

From Innovation Race to Algorithmic Dependence: AI Geopolitics and the Future of Emerging Economies

1st **Bharath Sangaveer Peruri**

*Mittal School of Business (of Aff.)
Lovely Professional University (of Aff.)
Phagwara, Punjab, INDIA
sangaveer.peruri@gmail.com*

2nd **Kritika Prasher**

*Mittal School of Business (of Aff.)
Lovely Professional University (of Aff.)
Phagwara, Punjab, INDIA
prasherkritika3@gmail.com*

Abstract—AI has stopped being just a technology story. It's a power story now. The US, China, and the EU are racing to lead it, and much of the rest of the world is navigating what that means for them — often on unfavorable terms.

The problem for Global South countries isn't getting access to AI. It's what that access costs. Building national infrastructure on imported systems you can't inspect, modify, or fully understand means you're not just buying software. You're exposing behavioral data, ceding influence over how decisions get made, and — gradually — losing the ability to chart your own course. Some researchers call this data colonialism. The label is contested, but the dynamic it describes is real.

This paper examines how that plays out across different countries and governance approaches, with India as the main case study. India is worth looking at closely: it has the technical depth and political will to push back against dependency, and it's actively doing so. The analysis pulls from policy documents and concrete examples, read through both realist and constructivist lenses — who holds structural power, and how shared narratives shape what countries think is possible.

The pattern that shows up repeatedly is this: states that control AI infrastructure set the rules and harvest the data. States that don't end up structurally subordinate — not through force, but through the slow accumulation of dependencies. The paper calls this an "algorithmic hierarchy." India's experience suggests that domestic investment and deliberate sovereignty-building can change the equation, though not easily.

The framework tying it all together — the "Algorithmic Dependence Model" — is the paper's main theoretical contribution. It links geopolitics, AI adoption, and dependency theory in a way that's meant to be practically useful for policymakers, not just analytically tidy.

Index Terms—AI Geopolitics, Algorithmic Dependence, Digital Sovereignty, Digital Colonialism, Global South, Strategic Technology Competition

I. INTRODUCTION

Artificial Intelligence (AI) which started as a mere system code now changed the entire way we see the reality and also the AI has redefined the 21st century. Today, the way AI was ahead of human capabilities as the way it is redefining the structures and creating help the way which will enhance the human kind the governments are trying to use AI in geo politics and make the most of it in ways we

cannot even imagine as this was helping them evolve as next generation updates for the workforces just like machine era. Rather than relying solely on territorial control or traditional military, modern global power-table is increasingly determined by algorithmic dependence, where state authority is decided through data accumulation, computational capacity and different algorithms. The global AI landscape is increasingly fragmented along political lines, driven by an intense strategic rivalry among major powers—primarily the United States, China, and the European Union. As these nations have one thing in common, technological supremacy and promote their divergent governance models, countries in the Global South face a new challenge of strategic vulnerability known as algorithmic dependence, which puts them in the shoes of dependent nations. When emerging economies import foreign AI systems—particularly for public security, border control, and civic administration the main problem they face is not getting access to source code which leads to dependence on them, algorithmic logic, or operational control. This situation leads to what is referred to as "digital colonialism" or "data colonialism," where developing countries are primarily seen as sources of data extraction, while powerful multinational corporations and leading nations retain authority over essential digital infrastructure. As a result, the dependence on AI systems that are sourced from outside sources poses a threat to worsen power imbalances, increase global inequality, and undermine emerging economies' algorithmic sovereignty and domestic autonomy. Most geopolitical research still centers on military force and territory. AI is reshaping what power means in practice, but the academic literature is slow to reflect that. There's plenty of writing on AI ethics, adoption, and policy — the problem is that it tends to stay in its lane, rarely connecting to the broader questions of structural inequality or what these dynamics actually mean for weaker states.

One gap stands out: very little work treats dependency theory, geopolitics, and AI adoption as parts of the same problem. They are. The US market-driven model, China's state-directed integration, and the EU's regulatory approach aren't just three different bets on how to govern technology.

They generate different power relationships — and for emerging economies trying to build digital infrastructure without surrendering autonomy, understanding those differences isn't academic. It's strategic.

This paper asks a fairly direct question: when countries in the Global South adopt foreign AI infrastructure, what do they give up — and do they fully understand the terms?

The research pursues four lines of inquiry:

- 1) How is AI redrawing global power, beyond the conventional measures of military and economic strength?
- 2) What makes imported AI systems structurally dependency-forming in their architecture, their licensing terms, their data flows?
- 3) Is there a framework that honestly connects geopolitics, adoption challenges, and dependency theory, rather than treating them as separate problems?
- 4) What have countries like India and Brazil actually done to assert digital sovereignty — through policy, local capacity, or regional alliances — and where has it moved the needle?

II. LITERATURE REVIEW

A. AI as a Geopolitical Asset

AI didn't stay in the lab. What began as a niche computational field — mostly academics, some private money, narrow applications — has turned into something states actively compete over. It now touches military systems, economic strategy, and the informal norms that quietly structure international relations.

This isn't just about governments adopting new tools. The nature of power itself is shifting. Countries that lead in AI can project influence in ways that don't show up in traditional measures — troop counts, GDP, nuclear arsenals. They shape how other societies are governed, how decisions get made, and increasingly, who gets to set the rules. That's a different kind of leverage than the 20th century was built around.

The race to lead in AI isn't really about efficiency. It's about who shapes the next era of global power — and who doesn't. Scholars have compared it to the space race, and the analogy holds in one uncomfortable way: it's being treated as zero-sum. Win, and you gain real leverage across economics, military capability, and political influence. Fall behind, and you lose ground in all three simultaneously. That logic is now driving state behavior in ways that go beyond research funding.

Governments are treating sovereign AI — the ability to develop, control, and deploy AI infrastructure domestically — as a national security priority. To get there, they're building up exactly the capacities you'd expect: surveillance infrastructure, data extraction pipelines, tighter control over information systems. The stated goal is independence from technological rivals. The side effects are worth watching closely.

The ability of AI to transform three fundamental areas of national power—military superiority, economic competitiveness, and political/ideological influence—determines its geopolitical weight as an asset (from *The Geopolitical Impact of Artificial Intelligence*).

1) **1. Reshaping Defense and Military Superiority:** Of all the ways AI is rewriting geopolitics, the military dimension is probably where the stakes are most immediate. Deterrence the strategic logic that kept great power conflict in check for decades depended on predictability, clear red lines, and the assumption that humans controlled the trigger. AI is quietly eroding all three.

- **Autonomous Weapons and Cyberwarfare:** Cyberwarfare has become faster and harder to contain. Machine learning lets attackers probe defenses at scale, adapt in real time, and hit with more precision than was possible even a few years ago. Defense has automated too, but it's largely reactive — and the offense tends to move first. Autonomous weapons are a different kind of problem, and arguably a harder one. A system that can select and engage targets without direct human input doesn't just shift battlefield dynamics. It breaks the chain of accountability that military and international law were built around. Who is responsible when an autonomous system kills the wrong people? That question is live, contested, and so far unanswered — while the weapons keep getting deployed.

- **Intelligence and Battlefield Awareness:** States are deploying AI across their militaries — drone swarm coordination, mass intelligence collection, predictive targeting, and command systems that route around traditional decision chains. DARPA's work on AI-driven command and control is a documented example: it's changed how quickly the US military can act, and who in the chain actually makes the call.

The nuclear arms race comparison comes up constantly in this literature. It captures something true — the competitive logic, the fear of falling behind, the sense that the technology is moving faster than the governance around it. Where it breaks down is that nuclear weapons were visible, countable, and eventually subject to treaties. AI capabilities are none of those things. That makes the race harder to slow, and harder to verify if anyone actually does.

2. Economic Competitiveness and Algorithmic Power

Economists classify AI as a general-purpose technology — the kind that rewires whole economies rather than improving one sector. The shift worth paying attention to is what counts as a strategic asset now. The last century ran on oil and industrial output. This one runs on data, algorithms, and whoever controls the platforms everything flows through.

The gains from that shift are not spreading evenly:

- Early movers in AI capture disproportionate productivity growth and, more importantly, start setting the architecture of global supply chains — which other countries then plug into on terms they didn't negotiate.
- AI concentrates fast. Network effects and economies of scale mean the technology heavily rewards whoever got there first, which is why most of the value has pooled in Silicon Valley and Shenzhen rather than diffusing

outward.

- Scholars call this "algorithmic power" — states and corporations that own the infrastructure don't just profit from it, they write the rules. Economic dependencies, trade conditions, and the norms governing digital exchange all flow from whoever controls the stack underneath.

3. Political Influence and International Relations Finally, AI has become a profound tool for political influence and "technological statecraft". AI-driven digital platforms and information systems are deployed as instruments of soft power, enabling governments to shape public discourse and expand their ideological reach globally. For example, the export of AI-enabled




Fig. 1. Investments in AI and Blockchain

surveillance systems (such as facial recognition and predictive policing) has become a foreign policy tool. China's "Digital Silk Road" provides turnkey surveillance infrastructures to nations in the Global South, bundling technology with diplomatic alignment and creating deep geopolitical dependencies (from *Algorithmic Sovereignty and the New Security Dependencies*).



Fig. 2. AI Strategies

B. Digital Divide & Technology Inequality

AI is moving fast, and the economic disruption is real. What's less discussed is how unevenly it's landing. The race for AI dominance isn't functioning as a global equalizer — it's accelerating the gap between the countries driving it and everyone else.

A small number of states control what actually matters: the frontier research, the physical compute infrastructure, the supply chains for critical hardware. That concentration doesn't stay contained. It generates structural dependencies — other countries building on foreign platforms, running on foreign chips, operating within technical and legal frameworks they had no hand in designing. The international system is being quietly reorganized around those fault lines.

