11

Operationalizing AI Sovereignty Through Agency, Interoperability, and Openness

Luca Belli

Artificial intelligence sovereignty has rapidly transitioned from a niche academic concept to a cornerstone of contemporary digital policy. Yet despite the ubiquity of digital sovereignty and AI sovereignty in current technological-policy debates, both goals remain conceptually fluid and operationally elusive. Governments frequently invoke the protection of sovereignty to legitimize public investments, justify regulatory guardrails, or navigate geopolitical rivalries, yet they often fail to articulate a concrete road map for its realization.

For the past four years, the CyberBRICS project of the Center for Technology and Society at FGV Law School, Rio de Janeiro, has applied the Key AI Sovereignty Enablers (KASE) framework to analyze how technological dependencies intersect with AI governance.[246] Building on this foundation, this article explores a practical path to operationalize the concept of AI sovereignty. The KASE framework focuses on the governance, regulatory, and industrial policy dimensions of eight fundamental layers composing an AI technology stack: data, software and algorithmic models, computational capacity, meaningful connectivity, energy, education and research, cybersecurity, and AI-safety frameworks.

Exercising the framework effectively will require reducing persistent tensions between access and control, as well as innovation and dependency, while emphasizing the role of cooperative institutional and technological arrangements in addressing structural constraints and enabling collective capacity building.[247] More broadly, it will require alleviating tension between aspirations for technological autonomy and the realities of global integration. Some scholars argue that total AI sovereignty is an illusion, either because it overlooks the physical infrastructure underpinning AI or confuses control over artifacts with genuine capability.[248] From this perspective, building resilience—the ability to withstand external shocks—is a more realistic objective. At the same time, private sector actors frame sovereign AI around infrastructure control and data localization, reflecting a broader tension between political desire, technical constraints, and commercial incentives.

In this context, AI sovereignty should be operationalized as a strategic capacity built through three interdependent pillars: agency, interoperability, and openness.[249] Grounded in these pillars, the KASE framework offers a blueprint for states to navigate global dependencies while retaining the power to shape their own technological futures. In short, this more productive approach could help countries move away from sovereignty as an abstract aspiration toward sovereignty as an operational capacity. This would offer a way to navigate dependence without denying it exists and to exercise influence without requiring full control. The analysis and recommended steps that follow apply the pillars across the key enabling conditions of AI systems.

For the purpose of the analysis, AI sovereignty is defined as the capacity of a given country to understand, develop, and regulate digital or AI systems, thus retaining self-determination, agency, and control over such systems.[250] This definition emphasizes that sovereignty discussions involve not only formal regulatory authority, but also—and, perhaps, chiefly—the technological, institutional, and scientific capabilities necessary to shape the development and governance of AI systems.[251]

Rethinking Sovereignty: From Control to Strategic Agency, Interoperability, and Openness

The traditional Westphalian view of sovereignty is rooted in the control of physical territory, resources, and independent decisionmaking; and it constitutes a pivotal principle of modern international law.[252] Transposed to the digital sphere, this logic manifests as the pursuit of full-stack autonomy, implying the domestic production of all the abovementioned Key AI Sovereignty Enablers. However, such ambitions clash with the realities of the AI value chain, which is defined by extreme specialization, immense capital requirements, and powerful network effects.[253] In this environment, self-sufficiency is a logistical impossibility for most countries.

A more useful framing is to reevaluate sovereignty through the lens of three pillars. The first is *agency*, the substantive capacity to exercise autonomous choice within a state of interdependence. Agency involves the capacity to autonomously assess what technological choices one should pursue to achieve one's right to self-determination. Concretely, it means being able to understand the function of AI systems, access critical technologies on fair terms, adapt systems to local linguistic or cultural needs, and exert influence over global governance standards in accordance with one's own values. Under these conditions, sovereignty is better understood as agency under constraint. The relevant question is not whether a state controls every layer of the stack, but whether it retains the capacity to make consequential choices: to switch providers, impose conditions, adapt technologies, and freely engage in national or international rule-making processes.

Agency, in this sense, is negotiated, contingent, and often collective. This perspective recognizes that cooperation across countries does not diminish sovereignty.[254] Rather, strategic partnerships often provide the leverage required to maintain autonomy, allowing countries to diversify and manage their dependencies on favorable terms.

