Working Paper

Two Structures of AI Power:
National Capacity, Relational Boundary Control, and Systemic Influence in the Emerging AI Order

Author: Shaoyuan Wu
ORCID: https://orcid.org/0009-0008-0660-8232
Affiliation: Global AI Governance and Policy Research Center, EPINOVA LLC
Date: September 1, 2026

T-Series Information:
The WP-T series develops general theoretical frameworks for analyzing power, networks, systems, boundaries, and structural transformation across multiple domains. It emphasizes transferable concepts that can support later applied research in AI, governance, geopolitics, infrastructure, and security.

Recommended Citation:
Wu, Shaoyuan. (2026). *Two Structures of AI Power: National Capacity, Relational Boundary Control, and Systemic Influence in the Emerging AI Order* (EPINOVA Working Paper No. EPINOVA–WP–T–2026–02). Global AI Governance and Policy Research Center, EPINOVA LLC. https://doi.org/10.67037/epinova.wp.t.2026.002

Executive Summary

Artificial intelligence (AI) competition is often measured through models, compute, investment, and adoption, yet these indicators conflate resources, relational leverage, and system-level outcomes. This article develops a two-structure framework that distinguishes AI National Power (ANP) from AI Boundary Power (ABP). ANP is the usable capacity of a national system to generate, mobilize, govern, and sustain AI-relevant resources. ABP is an actor-level relational capacity to alter the conditions under which another actor can obtain, transfer, integrate, deploy, or continue using an AI-relevant resource at a specified interface. The framework further separates boundary control from target-side boundary autonomy and from systemic influence as a later outcome. Through conceptual synthesis, operationalization, and mechanism illustrations involving advanced-computing export controls and European Union AI regulation, the article identifies effective control, interface criticality, enforceability, substitutability, and organizational conversion as key conditioning mechanisms. It argues that national depth and boundary control are distinct but potentially complementary forms of AI power. Their interaction helps explain why technologically capable states may remain externally constrained, why narrower actors can exercise consequential leverage, and why durable influence depends on more than possession of advanced AI.

Keywords: AI power; national power; boundary power; technological interdependence; AI sovereignty; structural power

1. Introduction

Artificial intelligence has become a central object of strategic competition, but the language of "AI power" is often asked to describe fundamentally different things. Model benchmarks, investment totals, compute capacity, research output, data-center construction, platform reach, regulatory influence, military integration, and public-sector adoption are all used as evidence of power. They do not, however, refer to the same unit of analysis or the same mechanism. A state may possess deep domestic AI capability while remaining dependent on external chokepoints. A regulator or platform may alter the behavior of foreign firms without possessing a comprehensive domestic AI stack. A supplier may control a critical interface while remaining vulnerable to upstream inputs controlled by another actor. These positions cannot be represented adequately as points on a single leaderboard.

GLOBAL AI GOVERNANCE RESEARCH CENTER

EPINOVA

Recent measurement and governance research has begun to move beyond model performance by examining infrastructure, investment, adoption, policy, and the distribution of power among AI actors. National-strategy and technology-sovereignty research likewise emphasizes that states differ in how they build, govern, and preserve agency under interdependence (Kim et al., 2025; Papyshev & Yarime, 2023; Radu, 2021; Stanford Institute for Human-Centered Artificial Intelligence, 2026). The unresolved analytical problem is that domestic capacity, external leverage, autonomy from dependence, and system-shaping influence are still frequently treated as interchangeable manifestations of a single variable.

This article asks a narrower question: how should AI power be conceptualized when national capability, control over consequential interfaces, target-side autonomy, and durable system-level effects do not coincide? It develops a two-structure framework that distinguishes AI National Power (ANP) from AI Boundary Power (ABP). ANP is a state-level capacity construct: the usable capacity of a territorially and politically constituted national system to generate, mobilize, govern, and sustain AI-relevant resources for specified functions. ABP is an actor-level relational construct: the effective capacity to alter the conditions under which another actor can obtain, transfer, integrate, deploy, or continue using an AI-relevant resource at a defined interface.

The framework adds two further distinctions. First, control over another actor's crossing conditions is not the same as the target's ability to switch, bypass, substitute, or internalize; the latter is termed boundary autonomy. Second, neither high ANP nor high ABP is defined as systemic influence. Systemic influence is a later outcome: a persistent and sufficiently broad change in feasible option sets, network structure, or governing rules within a pre-specified domain. This separation prevents resource possession, bilateral leverage, and system transformation from being counted as the same phenomenon.

The article is conceptual rather than a country ranking or causal test. It proceeds through concept formation, comparison with adjacent literatures, operationalization, falsifiable propositions, and two mechanism illustrations. The contribution is not a new universal index. It is an analytical architecture for specifying who possesses what kind of power, over whom, at which interface, through what mechanism, and with what observable consequences. This is especially relevant to international studies because AI power is simultaneously national, transnational, infrastructural, institutional, and relational.

2. Theoretical Positioning and Research Design

The first building block is a system-level account of AI capability. Earlier work in this research program treated model capability as only one component of consequential AI capacity and distinguished model capability, data quality, resource access, authority, real-world interfaces, and rule space (Wu, 2026b). That framework is useful at the level of an AI system or organizational deployment because it asks what a system can know, access, mobilize, do under authorization, connect to, and change. It should not, however, be mechanically aggregated into a national score. The national question is which capabilities a political system can reliably mobilize, coordinate, sustain, and protect under specified conditions.

The second building block comes from structural and network theories of power. Structural power emphasizes control over the frameworks within which others act. Network approaches distinguish positional centrality from the ability to program, switch, or exploit asymmetric networks. Weaponized interdependence shows that network hubs become coercively consequential when concentrated position is joined to jurisdiction and institutional capacity, while regulatory-power research shows how market access can externalize rules beyond territorial borders (Bradford, 2020; Castells, 2011; Farrell & Newman, 2019; Strange, 1988). These traditions already establish that power is not reducible to material inventories.

ABP therefore does not claim that earlier research overlooked access control. Its incremental value is narrower: it supplies a common interface-level vocabulary for AI systems in which strategically important boundaries are distributed across compute, data, models, platforms, cloud infrastructure, technical standards, procurement, finance, and security authorization. It also makes target substitutability part of the relation rather than an afterthought. The concept draws directly on a broader conditional-connectivity framework in which connectivity is actor-, function-, direction-, and time-specific and in which usable access depends on boundary conditions rather than topology alone (Wu, 2026a). In the present article that general boundary-control logic is adapted to AI-specific interfaces.

The third building block is the literature on technological and digital sovereignty. Sovereignty is treated here neither as autarky nor as a synonym for power. A state can have high ANP while remaining exposed to an externally controlled chip, cloud, model, or standards interface; another state can retain meaningful agency through diversified suppliers, trusted partnerships, interoperable architectures, and enforceable fallback arrangements (Edler et al., 2023; Tanner et al., 2026; Woods, 2023). The framework therefore treats sovereignty as a configuration of capability, dependable access, and dependence exposure.

Methodologically, the article uses concept reconstruction and mechanism-centered comparison. Each construct is assigned a unit of analysis, constituent mechanism, and distinct outcome set. ANP is evaluated at the state level; ABP at the controller-target-interface-time level; boundary autonomy at the target level; and systemic influence at the level of a pre-specified network, market, sector, or rule domain. The empirical illustrations are selected because they expose different forms of boundary control—licensing and denial in advanced compute, and market-access regulation in the European Union—rather than because they constitute a representative sample. Their purpose is to demonstrate coding logic and scope conditions, not to estimate average effects.