Artificial Intelligence and Global Power Dynamics makes a pointed argument: AI isn't a neutral technology that disadvantages some countries by accident. It actively marginalizes states that lack the infrastructure to develop or regulate it — and that's most of the Global South.

The mechanism is structural. Leading economies build the models, own the platforms, and set the technical standards. Emerging economies adopt what's available, on terms set elsewhere, feeding data into systems they don't control and can't meaningfully audit. Wealth and geopolitical leverage accumulate at one end. At the other, countries find themselves not gradually catching up to the digital economy but increasingly unable to participate in it on their own terms. That's a different problem than a simple development gap — and it doesn't self-correct.

Unequal Access to Capital, Infrastructure, and Education The divide comes down to resources — capital investment, physical infrastructure, and the depth of a country's technical workforce. These aren't evenly distributed, and AI readiness maps onto them almost exactly.

An IMF analysis of 125 countries put numbers to this. Assessing readiness across digital infrastructure, human capital, labor market policy, and innovation capacity, the results split cleanly along existing wealth lines: advanced economies are substantially better prepared than emerging markets or low-income countries. That finding alone isn't surprising. What's worth noting is the IMF's observation about the asymmetry of risk. Wealthier countries face real exposure — AI will disrupt their labor markets too — but they have the institutional infrastructure to manage it. Retraining programs, social safety nets, years of accumulated educational investment. Developing economies face similar disruption with none of those buffers in place. The technology arrives either way. The capacity to absorb it doesn't.

Start with the basics. Hundreds of millions of people across Africa don't have reliable electricity. Without that, broadband is out of reach and data centers don't enter the conversation. *The Bulletin of "Carol I" National Defence University* makes the point plainly: digital sovereignty requires functioning utilities first. Countries that can't keep the lights on reliably cannot build the local talent pipelines, startups, or educational ecosystems that indigenous AI development demands.

Country	GDP Growth Rate (%)	Digital Economy Size
United States	5.7	2,000
China	8.1	7,000
Germany	2.9	550
India	8.7	200
Brazil	4.5	150
Nigeria	3.4	40

That infrastructure gap has a downstream effect that's easy to predict and hard to escape. When you can't build your own systems, you buy someone else's. Across much of Africa, Latin America, and Southeast Asia, AI adoption means importing products and running workloads on cloud infrastructure based in the US or China. Maintenance, updates, and the algorithmic logic underneath — all of it lives outside the country using it, managed by providers operating under foreign laws and foreign commercial incentives. The phrase "digital colonialism" gets used to describe this arrangement. Whether or not the label fits, the dependency it points to is real and, for most of these countries, deepening rather than narrowing.

The term "data colonialism" is deliberately provocative, and that's the point. The scholars using it are drawing a structural parallel: powerful actors extract valuable resources from less powerful ones, and what gets left behind is dependency, not development. The resource has changed — it's behavioral data now, not land or labor — but the directionality is familiar. Populations across the Global South generate the data. Foreign tech companies harvest it, build with it, and sell the resulting systems back to the same populations, often at a price that deepens the original imbalance.

What this produces, over time, is a hierarchy that doesn't require coercion to maintain. Emerging economies get assigned roles — data providers, consumer markets, dependent operators of infrastructure they didn't design and can't meaningfully modify. The questions that matter most, what the systems optimize for, whose values they encode, who benefits from the insights they generate, get answered somewhere else, by someone else, for reasons that have little to do with the countries actually living inside these systems.

C. Dependency Theory in the Digital Age

Dependency theory was built to explain a recurring pattern in the global economy: resources and wealth flowing from poorer peripheral states to wealthy core ones through unequal trade and industrial structures. The periphery extracts and supplies. The core processes, captures the value, and sells back. What keeps the arrangement in place isn't force — it's the structural conditions that make alternatives difficult to build and easy to defer.

That framework was developed to describe industrial-era relationships, but it keeps showing up in analyses of the digital economy — and not by accident. The inputs have changed. Raw materials became data. Factory output became algorithmic infrastructure. But the flow of value, and who controls it, follows a recognizable path. *The Geopolitics of Artificial Intelligence* describes dependency theory as now

being "refracted through digital infrastructures rather than industrial production." That's a precise way of putting it. The underlying logic didn't become obsolete. It found new infrastructure to run on.

Map the current AI landscape onto dependency theory and the structure becomes hard to miss. The US and China and the tech multinationals operating out of both control the algorithms, the cloud infrastructure, and the hardware that everything runs on. That's the core. Developing nations in the Global South occupy the periphery, and their primary export has shifted from raw agricultural or industrial goods to something less visible but just as extractable: unregulated behavioral data generated by their own populations.

Algorithmic Sovereignty and the New Security Dependencies uses the term digital colonialism to describe the specific mechanism at work. Foreign powers provide the infrastructure — AI surveillance systems, smart city platforms, ready-to-deploy government tools — but ownership of what actually matters stays elsewhere. The data centers are foreign. The source code is foreign. The algorithmic parameters are set by people who answer to foreign governments and foreign shareholders. Data flows out, gets processed and monetized somewhere else, and what returns to the host country is a finished output with no transparency into how it was built, what it optimizes for, or whose interests shaped it. The country gets the product. Someone else keeps the value.

Algorithmic Dependence and Vendor Lock-In This is where algorithmic dependence becomes structural rather than incidental. Countries that lack the capital, technical workforce, and domestic capacity to build their own AI systems don't just start behind — they stay behind, and the gap compounds. Every system adopted from a foreign vendor creates a new dependency: on that vendor for maintenance, for updates, for security. There's no obvious exit ramp, and the vendors have little incentive to build one.

The dependency gets technically enforced through source code blackboxing and vendor lock-in. A government that deploys a foreign AI system gets access to what it produces, not to how it works. The underlying logic — the criteria, the weightings, the design choices baked into the algorithm — remains invisible. For routine applications, that's a governance problem. When those systems feed into national security decisions, it's something more serious: the actual reasoning architecture is sitting in a black box controlled by a foreign company operating under a foreign legal system, and the adopting government has no real way to audit, challenge, or override it. That's not dependence in a loose sense. It's a structural transfer of sovereignty. Ultimately, this digital dependency severely limits a nation's capacity to exercise sovereign control over its digital future. It fundamentally restricts domestic autonomy, as the rules, biases, and legal jurisdictions of the foreign tech providers override the host country's domestic legal frameworks, creating a new hierarchy of global power where algorithmic control dictates geopolitical subordination.

Attribute	Vendor A (e.g., China)	Vendor B (e.g., U.S.)	Vendor C (e.g., EU)
Source Code Access	No access; proprietary & encrypted	Partial access under strict NDA	Limited access via certified partners
Data Ownership Model	Vendor-controlled or co-owned	Government-owned, vendor-accessible	Locally-owned, with export restrictions
System Control	Remote auto-updates; no local veto	Negotiated patches; vendor-initiated	Local control with vendor collaboration
Cloud Location	Primarily offshore (vendor jurisdiction)	Mixed local/foreign data centers	EU-compliant local hosting mandates
Auditability	Closed-loop system, non-auditable	Third-party audit possible with vendor	Independent audits permissible by design
User Customization	Minimal; fixed use-case configurations	Medium; adaptable through APIs	High; modular and standards-driven
Jurisdictional Control	Foreign legal jurisdiction dominates	Shared legal responsibility	National prioritized sovereignty
Algorithm Transparency	None; model logic obscured	Partial explanations provided	Full explainability frameworks required

Fig. 3. Comparison

World Journal of Advanced Research and Reviews, 2025, 27(02), 162-180

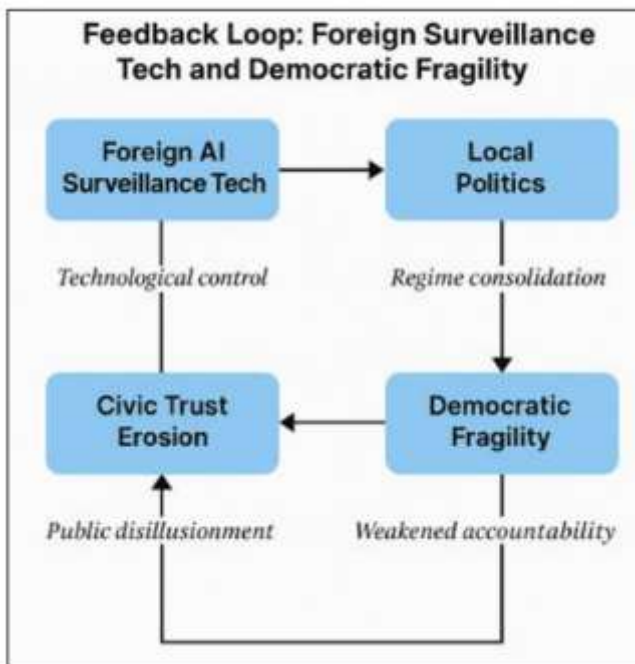


Fig. 4. Feedback Loop

D. Gaps in Existing Research

The literature on AI and international relations is expanding, but it has a structural problem. Classical geopolitics was built around territory, military force, and physical control of space. Technology studies tends to treat AI as an economic or regulatory question, examined within its own disciplinary boundaries. Neither is wrong, but they operate in parallel rather than in dialogue — and the gap between them is where a coherent analysis of AI’s geopolitical consequences should sit.

What falls into that gap is the systemic picture: how AI adoption shapes and is shaped by global power imbalances, how technology choices connect to questions of sovereignty, how decisions made in a handful of dominant states create structural conditions for everyone else. Current research tends

to treat AI policy and technology adoption as separate from those larger dynamics rather than embedded in them.