This means, however, that the capacity to exercise agency as a strategic objective relies on a second pillar, *interoperability*, serving as its primary technical enabler. In a fragmented digital landscape, the lack of interoperability acts as a mechanism of "vendor lock" confining applications into proprietary-walled gardens that preclude the ability to switch providers and stifle localized innovations. Conversely, when digital architecture is interoperable, it lowers switching costs and catalyzes competitive dynamics.[255] Interoperability spans data, software, and infrastructure, requiring common standards, compatible systems, and open architecture to permit the secure exchange of datasets, prevent technical silos, and ensure that sovereign cloud solutions are not tethered to a single vendor's road map.

However, for interoperability to be successful, one has to have alternative options to interoperate with. Hence, pursuing interoperability is as much a political and industrial choice as a technical one. Policies that mandate open standards and data portability are essential for preventing the unilateral imposition of rules by dominant external actors. Industrial policy facilitating the development of alternatives is essential to make sure that interoperability is more than a normative obligation.[256] Ensuring that standards are developed through inclusive, multistakeholder processes is therefore central to the broader pursuit of digital autonomy.

Thus, *openness* is the third pillar of AI sovereignty; it lowers barrier to entry and enables distributed experimentation and rapid scaling. Yet openness is not a neutral benefit. It can sometimes entrench structural concentration.[257] Open components are frequently integrated into proprietary stacks controlled by a few dominant firms, allowing them to externalize the costs of development while capturing the lion's share of the value.[258] Furthermore, a state that consumes open technologies without possessing the domestic human capital to modify or audit them risks remaining in a state of passive dependency.

The goal, therefore, is strategic openness.[259] This means investing in open infrastructure, such as digital public infrastructures, and ensuring that regulatory frameworks do not become a pretext for unregulated data extraction or unfair competitive practices. Openness must be treated as a strategic tool embedded within a larger industrial and governance strategy, rather than an end in itself.

Applying the Three Pillars to Avoid "Sovereignty-Washing"

Applying the pillars of agency, interoperability, and openness across the KASE framework's eight layers spanning the AI technology stack can bring meaningful change, but operationalizing the pillars requires a coordinated approach:[260]

1. **Data:** Agency depends on the ability to curate high-quality local datasets. Interoperability relies on shared standards. Openness is realized through the establishment of data stewardship institutions that manage data as a valuable resource to be handled in the public interest.

2. **Software and models:** Agency is the capacity to develop, fine tune, and deploy models. Interoperability ensures that these models function across platforms. Openness is secured by making reusable, transparent code available.

3. **Hardware and infrastructure:** Given the high barriers to entry in hardware, agency involves securing strategic positions in the global supply chain rather than full production. Interoperability ensures cross-system compatibility. Openness is reflected in the adoption of standardized hardware architecture.

4. **Meaningful connectivity:** Reliable, affordable, and universally accessible connectivity infrastructure is critical for developing, operating, and accessing AI systems. Hence, meaningful connectivity requires open and nondiscriminatory internet access, avoiding practices like "zero-rating plans" where private firms subsidize access to just a few apps or websites, limiting users' agency and capacity to freely access AI services.[261]

5. **Energy and mineral resources:** Agency is contingent on countries accessing critical resources so they can physically build sustainable digital infrastructure. Open energy markets and transparent supply chains prevent resource gatekeeping, ensuring interoperable networks can scale sustainably without creating new external dependencies.

6. **Capable and resilient humanware:**[262] The level of digital literacy and cognitive resilience in a society determines its capacity to exercise agency, understand interoperability, and work to promote openness. Only technological systems that retain a human element can actively shape technological evolution rather than passively react to it.

7. **Cybersecurity:** Preserves agency by safeguarding supply chain integrity and protecting domestic infrastructure from foreign interference or nonstate threats. Without robust cyber defenses, open and interoperable architectures risk becoming vectors for external manipulation and systemic instability.

8. **AI safety frameworks:** Preserve sovereign agency by ensuring model robustness, predictability, and alignment with domestic governance standards. Establishing rigorous safety standards enables states to confidently adopt open systems and foster integration without exposing critical infrastructure to operational risks or structural exploitation

Applying the three pillars to the KASE framework helps structure both the conceptual approach and operational strategy for AI sovereignty. It also counters the growing, imprecise use of the term "AI sovereignty," a term increasingly invoked to justify measures that have little to do with autonomy. Building off the pillars, the framework can be used to assess the proliferation of initiatives branded as "sovereign" that fail to build real agency, such as localized cloud infrastructures dependent on foreign providers or "national" AI systems built on externally controlled models.

Countries' claims to sovereignty should be evaluated against concrete criteria: Does the specific AI-sovereignty initiative increase its ability to make independent choices? Does it reduce lock-in? Does it preserve openness without creating new vulnerabilities? Any initiative that restricts a nation's ability to choose ultimately locks the country into a single provider or creates opaque dependencies that should be viewed with skepticism, regardless of its label.

