

# The Stack Capture Race

Flyxion

January 31, 2026

## Abstract

Contemporary discourse surrounding artificial general intelligence is dominated by speculative narratives that frame technological development as a race toward autonomous machine minds. This essay argues that such narratives obscure a more consequential and already ongoing transformation: the competition to control the layered infrastructure through which cognition, coordination, and decision-making are mediated. Rather than pursuing general intelligence in the abstract, major technology firms are engaged in a race to capture the cognitive stack, spanning physical computation, energy and resource extraction, data persistence, interface mediation, and narrative interpretation.

By situating recent developments in artificial intelligence alongside cloud infrastructure, data centers, semiconductor supply chains, and military integration, this paper reframes AI not as a primarily epistemic breakthrough but as a phase shift in industrial and geopolitical organization. The analysis shows how platform differentiation masks systemic convergence, producing interlocking forms of enclosure that reshape labor, memory, and institutional judgment. Particular attention is paid to the role of gaming platforms, efficiency-oriented research under constraint, and the material dependencies of large-scale computation.

The essay concludes by arguing that the defining political question raised by contemporary AI is not whether machines will become intelligent, but who controls the energy, memory, and decision infrastructures upon which social life increasingly depends. The elimination of ephemerality, rather than the emergence of artificial cognition, is identified as the central rupture introduced by the stack capture race.

# 1 Introduction: From Intelligence Narratives to Infrastructure Politics

Public and academic discourse surrounding artificial intelligence has increasingly converged on the concept of artificial general intelligence as its organizing horizon. Policy statements, corporate roadmaps, and popular commentary alike frame contemporary developments as steps along a speculative trajectory toward machines capable of autonomous reasoning, self-directed agency, and general problem solving. Within this narrative, the present appears as a transitional moment, valuable primarily insofar as it anticipates a future rupture. The legitimacy of analysis is thereby tethered to prediction: to forecasting when, how, and by whom such a system might emerge.

This orientation, however, has become increasingly inadequate as an explanatory framework. The most consequential transformations associated with contemporary artificial intelligence are no longer hypothetical or incipient. They are already embedded in the ordinary operation of communication systems, workplaces, logistics networks, research practices, and state institutions. These changes do not depend on the realization of artificial general intelligence, nor do they presuppose the existence of autonomous machine agents. They arise instead from the consolidation of control over the infrastructural layers through which human cognition and coordination are mediated.

This essay therefore proposes a shift in analytical focus away from intelligence as a speculative endpoint and toward infrastructure as a site of political struggle. Rather than asking whether current systems will become intelligent in a general sense, it examines how existing technologies are reorganizing the conditions under which knowledge is produced, memory is stored, decisions are made, and action is coordinated. The central claim advanced here is that the present moment is best understood as a race to capture the cognitive stack: a vertically integrated assemblage encompassing physical computation, energy and resource extraction, data persistence, interface design, and narrative interpretation.

Framed in this way, artificial intelligence appears less as a discrete technological breakthrough than as a catalyst accelerating long-standing tendencies toward infrastructural enclosure. Historical parallels can be found in earlier moments of industrial consolidation, including the expansion of rail and telegraph networks in the nineteenth century, the electrification of production and domestic life in the early twentieth century, and the standardization of operating systems and network protocols in the late twentieth century. In each case, control over seemingly technical layers of coordination produced durable asymmetries of power that reshaped economic and political life without requiring overt changes in formal governance.

What distinguishes the contemporary situation is the direct entanglement of these infrastructural layers with cognitive functions traditionally associated with human judgment. Systems for search, communication, document production, code development, and organizational coordination increasingly operate as memory prostheses and decision-support mechanisms. As these systems are integrated into cloud platforms and augmented with large-scale machine learning models, they acquire the capacity not merely to store or transmit information, but to summarize, prioritize, and narrate it. The resulting infrastructures do not simply assist cognition; they actively structure what can be known, recalled, and acted upon.

The competition among major technology firms is often described in terms of model capabilities, parameter counts, or benchmark performance. While such metrics are not irrelevant, they obscure the deeper convergence taking place at the level of infrastructure. Distinct platforms pursue different surface strategies, emphasizing epistemic authority, workflow integration, social mediation, or real-time interpretation. Yet these strategies are unified by a shared orientation toward persistence, enclosure, and dependence. The race is not to build an artificial mind that can replace human judgment, but to become the substrate within which judgment is exercised.

Understanding contemporary artificial intelligence in these terms also clarifies why debates over privacy, alignment, and safety frequently miss their target. These debates tend to assume discrete systems acting upon users from the outside, rather than persistent environments within which users act. The political stakes of the stack capture race lie less in the possibility of runaway intelligence than in the normalization of infrastructures that remember everything, forget nothing, and render judgment increasingly legible to institutional systems rather than human discretion.

The sections that follow elaborate this argument by examining the layered structure of the cognitive stack and the differentiated roles played by major platform actors. Particular attention is given to the material foundations of artificial intelligence, including semiconductor supply chains, data centers, energy systems, and extractive industries, as well as to the military and geopolitical dimensions of large-scale computation. By situating artificial intelligence within these broader contexts, the essay seeks to move beyond speculative futures and toward a forensic analysis of the infrastructural present.

## 2 From Prediction to Forensics

Much of the twentieth- and early twenty-first-century discourse on technological change has been oriented toward prediction. Futurism, scenario planning, and speculative ethics have treated emerging technologies as incipient forces whose primary significance lies in what they may become. Artificial intelligence has been no exception. From early expert systems to contemporary machine learning, debates have repeatedly centered on thresholds, timelines, and discontinuities, asking when machines might achieve human-level intelligence or surpass it.