There’s also a focus problem. The strategic competition between the US, China, and the EU generates most of the intellectual and political attention — it’s high-stakes, well-resourced, and easy to frame as a coherent narrative. That framing crowds out something important. Countries like India and Brazil aren’t observers of this competition. They’re subject to its consequences: foreign infrastructure they can’t audit, data flows they can’t control, digital systems encoding values and priorities set elsewhere. These are live, concrete problems — not edge cases. The fact that they get treated as secondary concerns in a literature organized around superpower rivalry says more about the literature’s blind spots than about the actual stakes involved.

One gap in the current literature is more fundamental than the others. Research describes the outcomes of AI dependency well enough — loss of autonomy, data extraction, structural subordination — but the causal mechanisms are underexplored. How, specifically, does importing foreign AI infrastructure translate into sovereign vulnerability? What are the technical, legal, and political pathways through which that dependency gets established and entrenched? Treating AI as simply a tool of economic or military competition misses this. It needs to be examined as a mechanism of structural reliance — one that functions through digital infrastructure the way earlier forms of economic imperialism functioned through trade dependency and industrial control.

The “Algorithmic Dependence Model” proposed in this paper is an attempt to build that causal account. It pulls together three bodies of thinking that rarely get integrated: the state-power analysis of classical geopolitics, the structural logic of dependency theory, and the infrastructural and governance concerns of technology studies. The aim isn’t theoretical synthesis for its own sake. It’s to produce a framework that can actually trace how dependence on opaque, externally controlled AI systems creates and deepens strategic vulnerability — and to do it in terms grounded in the realities facing emerging economies, not the dynamics of competition between states that already control the infrastructure in question.

III. CONCEPTUAL FRAMEWORK: THE ALGORITHMIC DEPENDENCE MODEL

The Algorithmic Dependence Model pulls together three bodies of thinking that are rarely integrated: classical dependency theory, contemporary state-formation frameworks, and the emerging literature on digital colonialism.

Dependency theory’s core claim is structural. In unequal global economic arrangements, value flows from peripheral developing states to wealthy core ones. The periphery provides the raw inputs — agricultural goods, natural resources, cheap labor. The core processes, manufactures, and captures the margin. What keeps the arrangement in place isn’t just power asymmetry in any given moment; it’s that peripheral states lack the domestic capacity to build alternatives, and the structure reproduces the conditions of its own continuation.

That pattern is legible in the current AI landscape, just running on different infrastructure. Raw materials have become behavioral data. Manufactured exports have become algorithmic systems and platform services. The flow of value still runs in the same direction. Contemporary geopolitical writing describes this as dependency theory refracted through digital infrastructure rather than industrial production — a precise formulation, because the underlying logic is genuinely the same. What changed is the medium, not the structure.

State-formation theory complicates the picture in a useful way. It argues that sustained external pressure — security competition, geopolitical rivalry, economic threat — forces states into a choice: build your own extractive, coercive, and informational capacities, or become structurally dependent on states that already have. The AI race is that kind of pressure, running at speed. States that can't develop sovereign capabilities fast enough don't simply lag. They get absorbed into infrastructure built by others, governed by others, and optimized for others — on terms they didn't negotiate and often can't renegotiate later.

The Algorithmic Dependence Model maps the specific causal pathway by which importing foreign AI technology undermines domestic autonomy in Global South countries. The process unfolds across four sequential stages — Infrastructure Dependency, Platform Dependency, Data Dependency, and Strategic Vulnerability — and the sequencing matters.

Dependency doesn't arrive all at once. It accumulates. Each stage produces conditions that make the next stage more likely and harder to exit. A country that starts by adopting foreign infrastructure finds itself, almost inevitably, dependent on the platforms built on top of it. Platform dependency shapes what data gets collected and where it flows. Data dependency, once established, translates directly into strategic vulnerability — the point at which a state's ability to make autonomous decisions about its own digital future is genuinely compromised. By then, the problem isn't a bad procurement decision that can be reversed. It's a structural condition that took years to build and would take sustained, deliberate effort to dismantle.

A. Stage 1: Infrastructure and Material Dependence

The dependency cycle starts before any algorithm gets written or any platform gets adopted. It starts with rock.

Rare earth elements — seventeen minerals essential for semiconductors, advanced computing hardware, and quantum processors — sit at the physical base of everything AI runs on. China controls roughly 60% of global extraction and 85% of refining capacity. Most discussions about AI geopolitics focus on data and algorithms. The hardware substrate those systems depend on gets far less attention, which is convenient for the one country that dominates it.

For emerging economies that lack domestic REE reserves or the manufacturing infrastructure to process them into chips, this creates a hard constraint that no policy document fixes easily. You cannot build what you cannot fabricate. The practical result is that Global South countries adopt foreign-built

digital infrastructure — cloud computing networks, telecommunications systems, turnkey platforms — supplied by US or Chinese multinationals because there isn't a realistic alternative. The infrastructure lands on their soil. The ownership stays elsewhere. Management, maintenance, upgrade cycles, architectural decisions — all of it controlled externally, by companies operating under foreign laws and answering to foreign governments.

Everything built on top of that arrangement inherits its dependency. The hardware layer isn't just the first stage of the model. It's the condition that makes every subsequent stage difficult to escape.

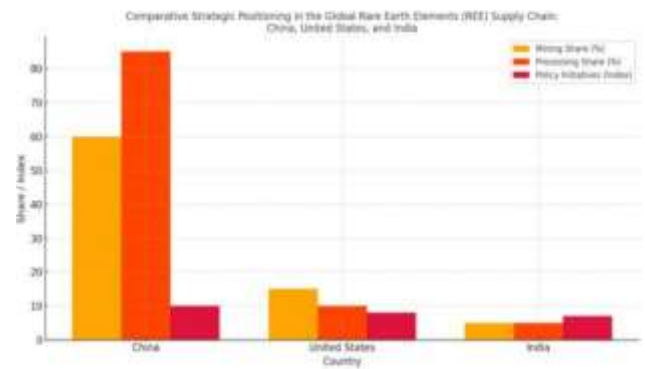


Fig. 5. Global REE

B. Stage 2: Platform Dependence

Once physical infrastructure is in place, dependency moves up the stack — and the software layer is where it starts to feel permanent.

Emerging economies that deploy foreign-developed AI for public administration — facial recognition systems, predictive policing tools, smart city platforms — aren't making a straightforward technology purchase. They're entering an operational relationship with terms that heavily favor the vendor. The systems get installed. The logic that drives them stays locked away.

Source code blackboxing is the technical name for this, and vendor lock-in is what it produces over time. Governments adopting these platforms are typically denied access to the underlying computational logic — the algorithmic parameters, the decision criteria, the weightings that determine outputs. What the system decides is visible. How it decides is not. For a predictive policing tool or a surveillance network feeding into law enforcement, that's not a minor transparency gap. It means the reasoning behind consequential public decisions is sitting in a black box controlled by a foreign company.

Architectural comparisons of AI surveillance systems show this opacity isn't an oversight — it's structural. Dependency gets enforced through a specific set of technical constraints:

- **Source code access:** Proprietary and encrypted. The governments deploying these systems have no access to the logic driving them — not during procurement, not during operation.

- **Algorithm transparency:** Nonexistent. Model logic is deliberately obscured, which means domestic authorities cannot audit how decisions get made, what criteria the system applies, or what biases it carries. They receive outputs. The reasoning stays hidden.
- **System update control:** Vendors push remote updates automatically, with no local veto. Governance protocols and security parameters can change without notice, set entirely by a company operating under foreign law and answering to foreign shareholders.

What this produces, in practice, is a quiet transfer of sovereign function. A government running predictive policing or civic surveillance on these platforms isn't really making the decisions those systems execute — it's ratifying them after the fact, with no real ability to inspect, challenge, or override the underlying logic. The populations subject to these systems are governed, in a meaningful sense, by infrastructure they have no democratic claim over and their governments have no genuine control of. That's not a procurement problem. It's a sovereignty problem dressed up as a technology contract.

Attribute	Vendor A (e.g., China)	Vendor B (e.g., U.S.)	Vendor C (e.g., EU)
Source Code Access	No access, proprietary & encrypted	Partial access under strict NDA	Limited access via certified partners
Data Model Ownership	Vendor-controlled or co-owned	Government-owned, vendor-accessible	Locally-owned, with export restrictions
System Control Updates	Remote auto-updates; no local veto	Negotiated patches; vendor-initiated	Local control with vendor collaboration
Cloud Hosting Location	Primarily offshore (vendor jurisdiction)	Mixed local/foreign data centers	EU-compliant local hosting mandates
Auditability	Closed-loop system, non-auditable	Third-party audit possible with vendor	Independent audits permissible by design
User Customization	Minimal, fixed use-case configurations	Medium, adaptable through APIs	High, modular and standard-driven
Jurisdictional Control	Foreign legal jurisdiction dominates	Shared legal responsibility	National sovereignty prioritized
Algorithm Transparency	None; model logic obscured	Partial explanations provided	Full explainability frameworks required

Fig. 6. Enter Caption

C. Stage 3: Data Dependence

The third stage follows from the second with a kind of structural inevitability. Foreign platforms embedded in a country's public infrastructure don't just sit there — they run continuously, and they generate data constantly.

The underlying logic was named by Shoshana Zuboff as surveillance capitalism: human experience gets converted into behavioral data, processed into predictive models, and monetized. In emerging economies operating on foreign AI infrastructure, that process runs the same way — except the value it produces leaves the country. Biometric records, raw surveillance footage, location data, sensitive personal information: collected inside the host country, transmitted to servers outside it, processed under legal frameworks the host country never agreed to and can't override.

