choice of open source is disadvantageous to maintaining its own technical advantage, but it is undoubtedly beneficial to boosting enthusiasm for open source, building the open-source community, and prospering the AI ecosystem… guiding more open, transparent, and shared AI open-source development will become our country's key move to break through 'small yard, high fence' (xiǎo yuàn gāo qiáng, 小院高墙) and 'decoupling and chain-severing'.

The release is acknowledged as technically costly to DeepSeek while politically valuable to China, showing the strategic calculation the official commons narrative simultaneously enacts. Xinhua extends the logic to governance: "*China should actively lead or participate in setting international open-source AI model technical standards… jointly build a global AI governance alliance*." Open source is thereby articulated as infrastructure for an alternative global AI order in which China participates as standard-setter rather than standard-receiver.

These official narratives reveal an important asymmetry. The US, holding key infrastructural chokepoints, reads open distribution as a threat to the exclusivity that supports strategic advantage. China, positioned more ambivalently within that hierarchy, reads the same distribution as loosening the effects of exclusion. The same release appears threatening or emancipatory, not because the technical object changes but because the formations encountering it occupy different positions in relation to compute, standards, and institutional power in the AI stack (Widder et al., 2024).

The Zhihu corpus introduces a third strategic reading that views the open-source release as a denial-of-revenue weapon, strategically valuable precisely because it is ungovernable. One post articulates:

> China doesn't need to make money; what matters is that America doesn't make money. DS is calling on the people of the world to resist American technological hegemony together… Chinese leadership would be laughing in their dreams.

Another post identifies the ungovernable character of the instrument as its political substance:

> If Trump wants to block OpenAI, he could do it very easily. Now that DeepSeek is open source, go ahead and try to block it… this equals unstoppable AI dumping. A transparent strategy: there's absolutely nothing you can do.

The Zhihu strategic reading simultaneously validates the US security chain's concern that Chinese open-source *is* a strategic instrument against US interests and exposes its operational futility, since the instrument works precisely because it cannot be stopped.

Another Zhihu post identifies the dissolution of the investment narrative sustaining the scale imaginary:

> Fortunately, OpenAI went closed-source… After DeepSeek went open-source, with less than one-tenth of OpenAI's cost, achieving comparable or even better results—how is the OpenAI-Nvidia story supposed to be told now?

Because R1's weights are openly downloadable, users and organizations can deploy the model on their own hardware (via inference engines like vLLM and llama.cpp—software that loads the weights and executes the model on local servers or even consumer GPUs) rather than paying for access through OpenAI's cloud API. This means the "OpenAI-Nvidia story"—an investment thesis built on AI's structural dependence on centralized, proprietary, cloud-delivered cloud infrastructure—is disrupted not through argument but through the material practice of every user who chooses local deployment over API access. The open-source release is politically consequential not only for what it symbolizes but for how it alters the conditions of technical circulation (Kapoor et al., 2024).

Some Reddit users make this material character explicit:

> What they're actually doing is better than dropping the full stack… They're working on getting the optimizations ported to popular open source inference engines. This means we're getting DS optimizations in vLLM, and likely Llama.cpp, kobold, etc., as well.

When the model weight enters shared infrastructure through local deployment, it is constituted as a technical artefact whose Chinese origin is retained in historical attribution but recedes in current use. This is what the "unstoppable AI dumping" that Zhihu posts identify as a transparent strategy, whose power lies in its unblockability, and what US official discourse frames as a governance emergency. Both formations recognize the same material fact (open model weights cannot be effectively prohibited once released) yet reach opposite assessments of what follows. In practitioner communities, use becomes a mode of articulation in its own right: local deployment, optimization, and integration into open-source stacks do not simply respond to the released object but generate new deployments, technical dependencies, and governance dilemmas.

Distillation and the asymmetric governance of common practice

If the openness debate exposes how the same release can be articulated as commons or threat depending on structural position, the distillation controversy reveals how the same practice can be articulated as ordinary technique, or as theft depending on who performs it, or as evidence of double standards. The disagreement is not about whether the practice exists but about how it is classified and governed.

In the US official discourse, distillation is first described neutrally in the *Protecting Our Edge* hearing:

> This is a very common technique that a lot of US AI developers will use as well. Essentially, it's using a more advanced AI model to get some answers and to use those as part of your training data.

However, it is reconstituted in the investigative report as a criminal act:

> DeepSeek personnel infiltrated US AI models and fraudulently evaded protective measures under aliases… using a sophisticated network of international banking channels.