This predictive orientation has increasingly lost its explanatory power. The most consequential effects of artificial intelligence are no longer deferred to an imagined future but are already observable in institutional practice. Systems for hiring, evaluation, logistics, content moderation, and decision support operate today with measurable consequences for labor, governance, and social coordination. In this context, continued emphasis on speculative endpoints risks obscuring the mechanisms through which power is currently exercised.

A forensic approach offers an alternative analytic stance. Rather than asking what artificial intelligence will become, it asks what it has already done and how those effects are produced. This shift entails attending to infrastructures, incentives, and defaults rather than capabilities alone. It treats artificial intelligence not as an autonomous agent in waiting, but as a set of techniques embedded within organizational systems that redistribute authority, memory, and responsibility.

Adopting a forensic perspective also alters the role of critique. The task is no longer to an-

ticipate hypothetical harms, but to reconstruct causal pathways linking technical design to social outcome. Such reconstruction requires attention to mundane details of implementation, including data retention policies, interface affordances, and energy provisioning. By foregrounding these factors, the essay situates artificial intelligence within a longer history of infrastructural power, preparing the ground for an analysis of stack capture as an already unfolding process rather than a future contingency.

### 3 Stacks, Layers, and the Politics of Mediation

The language of stacks and layers has long been used within computing to describe hierarchical organization, from hardware abstraction to network protocols and application software. In technical contexts, such layering is often presented as a neutral design principle, enabling modularity, interoperability, and efficiency. Yet when extended beyond engineering practice, layered architectures acquire political significance. Decisions about where boundaries are drawn, which layers are exposed, and who controls their interfaces shape how power is distributed within complex systems.

In social and institutional settings, mediation is rarely neutral. Every layer that intervenes between intention and action introduces constraints, defaults, and priorities. Historically, many such mediations were fragmented across institutions, professions, and material practices. Memory resided in archives, judgment in human deliberation, coordination in bureaucratic routines. The contemporary reorganization of these functions into integrated digital systems collapses previously distinct layers into unified platforms.

The concept of a cognitive stack captures this consolidation without presupposing a singular technological essence. It refers instead to the alignment of multiple layers of mediation under common ownership or governance, producing cumulative effects that exceed the sum of their parts. Control over any single layer may appear limited, but control over their integration enables durable forms of influence that operate below the threshold of explicit decision.

Understanding stacks as political formations rather than technical conveniences is essential for interpreting the current moment. The competition among platforms is not merely over services or features, but over the architecture through which mediation itself occurs. This perspective allows the analysis to move beyond surface-level debates about tools and toward the deeper question of how layered control reshapes cognition, coordination, and accountability.

### 4 Cognitive Stacks and the End of Interface Neutrality

The concept of a cognitive stack provides a means of describing the layered organization through which contemporary systems mediate perception, memory, coordination, and action. The term is not intended as a metaphor but as an analytic abstraction that foregrounds vertical integration. At its base lie physical substrates such as energy generation, mineral extraction, semiconductor fabrication, and data center construction. Above these layers sit computational infrastructures, including cloud platforms, scheduling systems, and large-scale machine learning models. At the uppermost levels are interfaces through which users encounter these systems, encompassing com-

munication tools, productivity software, recommendation systems, and narrative synthesis engines. What distinguishes the present configuration is not the existence of these layers individually, but their increasing consolidation under unified institutional control.

Historically, interfaces were treated as neutral conduits between human intention and technical execution. Word processors, operating systems, and network protocols were understood as tools whose political significance derived primarily from access rather than behavior. This assumption no longer holds. Contemporary interfaces do not merely transmit commands or display results; they actively participate in cognitive processes by shaping attention, prioritizing information, and rendering certain actions easier or more difficult than others. As a result, the interface becomes a site where power is exercised not through prohibition, but through affordance.

This transformation is closely tied to the integration of persistent data storage with machine learning systems capable of summarization and inference. When communication histories, documents, code repositories, and audiovisual records are retained indefinitely and made available for algorithmic processing, interfaces acquire temporal depth. They no longer represent only the present state of a task but implicitly reference past actions, prior decisions, and inferred intentions. In such environments, memory ceases to be a personal or collective practice and becomes an infrastructural default.

The end of interface neutrality is most evident in the way contemporary systems collapse distinctions between assistance and judgment. Recommendation engines suggest not only what to read or watch, but what to emphasize, revise, or omit. Writing and coding tools propose completions that implicitly define norms of adequacy and correctness. Organizational platforms reconstruct decision histories in ways that privilege traceable outcomes over informal deliberation. These interventions do not require explicit coercion to be effective. They operate by subtly reconfiguring the conditions under which judgment is exercised, rendering some possibilities salient while relegating others to obscurity.

This shift has significant political implications. When interfaces are no longer neutral, control over their design and operation becomes a form of governance. Decisions about what is surfaced, archived, or synthesized shape institutional memory and influence future action. Unlike traditional forms of authority, which are often legible and contestable, interface-mediated power is distributed, incremental, and difficult to localize. Its effects accumulate over time, producing path dependencies that are experienced as convenience rather than constraint.