The extraterritorial dimension is what closes the trap. Data generated by citizens in a Global South country, sitting on servers in the US or China, falls under US or Chinese jurisdiction. The US CLOUD Act and China's Cybersecurity Law both give their respective governments compulsory access

to data held on domestic servers — regardless of where the people who generated that data actually live. The host country's own privacy legislation simply doesn't reach. Its citizens are subject to foreign legal exposure through infrastructure their own government procured and deployed.

The consequences compound. The host country can't protect its citizens' data. It can't use that data to train domestic AI models — the resource that would eventually reduce dependency gets extracted before it can be retained. And it can't enforce its own digital rights within its own borders. What gets built, slowly and almost invisibly, is a relationship in which a country's population generates value that accrues entirely elsewhere. The digital colonialism framing isn't rhetorical. It's a description of the actual transaction.

D. Stage 4: Strategic Vulnerability

Infrastructure dependency, platform dependency, and data dependency don't stay separate. They accumulate — and what they accumulate into is the fourth stage: strategic vulnerability, the point at which algorithmic sovereignty has been functionally dismantled.

Algorithmic sovereignty is the capacity to develop, deploy, and regulate the systems shaping your country's security and civic life. By this stage, that capacity exists mostly on paper. The systems are running. Control over them isn't.

The technical exposure is concrete. Critical defense and surveillance infrastructure operating on externally hosted, blackboxed platforms carries an inherent vulnerability: the vendor controls it remotely, and the host government cannot verify what access exists or under what conditions it might be used. During diplomatic tensions or outright conflict, a foreign state or corporation with backend access to these systems could degrade, manipulate, or shut them down. This isn't a hypothetical edge case. It's a structural feature of the dependency relationship, present from the moment the contract was signed.

The political damage is slower but reaches further. Foreign technical systems embedded in domestic governance don't just create security risks — they create accountability gaps that domestic institutions can't close. Local laws get displaced by the operational defaults of foreign platforms. Judicial oversight can't reach systems whose logic is proprietary and whose servers sit in another jurisdiction. Democratic accountability doesn't collapse dramatically; it just stops applying to the decisions that increasingly matter. Foreign algorithms govern hiring, policing, resource allocation, and civic surveillance — and there is no meaningful domestic recourse for the people subject to them.

What this produces is a cycle that compounds itself. Each layer of dependency makes the next layer more entrenched and the exit more costly. Sovereignty doesn't get seized. It gets contracted away, gradually, through procurement decisions that looked like technology choices and turned out to be something closer to constitutional ones.

IV. METHODOLOGY

A. Research Design

The research is qualitative and draws across comparative geopolitics, international political economy, and technology governance studies. That interdisciplinary reach isn't methodological decoration — the questions this paper asks don't sit neatly inside any single discipline, and answering them requires analytical tools that can move between state behavior, economic structure, and technical infrastructure without losing the thread connecting them.

The theoretical framework is intentionally pluralistic, because the phenomenon being studied is genuinely multi-layered. Realism provides the lens for resource competition and national security — the zero-sum logic driving states to accumulate AI capabilities and deny them to rivals. Techno-nationalism captures something realism alone misses: the ideological dimension of sovereign AI ambitions, the way states frame technological self-sufficiency as an identity project, not just a strategic one. Liberal institutionalism adds the cooperative layer — the multilateral frameworks, regional alliances, and governance institutions through which states sometimes manage competition rather than simply prosecuting it.

Using all three isn't a compromise between competing theories. It's a recognition that the modern AI ecosystem is simultaneously a security competition, an economic rivalry, and a governance problem — and collapsing it into any one of those frames produces a distorted picture of what's actually happening.

B. Data Sources

The evidence base pulls from three categories of sources, reviewed systematically across primary and secondary materials:

- **AI and economic reports:** The IMF's AI Preparedness Index, Stanford's annual AI Index, and McKinsey Global Institute assessments on generative AI supply the quantitative grounding — capability gaps, adoption rates, readiness differentials across income groups. These aren't treated as neutral data but as documents that themselves reflect particular framings of what AI preparedness means and who it's measured for.
- **Government policies and strategic directives:** National AI strategies and legislative records from 2020 to 2025 anchor the policy analysis. The EU AI Act, China's New Generation AI Development Plan, US Executive Orders on AI, and India's Critical Minerals Mission frameworks are examined for their stated objectives, but also for their structural assumptions — what kind of AI future each document is trying to build, and whose interests that future serves.
- **Industry data and grey literature:** Corporate disclosures, white papers, and think tank reports fill in what official policy documents rarely address directly: how vendor lock-in actually works in practice, what data ownership models look like in contractual terms, and how

cloud infrastructure architectures create dependencies that persist long after the initial procurement decision. These sources require more critical handling than government documents — they carry their own institutional interests — but they're often the most precise account of operational mechanics available.

C. Approach

The analytical work proceeds on two levels:

Comparative analysis examines the three dominant AI powers — the US, China, and the EU — not just as policy case studies but as competing models of how AI development and governance should be organized. The US approach bets on market competition and private sector leadership. China's model centers state direction and strategic industrial policy. The EU has staked its position on regulatory standard-setting as a form of geopolitical influence. Mapping where these models converge, where they diverge, and what each one produces for countries downstream is essential to understanding the structural environment that Global South states are navigating.

Case study — India brings the framework down to ground level. India is the paper's primary case for how a developing middle power actually manages AI dependency in practice — and it's a more interesting case than most, because India is actively trying to change its position rather than simply accepting it. The country holds domestic REE deposits and has launched policy frameworks to develop them. It has invested in indigenous computing infrastructure. It's engaged in plurilateral arrangements like the US-India Initiative on Critical and Emerging Technologies (iCET) — using strategic alignment selectively, without full commitment to either dominant bloc. The case study examines what these moves have realistically achieved, where structural constraints have limited them, and what the Indian experience suggests about the options genuinely available to countries trying to shift from technology dependency toward something resembling co-creation. India doesn't resolve the problem the paper is analyzing. But it shows what serious engagement with that problem looks like from the inside.

D. Limitations

Two limitations deserve honest acknowledgment rather than boilerplate qualification.

Secondary data reliance is a genuine constraint, not just a methodological footnote. The analysis works from government reports, peer-reviewed literature, and consulting assessments — which means it can describe the outcomes and structures of AI dependency with reasonable confidence, but can't fully penetrate the technical mechanics underneath. The blackboxing that this paper identifies as a sovereignty problem turns out to also be a research problem: proprietary algorithms operated by US and Chinese multinationals are not available for independent inspection, and source code opacity forecloses the kind of detailed technical analysis that would strengthen some of the paper's claims. The limitation is structural — it reflects the same concentration of informational power the

paper is critiquing — but it still shapes what can and can't be concluded with confidence.

The pace of change in AI creates a different kind of problem. Market positions that look stable today can shift within a year. Regulatory frameworks under development may look quite different once implemented. Geopolitical alignments — particularly in the US-India-China triangle — are actively in flux. Some of the specific dynamics analyzed here will have moved by the time this paper reaches readers. The Algorithmic Dependence Model is designed as a durable analytical framework, not a snapshot of current events, but applying it requires ongoing updating — and readers should treat the specific case material as illustrative of structural patterns rather than as a fixed empirical record.

V. GLOBAL AI POWER STRUCTURE

AI's rapid development hasn't distributed power more evenly across the international system. If anything, it's done the opposite. The states and corporations that arrived at this technology first — with the capital, the infrastructure, and the institutional capacity to scale it — have used that head start to entrench advantages that compound over time. The global AI power structure isn't pluralizing. It's concentrating around a small number of actors who control what everyone else depends on.

Four dimensions define that concentration in concrete terms: who funds frontier research and at what scale, who controls the hardware and semiconductor supply chains that AI runs on, who sets the regulatory frameworks that govern how AI gets built and deployed globally, and how deeply AI has been woven into military infrastructure and national security architecture. None of these dimensions operates independently — they reinforce each other, and together they determine which states shape the digital future and which ones navigate it on terms they didn't negotiate.

A. AI Investment Concentration

AI development at the frontier has become extraordinarily capital-intensive — and the scale required to stay competitive has effectively narrowed the race to two serious contenders. Both the United States and China treat AI leadership as a core national security and economic imperative, not a policy option to be weighed against alternatives.

The investments reflect that framing. The US is mobilizing its private sector and venture capital networks behind projects like the \$500 billion Stargate initiative — a four-year effort to build the largest AI infrastructure ever attempted, combining federal direction with corporate scale. China's response is more directly state-administered: a roughly 1 trillion yuan (\$138 billion) AI Industry Development Action Plan financing national champions like Baidu, Alibaba, and DeepSeek while expanding domestic accelerator infrastructure. The architecture differs. The intent — to dominate — is the same.

The EU's €200 billion commitment to AI gigafactories and India's €1.25 billion computing infrastructure investment signal serious engagement, but the gap in scale is hard to paper

over. These are meaningful numbers in isolation. Against the US-China figures, they represent something closer to competitive participation than competitive parity. What's taking shape, across all these investment flows, is a global hierarchy in which frontier AI innovation and the wealth it generates concentrate heavily in two countries — and the rest of the world, including the EU and India, navigates an order built around infrastructure they didn't build and decisions they didn't make.

B. Compute & Semiconductor Control

The AI power structure doesn't float above physical reality. It runs on hardware — microprocessors, semiconductors, cloud infrastructure — and the geopolitics of who controls that hardware is at least as consequential as the geopolitics of who writes the software.

The dependencies are asymmetric and poorly understood outside specialist circles. China controls roughly 60% of global rare earth extraction and 85% of refining capacity. Those minerals go into every chip, processor, and advanced computing system the AI industry depends on. Further up the supply chain, sophisticated microprocessor manufacturing is concentrated in two places: TSMC in Taiwan and Samsung in South Korea together account for nearly 75% of global chip production contracts. That means the physical continuity of the global AI ecosystem runs through two countries sitting in one of the more geopolitically volatile regions on the planet. That's not a risk that gets managed away with supply chain diversification announcements — it's structural, and it's not changing quickly.