3. Analytical Architecture: Four Objects That Should Not Be Collapsed

The framework separates four analytical objects. The first is AI-system capability: the functional capacity of a specific model, agent, platform, or organizational deployment under defined permissions and resources. The second is ANP: usable AI capacity organized at the national-system level. The third is ABP: effective relational control over the crossing conditions faced by another actor. The fourth is systemic influence: durable and sufficiently broad changes that follow from the interaction of capacities, boundary relations, adaptation, and network structure.

ANP(state, time) is shorthand for the usable national AI capacity of a specified state at a specified time. ABP(controller → target | interface, time) identifies a relation rather than an actor-level stock. The notation is descriptive, not an estimand. Its purpose is to force the analyst to identify the controller, target, interface, and observation window before constructing a score or making a causal claim.

AI National Power is defined as the usable capacity of a territorially and politically constituted national system to generate, sustain, mobilize, govern, and convert AI-relevant resources into economic, administrative, scientific, social, and strategic effects. Firms, universities, laboratories, data centers, and other organizations contribute to ANP only to the extent that their capabilities are available to the national system under the conditions being studied. Presence is therefore not equivalent to mobilizability. A foreign-owned cloud region may strengthen local hosting while remaining dependent on remote administration, externally controlled software, imported accelerators, or contractual restrictions. Conversely, overseas research and corporate networks may contribute to usable national capability even when the underlying assets are not physically domestic.

AI Boundary Power is defined as the effective capacity of an actor to change the conditions under which another actor can obtain, transfer, integrate, deploy, or continue using an AI-relevant resource or function at a specified interface. Change includes granting, denying, conditioning, pricing, throttling, authenticating, certifying, making interoperable, or restructuring access. The definition centers control over another actor's crossing conditions. A user's ability to traverse an interface is not itself boundary power.

Boundary autonomy captures the target-side capacity to reduce exposure to another actor's boundary control. It includes the ability to switch providers, multi-home, bypass, substitute, redesign, localize, or internalize. Boundary autonomy can reduce vulnerability without making the target a gatekeeper. It becomes bargaining power only when credible exit or substitution changes the controller's behavior.

Systemic influence is reserved for persistent and sufficiently broad changes in other actors' feasible option sets, network connections, or governing rules within a pre-specified domain. Persistence requires effects that survive beyond a single transaction or short interruption. Breadth requires effects that extend beyond the initiating bilateral relation to a class of participants, multiple network ties, or rules governing future participation. A durable bilateral contract change may demonstrate consequential ABP without demonstrating systemic influence unless it also alters the conditions facing a wider actor set.