The notion of a cognitive stack thus captures a structural inversion in the relationship between humans and their tools. Rather than external instruments used episodically, digital systems increasingly function as environments within which cognition unfolds. The race to capture the stack is therefore not a competition over superior intelligence in isolation, but over the conditions under which intelligence, human or artificial, is exercised. To control the stack is to control not what people think, but how thinking itself is scaffolded, remembered, and rendered actionable.

## 5 Enclosure Beyond Land and Labor

The concept of enclosure originates in the transformation of land from shared or customary use into exclusive private ownership. Historically associated with agrarian capitalism, enclosure restructured social relations by restricting access to resources essential for subsistence and participation. Over time, the logic of enclosure extended beyond land to encompass labor, knowledge, and communication, adapting to new technological and economic conditions.

In contemporary digital contexts, enclosure operates not through fences or legal title alone, but through infrastructural dependency. Access to platforms, networks, and services is formally open yet substantively constrained by terms of use, technical standards, and economic necessity. Participation becomes conditional on compliance with systems that mediate interaction, store memory, and structure visibility. The resulting form of enclosure is less overt than its historical predecessors, but no less consequential.

Cognitive enclosure represents a further extension of this logic. When systems mediate not only access to resources but the processes of thinking, remembering, and deciding, enclosure reaches into domains previously governed by informal norms and human discretion. The privatization of mediation itself alters the conditions under which agency is exercised, rendering cognition dependent on infrastructures that are neither transparent nor democratically accountable.

Situating artificial intelligence within this lineage clarifies why its political significance cannot be reduced to questions of capability or intent. The enclosure at stake is not of intelligence as such, but of the means through which intelligence is expressed and organized. The stack capture race thus appears as a contemporary form of enclosure, one that operates through integration rather than exclusion and through convenience rather than coercion.

## 6 Platform Differentiation Within a Shared Enclosure

The competition among major technology platforms is often described in terms of rivalry, disruption, or leadership in artificial intelligence research. Such descriptions emphasize surface-level differentiation while obscuring a deeper structural convergence. Although firms pursue distinct strategies and occupy different cultural positions, they increasingly participate in a shared project of infrastructural enclosure. Each seeks to dominate particular layers of the cognitive stack, not in isolation, but in ways that reinforce the stability of the overall system.

This differentiation is most visible at the level of interfaces and user-facing services. Some platforms emphasize epistemic mediation, positioning themselves as authoritative sources of knowledge, search, and synthesis. Others prioritize workflow integration, embedding themselves within the routines of professional labor and institutional coordination. Still others focus on social mediation, capturing attention, affect, and behavioral feedback at scale, or on real-time narrative interpretation that claims immediacy and relevance in moments of uncertainty. These orientations appear competitive insofar as they vie for user loyalty and market share, yet they are complementary in their infrastructural effects.

What unites these approaches is a shared commitment to persistence. Communication is archived by default, drafts are versioned indefinitely, interactions are logged, and histories are

rendered searchable and analyzable. The accumulation of data across time allows platforms to reconstruct not only what occurred, but how it unfolded, enabling retrospective interpretation and optimization. In this sense, differentiation at the interface level masks homogeneity at the temporal level. Regardless of whether a platform presents itself as a productivity tool, a knowledge assistant, or a social space, it participates in the same expansion of institutional memory.

This convergence is further reinforced by the reliance of ostensibly competing platforms on common infrastructural foundations. Cloud computing services, semiconductor supply chains, content delivery networks, and energy-intensive data centers form a shared substrate upon which higher-level services depend. Even where platforms compete in the provision of these services, they do so within a narrow band of technical and economic constraints that favor scale, capital intensity, and long-term lock-in. As a result, competition tends to stabilize rather than destabilize the underlying system.

The appearance of rivalry thus serves an important legitimating function. By framing developments as a contest among firms, attention is directed toward questions of innovation, performance, and consumer choice. Less visible are the collective consequences of enclosure, including the normalization of persistent surveillance, the erosion of informal discretion, and the gradual transfer of judgment from human actors to infrastructural systems. These outcomes are not the result of collusion in the narrow legal sense, but of alignment around a shared model of value extraction grounded in cognitive mediation.

Understanding platform differentiation as occurring within a shared enclosure helps clarify why regulatory interventions often struggle to gain traction. Antitrust frameworks designed to address monopolistic behavior at the firm level are poorly suited to addressing power exercised through layered infrastructure. Even where individual platforms are constrained or fragmented, the broader logic of the cognitive stack remains intact. Control is distributed across interfaces, services, and physical substrates in ways that resist simple attribution.

The stack capture race therefore does not produce a single winner who dominates all layers of cognition. Instead, it yields a stable configuration in which multiple actors control interdependent segments of the stack, each reinforcing the others' relevance. This distributed form of dominance is less visible than traditional monopoly, but more resilient. It ensures that the enclosure of cognition proceeds incrementally, experienced by users as convenience, integration, and inevitability rather than as overt coercion.

## 7 Gaming, Pleasure, and the Normalization of Enclosure

The incorporation of gaming platforms into contemporary technology ecosystems reveals an often overlooked pathway through which enclosure becomes socially acceptable. Unlike enterprise software or administrative systems, games are adopted voluntarily, associated with leisure rather than obligation, and experienced as spaces of agency and play. These characteristics make gaming an ideal environment for experimenting with persistent identity, behavioral telemetry, and real-time moderation at scale. Long before similar mechanisms were introduced into workplaces and educational settings, gaming platforms had already normalized many of the practices that now underpin

cognitive enclosure.