Both dominant powers have drawn the obvious conclusion and started using hardware as a direct instrument of statecraft. The US has imposed progressively tighter export restrictions on high-end processors, specifically aimed at denying China access to the chips needed for frontier AI and military applications. China's Made in China 2025 strategy is the countermove: a sustained push toward domestic semiconductor self-sufficiency designed to reduce the exposure that US export controls are trying to exploit. Neither approach has fully achieved its objective. Both have confirmed that hardware is now a front line in the broader AI competition, not a background condition of it.

C. AI Regulation as Strategy

AI governance has become a geopolitical contest in its own right. The rules shaping how AI gets built, deployed, and regulated aren't emerging from multilateral negotiation — they're being driven by three competing models, each reflecting a fundamentally different answer to the question of what AI is actually for.

The US model — market-led, competition-driven: Washington treats AI as an instrument of economic and military primacy. Regulation stays deliberately flexible — not from neglect, but from a strategic judgment that constraining the private sector engines driving US AI leadership would hand advantage to rivals. Silicon Valley sets the pace; the state backstops and directs at the margins. The implicit theory is

that global influence follows technological dominance, and that dominance requires speed more than rules.

The Chinese model — state-directed, surveillance-integrated: Beijing has embedded AI into the apparatus of domestic governance — social credit systems, facial recognition infrastructure, smart city platforms, national economic planning. The New Generation AI Development Plan 2030 formalizes this integration. Internationally, China exports the same ready-built surveillance architecture to Global South countries, often through Belt and Road-adjacent arrangements. The strategic logic is explicit: create digital dependency early, and geopolitical alignment tends to follow. Infrastructure is the leverage.

The EU model — regulatory standard-setting as soft power: Brussels has made a different bet. Through the AI Act and GDPR, the EU has staked its position on human-centric design, algorithmic accountability, and enforceable privacy protections. The power move isn't speed or infrastructure — it's market size. By requiring any multinational operating in Europe to meet EU standards, Brussels effectively exports its regulatory framework globally, setting a baseline that companies and governments outside Europe have to consciously decide whether to meet or reject. Normative influence, exercised through legal architecture.

D. AI in Defense & Cybersecurity

Military AI is where the geopolitical competition becomes most visceral. The integration of AI into defense and intelligence systems is already reshaping how states calculate risk, project force, and define the threshold between peace and conflict — not as a theoretical future scenario, but as a present operational reality.

Autonomous weapons systems sit at the center of this. Militaries are pouring investment into platforms that can identify and engage targets without direct human input. The stated rationale is speed — compressing decision timelines to gain tactical advantage. What gets less attention is what happens to accountability when a system kills the wrong people, or to escalation dynamics when both sides are running automated responses faster than any human can intervene.

Cyberwarfare has quietly transformed alongside this. Machine learning now enables state-sponsored attacks on critical infrastructure that are more precise, more adaptive, and harder to attribute than anything available a decade ago. The same underlying capabilities drive large-scale deepfake production and disinformation operations — tools that sit in a grey zone between intelligence activity and active conflict, and are deliberately kept there. On the defensive side, AI allows militaries to process intelligence at volumes no human analyst could manage, shifting postures from reactive to predictive. That sounds like progress until automated threat assessment starts triggering responses before any human has examined the underlying judgment.

What's being eroded, across all of these developments, is the conceptual infrastructure that war and peace were organized around. The distinction between a cyberattack and an act of

war, between surveillance and aggression, between a human decision and an automated execution — all of it is becoming genuinely harder to locate. That ambiguity isn't a side effect of military AI. For some actors, it's a feature.

VI. EMERGING ECONOMIES: ADOPTION CHALLENGES

While the US, China, and the EU compete over who leads the AI era, most of the world is dealing with a more immediate problem: how to avoid being permanently locked into an order they had no hand in designing.

For countries in the Global South, the AI revolution isn't experienced as a race for global influence. It's experienced as a series of constrained choices made under financial pressure, institutional limitations, and structural deficits that don't resolve quickly. Capital is scarce. Technical workforces take decades to build. Basic infrastructure — reliable power, broadband connectivity, computing capacity — remains inadequate in ways that cap what's achievable regardless of how ambitious the policy agenda is.

What makes these obstacles more than a standard development story is where they lead. A country that can't build, finance, or operate its own AI ecosystem doesn't simply go without — it imports. And what's available to import are ready-built systems from foreign vendors: opaque, externally controlled, and structured in ways that create dependencies the purchasing country rarely fully understands at the time of procurement. That's the entry point into the Algorithmic Dependence Model. Infrastructure reliance creates platform reliance. Platform reliance shapes data flows. Data dependency erodes the capacity for autonomous decision-making. Strategic vulnerability is the endpoint — not because anyone planned it that way, but because the structural conditions of the international AI order make it the path of least resistance.

A. Infrastructure Constraints

The most fundamental barrier to AI adoption in the Global South isn't a shortage of ambition or policy frameworks. It's the absence of the physical conditions AI requires to function. Advanced data centers need reliable power at scale. Connectivity demands high-speed broadband infrastructure. Neither exists at adequate levels across much of the region. Around 592 million people in Africa lack dependable electricity access — a figure that renders conversations about AI readiness somewhat abstract for the communities behind it.

The hardware layer compounds the problem. AI runs on semiconductors and microprocessors manufactured from rare earth elements, and that supply chain is concentrated in ways that leave most emerging economies with no realistic domestic path to self-sufficiency. China controls approximately 60% of global REE extraction and 85% of refining capacity. Countries without domestic deposits and without the industrial infrastructure to process raw minerals into usable components can't build the hardware substrate that indigenous AI development depends on. Importing finished chips and systems is the only practical option — which means the physical foundation of a country's digital infrastructure is owned and controlled

elsewhere before a single line of policy is written. Everything built on top of that arrangement inherits the dependency.

Link to Framework: Stage 1 (Infrastructure Dependence) The infrastructure deficits described above aren't just a development problem — they're the mechanism that triggers the first stage of the Algorithmic Dependence Model. Countries that lack the capital and technical capacity to build domestic hardware networks, manufacturing infrastructure, or computing facilities face a binary choice that isn't really a choice: go without, or adopt what foreign vendors are offering. Most adopt. Cloud-hosted systems and telecommunications frameworks built by US and Chinese multinationals fill the gap because nothing domestic exists to fill it. At the point of adoption, this looks like a pragmatic response to a resource constraint. What it establishes, structurally, is something more consequential: a country's core digital infrastructure — the servers it runs on, the telecommunications architecture connecting it, the systems managing its data — sitting outside its borders, owned by foreign corporations, governed by foreign law, and maintained on terms the host country didn't negotiate from a position of strength.

This is Infrastructure Dependence in concrete terms. It's not a metaphor for vulnerability — it's the literal condition of having your country's digital foundation controlled by someone else. And because every subsequent layer of AI adoption gets built on top of this foundation, the dependency established here doesn't stay contained to infrastructure. It propagates upward through the entire stack.

B. Capital & Investment Gaps

Sovereign AI development isn't expensive the way most infrastructure projects are expensive. The costs are front-loaded, continuous, and scale with competition — meaning the gap between what dominant powers spend and what emerging economies can mobilize keeps widening even when developing countries are investing seriously.

The numbers make this concrete. The US is committing \$500 billion to the Stargate Project. China's AI industry development plan runs to roughly 1 trillion yuan, approximately \$138 billion, directed through state channels toward national champions and domestic accelerator infrastructure. Both figures reflect political decisions to treat AI investment as a national security expenditure — which means they're backed by the full resource mobilization capacity of major states, not just sectoral budgets.

The IMF's AI Preparedness Index documents what that disparity produces on the ground. Assessing readiness across digital infrastructure, human capital, innovation capacity, and regulatory frameworks across 125 countries, the findings split along predictable lines: advanced economies are significantly better prepared for AI adoption than emerging market economies or low-income countries. The index is useful not just as a ranking but as a map of where the structural constraints actually sit. For most of the countries at the bottom of it, the binding constraint isn't political will or strategic vision. It's capital — and the absence of capital at this stage is

precisely what pushes countries toward the ready-built foreign solutions that start the dependency cycle.

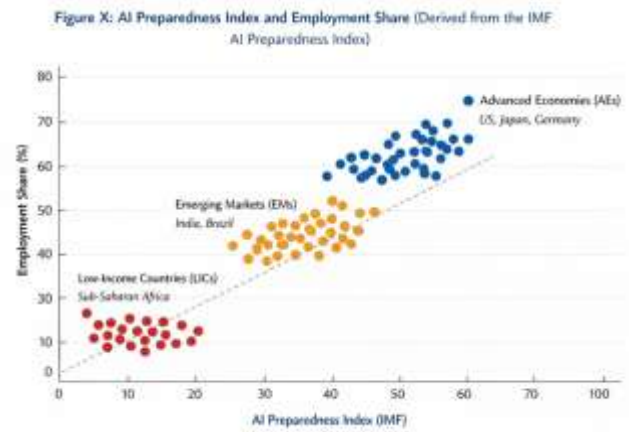


Fig. 7. Enter Caption

Link to Framework: Transition to Stage 2 (Platform Dependence) The capital gap has a direct and predictable consequence. Countries that can't fund domestic AI development don't simply go without — they look for what they can afford. China's Digital Silk Road has positioned itself precisely to meet that demand, offering turnkey AI surveillance systems paired with concessional financing that brings acquisition costs within reach for governments operating under severe fiscal constraints. For a country that needs surveillance infrastructure and has limited options, a subsidized, fully operational foreign system is a genuinely difficult offer to decline.