The escalation from API usage (a technical act) to infiltration (a security violation) to international fraud (a criminal act) is achieved through the accumulation of articulatory steps rather than through any change in the technical facts.

The Zhihu community's response does not deny the distillation but recontextualizes it through comparative moral analysis:

> Standing at the moral high ground of the open-source spirit: OpenAI privately scanned data from the internet, took it out and sold it for $200 a month, while DeepSeek also scanned it but completely open-sourced it for free, only 'allegedly' using some of OpenAI's outputs. Which of these two is more controversial, which is on the moral high ground — anyone with eyes can see.

This is not a defense of distillation as such but moral inversion: if OpenAI's monetization of freely contributed internet data is the prior violation, DeepSeek's replication of OpenAI's outputs is a secondary act whose ethical status cannot be adjudicated without the prior violation being named. The framing articulates distillation into an equivalence chain between OpenAI's training data acquisition and DeepSeek's distillation, making the asymmetric criminalization visible as an articulatory choice rather than a principled legal distinction.

Reddit's practitioner communities make similar points:

> 'Illegal distillation'… Everyone's doing it; there are one million datasets on HuggingFace, most generated with GPT. There are papers on it; it's common practice. OpenAI can at most block their accounts for TOS violations. Microsoft does it too, for the Phi series of models.

This is the discursive-material move that exposes the selectivity of the security logic of the US official narrative: the same practice by a US company is a standard development technique; by a Chinese company, it is theft. Model-development practices are not governed uniformly; instead, they are filtered through geopolitical assumptions about who counts as the legitimate innovator and who appears as the suspect imitator. Our findings extend Beer and Jacobsen's (2026) argument about the structuring potential of distillation by showing how such developments re-enter later rounds of political and institutional argument as evidence—the recursive dynamic the framework of discursive-material articulation aims to capture.

Discussion: From the DeepSeek moment to a research agenda for AI (geo)politics

The DeepSeek case shows the discursive-material articulation of AI infrastructure through three circuits—sovereignty, scale, and openness—each linking specific material referents to distinct discursive logics. Thus, articulation operates semiotically and infrastructurally simultaneously. AI infrastructure is not a neutral technical substrate prescribed or later interpreted through

discourse but is constituted through recursive relations between material arrangements and meaning-making practices that our framework delineates.

The discursive-material articulation is also shaped by a political economy specific to each circuit and is not equally available to all actors. What can be said persuasively depends partly on a formation's position in the AI stack and on the scale of its sunk material investment in a particular infrastructural trajectory. Infrastructure thus does not simply materialize imaginaries; it can also encumber them, creating rigidities that shape what actors can plausibly say about the systems they seek to resist or have helped build. The material dimension is, correspondingly, encountered unevenly. Communities positioned closer to the technical substrate often can articulate more granular accounts of friction, dependency, and constraint, while communities at a greater distance are more likely to reproduce higher-level official frames. This clarifies a form of political agency that is not fully counter-hegemonic yet still consequential: technically-grounded skepticism, which does not reject dominant frames outright while limiting their closure by preserving material specificity.

These findings open at least three directions for future research. First, future work should explore the situated geopolitics of AI discourse. The DeepSeek case shows that the "same" infrastructural object carries strikingly different political identities in different discursive arenas, inviting a more thorough comparative geopolitics of AI that resists US-centrism and attends to how computational autonomy is articulated within specific national and regional configurations, rather than read off a single dominant frame (cf. Bareis and Katzenbach, 2022; Rao, 2023). Second, the discursive-material entanglement is at the heart of critical AI studies: as recent critical AI scholarship rightfully turns toward materiality after years dominated by discursive and ethical critique, our discursive-material framework cautions against treating materiality as a new analytic privilege and instead emphasizes their mutual articulation. Discursive contestation, such as the "stochastic parrots" argument (Bender et al., 2021) and the destabilization of AI's claimed thinghood (Suchman, 2023), and material contestation, such as data center protests, AI worker organizing, infrastructural sabotage, are not parallel but recursively coupled: discursive contestation creates the interpretive conditions in which material contestation becomes legible as political resistance, while material contestation generates the empirical material that discursive critique can mobilize. Finally, whereas we examine two communities at one moment, the conditions under which technically-grounded skepticism and other partial dissent move from localized frictions to durable infrastructural politics over time merit further examination.