Online games require continuous identity management, the maintenance of persistent worlds, and the coordination of large populations across time and space. Achievements, rankings, cosmetic differentiation, and social affiliations transform identity into a legible and comparable construct, while telemetry systems record actions with fine granularity. Crucially, these features are not experienced as surveillance. They are framed as feedback, progression, and community participation. The result is a form of consent that is affective rather than contractual, grounded in enjoyment rather than compliance.

This affective normalization has significant consequences for how similar mechanisms are later received in non-entertainment contexts. Systems that track productivity, log communication, or reconstruct decision histories draw on design principles first refined in gaming environments. Badges, performance metrics, and dashboards echo achievement systems, while persistent profiles and social graphs mirror player identities and guild structures. What changes is not the underlying logic, but the context in which it is applied. Control mechanisms that would provoke resistance if introduced directly into workplaces become acceptable when presented as extensions of familiar, pleasurable interactions.

Gaming also provides a controlled environment in which platforms can study human adaptation to algorithmic mediation. Player behavior offers rich data on motivation, frustration, cooperation, and compliance under rule-bound conditions. This knowledge is transferable to broader domains of social coordination, informing the design of interfaces that guide behavior without explicit instruction. The capacity to tune feedback loops, adjust incentives, and manage conflict at scale is developed not through abstract theory, but through sustained experimentation within entertainment systems.

The significance of gaming for the stack capture race lies not in its cultural content, but in its infrastructural role as a rehearsal space. It demonstrates how persistent environments can be made desirable, how monitoring can be rendered invisible, and how enclosure can be experienced as empowerment. When similar techniques are deployed in professional or civic contexts, they encounter a population already acclimated to their logic. The boundary between play and work, once a site of distinction, becomes a vector for the migration of control.

By the time cognitive mediation is introduced into domains such as writing, coding, or organizational coordination, the underlying practices have already been socially legitimized. The novelty lies not in the mechanisms themselves, but in their repositioning as tools for efficiency and insight. Gaming thus occupies a critical but understated position within the contemporary technological landscape. It is the domain in which enclosure is perfected under conditions of pleasure, preparing the ground for its extension into the rest of social life.

## 8 Compute as Heavy Industry: GPUs, Data Centers, and Physical Constraint

The abstraction of artificial intelligence as a purely informational or computational phenomenon obscures the extent to which contemporary systems are grounded in heavy industrial infrastructure.

Despite the language of clouds and virtuality, large-scale machine learning depends on material arrangements whose scale and rigidity more closely resemble twentieth-century manufacturing than software development. The expansion of artificial intelligence must therefore be understood as a process of reindustrialization, one that reintroduces familiar constraints of energy, land, labor, and capital under the guise of digital innovation.

At the center of this transformation lies the graphical processing unit, which has become the decisive bottleneck for advanced computation. High-performance GPUs function as the primary means through which machine learning models are trained and deployed at scale. Their production requires complex global supply chains, advanced fabrication facilities, and long-term capital commitments that sharply limit entry. As a result, access to computational capacity is no longer primarily a function of technical ingenuity but of geopolitical alignment, purchasing power, and infrastructural foresight. The concentration of GPU manufacturing and design thus operates as a form of cognitive chokepoint, regulating not only who can build advanced systems but what kinds of systems are economically viable.

The role of data centers further illustrates the industrial character of contemporary computation. Far from lightweight server farms, modern AI-oriented data centers are energy-intensive facilities designed to sustain continuous high-load operation. They require proximity to reliable baseload power, abundant water resources for cooling, and favorable regulatory environments. Their construction reshapes local geographies, drawing electricity away from surrounding communities, stressing municipal infrastructure, and introducing new dependencies between technology firms and regional governments. These facilities are not easily repurposed or relocated, embedding artificial intelligence development within fixed physical sites that carry long-term political and environmental consequences.

Energy emerges as a particularly salient constraint within this system. The computational workloads associated with training and operating large models demand uninterrupted power at scales that exceed the capacity of many renewable systems operating alone. As a result, artificial intelligence development has reinforced the importance of fossil fuels, nuclear energy, and other forms of dispatchable generation capable of meeting constant demand. This dynamic complicates narratives that position digital technologies as inherently dematerialized or environmentally progressive. Rather than displacing extractive industries, artificial intelligence extends their relevance by creating new forms of demand that align with existing energy infrastructures.

The material requirements of computation also extend upstream into mining and resource extraction. Semiconductor manufacturing relies on high-purity silicon, rare earth elements, copper, and specialized chemicals, all of which are sourced through environmentally intensive processes. The expansion of data center infrastructure amplifies demand for these materials, linking cognitive technologies to global extractive regimes that disproportionately affect regions already subject to ecological and political vulnerability. In this sense, the cognitive stack rests upon a geological foundation, binding abstract processes of inference and optimization to the physical transformation of landscapes.

Transport and logistics further constrain the development of large-scale computation. High-end chips are produced in limited locations and must be shipped securely and rapidly to data centers

distributed across the globe. Disruptions to shipping routes, trade relations, or manufacturing capacity therefore have immediate implications for artificial intelligence development. The sensitivity of these systems to geopolitical instability underscores the extent to which computation has become entangled with national security concerns, trade policy, and strategic resource management.

Taken together, these factors undermine the notion that artificial intelligence development is governed primarily by innovation in algorithms or architectures. While such innovations remain important, they operate within a narrow corridor defined by material feasibility. The economics of compute privilege scale, continuity, and capital intensity, favoring actors capable of sustaining long-term infrastructural investment. This reality helps explain why the stack capture race converges on a small number of dominant platforms, despite ongoing experimentation and competition at the model level.