The adoption decision looks like a practical financial choice. What it creates is Platform Dependency. These systems arrive closed-source, with no access to underlying code and no transparency into algorithmic logic. Vendor lock-in isn't a side effect of the contract — it's built into the product architecture. A government that deploys one of these platforms to fill a budget gap finds itself, months or years later, unable to audit what the system is actually doing, unable to modify it to reflect domestic legal standards or policy priorities, and unable to migrate away from it without dismantling the institutional processes that have been built around it. The switching costs grow with every passing year of integration.

The initial adoption was driven by the absence of alternatives. The dependency that follows is maintained by the accumulation of sunk costs, institutional inertia, and the technical barriers that closed-source architecture deliberately creates. Affordability opened the door. The architecture makes sure it stays open only in one direction.

C. Talent & Brain Drain

Even setting aside infrastructure deficits and capital constraints, a third structural barrier operates independently of both: human capital. Building, deploying, and governing sophisticated AI systems requires deep technical expertise — and that expertise is exactly what the global technology industry is most aggressively pulling out of the Global South.

The mechanism is straightforward. Advanced economies offer salaries, research environments, and career trajectories that developing countries genuinely cannot match. Skilled engineers, data scientists, and AI researchers weigh their options and leave — not as a political statement, but as a rational response to a dramatically asymmetric opportunity structure. The individual decision is entirely understandable. The aggregate effect is that the countries most in need of technical capacity are systematically losing the people who would build it.

What remains is a hollowed-out talent pipeline: weakened universities, thinned research communities, and public institutions trying to govern AI systems they lack the in-house expertise to fully understand. The IMF’s AI preparedness findings make this concrete — developing nations lack not just current technical workers but the underlying conditions that produce them over time: accumulated educational investment, labor market structures that reward technical skill, and social safety nets that allow people to take the risks that advanced technical careers require. Closing that gap isn’t a matter of training programs or scholarship initiatives. It reflects decades of structural underinvestment, and it compounds every year that the brain drain continues.

Link to Framework: Deepening Stage 2 (Platform Dependence) The talent deficit does something specific to the dependency dynamic: it makes it self-sealing. Countries that lack the domestic technical capacity to audit or modify foreign AI systems aren’t just dependent on vendors at the point of procurement — they remain dependent for the entire operational life of the system. Routine maintenance, camera recalibration, algorithmic retraining, system updates — all of it flows through international contractors because there’s no domestic alternative. The vendor relationship doesn’t conclude after installation. It becomes permanent infrastructure.

The accountability implications are serious and underappreciated. Black-boxed source code isn’t just an inconvenience for technically capable governments — for institutions without the in-house expertise to read it, it’s a complete barrier. Domestic authorities running foreign AI systems for predictive policing, border surveillance, or public administration cannot independently verify whether the algorithms contain embedded biases, cannot confirm how decisions are being generated, and cannot assess whether the system is behaving as contracted. They see what the system produces. The reasoning behind it is inaccessible — not because anyone is actively concealing it in the moment, but because the combination of proprietary architecture and absent domestic expertise makes scrutiny structurally impossible.

This is what Platform Dependency looks like when the human capital layer is also missing. It’s no longer just a procurement problem or a contractual arrangement. It’s a condition in which local institutions govern their populations through systems they fundamentally cannot understand, cannot challenge, and — as long as the talent pipeline remains thin — cannot replace.

Dimension	Observation in Global South	Implications for Sovereignty
Operational Control	Foreign vendors manage updates and data storage.	Limits domestic technical capacity; weakens state control over auditing.
Legal Oversight	Lack of AI-specific legislative frameworks and judicial redress.	Legal opacity enables executive overreach and misuse of surveillance data.
Civic Participation	Public consultations absent in procurement processes.	Deepens mistrust, reduces legitimacy, and increases societal resistance.
Transparency	Algorithmic decisions not subject to audit or explanation.	Citizens have no means to challenge wrongful surveillance outcomes.

Fig. 8. Governance Issues

D. Policy & Governance Issues

AI technology moves faster than the regulatory capacity to govern it — and in much of the Global South, the gap between the two is wide enough to drive serious governance failures through.

Surveillance platforms frequently get deployed reactively: a spike in urban crime, a period of political instability, external security pressure. The decision to adopt happens quickly, often under urgency framing, and the institutional infrastructure for oversight — procurement rules, parliamentary review, civil society consultation — gets bypassed or deferred. The technology arrives. The governance framework doesn’t.

What fills that vacuum is often what researchers call elite capture or policy co-optation. Security ministries and military leadership negotiate bilateral contracts with foreign technology companies directly, outside normal procurement channels, under national security cover that insulates the arrangement from public scrutiny. Civil society doesn’t have a seat at the table. Opposition parties don’t either. The result is critical public surveillance infrastructure governed by political relationships rather than legal accountability — with foreign contractors who have their own strong reasons to keep the terms opaque. The legal consequences are more structural than they appear. Imported AI systems were designed for different legal contexts, and their operational defaults don’t automatically conform to local constitutional standards. When a foreign surveillance platform’s built-in parameters conflict with domestic privacy protections or due process requirements, the platform tends to win — not because anyone has legally overridden the constitution, but because the system simply operates according to its own logic and no one in the host government has the technical capacity or institutional authority to override it. Foreign architecture quietly displaces domestic law. That’s legal extraterritoriality, and it happens without a single court case or treaty negotiation — just a procurement decision and a gap in regulatory capacity that was never closed. **Link to Framework: Stage 4 (Strategic Vulnerability)** Elite capture and regulatory failure don’t stay contained to the governance layer. They feed directly into the model’s final stage Strategic Vulnerability where the accumulated dependencies become a political condition rather than just a technical one.

Foreign surveillance infrastructure operating without meaningful legal oversight produces democratic damage that’s hard

to see in any single moment but accumulates steadily. Citizens governed by systems they can't scrutinize, making decisions through processes no domestic institution can audit, start to lose confidence in the state's ability — or willingness — to act in their interest. That erosion of trust isn't abstract. It shows up in institutional legitimacy, in civic participation, in the widening distance between formal democratic structures and the actual exercise of power.

The political elites who facilitated the original arrangements have little incentive to reverse them. The same opacity that protects the foreign vendor protects the ministers who signed the contracts. Accountability requires transparency, and transparency would expose both.

What gets consolidated at this stage is a governing arrangement where the decisions that matter most — who gets flagged as a threat, how risk gets calculated, what triggers surveillance or intervention — are effectively made by algorithmic systems designed in another country, owned by a foreign corporation, and constrained by no domestic law that can actually reach them. The host state's legal framework persists formally. In operational reality, it has been subordinated to the defaults of foreign technology. This isn't democratic backsliding in the ordinary sense of weakening institutions or concentrating power in a strongman. It's a quieter and in some ways harder problem: a state that governs through infrastructure it doesn't control, applying logic it can't inspect, on behalf of interests that aren't its own.

E. Data Colonialism Risks

Data is the primary resource of the AI economy — and for countries in the Global South operating on foreign cloud infrastructure, it's a resource that flows outward by design.

The mechanism requires no malicious intent to operate. Emerging economies running surveillance systems, smart city platforms, or public administration tools hosted by Amazon Web Services, Huawei, or comparable foreign providers transmit citizens' behavioral data, biometric records, and raw surveillance footage to servers outside their borders as a routine condition of using those systems. The data leaves. It gets processed elsewhere. The analytical value it generates — the trained models, the behavioral insights, the predictive capabilities built from it — accrues to the vendor and, through the vendor, to the dominant states whose legal systems govern its use. The country that produced the data ends up with the outputs it paid for. The underlying resource is gone.

The digital colonialism framing is contested in academic literature, but what it points at is real. Historical colonialism extracted physical resources from occupied territories through direct control. This arrangement works differently — through contracts, cloud architecture, and data flows rather than territorial occupation — but the directional logic is structurally similar. Populations in Africa, Latin America, and Southeast Asia generate behavioral surplus that dominant states and multinational corporations capture, process, and monetize. The raw material comes from the periphery. The value concentrates at the core.

The legal layer closes whatever escape route domestic policy might otherwise offer. Data sitting on foreign servers falls under the jurisdiction of the country where those servers are located. The US CLOUD Act and China's Cybersecurity Law both give their respective governments compulsory access to data held by domestic companies — regardless of where the people who generated that data live, and regardless of what privacy protections the host country's own laws provide. A government in sub-Saharan Africa or Southeast Asia that has built its public infrastructure on foreign cloud services has no reliable mechanism to prevent a foreign intelligence agency from accessing its citizens' most sensitive information. Its sovereignty over that data ends at the border of the server farm.

Link to Framework: Stage 3 (Data Dependence) This is Data Dependence — and it's the stage where the dependency cycle stops being reversible through ordinary policy measures.

The core problem is resource capture. Data generated by citizens of emerging economies — behavioral records, biometric profiles, interaction patterns — is the raw material from which AI systems get trained and improved. When that data flows permanently to foreign servers, processed and monetized under foreign legal frameworks, the host country loses more than privacy control. It loses the resource base that would otherwise support domestic AI development. You cannot train indigenous models on data you don't have access to. You cannot build sovereign AI capability on a resource that's been extracted before you can retain it.

What this produces isn't just technical dependency — it's a structural geopolitical position that compounds with time. Every year that citizens' data feeds foreign AI development rather than domestic capability is a year the gap between the host country and dominant AI powers widens. National data sovereignty — the ability to govern how your population's information gets collected, stored, processed, and used — has been functionally dismantled, not through any single decision that could be identified and reversed, but through the accumulated logic of infrastructure choices that each looked, at the time, like pragmatic responses to resource constraints.