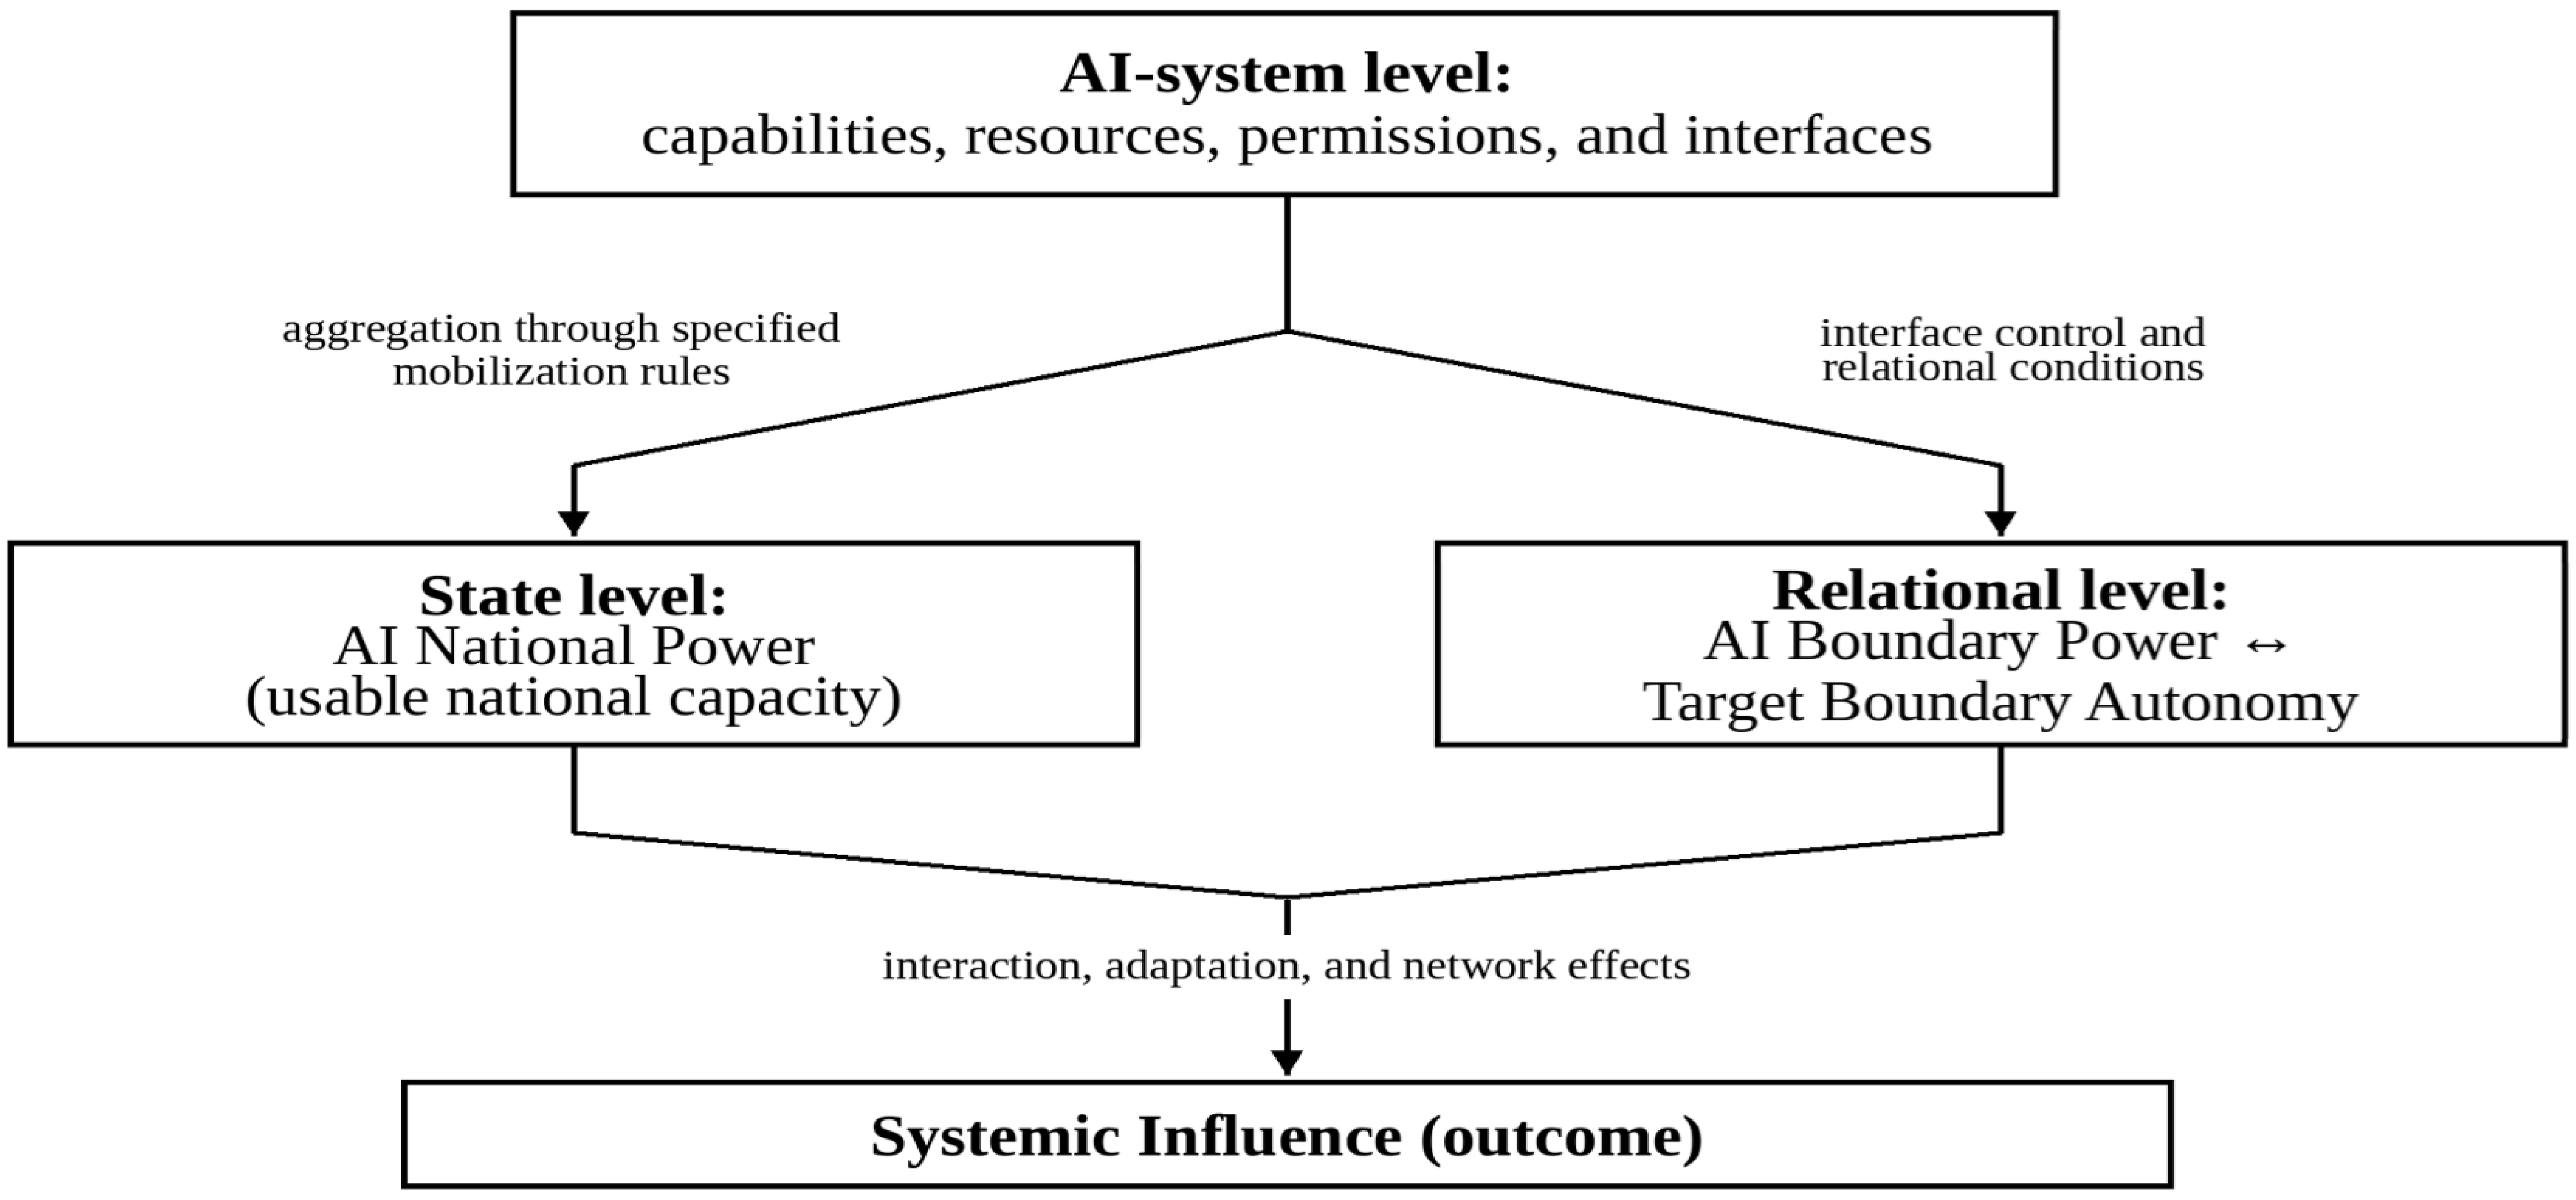

Figure 1. Analytical Architecture of AI Power

Note: ANP is a state-level capacity construct; ABP is a controller-target-interface relation; boundary autonomy is target-side; systemic influence is a later outcome and is not guaranteed by any single configuration.

4. AI National Power: Resources, Conversion, and Conditioning

ANP should not be constructed as an undifferentiated list of everything that might matter to AI. A more defensible architecture separates resource bases, conversion mechanisms, and conditioning variables. The distinction matters because a state can possess large resource stocks without converting them into usable capability, while political, linguistic, or normative conditions can amplify some uses and constrain others.

Territorial-infrastructural capacity includes the physical conditions needed to host and sustain AI: electricity, land, water, data centers, fiber, cloud regions, semiconductor and logistics access, cooling, maintenance, and physical resilience. AI is digital in output but materially embedded in operation. Energy demand and the geography of compute therefore affect where AI capacity can be built, operated, and sustained (International Energy Agency, 2025; Lehdonvirta et al., 2025). Territory matters not as land area in the abstract but as a bundle of infrastructures, jurisdictional competencies, and bottlenecks.

Human-epistemic capacity concerns researchers, engineers, universities, R&D institutions, expert communities, specialized data, and the educational systems that reproduce them. Economic-industrial capacity concerns capital, firms, semiconductor and software participation, manufacturing, startup formation, procurement resources, and demand. These are resource bases, not outcomes. Their strategic value depends on whether the national system can organize them into usable capability.

Political-administrative mobilization is the principal conversion mechanism. It includes coordination capacity, lawful authority, procurement competence, implementation institutions, cyber and physical security, and the ability to mobilize resources under crisis or strategic pressure. The distinction resembles the broader idea of infrastructural state power: the capacity to coordinate and implement across a territory, rather than merely to possess formal authority (Mann, 1984). For AI, the crucial question is whether dispersed public and private resources can be called upon, integrated, and sustained for a defined function.

Cultural-linguistic fit and normative-strategic orientation are better modeled as conditioning variables than as additive resources. Language coverage, domain vocabulary, local knowledge, representation, and social legitimacy can raise usability in one context while creating constraints in another. Research on Indigenous languages and AI illustrates why representational inclusion and data sovereignty cannot be reduced to the size of a language market (Kwok, 2026). Similarly, stronger privacy rules, openness requirements, security exceptions, state control, or market competition can expand some capabilities while limiting others. National AI strategies reveal different combinations of development, control, and promotion rather than a single monotonic scale of governance strength (Papyshev & Yarime, 2023; Radu, 2021).