Understanding compute as heavy industry also clarifies the political stakes of artificial intelligence. Control over computational infrastructure translates into control over the tempo and direction of cognitive mediation itself. Decisions about where data centers are built, how energy is sourced, and which applications are prioritized shape not only technological outcomes but social and institutional ones. The apparent neutrality of computation masks a series of choices that distribute costs and benefits unevenly across populations and regions.

The stack capture race thus unfolds not only in laboratories and software repositories but across power grids, supply chains, and extraction sites. Artificial intelligence, far from heralding a post-industrial future, reasserts the centrality of material constraint in shaping technological possibility. Any analysis that neglects this dimension risks mistaking surface-level innovation for structural transformation.

## 9 Energy, Extraction, and the Return of Resource Geopolitics

The material demands of large-scale computation draw artificial intelligence directly into the domain of energy politics and resource extraction, reviving patterns more commonly associated with earlier industrial transitions. Although contemporary AI is often framed as a clean or dematerialized technology, its expansion intensifies dependence on energy systems characterized by geographic fixity, long investment horizons, and geopolitical contestation. In this respect, the stack capture race does not transcend the political economy of resources but reinscribes it in a new register.

The defining feature of computational energy demand is its constancy. Training and inference workloads require uninterrupted power delivery over extended periods, rendering them poorly suited to energy systems based on intermittency alone. While renewable sources play an increasing role in marginal capacity, they are frequently supplemented by fossil fuels, nuclear generation, or other dispatchable sources capable of maintaining stable output. This requirement has led technology firms to pursue long-term power purchase agreements, direct investment in generation facilities, and partnerships with energy producers whose assets were previously associated with declining or transitional sectors. Artificial intelligence thus acts as a stabilizing force for incumbent energy regimes even as it is rhetorically associated with technological futurism.

These energy dependencies introduce new forms of spatial concentration. Data centers cluster

near sources of cheap and reliable power, reshaping regional economies and infrastructure priorities. Local governments compete to attract such facilities through tax incentives and regulatory concessions, often accepting long-term environmental and fiscal costs in exchange for short-term investment. The resulting arrangements resemble earlier forms of industrial zoning, in which communities are bound to particular extractive or energy-intensive activities with limited capacity for exit. In this way, cognitive infrastructure becomes a driver of place-based inequality, linking abstract processes of computation to concrete territorial commitments.

Upstream from energy generation lies the expanding demand for raw materials. Semiconductor fabrication, power transmission, and cooling infrastructure require vast quantities of copper, aluminum, rare earth elements, and specialized minerals. The extraction and processing of these materials are geographically uneven and environmentally destructive, frequently concentrated in regions with limited regulatory oversight or political leverage. As demand grows, so too does pressure on mining rights, water access, and land use, intensifying conflicts that echo those associated with earlier waves of industrialization. Artificial intelligence, far from dissolving material scarcity, redistributes it along new vectors.

These dynamics have clear geopolitical implications. States increasingly recognize computational capacity as a strategic asset, tied not only to economic competitiveness but to military and intelligence capabilities. Control over energy supplies, mineral resources, and manufacturing capacity becomes inseparable from control over cognitive infrastructure. Export controls, sanctions, and trade restrictions aimed at limiting access to advanced computation reveal the extent to which artificial intelligence is embedded in broader struggles over technological sovereignty. Such measures do not halt development but redirect it, reinforcing regional blocs and alternative supply chains.

The reemergence of resource geopolitics challenges narratives that treat artificial intelligence as a purely informational domain governed by market competition and innovation alone. Instead, it situates AI within a familiar pattern in which technological advances amplify the strategic importance of energy and materials. The difference lies in the object of control. Whereas previous regimes centered on transportation, manufacturing, or fuel, the current configuration centers on the capacity to mediate cognition itself. Energy and extraction are no longer merely inputs to production but prerequisites for participation in the informational order.

By embedding artificial intelligence within these material and geopolitical constraints, the stack capture race acquires a durability that exceeds any particular model or platform. Investments in energy infrastructure and resource extraction create sunk costs and long-term dependencies that shape technological trajectories for decades. The future of cognitive mediation is thus bound not only to advances in algorithms but to decisions made in mines, power plants, and regulatory chambers. Recognizing this continuity is essential to understanding artificial intelligence not as a rupture from industrial history, but as its latest and most abstract expression.

## 10 Military Entanglement and the Dual-Use Stack

The convergence of artificial intelligence, large-scale computation, and infrastructural consolidation renders any strict separation between civilian and military applications increasingly untenable.

From its material foundations to its highest-level interfaces, the contemporary cognitive stack is structured in ways that align naturally with defense and security objectives. This alignment does not require explicit militarization to be effective. It arises from the shared requirements of decision-making under uncertainty, large-scale coordination, simulation, and logistical optimization. As a result, military integration is not an external imposition on the stack capture race, but an intrinsic dimension of it.

At the infrastructural level, cloud computing platforms and data centers are designed to support high-reliability workloads, secure data handling, and rapid scalability. These characteristics are equally valuable for enterprise productivity and for intelligence analysis, battlefield simulation, and command-and-control systems. Contracts between technology firms and defense agencies formalize this overlap, but the underlying compatibility precedes any particular agreement. The same systems that store organizational memory or optimize supply chains can be repurposed to model force deployment, analyze surveillance data, or coordinate autonomous systems. The distinction between civilian and military use thus becomes a matter of access control rather than architectural difference.