The cycle closes here. The dependency began with the absence of infrastructure. It was deepened by platform adoption and governance failure. It completes itself at the data layer, where the resource that could theoretically enable a path toward independence is the same resource being systematically extracted — leaving the country more dependent after each iteration than it was before.

VII. CASE STUDY: INDIA

India serves as this paper's central case study not because it's a typical example of AI dependency in the Global South, but because it isn't.

Most dependency narratives feature a relatively passive subject a country that adopts foreign infrastructure because it has no realistic alternative and lacks the capacity to push back meaningfully. India complicates that picture. It has a large and technically sophisticated workforce, established IT

export infrastructure, and enough strategic weight to negotiate with dominant AI powers rather than simply accepting their terms. At the same time, it faces structural constraints that its technical strengths don't resolve: critical hardware import dependency, an underdeveloped domestic semiconductor industry, and an AI ecosystem that still runs substantially on foreign platforms and cloud services it doesn't control.

The coexistence of genuine capability and genuine vulnerability is what makes India worth examining closely. It shows what the dependency problem looks like for a country that's actively trying to solve it — what policy levers are available to a middle power with real assets, where structural limits bite regardless of political will, and how wide the gap remains between articulating sovereign AI ambitions and actually building the infrastructure to back them up. For other developing countries weighing similar strategies, India's experience is more instructive than a simpler success or failure story would be.

A. Strengths: Talent Pool, IT Ecosystem, and Digital Public Infrastructure

India's positioning in the geopolitical AI landscape is anchored by three formidable domestic strengths: a massive demographic dividend of digital talent, a thriving IT innovation ecosystem, and a pioneering approach to Digital Public Infrastructure (DPI).

Digital Public Infrastructure (DPI) India's digital infrastructure story doesn't map onto either dominant model. The US built its digital economy on private sector platforms owned by a handful of corporations. China's digital infrastructure is state-controlled and closed. India has done something architecturally distinct: built open, state-backed digital public goods that sit outside both frameworks and remain, at least in design principle, accessible rather than proprietary.

The scale of what exists is striking. Aadhaar is the world's largest biometric identification system — over a billion enrollments, used as the authentication backbone for government services, financial access, and welfare delivery. UPI has become one of the most actively used digital payments infrastructures anywhere, processing transaction volumes that have fundamentally changed how money moves across the country. The Digital India Mission has pushed e-governance and digital literacy initiatives across a population and geographic scale that makes the logistical achievement alone significant.

Taken together, these platforms represent something strategically important: a pre-existing, domestically owned data infrastructure that's already operating at AI-relevant scale. For a country trying to reduce dependency on foreign platforms, having functional domestic alternatives isn't a minor advantage — it's the difference between building sovereign AI capability on your own foundation and building it on someone else's. India has the foundation. The question the rest of this case study examines is how effectively it's being used, and where structural constraints still bind despite it.

IT Ecosystem and Talent Pool India's AI ecosystem has genuine depth behind the headline numbers. Decades of

Country	City	No of Startups	Investment in Tech (2023)
India	Bangalore	10,000	\$15 billion
India	Hyderabad	7,000	\$12 billion
Israel	Tel Aviv	6,000	\$10 billion

IT outsourcing built something more durable than a service industry — it produced large pools of technical expertise, established working relationships with global technology companies, and created the engineering culture that now feeds directly into AI development. Bangalore and Hyderabad have evolved from delivery centers into active technology clusters, drawing venture capital and international R&D investment at scale.

The startup layer is substantial. Thousands of AI companies operate across sectors — enterprise software, healthcare, agriculture, financial services — supported by a talent pipeline that remains one of India's most significant strategic assets. The government has recognized this and leaned into it, backing AI development through policy frameworks and public investment on the understanding that the economic spillovers extend well beyond the technology sector.

Market projections put India's AI sector at roughly €17 billion by 2027, growing at 25 to 35 percent annually, with generative AI investment accelerating that trajectory. These are real numbers reflecting real activity. But commercial scale and sovereign capability aren't the same thing. A thriving startup ecosystem that runs on AWS, builds on OpenAI's APIs, and depends on imported Nvidia chips is growing within a dependency structure, not outside it. The case study that follows is partly about whether India's innovation base is translating into the kind of infrastructure control — domestic semiconductors, indigenous foundation models, sovereign cloud capacity — that would actually change its structural position, or whether it's building impressive capability on a foundation it still doesn't own.

B. Weaknesses: Compute Dependence and Chip Manufacturing Absence

Despite its software prowess and vast talent pool, India faces severe vulnerabilities at the foundational hardware and material layers of the AI supply chain.

Absence of Indigenous Chip Manufacturing India's most significant structural vulnerability is one no amount of software talent resolves: it cannot manufacture advanced semiconductors.

Frontier chip fabrication requires multibillion-dollar facilities, extraordinarily precise supply chains, and manufacturing expertise that accumulates over decades rather than policy cycles. TSMC in Taiwan and Samsung in South Korea together handle roughly 75% of global chip production contracts — a concentration that reflects how long and how specifically both countries invested in building that capacity. India is not in that picture at the advanced node level, and the distance between where it is and where it would need to be is measured in decades of investment and industrial development, not years.

The practical consequence is straightforward. Every advanced processor powering India's AI applications — in its data centers, government platforms, military systems, and startup ecosystem — is imported. The physical hardware that India's digital economy runs on is manufactured elsewhere, shipped in, and dependent on supply chains India doesn't control and can't guarantee under geopolitical pressure. That's not a secondary concern sitting beneath more visible strategic vulnerabilities. It's the foundational one — and it persists regardless of how capable India's engineers are or how ambitious its AI policy frameworks become. Sophisticated software built on imported hardware is still built on hardware someone else controls.

Compute and Material Dependence The semiconductor deficit sits inside a wider hardware problem. AI systems at scale depend on rare earth elements — the minerals that go into GPUs, processors, and the full stack of advanced computing hardware. China controls approximately 60% of global REE mining and 85% of processing capacity. India's share is around 5%. That gap isn't just a trade imbalance — it means India cannot independently source the raw materials its AI hardware requires, which makes every layer of physical infrastructure built on top of it dependent on supply chains it doesn't control.

The compute situation reflects the same structural reality. Large AI models require substantial GPU arrays and the data center capacity to run them. The Indian government has acknowledged this directly — the India AI Mission commits \$1.25 billion to building domestic compute infrastructure targeting 10,000+ GPUs. That's a serious investment and a meaningful policy signal. It's also a fraction of what frontier AI development requires, and it addresses a gap that currently leaves India's AI workloads running predominantly on Amazon Web Services and Google Cloud. Foreign platforms, foreign servers, foreign legal jurisdiction — with all the data sovereignty implications that follow.

The India AI Mission represents genuine strategic intent. But intent and structural position aren't the same thing. Until domestic compute capacity, REE processing, and semiconductor manufacturing reach a scale that actually reduces import dependency rather than just reducing its rate of growth, India sits firmly within the Infrastructure Dependence stage of the model — not as a passive victim of the dynamic, but as a country actively working against structural constraints that remain, for now, largely intact.

C. *Opportunities: AI for Public Services and Global South Leadership*

India is strategically leveraging its growing technological capacity to address domestic challenges while positioning itself as an international leader for emerging economies.

AI for Public Services and Inclusive Growth India's domestic AI agenda starts from a different set of imperatives than the US or Chinese models. With 1.4 billion people and persistent structural gaps in healthcare access, welfare delivery, and financial inclusion, the most pressing domestic case for

AI isn't market dominance — it's whether the technology can help the state function more effectively for the majority of its population that existing systems consistently underserve.

The policy framework "Leveraging AI for Inclusive Growth" captures this orientation. In practice, it means AI-driven tools being deployed for coordination between government departments, urban planning optimization, and environmental impact modeling for zoning decisions. These aren't headline-generating applications. They're attempts to solve the unglamorous but consequential problem of making a large, complex state work better at scale.

The underlying opportunity is real. India's Digital Public Infrastructure — Aadhaar, UPI, the e-governance stack — already operates at a reach that most countries can't replicate. It touches nearly the entire population and generates data flows at a scale that's architecturally compatible with AI-driven service delivery. Integrating AI with that existing foundation creates a plausible path to transforming how healthcare, financial services, and welfare programs reach communities that have historically fallen through the gaps. The question isn't whether the infrastructure could support that transformation — it demonstrably could. The question is whether India can govern the AI systems doing the delivering in ways that keep them accountable to the people they're supposed to serve, rather than to the foreign vendors or domestic elites who control the technical architecture underneath.

Global South Leadership and Plurilateral Diplomacy India's response to the structural pressures of AI geopolitics isn't to pick a side — it's to avoid being fully captured by either. The approach gets described in policy circles as technological statecraft: engaging with multiple frameworks simultaneously, building leverage through selective participation rather than wholesale alignment, and treating strategic autonomy as a goal worth paying diplomatic costs to maintain.

The plurilateral engagements are the practical expression of this. The US-India Initiative on Critical and Emerging Technologies (iCET) pairs science diplomacy with industrial cooperation on REE supply chains and AI hardware — reducing India's exposure to Chinese-controlled inputs while stopping well short of the kind of dependency on American technology that would simply recreate the problem in a different direction. The Minerals Security Partnership and the Quad's technology working group operate on the same logic: use multilateral structures to diversify the dependency landscape rather than consolidate it around a single dominant partner.

The more ambitious claim India is making is that it represents something the Global South actually needs — a democratic, open-source alternative to China's Digital Silk Road model, which offers infrastructure on terms that this paper has documented in detail. India is positioning itself as a different kind of partner: one that promotes rule-based resource governance and open digital architecture rather than dependency-generating turnkey systems. The pitch has genuine appeal in parts of Africa and Latin America that have experienced the Digital Silk Road's terms firsthand.