Operationalization should therefore begin with a specified sector and time horizon. Analysts should inventory relevant resource bases; code location, ownership, jurisdiction, and mobilizability separately; identify conversion institutions; and then measure the usable capability that results. The dependent construct is not the number of assets inside national borders but the level of capability the national system can actually call upon and sustain at a specified quality, cost, and time window.

Table 1. Role-Based Architecture of AI National Power

Role	Component	Illustrative Content	Analytical Question
Resource Base	Territorial-infrastructural capacity	Electricity, land, water, data centers, fiber, cloud regions, semiconductor/logistics access, physical resilience.	Can relevant AI capability be hosted, powered, connected, secured, and scaled?
Resource Base	Human-epistemic capacity	Researchers, engineers, universities, R&D institutions, expert communities, specialized data, education and training.	Can the system generate, reproduce, and retain knowledge and expertise?
Resource Base	Economic-industrial capacity	Capital, firms, manufacturing, semiconductors, software/services, demand, startup formation, procurement resources.	Can AI capability be financed, produced, commercialized, and sustained?
Conversion Mechanism	Political-administrative mobilization	Coordination capacity, lawful authority, procurement competence, cybersecurity, implementation institutions, crisis mobilization.	Can dispersed resources be converted into usable national capability?
Conditioning Variable	Cultural-linguistic fit	Language coverage, local knowledge, institutional vocabularies, representational inclusion, social legitimacy.	How well does available AI capability fit the society and institutions in which it is used?
Conditioning Variable	Normative-strategic orientation	Privacy-security balance, openness, state-market relations, risk tolerance, autonomy goals, legitimacy expectations.	Which uses are authorized, prioritized, resisted, or redirected?

Note: Components are separated by analytical role to avoid treating every national attribute as an additive resource

5. AI Boundary Power: Interface Control, Criticality, and Autonomy

A boundary is any consequential threshold that distinguishes permitted from non-permitted movement between functional domains. In AI systems it may be physical, legal, technical, contractual, institutional, or organizational. Examples include an export license for an accelerator, an API key for a model, a cloud contract, an app-store approval, a model license, a data-localization rule, a procurement certification, a security clearance, or an interoperability standard.

ABP is inherently relational. The same controller can possess substantial power over one target and little over another because the targets differ in dependence and substitutes. A chip restriction may be severe for an actor without comparable accelerators and modest for another with domestic alternatives. A platform's conditions may bind a developer locked into its APIs but matter less to a firm capable of migrating workloads. The relevant observation is therefore the controller-target-interface-time relation rather than a generic score for the controller.

Three substantive axes organize boundary observations. The object or interface identifies what crosses the boundary: data, compute, models, platforms and infrastructure, knowledge and talent, or capital and finance. The control mechanism identifies how crossing conditions change: licensing, pricing, authentication, filtering or throttling, interoperability, certification and standards, procurement and market access, or denial and termination. The application domain identifies the consequence environment: commercial and consumer use, public administration, science and education, critical infrastructure, or defense and intelligence. Controller, target, jurisdiction, duration, and direction are recorded as relational metadata rather than treated as additional dimensions.

Formal authority alone does not establish consequential boundary power. Four conditions are especially important. First, the controller must possess actual discretion over the relevant access or terms. Second, the interface must be critical to the target's intended function. Third, substitutes must be sufficiently costly, slow, inferior, or unavailable for the condition to matter. Fourth, the controller must be able to enforce the condition at an acceptable political, legal, technical, or commercial cost. ABP therefore rises as effective control, criticality, and enforceability increase and as target substitutability decreases.

This formulation sharpens the relationship between boundary power and conditional connectivity. A node can be topologically central yet exercise little leverage if it lacks authority over the relevant crossing conditions or if targets can substitute quickly. Conversely, a less central actor can exercise strong relational power when it controls a narrow but indispensable interface with long substitution times. The distinction between centrality and control is therefore not semantic; it produces different expectations about observable effects (Wu, 2026a).

Counterpower operates through substitution, multi-homing, localization, redesign, alternative standards, political retaliation, or construction of new networks. These responses are consistent with network theories that emphasize switching and reprogramming and with work showing that asymmetric-network leverage depends on jurisdictional and institutional conditions (Castells, 2011; Farrell & Newman, 2019). ABP should consequently be treated as dynamic. A powerful boundary can erode, migrate to an adjacent layer, or dissipate as technical and political conditions change.

Table 2. Analytical Axes and Relational Metadata for AI Boundaries

Axis	Categories	Examples	Purpose
Object or Interface	Data; compute; models; platforms/infrastructure; knowledge/talent; capital/finance	Training data, accelerators, model weights/APIs, cloud/app ecosystems, tacit expertise, financing/payment rails.	Identifies what is crossing or being connected.
Control Mechanism	Licensing; pricing; authentication; filtering/throttling; interoperability; certification/standards; procurement/market access; denial/termination	Export licenses, API quotas, identity requirements, compatibility rules, audit certification, procurement eligibility.	Identifies how crossing conditions are altered.
Application Domain	Commercial/consumer; public administration; science/education; critical infrastructure; defense/intelligence	Enterprise AI, government services, research systems, grids, command-and-control.	Identifies the consequence environment and authority level.
Relational Metadata	Controller; target; interface; jurisdiction; duration; direction	State-to-firm chip licensing; platform-to-developer API access; regulator-to-provider market-entry condition.	Records the scope needed to interpret an ABP observation.

Note: Relational metadata define the scope of an ABP observation and should not be treated as additive dimensions.

6. State Configurations and Mechanism-Based Propositions

For states only, ANP and state-level ABP can be displayed as a two-dimensional analytical space. The matrix is not a global ranking. It distinguishes dependent or peripheral configurations (low ANP, low state ABP), domestically deep configurations (high ANP, low state ABP), boundary-positioned configurations (low ANP, high state ABP), and comprehensive configurations (high ANP, high state ABP). The European Union, firms, platforms, and standards bodies are not assigned ANP because ANP is reserved for states; their relational power can still be analyzed through ABP.

A dependent or peripheral state has limited usable national capacity and limited ability to condition the external interfaces on which it depends. Its strategic problem is not simply catching up in model performance but reducing asymmetric dependence across compute, cloud, data, procurement, skills, and standards while retaining useful external access. A domestically deep state possesses substantial research, infrastructure, industrial, or administrative capacity but limited external interface control. It may be resilient and technically capable without being able to reshape the surrounding network. A boundary-positioned state controls one or more consequential interfaces, such as energy-backed compute sites, specialized manufacturing inputs, regulatory markets, logistics routes, or trusted hosting, without possessing broad national AI depth. Such leverage can be substantial but narrow and vulnerable to substitution. Comprehensive configurations combine broad national capacity with control over several consequential external interfaces, creating favorable conditions for wider influence without guaranteeing it.

The framework generates five propositions.