Deepseek R1 is now a year and a half old—a fossil in the discursive field of AI. But the spectre of Chinese models roars. Usage of Chinese open-weight models—MiMo by Xiaomi, Hy by Tencent, and R1's descendants by DeepSeek—continues to dwarf American counterparts (OpenRouter, 2026), as US business press posits "the end of America's dominance in the AI race" (Knutson, 2026). Discourse that Chinese companies circumvent "loopholes" in chip export bans via remote

access, reap government subsidies (Knutson, 2026), or distill and "reverse engineer all the secret ingredients" of American models echoes a long history of US industry calling Chinese goods "cheap knockoffs." Yet, China is increasingly regarded as a technological trailblazer, challenging U.S. dominance over markets and imaginaries alike (Pollio, 2024). In all of this are competing discursive claims about material infrastructure as it does or should exist; and more widely, about an AI future continuously reimagined amid geopolitical anxiety. DeepSeek, as both discourse and thing, has only just cracked this open.

Acknowledgements

We are grateful to Inha Cha, Dawn Nafus, and Lucy Suchman and for their helpful feedback on earlier versions of this draft.

References

Allen, G. C., Benson, E., & Reinsch, W. A. (2022, November). *Improved Export Controls Enforcement Technology Needed for U.S. National Security*. https://www.csis.org/analysis/improved-export-controls-enforcementtechnology-needed-us-national-security

Bareis, J., & Katzenbach, C. (2022). Talking AI into Being: The Narratives and Imaginaries of National AI Strategies and Their Performative Politics. *Science, Technology, & Human Values*, *47*(5), 855–881. https://doi.org/10.1177/01622439211030007

Beckert, J. (2016). *Imagined Futures*. Cambridge, MA: Harvard University Press.

Beer, D., & Jacobsen, B. N. (2026). Artificial intelligence, conceptions of distillation and the reframing of reasoning. *Big Data & Society*, *13*(2), 20539517261442383. https://doi.org/10.1177/20539517261442383

Bender, E. M., Gebru, T., McMillan-Major, A., et al. (2021, March). On the dangers of stochastic parrots: Can language models be too big?🦜. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* (pp. 610-623).

Bennani-Taylor, S. (2024). Infrastructuring AI: The stabilization of "artificial intelligence" in and beyond national AI strategies. *First Monday*. https://doi.org/10.5210/fm.v29i2.13568

Bennett, J. (2001). *The enchantment of modern life: Attachments, crossings, and ethics*. Princeton University Press.

Besiroglu, T., Bergerson, S. A., Michael, A., et al. (2024). *The Compute Divide in Machine Learning: A Threat to Academic Contribution and Scrutiny?* (arXiv:2401.02452). arXiv. https://doi.org/10.48550/arXiv.2401.02452

Borup, M., Brown, N., Konrad, K., et al. (2006). The sociology of expectations in science and technology. *Technology Analysis & Strategic Management*, *18*(3–4), 285–298. https://doi.org/10.1080/09537320600777002

Brown, N., Rappert, B., & Webster, A. (2016). Introducing Contested Futures: From Looking into the Future to Looking at the Future. In *Contested Futures: A Sociology of Prospective Techno-Science* (pp. 3–20). New York: Routledge.

Bueger, C., Liebetrau, T., & Stockbruegger, J. (2023). Theorizing Infrastructures in Global Politics. *International Studies Quarterly*, *67*(4), sqad101. https://doi.org/10.1093/isq/sqad101

Carpentier, N. (2017). *The Discursive-Material Knot*. Peter Lang US. https://doi.org/10.3726/978-1-4331-3754-9

Cave, S., & ÓhÉigeartaigh, S. S. (2018). An AI Race for Strategic Advantage: Rhetoric and Risks. *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 36–40. https://doi.org/10.1145/3278721.3278780

Chen, Q., & Zhang, Y. (2026, April 14). *China activates largest scientific AI computing cluster in Zhengzhou*. https://www.globaltimes.cn/page/202604/1358913.shtml

Chiang, M.-H. (2023). Taiwan Semiconductor Manufacturing Company: A Key Chip in the Global Political Economy. *East Asian Policy*, *15*(01), 36–46. https://doi.org/10.1142/S179393052300003X

Collier, S. J. (2011). *Post-Soviet social: Neoliberalism, social modernity, biopolitics*. Princeton University Press.

DeepSeek-AI, Liu, A., Feng, B., et al. (2025). *DeepSeek-V3 Technical Report* (arXiv:2412.19437). arXiv. https://doi.org/10.48550/arXiv.2412.19437

DeLuca, K. (1999). *Articulation Theory: A Discursive Grounding for Rhetorical Practice*. *32*(4).

du Gay, P., Hall, S., Janes, L., et al. (1997) *Doing Cultural Studies: The Story of the Sony Walkman*. Thousand Oaks, CA: Sage.