Cooperative Digital Sovereignty: Institutional Pathways Toward an AI Network of Networks

Individual states, particularly those outside major tech hubs, often lack the resources and scale to build AI ecosystems independently. This is not an insurmountable obstacle, but it necessitates adopting a proactive approach through "cooperative digital sovereignty" that is grounded in international collaboration on research and development, cross-border data flows, and global supply chains.[263] Crucially, a decentralized, autonomous, yet cooperative model of digital sovereignty, rooted in the internet's "network of networks" ethos, would leverage interoperability and openness to expand agency, while preserving a wide array of technological options.[264]

Cooperative arrangements can offer more viable technological alternatives. Shared infrastructures, aligned regulatory frameworks, and joint capacity-building initiatives can reduce duplication while increasing leverage. By pooling resources, aligning normative regimes, and building open digital infrastructures, states can achieve economies of scale that would be impossible in isolation. To facilitate this, intergovernmental organizations like the United Nations (UN) and the Digital Cooperation Organization (DCO), alongside development banks such as the Inter-American Development Bank (IDB), the New Development Bank (NBD), the Asian Development Bank (ADB), and the African Development Bank (AfDB), have a crucial role to play, providing platforms for standard setting and resource coordination.[265] Such a cooperative model would allow member states to preserve their strategic autonomy while counterbalancing the influence of dominant tech-states or corporations.

This scenario requires a high level of trust and complex diplomatic and resource alignment; the difficulty lies not in identifying areas for cooperation but in institutionalizing cooperation that will remain durable under conditions of unequal capabilities and geopolitical competition.

Divergent national interests, asymmetries in technological capacity, and incentives toward concentration make collective action difficult to sustain.[266] Yet, the alternative scenario in which digital dependence deepens and access to compute, models, infrastructure, and standards becomes concentrated in a small number of foreign firms and jurisdictions will impose greater long-term costs on economic resilience and political agency.[267]

The crux of the challenge is clarifying what sovereignty means when applied to the AI technology stack. The objective cannot be complete technological self-sufficiency. Progressive, frontier AI, for instance, increasingly depends on high capital expenditures, access to advanced semiconductors, hyperscale computing infrastructure, large proprietary datasets, and network effects that reinforce incumbent advantage. Under current conditions, few states can do this, and regulatory interventions alone cannot eliminate these structural asymmetries.

Instead, sovereignty should be understood as the capacity to organize and govern strategic interdependence. States do not have to reproduce the technological model of dominant AI powers; they can pursue agency by coordinating resources across selected parts of the innovation system while maintaining diversified external linkages elsewhere.

This implies three complementary strategies. First, states should leverage industrial policy, directing public investment to support national research laboratories, domain-specific model development, public-interest datasets, sovereign cloud layers, and specialized compute infrastructure for research and strategic sectors. The objective is to foster AI development by leveraging national assets to meet real national needs.[268]

Second, states should design regulatory frameworks that reduce structural dependence by enabling interoperable and open technological ecosystems. Regulatory regimes matter not because they neutralize incumbent advantages but because they shape market architecture: By mandating interoperability, portability, open standards, transparent procurement, and fair access to critical infrastructure, they can lower switching costs and create space for domestic firms and public institutions to participate.

Lastly, states should support multistakeholder and multilateral governance that will distribute agenda-setting authority across governments, firms, technical communities, and public institutions rather than allowing standards and infrastructures to become privately governed. Multilateral cooperation can pool capabilities that no single state can efficiently develop.[269] Shared compute facilities, regional research networks, federated data infrastructures, coordinated procurement, and joint investment vehicles allow states to achieve scale without requiring national self-sufficiency.

This approach recognizes that frontier AI competition is only one dimension of technological power. Countries can still exercise meaningful influence by developing capabilities in application layers, domain-specific models, public-sector AI, specialized hardware niches, standards development, open-source ecosystems, and strategic deployment capacities.

Importantly, the above measures should also help counter weaponized interdependence. When critical AI resources are diversified across providers, governed through interoperable standards, embedded in regional partnerships, and supplemented by domestic adaptation, external actors face greater costs in converting economic dependence into political leverage. Sovereignty therefore emerges from ensuring that dependencies remain plural, negotiable, and institutionally governed.