- **P1—Substitutability conditions restrictive boundary effects:** for comparably exposed targets facing restrictions of similar severity and functional scope, higher pre-intervention switching costs and longer substitution times should be associated with larger subsequent losses of deployment capacity. Comparable losses across targets with sharply different substitution conditions would weaken the mechanism.
- **P2—Organizational conversion conditions national capacity:** holding technical resource stocks approximately constant, states with stronger lawful coordination, procurement, implementation, and mobilization capacity should convert a larger share of those resources into usable sectoral AI capability. If resource stocks predict usable capability equally well regardless of conversion institutions, the mechanism is weakened.
- **P3—National depth and state boundary power are complementary but not jointly sufficient:** among states compared over the same sector, target set, and time window, the marginal association between ANP and the breadth and durability of external influence should increase as effective state ABP rises. High ANP and high ABP at an initial time do not guarantee systemic influence because later substitution, coalition change, enforcement costs, or network redesign may erode the relation.
- **P4—Boundary autonomy reduces vulnerability:** among targets facing restrictions of similar severity, greater pre-intervention switching, bypass, substitution, or internalization capacity should be associated with smaller subsequent losses of access or deployment capacity. Whether this autonomy also generates bargaining power is a separate empirical question requiring observed changes in controller behavior.
- **P5—Boundary erosion can produce migration or dissipation of power:** when substitution reduces control at one AI layer, leverage should move to an adjacent interface only if that interface remains scarce and consequential. Otherwise, overall boundary power should decline rather than automatically migrate. Evidence of systematic migration regardless of adjacent scarcity would weaken the proposition.