Machine learning models further blur this boundary by functioning as general-purpose tools for pattern recognition, prediction, and planning. Techniques developed for commercial applications such as recommendation, language processing, or image classification transfer readily to military contexts. Training data and deployment environments may differ, but the core capabilities remain aligned. This dual-use character complicates ethical and regulatory responses, as restrictions aimed at preventing military misuse often target downstream applications rather than upstream infrastructure. Yet it is precisely this infrastructure that confers strategic advantage.

The role of simulation merits particular attention. Artificial intelligence enables the construction of detailed virtual environments in which scenarios can be explored, strategies tested, and outcomes evaluated without immediate real-world consequences. Such simulations are invaluable for military planning, where uncertainty and risk are endemic. At the same time, similar techniques underpin commercial forecasting, logistics optimization, and organizational decision support. The increasing fidelity of these simulations, supported by advances in compute and data integration, strengthens the appeal of AI as a decision-support technology across domains, reinforcing the incentives for infrastructural consolidation.

The military relevance of the cognitive stack also extends to its physical substrate. Data centers, energy supplies, and semiconductor fabrication facilities become strategic assets whose protection and continuity are matters of national security. Disruptions to power grids, supply chains, or manufacturing capacity have cascading effects on computational capability. As a result, the geography of artificial intelligence infrastructure acquires defensive significance, influencing planning around resilience, redundancy, and territorial control. These considerations further entrench AI within state-level security frameworks.

The integration of artificial intelligence into military contexts also feeds back into civilian development. Defense funding, security requirements, and operational demands shape research priorities and deployment standards. Emphasis on reliability, robustness, and scalability reinforces design choices that favor large, centralized systems over smaller, decentralized alternatives. In this way, military entanglement contributes to the stabilization of the stack capture regime, aligning technolo-

logical evolution with institutional preferences for control and predictability.

Understanding this dual-use dynamic is essential for assessing the political implications of artificial intelligence. The stack capture race does not merely reflect commercial competition or technological ambition; it participates in a broader reconfiguration of power in which cognitive mediation becomes a strategic resource. Military applications do not represent an aberration from this process but a continuation of it under conditions of heightened stakes. Any attempt to evaluate the social consequences of artificial intelligence must therefore grapple with its role within contemporary security architectures, rather than treating militarization as a peripheral concern.

## 11 Constraint, Efficiency, and the Limits of Abundance

The prevailing narrative of artificial intelligence development assumes a direct relationship between progress and computational abundance. Larger models, greater quantities of data, and more powerful hardware are treated as the primary drivers of advancement. Within this framework, dominance is secured through access to scale, and constraints are viewed as obstacles to be overcome. Yet recent developments challenge this assumption by demonstrating that significant capabilities can emerge under conditions of limitation, calling into question the inevitability of compute-centric dominance.

Research conducted under constraint foregrounds efficiency as a primary design principle. Rather than maximizing parameter counts or training time, such approaches emphasize algorithmic refinement, architectural parsimony, and careful allocation of computational resources. These strategies do not merely compensate for scarcity; they reshape the character of the resulting systems. Models developed in constrained environments often exhibit different trade-offs, prioritizing reasoning depth, adaptability, or interpretability over raw throughput. In doing so, they reveal that abundance is not the sole pathway to effective cognitive mediation.

The significance of this shift extends beyond technical performance. Constraints imposed by export controls, sanctions, or limited access to advanced hardware alter the geopolitical landscape of artificial intelligence development. They incentivize alternative trajectories that reduce dependence on centralized supply chains and challenge the assumption that leadership in AI is inseparable from control over the most advanced manufacturing capabilities. Such developments complicate efforts to regulate or contain technological diffusion through hardware restrictions alone.

Efficiency-oriented research also exposes vulnerabilities in the prevailing infrastructure model. Systems optimized for extreme scale entail high fixed costs, energy consumption, and environmental impact. Their economic viability depends on sustained demand and continuous expansion, conditions that may not hold indefinitely. By contrast, approaches that achieve comparable functionality with reduced resources suggest the possibility of more distributed and resilient configurations. These alternatives threaten not only established narratives of progress but the business models predicated on perpetual infrastructural growth.

The emergence of effective systems under constraint thus has strategic implications for the stack capture race. It demonstrates that control over cognitive mediation cannot be secured solely through accumulation of hardware and energy resources. Algorithmic ingenuity and organizational

adaptation remain potent forces, capable of shifting competitive balances and undermining chokepoints. At the same time, such developments do not negate the importance of infrastructure; they reframe it as one component within a more complex landscape of capability.

By highlighting the limits of abundance, efficiency-oriented approaches invite a reconsideration of what constitutes leadership in artificial intelligence. Rather than equating progress with scale, they foreground questions of sustainability, accessibility, and adaptability. In doing so, they open conceptual space for alternatives to the dominant model of stack capture, even as they operate within the constraints imposed by existing political and economic structures. The tension between abundance and efficiency thus becomes a fault line within the broader transformation of cognitive infrastructure, with implications that extend well beyond the technical domain.

## 12 Grassroots Resistance and the Limits of Enclosure

The expansion of cognitive infrastructure has not proceeded without resistance. Alongside the consolidation of platforms, data centers, and energy systems, a heterogeneous array of grassroots organizations has emerged to contest the social, ecological, and political consequences of technological enclosure. These movements differ widely in origin, scale, and ideology, yet they share a common refusal to accept the inevitability of stack capture as a neutral or purely technical process. Their interventions expose the fact that the cognitive stack is not merely built but negotiated, often contentiously, within specific social and territorial contexts.