Edwards, P. N. (1996). *The closed world: Computers and the politics of discourse in Cold War America*. MIT Press.

Enright, T. (2022). The Infrastructural Imagination. *Journal of Urban Technology*, *29*(1), 101–107. https://doi.org/10.1080/10630732.2021.2001715

Fairclough, N. (2023). Critical discourse analysis. In *The Routledge handbook of discourse analysis*. Routledge.

Franzke, A.S., Bechmann, A., Zimmer, M., et al. (2020). *Internet research: Ethical guidelines 3.0*. https://aoir.org/reports/ethics3.pdf

Goldman Sachs. (2024, June 27). Gen AI: Too much spend, too little benefit? Top of Mind. https://www.goldmansachs.com/insights/top-of-mind/gen-ai-too-much-spend-too-little-benefit

Guenduez, A. A., & Mettler, T. (2023). Strategically constructed narratives on artificial intelligence: What stories are told in governmental artificial intelligence policies? *Government Information Quarterly*, *40*(1), 101719. https://doi.org/10.1016/j.giq.2022.101719

Guo, D., Yang, D., Zhang, H., et al. (2025). DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*, *645*(8081), 633–638. https://doi.org/10.1038/s41586-025-09422-z

Hall, S. (1986). Gramsci's Relevance for the Study of Race and Ethnicity. *Journal of Communication Inquiry*, *10*(2), 5–27. https://doi.org/10.1177/019685998601000202

Hall, Stuart. 1997. "The Work of Representation", in S. Hall (ed.) Representation: Cultural Representations and Signifying Practices, London: Sage and The Open University.

Harvey, P., & Knox, H. (2012). The Enchantments of Infrastructure. *Mobilities*, *7*(4), 521–536. https://doi.org/10.1080/17450101.2012.718935

Hauge, J., Houtzager, B., & Hörmann, A. J. (2025). The new economic nationalism: Industrial policy and national security in the United States, China, and the European Union. *Geoforum*, *166*, 104382. https://doi.org/10.1016/j.geoforum.2025.104382

Henderson, P., & Lemley, M. A. (2024). The Mirage of Artificial Intelligence Terms of Use Restrictions. *arXiv.Org*. Publicly Available Content Database; SciTech Premium Collection (3143056193). https://ezproxy.lib.utexas.edu/login?url=https://www.proquest.com/working-papers/mirage-artificial-intelligence-terms-use/docview/3143056193/se-2?accountid=7118

Jevons, W. S. 1865. *The Coal Question*. London: MacMillan.

Kapoor, S., Bommasani, R., Klyman, K., et al. (2024). *On the Societal Impact of Open Foundation Models* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2403.07918

Konrad, K., Van Lente, H., Groves, C., et al. (2016). Performing and governing the future in science and technology. In *The handbook of science and technology studies* (pp. 465–493).

Knutson, R. (Host). (2026, July 28). *Chinese AI is spooking Silicon Valley* [Audio podcast episode]. In *The Journal*. The Wall Street Journal. https://www.wsj.com/podcasts/the-journal/chinese-ai-is-spooking-silicon-valley/f8c91c99-bbdf-48bf-982a-65132fee8bf1

Laclau, E. (1990). *New Reflections on the Revolution of Our Time*. London: Verso.

Laclau, E., & Mouffe, C. (1985). *Hegemony and socialist strategy: Towards a radical democratic politics*.

Larkin, B. (2013). The Politics and Poetics of Infrastructure. *Annual Review of Anthropology*, *42*(1), 327–343. https://doi.org/10.1146/annurev-anthro-092412-155522

Larkin, B. (2018). Promising forms: The political aesthetics of infrastructure. In *The promise of infrastructure* (pp. 175–202). Duke University Press.

Lindtner, S. M. (n.d.). *Prototype Nation: China and the Contested Promise of Innovation*. Princeton University Press. (Original work published 2020)

Liu, Z., Li, H., Chen, A., et al. (2024). Understanding Public Perceptions of AI Conversational Agents: A Cross-Cultural Analysis. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–17. https://doi.org/10.1145/3613904.3642840

Mager, A., & Katzenbach, C. (2021). Future imaginaries in the making and governing of digital technology: Multiple, contested, commodified. *New Media & Society*, *23*(2), 223–236. https://doi.org/10.1177/1461444820929321

McInerney, K. (2024). Yellow Techno-Peril: The 'Clash of Civilizations' and anti-Chinese racial rhetoric in the US–China AI arms race. *Big Data & Society*, *11*(2), 20539517241227873. https://doi.org/10.1177/20539517241227873

Monsees, L. (2025). The paradox of semiconductors—EU governance between sovereignty and interdependence. *Cambridge Review of International Affairs*, *38*(1), 3–21. https://doi.org/10.1080/09557571.2024.2405915

Mosco, V. (2005). *The Digital Sublime (MIT Press): Myth, Power, and Cyberspace*. MIT Press.