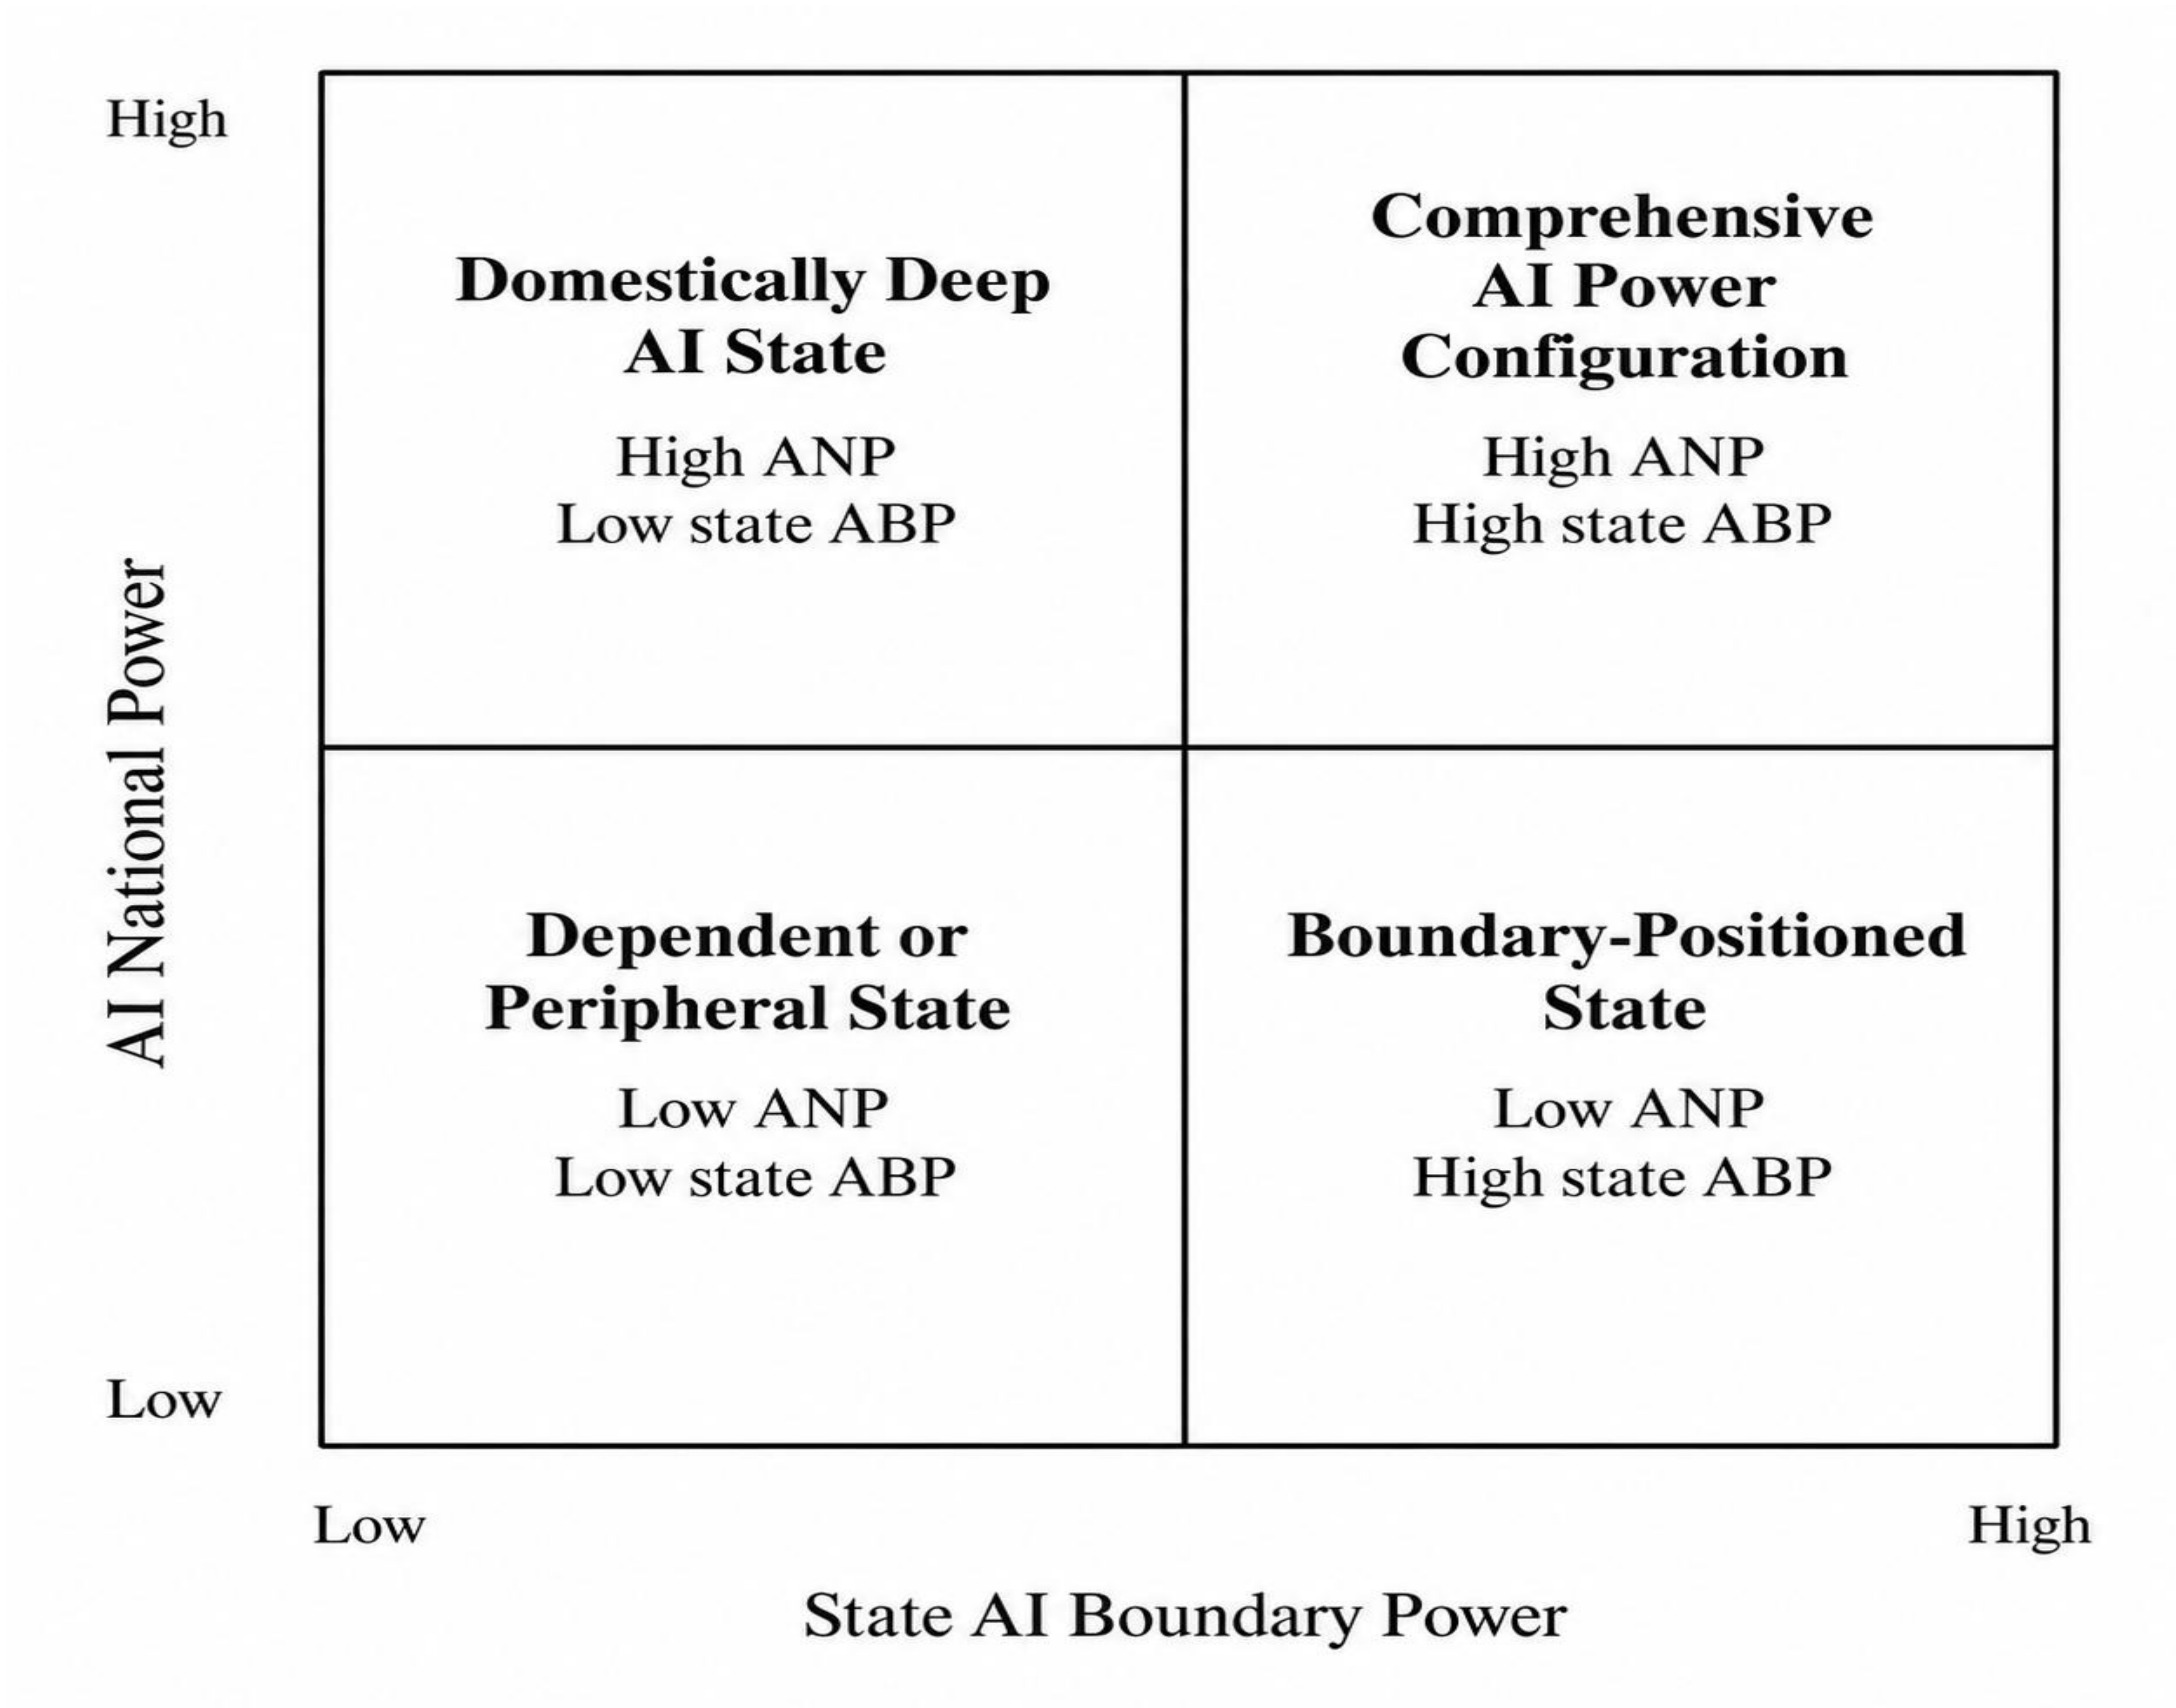

Figure 2. State Configurations in the National-Boundary Power Space

Note: The matrix applies only to states and is an analytical map, not a ranking. High ANP and high state ABP facilitate but do not define systemic influence.

7. Operationalization and Measurement Strategy

A defensible measurement strategy begins by assigning each variable one role. Resource inputs should not be reused as outcomes; conversion mechanisms should not be counted as both causes and results; and systemic influence should be observed separately from the evidence used to establish ANP or ABP. This prevents circularity, especially in studies that might otherwise infer power from an intervention and then use the same intervention as proof of its outcome.

Existing datasets can supply partial ANP indicators. The AI Index contains measures of research, model development, investment, infrastructure, and adoption; OECD work provides a method for measuring domestic public-cloud compute availability; and UN Trade and Development identifies infrastructure, data, and skills as important leverage points for inclusive AI development (Lehdonvirta et al., 2025; Stanford Institute for Human-Centered Artificial Intelligence, 2026; UN Trade and Development, 2025). These sources should be organized by sector and time rather than summed mechanically into a universal score.

A practical ANP study can proceed in four stages: specify the sector or national function; inventory relevant resource bases; estimate mobilizability and organizational conversion; and measure the usable capability that results under a stated cost, quality, and time window. If the sector is public administration, for example, procurement competence and institutional data access are conversion mechanisms, while the ability to sustain specified AI-enabled administrative functions is part of usable ANP. Later foreign spillovers are not.

ABP measurement begins with a boundary map rather than a country score. Researchers should identify the controller, target, interface, direction, and time window; document legal or technical discretion; estimate the share and functional criticality of affected flows; measure switching costs and alternative routes; and observe whether access conditions can actually be changed. Market concentration and network betweenness can predict potential leverage but do not substitute for evidence of control.

Event studies can exploit changes in export controls, platform access, licensing rules, cloud contracts, procurement standards, or certification requirements. Candidate outcomes include prices, delays, rerouting, redesign, supplier switching, localization, or loss of deployment capability. Matched comparisons can examine similarly exposed targets with different substitute availability, directly testing P1 and P4.