Indigenous communities occupy a particularly salient position within this landscape of resistance. The material requirements of artificial intelligence infrastructure, including energy generation, mineral extraction, water access, and land use, frequently intersect with Indigenous territories and treaty lands. Opposition to data center construction, mining operations, and energy projects is therefore not solely environmental but epistemic and political. These struggles assert alternative conceptions of stewardship, temporality, and collective memory that conflict directly with the logics of persistent extraction and infrastructural permanence. In resisting enclosure, Indigenous movements challenge not only specific projects but the underlying assumption that cognition can be industrialized without regard to place, history, or relational obligation.

Across Europe, resistance has often taken a different but complementary form. Regulatory activism, labor organization, and digital rights movements have sought to limit the power of large platforms through legal, institutional, and collective mechanisms. Concerns over data protection, algorithmic governance, workplace surveillance, and environmental impact have translated into policy debates and public campaigns that foreground the social costs of technological integration. While such efforts operate within existing political frameworks, they nonetheless articulate a countervailing vision in which cognitive infrastructure is subject to democratic oversight rather than market inevitability.

Grassroots opposition is not limited to formal activism or regulation. Informal practices of refusal, adaptation, and withdrawal also play a role. Communities experiment with alternative platforms, localized infrastructures, and modes of coordination that privilege ephemerality, autonomy, or mutual aid over efficiency and scale. These efforts are often fragile and uneven, constrained

by the pervasive reach of dominant systems. Yet they demonstrate that enclosure is neither total nor uncontested. The persistence of such practices indicates that cognitive mediation remains a site of struggle rather than a settled condition.

Importantly, resistance does not always take the form of outright rejection. In many cases, it involves selective engagement, negotiation, or the imposition of conditions on participation. Indigenous agreements over land use, European regulatory compromises, and community-level bargaining over infrastructure projects illustrate how the expansion of the cognitive stack is shaped by ongoing conflict. These interactions complicate narratives that portray technological development as a unidirectional force imposed from above. Instead, they reveal a dynamic process in which power is exercised, resisted, and reconfigured through social action.

The existence of grassroots resistance underscores a central claim of this essay: that the stack capture race is not merely a technological competition but a political transformation with uneven effects. Efforts to contest enclosure draw attention to values marginalized by dominant platforms, including ecological sustainability, collective memory, and the right to refuse persistent mediation. While such movements face significant structural disadvantages, their presence challenges the assumption that the future of cognition is already determined. In doing so, they reassert the possibility that alternative arrangements, however constrained, remain conceivable within an increasingly enclosed world.

### 13 Stack Capture as Regime Transition

Taken together, the developments examined in the preceding sections suggest that the contemporary transformation associated with artificial intelligence is best understood not as a sectoral shift or a technological cycle, but as a regime transition. The concept of regime transition is used here to denote a reorganization of social, economic, and political relations that alters the underlying conditions of coordination rather than merely introducing new tools within an existing order. In this sense, the stack capture race marks a transition comparable in scope to earlier reorganizations associated with industrialization, electrification, or the rise of networked computation.

What distinguishes a regime transition from incremental change is the stabilization of new dependencies. As cognitive mediation becomes embedded within infrastructural layers that are capital intensive, energy dependent, and territorially fixed, alternatives become progressively harder to sustain. Individuals, institutions, and states adapt their practices to the affordances of these systems, reinforcing their centrality even in the absence of explicit coercion. Over time, what began as optional assistance becomes a baseline expectation, and withdrawal comes to appear as dysfunction rather than choice.

This process is reinforced by the interdependence of the actors involved. No single platform controls the entire cognitive stack, yet each relies on the others to sustain the overall configuration. Cloud providers depend on energy and extraction regimes; interface platforms depend on cloud infrastructure; model developers depend on both. Military and state actors rely on the same systems for planning, logistics, and intelligence, further entrenching their strategic importance. The result is a distributed form of power that resists disruption precisely because it lacks a single

point of control.

Regime transition is also evident in the changing nature of governance. Decisions that once occurred through explicit institutional processes are increasingly mediated by infrastructural defaults. Questions of prioritization, relevance, and accountability are resolved through interface design, data retention policies, and algorithmic synthesis rather than through deliberation or rule-making alone. Authority migrates from visible decision-makers to the architects and operators of systems that structure decision space itself. This shift does not abolish formal governance, but it renders it reactive, operating within constraints established elsewhere.

The temporal dimension of this transition is particularly significant. Persistent storage, comprehensive logging, and retrospective analysis alter the relationship between action and judgment. Past behavior becomes permanently available for reinterpretation, reducing the role of context, intention, and forgetting in social evaluation. Institutions gain the capacity to reconstruct narratives of performance and responsibility with unprecedented granularity, while individuals lose the protective ambiguity that once accompanied memory's fragility. This asymmetry contributes to a rebalancing of power that favors organizations over persons and systems over discretion.

Importantly, the stack capture regime does not depend on the realization of artificial general intelligence. Its stability derives from infrastructural integration rather than cognitive autonomy. Even modest forms of machine learning, when embedded within persistent environments and coupled to large-scale data retention, suffice to reshape social relations. The promise of more advanced intelligence serves primarily to justify continued investment and expansion, not to define the regime's functional core.