KASE Operational Steps: From Concept to Implementation

Operationalizing the KASE framework's layers requires a deliberate sequencing of actions. Sovereignty cannot be engineered in the abstract; it must be built through institutional design, regulatory calibration, sustained investment, and solid capacity building. A first step lies in establishing multistakeholder governance mechanisms capable of providing visibility on the complexity of all the layers and coordinating public and private actors that play key roles in each layer. Without such coordination, policy will remain fragmented and reactive at best.

Second, regulations must be recalibrated to prioritize interoperability, transparency, and contestability in any given layer. This includes enforcing data portability, mandating auditability for high-risk AI systems, and ensuring that standards-setting processes remain open and inclusive. Regulation, in this sense, is not constraining; instead, it becomes constitutive of digital ecosystems.

Third, industrial policy must move beyond generalized support toward targeted capability development. This entails identifying strategic entry points within different layers—whether in data infrastructures, model adaptation, public compute capacity or specific segments of the hardware value chain—and concentrating resources where meaningful benefits and durable advantages can be developed.

Such resource concentration should not aim to replicate the full technological ecosystems of leading AI powers. Frontier AI requires scale, capital, and cumulative advantages that few countries possess. Instead, states should seek strategic scale by focusing on capabilities that preserve agency, reduce vulnerability, and create leverage within global value chains. Crucially, this often requires cooperative approaches, such as shared infrastructure, coordinated investment, joint procurement, and common standards, to shape interdependence rather than passively accept it.

Fourth, capacity building must be treated as a long-term commitment rather than an auxiliary measure. Public administrations require technical expertise to govern effectively; educational systems must align with emerging skill demands; and cross-border talent mobility should be facilitated rather than restricted.

Fifth, cooperative mechanisms must be operationalized through shared (digital) infrastructures and coordinated investment frameworks. Joint procurement of compute resources, federated data spaces, and regional AI research initiatives can create economies of scale while preserving national agency. Moreover, interdependency in mutualized digital infrastructures has the potential to strongly reduce geopolitical frictions, creating a considerable incentive for cooperation.

Finally, implementation must be accompanied by continuous monitoring and adaptive governance, grounded in situational awareness.[270] Sovereignty is not a fixed endpoint but an evolving condition. A large number of technology scenarios must be continuously monitored and contextualized to understand and mitigate risks and take advantage of opportunities. From this perspective, metrics, audits, and iterative policy adjustments are essential to ensure that strategies remain aligned with technological and geopolitical developments.

Being Pragmatic About AI Sovereignty

In its most idealistic form, AI sovereignty promises a level of control that the modern, interconnected world simply cannot provide. Yet abandoning the concept altogether would ignore legitimate concerns about dependency, control, and distribution of value. Hence, the solution lies in a pragmatic redefinition.

By grounding sovereignty in strategic capacity for agency, interoperability, and openness, policymakers can move beyond rhetoric to achieve substantial, yet incremental, change. The goal of the proposed framework is not to eliminate the difficult trade-offs inherent in digital policy, but rather to provide a clearer lens through which to evaluate them. It is important to remember that in a geopolitically contested and complex digital landscape, true sovereignty is neither found in isolation nor in confrontation, but in the deliberate and strategic management of interdependence.

As emphasized above, this pragmatic imperative is dictated by the material realities of the AI value chain. Defined by extreme specialization, immense capital expenditures, and entrenched network effects, frontier AI development depends on access to tightly concentrated bottlenecks, from advanced semiconductors to hyperscale compute infrastructure. In an ecosystem where incumbent advantages continuously compound, attempting to achieve complete autarky is not only economically unviable, but becomes structurally impossible.

Consequently, cultivating meaningful agency, openness, and interoperability within this concentrated landscape demands a holistic, multipronged strategic framework. Statutory oversight alone is insufficient; while adopting comprehensive legislation and ensuring its rigorous enforcement remain critical baselines, true governance requires proactive industrial policy to create the economic incentives necessary for market alignment and distributed

innovation. Furthermore, bridging the gap between structural dependency and strategic autonomy relies on establishing sound multistakeholder governance mechanisms. By deliberately fostering synergy among researchers, developers, regulators, and commercial entities, policymakers have a chance to build a collaborative ecosystem where public interest and market realities align, thus ensuring that sovereignty is achieved not through isolationist autarky, but through active, value-driven participation in the global AI value chain.

Hence, a more modest and ultimately more fruitful approach is to treat sovereignty as a matter of structured agency. Indeed, policymakers should not seek absolute self-sufficiency but instead the capacity to move strategically within constraints. The path forward will be demanding, incremental, and necessarily imperfect. But it is also, under present conditions, the only credible path to retaining sovereignty on the AI stack.