Systemic influence should be coded only when both persistence and breadth are demonstrated within a pre-specified domain. Candidate indicators include durable standards adoption, repeated product-architecture changes across firms, creation of alternative supply chains, persistent network rewiring, or institutional adoption of new participation rules. The evidence used to establish an ABP relation should be temporally or analytically distinct from the later evidence used to establish systemic influence. A strong boundary intervention that produces only a short, bilateral disruption is not yet a system-level outcome.

Table 3. Measurement Roles and Illustrative Indicators

Role	Examples	Possible Indicators	Do Not Code As
National Resource Inputs	Infrastructure; expertise; industrial capacity	AI-capable compute availability; grid reliability; R&D labor; research institutions; semiconductor/software participation; investment capacity.	Systemic influence or regulatory spillover.
National Conversion Mechanisms	Coordination; procurement; lawful authority; implementation capacity	Procurement competence; institutional coordination; continuity planning; administrative capacity; cyber/physical resilience.	Public adoption itself when adoption is the outcome being explained.
Conditioning Variables	Cultural-linguistic fit; normative-strategic orientation	Language coverage; domain datasets; representational fit; policy orientation; legitimacy measures.	Monotonic “more is better” scores without a mechanism.
Construct: Usable ANP	Mobilized and sustainable national AI capability in a specified sector or function	Deployable capacity or task throughput; resources lawfully and organizationally callable; sustained availability; redundancy.	Raw resource inventories or later external effects.
ABP Relation	Effective control; criticality; enforceability; low substitutability	Documented access-control discretion; affected-flow criticality; switching cost/time; alternatives; enforcement evidence.	Ownership, market share, centrality, or later systemic outcomes by themselves.
Boundary Autonomy	Target switching, bypass, substitution, internalization	Supplier diversity; portability; interoperability; local fallback; time to substitute.	Power over third parties unless bargaining effects are observed.
Outcome: Systemic Influence	Durable and broad changes in options, networks, or rules	Persistent rerouting; standards adoption; network rewiring; repeated architecture changes; durable compliance spillovers.	Evidence already used to establish ANP or ABP in the same observation window.

Note: The measurement architecture separates inputs, mechanisms, constructs, target autonomy, and later outcomes to reduce circularity.

8. Mechanism Illustrations

The following illustrations demonstrate how the constructs can be separated empirically. They do not rank actors or establish the general validity of the propositions.

8.1 US Advanced-Computing Export Controls

US controls on advanced computing semiconductors provide a clear example of a compute boundary. In October 2023, the Bureau of Industry and Security updated restrictions on advanced-computing semiconductors, semiconductor-manufacturing equipment, and related supercomputing end uses (Bureau of Industry and Security, 2023). The mechanism is not simply US technological leadership or the centrality of American firms. It is the combination of legal jurisdiction, licensing, product coverage, and enforcement that changes the conditions under which affected actors can acquire specified computing resources.

In ABP terms, the controller is the US government; the immediate targets include regulated exporters and affected end users; the object is advanced compute; and the mechanism is licensing and denial. The strength of the relation depends on the criticality of the covered products, the scope of jurisdiction, enforceability, and the speed and quality of target-side substitutes. This is why national capability and boundary power must be separated: the restriction acts through a relation, not merely through the stock of US AI assets.

8.2 Adaptation and Rule Expansion

The same episode shows why formal restrictions do not produce a fixed amount of power. NVIDIA’s fiscal-year 2025 Form 10-K reported that changing US restrictions affected a broad range of data-center products and prompted the company to develop products for the Chinese market that remained within the thresholds then in force (NVIDIA Corporation, 2025a). This demonstrates supplier adaptation to a regulatory boundary; it does not establish equivalent substitution for Chinese end users or show that US leverage had materially declined. The boundary subsequently tightened. In April 2025, US authorities extended licensing requirements to NVIDIA’s H20 and other integrated circuits meeting specified technical conditions for China and designated destinations. NVIDIA later reported that the licensing requirement would remain in force indefinitely (NVIDIA Corporation, 2025b).

The sequence illustrates P5. Product redesign may reduce the bite of one threshold, after which the controller may expand or relocate the boundary. Whether power actually migrates depends on the scarcity of the adjacent interface, target-side substitution, technical equivalence, and enforcement. It may instead remain with the original controller or dissipate. The downstream user effect should also be separated from the legal relation: the government directly conditions exporters, while consequences for end users are mediated through supply availability and substitutes.

8.3 European Union AI Regulation

The European Union (EU) illustrates a different form of boundary power. Because the EU is not a state, the framework does not assign it ANP. It can nevertheless exercise jurisdictional ABP by conditioning access to its market on compliance with Union rules. Article 53 of Regulation (EU) 2024/1689 links market participation to a set of obligations for providers of general-purpose AI models, including technical documentation, transparency toward downstream providers, compliance with Union copyright requirements, and disclosure concerning training content. Qualifying free and open-source models receive differentiated treatment under specified conditions (European Parliament & Council of the European Union, 2024).

The relevant obligations entered into application on August 2, 2025, while the Commission’s enforcement powers for general-purpose AI obligations entered into application on August 2, 2026; providers of models placed on the market before August 2, 2025 have a later compliance deadline (European Commission, 2026; European Parliament & Council of the European Union, 2024). The mechanism resembles the Brussels Effect: access to a valuable market can induce firms to restructure compliance practices beyond the regulator’s production base (Bradford, 2020). Yet the existence of the rule is not proof of systemic influence. Globalized compliance architectures would widen the effect; product segmentation, market exit, or low-cost substitution could narrow it. The illustration therefore demonstrates jurisdictional ABP while leaving system-level spillover as an empirical outcome to be measured later.

9. Implications for International Order and Middle-Power Strategy

The framework reframes several debates about the emerging AI order. First, national AI strength is not synonymous with technological self-sufficiency. A state can possess extensive domestic capability while remaining exposed to externally controlled interfaces, and it can preserve substantial agency without reproducing every layer of the AI stack. Managed interdependence, supplier diversification, interoperable standards, trusted alliances, and enforceable fallback arrangements can therefore increase autonomy even when full-stack sovereignty is infeasible (Tanner et al., 2026).

Second, AI returns attention to territory without restoring a purely territorial conception of power. Data-center clusters, energy systems, semiconductor manufacturing, fiber and undersea cables, ports, industrial zones, and jurisdictional boundaries all condition where AI can be developed and operated. Yet location is not control. The relevant question is whether geographically situated assets are usable, governable, resilient, and mobilizable under the scenario being studied.

Third, middle powers need not imitate frontier states across the entire stack. A portfolio strategy can distinguish capabilities that require domestic control from dependencies that can be diversified or governed contractually. Energy-backed compute, specialized semiconductor functions, trusted hosting, language technologies, public digital infrastructure, financial channels, standards expertise, and logistics can each create selective national capacity or boundary leverage. The risk is to mistake a narrow gatekeeping position for comprehensive power or to ignore how quickly substitutes may erode it.

Fourth, lower-capacity states make consequential choices through boundary architecture even when they do not own frontier models. Cloud procurement, digital identity, model APIs, interoperability requirements, data governance, public-sector contracts, and technical standards determine who can enter administrative systems, who controls updates and data flows, and how costly provider substitution will be. Infrastructure, data, and skills remain central to inclusive AI development, but the governance of interfaces determines whether access also preserves future autonomy (UN Trade and Development, 2025).