Understanding stack capture as a regime transition clarifies both its resilience and its vulnerability. Like earlier infrastructural regimes, it is difficult to dismantle once established, yet it remains contingent on material, political, and social conditions. Energy constraints, ecological limits, geopolitical conflict, and organized resistance all impose pressures that can redirect or destabilize its trajectory. The regime's apparent inevitability thus reflects its current alignment with dominant economic and institutional interests rather than any intrinsic necessity.

The final section turns to the normative implications of this transition, focusing on the erosion of ephemerality and the politics of memory. If stack capture represents a reorganization of how cognition is mediated and remembered, then the central question is no longer whether machines will think, but how societies will preserve the conditions for judgment, forgiveness, and refusal within systems designed to remember indefinitely.

## 14 Conclusion: Ephemerality, Memory, and the Politics of Forgetting

The analysis developed in this essay has sought to reframe contemporary debates about artificial intelligence by shifting attention away from speculative futures and toward the infrastructural present. Rather than interpreting current developments as steps along a linear path toward artificial general intelligence, the essay has argued that the defining transformation of the present moment lies in the consolidation of control over the cognitive stack. This consolidation reshapes how knowledge

is produced, how decisions are made, and how social action is coordinated, independent of whether machines ever attain autonomous intelligence.

At the center of this transformation lies a reconfiguration of memory. Persistent storage, comprehensive logging, and algorithmic synthesis have converted memory from a fragile, negotiated human practice into an infrastructural default. What is remembered is no longer primarily a matter of personal recollection or collective narration, but of system design and institutional retention. This shift alters the moral and political texture of social life. Forgetting, once an ordinary feature of human interaction and a precondition for forgiveness, discretion, and change, becomes increasingly difficult to justify within environments optimized for recall.

The erosion of ephemerality has consequences that extend beyond privacy in its conventional sense. It transforms accountability by enabling retrospective reconstruction of action divorced from context, intention, or situational constraint. It privileges legibility over judgment and traceability over understanding. In such environments, individuals are rendered permanently answerable to systems that remember more than they can interpret, while institutions gain the capacity to narrate events with an authority that exceeds lived experience. The imbalance this produces is structural rather than accidental, arising from the asymmetry between human finitude and infrastructural persistence.

The stack capture race intensifies this asymmetry by aligning cognitive mediation with capital-intensive, energy-dependent, and militarily entangled infrastructures. These arrangements favor scale, continuity, and enclosure, crowding out alternatives that depend on locality, discretion, or voluntary forgetting. Resistance movements, regulatory efforts, and efficiency-oriented research demonstrate that this trajectory is neither uncontested nor inevitable, yet they operate within a landscape increasingly shaped by infrastructural defaults that resist reversal. The struggle over cognitive mediation thus becomes a struggle over the conditions of social possibility rather than over any single technology.

Recognizing stack capture as a regime transition clarifies the stakes of contemporary technological change. The question is not whether artificial intelligence will surpass human intelligence, but whether societies will retain the capacity to define the terms under which cognition is mediated, remembered, and judged. This is a political question that cannot be resolved through technical optimization alone. It demands engagement with energy systems, extractive practices, labor relations, governance structures, and cultural norms surrounding memory and responsibility.

The future of artificial intelligence, understood in this light, is inseparable from the future of forgetting. Preserving spaces of ephemerality, ambiguity, and refusal may prove more consequential than achieving ever-greater computational power. If cognition is increasingly mediated by infrastructures designed to persist, then the defense of human judgment will depend not on resisting intelligence itself, but on contesting the conditions under which intelligence is made permanent. The stack capture race thus confronts societies with a choice that precedes any question of machine autonomy: whether to accept a world in which nothing is allowed to fade, or to insist that forgetting remains a constitutive element of freedom.

## References

- [1] B. H. Bratton. *The Stack: On Software and Sovereignty*. MIT Press, Cambridge, MA, 2016.
- [2] P. N. Edwards. *A Vast Machine: Computer Models, Climate Data, and the Politics of Global Warming*. MIT Press, Cambridge, MA, 2010.
- [3] S. Zuboff. *The Age of Surveillance Capitalism*. PublicAffairs, New York, 2019.
- [4] C. Doctorow. *The Internet Con: How to Seize the Means of Computation*. Verso, London, 2023.
- [5] K. Crawford. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press, New Haven, 2021.
- [6] M. Mazzucato. *The Value of Everything: Making and Taking in the Global Economy*. PublicAffairs, New York, 2018.
- [7] L. Winner. Do artifacts have politics? *Daedalus*, 109(1):121–136, 1980.
- [8] M. Heidegger. The question concerning technology. In *The Question Concerning Technology and Other Essays*. Harper & Row, New York, 1977.
- [9] F. Pasquale. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press, Cambridge, MA, 2015.
- [10] P. Virilio. *Speed and Politics: An Essay on Dromology*. Semiotext(e), New York, 1986.
- [11] J. C. Scott. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press, New Haven, 1998.
- [12] G. C. Bowker and S. L. Star. *Sorting Things Out: Classification and Its Consequences*. MIT Press, Cambridge, MA, 1999.
- [13] H. Arendt. *The Human Condition*. University of Chicago Press, Chicago, 1958.
- [14] B. Latour. *We Have Never Been Modern*. Harvard University Press, Cambridge, MA, 1993.
- [15] D. Graeber. *Bullshit Jobs: A Theory*. Simon & Schuster, New York, 2018.