Mukerji, C. (2009). *Impossible engineering: Technology and territoriality on the Canal du Midi*. Princeton University Press.

Muldoon, J., Valdivia, A., & Badger, A. (2026). The politics of artificial intelligence supply chains. *AI & SOCIETY*, *41*(2), 1175–1187. https://doi.org/10.1007/s00146-025-02625-y

Musiani, F. (2022). *Infrastructuring* digital sovereignty: A research agenda for an infrastructure-based sociology of digital self-determination practices. *Information, Communication & Society*, *25*(6), 785–800. https://doi.org/10.1080/1369118X.2022.2049850

OpenRouter. (2026.). *LLM rankings*. Retrieved August 6, 2026, from https://openrouter.ai/rankings

Pan, Z., & Kosicki, G. M. (1993). Framing analysis: An approach to news discourse. *Political Communication*, *10*(1), 55–75. https://doi.org/10.1080/10584609.1993.9962963

Pollio, A. (2024). *Silicon elsewhere: Nairobi, global China, and the promise of techno-capital*. (pp. 51–65). University of California Press.

Rao, Y. (2023). Discourse as infrastructure: How “New Infrastructure” policies re-infrastructure China. *Global Media and China*, *8*(3), 254–270. https://doi.org/10.1177/20594364231198605

Repnikova, M. (2017). *Media politics in China: Improvising power under authoritarianism*. Cambridge University Press.

Sastry, G., Heim, L., Belfield, H., et al. (2024). *Computing Power and the Governance of Artificial Intelligence* (arXiv:2402.08797). arXiv. https://doi.org/10.48550/arXiv.2402.08797

Star, S. L., & Griesemer, J. R. (1989). Institutional Ecology, `Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science*, *19*(3), 387–420. https://doi.org/10.1177/030631289019003001

Star, S. L., & Ruhleder, K. (1994). Steps towards an ecology of infrastructure: Complex problems in design and access for large-scale collaborative systems. *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work - CSCW '94*, 253–264. https://doi.org/10.1145/192844.193021

Suchman, L. (2023). The uncontroversial 'thingness' of AI. *Big Data & Society*, *10*(2), 20539517231206794.

Tan, J., & Liu, T. (2026). Working closely with data: Aspirations and constraints of Chinese data scientists. *Big Data & Society*, *13*(1), 20539517261424165. https://doi.org/10.1177/20539517261424165

Taylor, C. (2004). *Modern social imaginaries*. Duke University Press.

Trockman, A., & Savani, Y. (2026, February 19). *Unexpected externalities of distillation*. Antidistillation. https://antidistillation.com/blog/unexpected-externalities-of-distillation/

Van Der Vlist, F., Helmond, A., & Ferrari, F. (2024). Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. *Big Data & Society*, *11*(1), 20539517241232630. https://doi.org/10.1177/20539517241232630

van Lente, H. (2016). Forceful Futures: From Promise to Requirement. In *Contested Futures: A Sociology of Prospective Techno-science* (pp. 43–64). New York: Routledge.

Von Blomberg, M., & Liu, C. (2025). Modern social imaginaries. *Information, Communication & Society*, 1–17. https://doi.org/10.1080/1369118X.2025.2565335

Wang, Y., Kordi, Y., Mishra, S., et al. (2022). *Self-Instruct: Aligning Language Models with Self-Generated Instructions* (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2212.10560

Widder, D. G., & Hicks, M. (2024). *Watching the Generative AI Hype Bubble Deflate* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2408.08778

Widder, D. G., Whittaker, M., & West, S. M. (2024). Why 'open' AI systems are actually closed, and why this matters. *Nature*, *635*(8040), 827–833. https://doi.org/10.1038/s41586-024-08141-1

Yuan, J. E. (n.d.). *DeepSeek and the Shifting AI Landscape in China and the US*.

Zhang, B. (2025). Affective counterpublics under constraint: Emotion, platform governance, and AI ethics discourse on Chinese social media. *AI & SOCIETY*. https://doi.org/10.1007/s00146-025-02437-0