10. Limitations and Research Agenda

The framework has five principal limitations. First, ANP remains broad. Empirical studies must specify the sector, scenario, and time horizon because the capabilities relevant to public administration differ from those relevant to frontier-model training, critical infrastructure, or defense. Second, ABP is intentionally relational and data intensive. A single actor-level score can conceal major differences across targets and interfaces, so aggregation should follow rather than precede target-specific coding.

Third, control is often shared. States, firms, standards bodies, cloud providers, and infrastructure operators can possess overlapping legal, technical, contractual, or operational authority. Joint control and conflict among controllers should be modeled explicitly. Fourth, cultural-linguistic fit and normative orientation are not universal positive resources; their effects can be nonlinear, domain specific, and politically contested. Fifth, the mechanism illustrations are not causal tests. They show what should be observed when distinguishing authority, effective control, target substitution, and durable outcomes.

Future work should therefore combine network analysis, matched comparisons, event studies, and longitudinal process tracing. The most useful tests will examine cases in which resource stocks or network positions are similar but effective boundary control or substitution conditions differ. Such designs can evaluate whether the two-structure framework explains variation that would be missed by national-capacity rankings or centrality measures alone.

Conclusion

AI power is not one thing. The depth of a state’s usable AI capacity and an actor’s relational control over consequential interfaces are distinct analytical dimensions. AI National Power concerns what a national system can generate, mobilize, govern, and sustain. AI Boundary Power concerns whether an actor can effectively alter the conditions under which another actor obtains, transfers, integrates, deploys, or continues using AI-relevant resources. Boundary autonomy concerns the target’s ability to avoid or absorb such control. Systemic influence concerns the durable and sufficiently broad changes that follow within a pre-specified domain.

This separation helps organize several puzzles without assuming their answers. Technologically capable states can remain externally constrained. Smaller states and specialized non-state actors can exercise consequential leverage through narrow but critical interfaces. Regulators can reshape market-entry conditions without possessing national AI capacity. Territorial infrastructure remains important, but location must be distinguished from ownership, jurisdiction, mobilizability, and control. Culture and language matter through fit, representation, and diffusion, while governance orientation conditions which capabilities are politically usable.

The central proposition is conditional rather than categorical: national depth and boundary control are distinct but potentially complementary sources of AI power. Their translation into durable system-shaping influence depends on organizational conversion, effective control, interface criticality, target substitutability, enforceability, adaptation, and network structure. The emerging AI order will therefore be shaped not only by who develops advanced intelligence, but also by who can mobilize it, who controls the interfaces through which it travels, and who can build credible alternatives when those interfaces close.

Working Paper

References

Bradford, A. (2020). *The Brussels effect: How the European Union rules the world*. Oxford University Press. https://doi.org/10.1093/oso/9780190088583.001.0001

Bureau of Industry and Security. (2023, October 17). *Commerce strengthens restrictions on advanced computing semiconductors, semiconductor manufacturing equipment, and supercomputing items to countries of concern*. U.S. Department of Commerce. https://www.bis.gov/press-release/commerce-strengthens-restrictions-advanced-computing-semiconductors-semiconductor-manufacturing-equipment

Castells, M. (2011). A network theory of power. *International Journal of Communication, 5*, 773–787. https://ijoc.org/index.php/ijoc/article/view/1136

Edler, J., Blind, K., Kroll, H., & Schubert, T. (2023). Technology sovereignty as an emerging frame for innovation policy: Defining rationales, ends and means. *Research Policy, 52*(6), 104765. https://doi.org/10.1016/j.respol.2023.104765

European Commission. (2026). *Guidelines for providers of general-purpose AI models*. Shaping Europe's Digital Future. https://digital-strategy.ec.europa.eu/en/policies/guidelines-gpai-providers

European Parliament & Council of the European Union. (2024). *Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Farrell, H., & Newman, A. L. (2019). Weaponized interdependence: How global economic networks shape state coercion. *International Security, 44*(1), 42–79. https://doi.org/10.1162/isec_a_00351

International Energy Agency. (2025). *Energy and AI*. https://www.iea.org/reports/energy-and-ai

Kim, R. M., Kuehnert, B., Lazar, S., Singh, R., & Heidari, H. (2025). The AI Power Disparity Index: Toward a compound measure of AI actors' power to shape the AI ecosystem. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, 8*(2), 1453–1464. https://doi.org/10.1609/aies.v8i2.36645

Kwok, J. C. Y. (2026). Data colonialism and Indigenous languages in AI: A critical review of existing initiatives and their struggles with data sovereignty. *AI & Society, 41*, 6719–6729. https://doi.org/10.1007/s00146-026-03091-w

Lehdonvirta, V., Wu, B., Hawkins, Z. J., Caira, C., & Russo, L. (2025). *Measuring domestic public cloud compute availability for artificial intelligence* (OECD Artificial Intelligence Papers No. 49). OECD Publishing. https://doi.org/10.1787/8602a322-en

Mann, M. (1984). The autonomous power of the state: Its origins, mechanisms and results. *European Journal of Sociology, 25*(2), 185–213. https://doi.org/10.1017/S0003975600004239

NVIDIA Corporation. (2025a). *Annual report (Form 10-K), fiscal year ended January 26, 2025*. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1045810/000104581025000023/nvda-20250126.htm

NVIDIA Corporation. (2025b). *Current report (Form 8-K), event reported April 9, 2025*. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1045810/000104581025000082/nvda-20250409.htm

Papyshev, G., & Yarime, M. (2023). The state's role in governing artificial intelligence: Development, control, and promotion through national strategies. *Policy Design and Practice, 6*(1), 79–102. https://doi.org/10.1080/25741292.2022.2162252

Radu, R. (2021). Steering the governance of artificial intelligence: National strategies in perspective. *Policy and Society, 40*(2), 178–193. https://doi.org/10.1080/14494035.2021.1929728

Stanford Institute for Human-Centered Artificial Intelligence. (2026). *The AI Index 2026 annual report*. Stanford University. https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf

Strange, S. (1988). *States and markets*. Pinter.

Tanner, B., Kerry, C. F., Wyckoff, A. W., Kyosovska, N., Renda, A., & Tabassi, E. (2026). *Is AI sovereignty possible? Balancing autonomy and interdependence*. Brookings Institution & Centre for European Policy Studies. https://www.brookings.edu/articles/is-ai-sovereignty-possible-balancing-autonomy-and-interdependence/

UN Trade and Development. (2025). *Technology and innovation report 2025: Inclusive artificial intelligence for development*. United Nations. https://unctad.org/publication/technology-and-innovation-report-2025

Woods, A. K. (2023). Digital sovereignty + artificial intelligence. In A. Chander & H. Sun (Eds.), *Data sovereignty: From the digital silk road to the return of the state* (pp. 115–136). Oxford University Press. https://doi.org/10.1093/oso/9780197582794.003.0006

Wu, S. (2026a). *Beyond centrality: Conditional connectivity and boundary control in interdependent networks* (EPINOVA Working Paper No. EPINOVA–WP–T–2026–01). Global AI Governance and Policy Research Center, EPINOVA LLC. https://doi.org/10.67037/epinova.wp.t.2026.01

Wu, S. (2026b). *Beyond model capability: A system-level framework for AI power* (Policy Brief No. EPINOVA–2026–PB–58). Global AI Governance and Policy Research Center, EPINOVA LLC. https://doi.org/10.67037/epinova.pb.2026.058