But a credible alternative model requires more than a

different set of values. It requires demonstrated capability — domestic semiconductor production, sovereign cloud infrastructure, indigenous AI systems that other countries could actually adopt without recreating the dependency structures they're trying to escape. India is working toward that. It isn't there yet. The gap between the strategic positioning and the underlying capability is real, and how India closes it over the next decade will determine whether its claim to Global South technology leadership is substantive or aspirational.

D. Threats: Dependency on Foreign AI Systems

Despite proactive policy maneuvering, India remains highly susceptible to the geopolitical threats accompanying imported technology.

Data Colonialism and Algorithmic Opacity India's integration with foreign technology platforms is deep enough that data colonialism isn't a hypothetical risk — it's a description of current operating conditions. The behavioral data of over a billion people flows through systems run by multinational corporations under legal jurisdictions that aren't India's. It gets processed abroad, feeds foreign AI development, and generates commercial value that accrues to foreign companies. The economic asset sitting inside India's population scale is being extracted at the platform layer, and the decisions those systems inform remain largely opaque to Indian authorities.

The vulnerability is sharpest in sensitive sectors. Telecommunications networks, smart city infrastructure, and border surveillance systems running on closed-source foreign platforms embody exactly the conditions the Algorithmic Dependence Model traces: vendor lock-in with no realistic exit, algorithmic logic that domestic institutions can't audit, and system defaults that subordinate local legal standards to foreign corporate architecture in practice even when domestic law nominally governs. If India's domestic computing ambitions stall — if the India AI Mission's GPU targets don't translate into genuine sovereign capability — continued dependence on foreign turnkey systems in these sectors isn't a manageable inconvenience. It's a structural national security exposure.

India's push for data localization requirements — mandating that sensitive personal data be stored and processed domestically — is a direct response to this. The policy direction is sound. The harder questions are enforcement: whether domestic infrastructure exists at the scale needed to actually host what localization rules require, whether the regulations apply with the same force to the sectors where vulnerability is highest, and whether a localization mandate alone addresses the deeper problem of foreign algorithmic control over systems that happen to store their data locally. Keeping data inside the border doesn't automatically return control over what gets done with it.

E. India as a "Swing State" in AI Geopolitics

India's position in the global AI competition has shifted enough that "emerging economy navigating dependency" no longer fully captures it. India has become a swing actor — a state whose strategic alignment decisions carry genuine weight

in how the broader AI order develops, not just how India itself fares within it.

The policy approach that produced this position is worth examining on its own terms rather than just labeling it. India is simultaneously running techno-nationalist plays — domestic computing investment, data localization mandates, indigenous platform development — and realist ones, through the Critical Minerals Mission's effort to secure material sovereignty over the hardware inputs where dependency originates. At the same time, it's using liberal institutionalist frameworks — iCET, the Minerals Security Partnership, the Quad technology working group — to build strategic corridors that diversify its dependencies without consolidating them around a single dominant partner. These aren't contradictory moves. They're a coherent strategy for a country that has enough leverage to play multiple frameworks simultaneously and enough strategic discipline to avoid being captured by any of them.

Whether this translates into durable algorithmic sovereignty depends on questions that haven't been answered yet. Domestic semiconductor manufacturing remains underdeveloped. The gap between India's compute ambitions and its current infrastructure reality is significant. Data localization rules are only as effective as the domestic capacity to enforce and absorb them. The structural vulnerabilities documented earlier in this paper don't disappear because India's strategic posture is sophisticated.

What India demonstrates, though, is that the dependency cycle the Algorithmic Dependence Model describes isn't deterministic. Middle powers with real assets — engineering depth, data scale, demographic weight, and the political will to use them strategically — can push back against the structural conditions that trap smaller, less resourced economies. The outcome isn't guaranteed. But the possibility is real, and for other developing countries watching how India navigates this, that matters more than whether India fully succeeds.

VIII. DISCUSSION

A. Historical Parallels: From Coal and Oil to Algorithms and REEs

Every major shift in global power has a technology story underneath it. That's not a controversial claim — it's what the historical record shows repeatedly. Geography and natural resources establish the structural conditions states operate within. But the states that have actually reshaped those conditions — conquered distances that were supposed to be prohibitive, built empires that shouldn't have been sustainable, displaced powers that looked entrenched — have done it by developing or adopting technologies that others couldn't match or didn't see coming.

AI sits in that lineage. Understanding the scale of disruption it represents requires placing it inside that longer pattern rather than treating it as unprecedented. The geopolitical consequences of disruptive technology aren't new. What changes each time is which technology, which states, and which existing order gets reorganized around the new capability.

The Migration of Economic Hegemony The connection between technology and global power becomes most legible when you trace where economic hegemony has actually sat — and what shifted it each time.

The Industrial Revolutions are the standard reference, and for good reason. Steam engines, railroads, and the telegraph transformed more than logistics. They gave European states the administrative reach and military projection capacity to build and sustain colonial empires at a scale that earlier technologies couldn't support. The infrastructure of empire the ability to move troops, extract resources, transmit orders, and administer vast territories depended on technological edges that geography alone couldn't provide.

Jacques Attali's historical analysis traces the same pattern further back and with more precision. Global economic gravity, in his account, moved through a series of technology-driven transitions: Venice built its position on maritime innovation; Amsterdam on the mass production of the fluyt ship and the commercial infrastructure organized around it; London on steam power and the manufacturing revolution that followed. These weren't smooth handoffs. Each represented a relatively sharp reorganization of which state's technological advantage translated into economic dominance and political influence and each left the previous hegemon managing decline rather than leading the next era.

The center of gravity then crossed the Atlantic in successive waves. The internal combustion engine, the electric motor, the microprocessor, and the internet each reorganized economic geography within the United States and entrenched American dominance globally. The technology changed with each wave. The underlying dynamic — disruptive capability concentrating power in whoever develops and deploys it first stayed constant. AI, quantum computing, and blockchain are now generating the same kind of instability that each previous wave produced. The existing distribution of power isn't fixed. The states moving fastest on these technologies are accumulating advantages that compound. And the states that miss the transition, as history suggests, tend not to recover their previous position.

The Hydrocarbon Era and the Power of Monopolies The 20th century's clearest example of strategic resource power is also its most instructive analogy for what's happening now with technology supply chains.

Oil didn't just fuel industrial economies it structured global power around whoever controlled its extraction and distribution. The Middle East's geopolitical landscape was largely a product of hydrocarbon geography. OPEC's formation made that leverage collective and explicit. Major powers financed regional clients, built foreign policy around access guarantees, and treated petroleum security as a core national interest. The resource was economically valuable, but its deeper strategic significance was that it was weaponizable a dependency so fundamental that threatening it could coerce governments that no conventional military pressure could easily reach.

The 1973 embargo showed exactly how this worked. Imposed in response to Western support for Israel during the Yom Kippur War, it hit economies that had built themselves around

the assumption of uninterrupted oil access with no serious contingency for its sudden interruption. The damage was rapid and substantial industries stalled, governments panicked, and the political costs of Middle East policy became viscerally concrete for publics that had never thought much about where their energy came from. The embargo lasted months, not years. It didn't need to last longer. The demonstration was the point. The technology supply chain of the 2020s has the same structural properties. Semiconductor manufacturing concentrated in Taiwan and South Korea, rare earth processing dominated by China, cloud infrastructure controlled by a handful of US and Chinese corporations — these create chokepoints with the same coercive potential that petroleum concentration demonstrated in 1973. The US export restrictions on advanced chips to China, and China's periodic signals about rare earth access, are both drawing on the same playbook OPEC wrote. The resource is different. The underlying logic of dependency-as-leverage is identical.

The New Material Substrate: Data and Rare Earth Elements (REEs) The physical foundation of AI power looks less like a digital revolution and more like a resource extraction story the 20th century would recognize. Data is the primary currency of the AI economy — but data runs on hardware, hardware runs on semiconductors, and semiconductors require 17 critical minerals that don't distribute themselves conveniently across the globe. The digital economy has a material substrate, and who controls that substrate shapes everything built on top of it.

China's position in that substrate is the contemporary equivalent of what Deng Xiaoping identified in 1992: "The Middle East has oil, China has rare earths." The observation has aged into something closer to a strategic doctrine. China controls approximately 60% of global REE extraction and 85% of refining capacity — figures that have been built deliberately over decades of industrial policy, not inherited by accident. That concentration gives Beijing the same structural leverage over technology supply chains that OPEC demonstrated over energy markets: the ability to impose costs on dependent economies through supply restriction rather than military confrontation.

This isn't theoretical. In 2010, during a diplomatic dispute with Japan over contested islands, China temporarily halted rare earth exports — a targeted disruption that hit Japanese electronics and automotive manufacturers with immediate and measurable force. The restriction lasted months and was eventually walked back under international pressure. But the strategic demonstration had already been made: REE dependency created a vulnerability that could be activated at will, and the economies that had built supply chains around uninterrupted Chinese supply had no short-term alternative. Every government that watched that episode adjusted its threat assessment accordingly. The 1973 OPEC embargo had taught the same lesson about oil. China's 2010 move taught it again about the minerals that the next technological era runs on.

Cold War Echoes: The AI Arms Race and the Space Race The rare earth competition is one layer of the historical parallel. The broader AI race has the structural shape of Cold

War strategic competition — and that resemblance isn't accidental. Political elites on both sides are consciously invoking it, which itself shapes how governments respond and what domestic constituencies they're trying to mobilize.

The Space Race comparison is the most frequently drawn, and it captures something real. The original Space Race wasn't about scientific curiosity or human exploration — it was about demonstrating technological and industrial capability in ways that carried direct military implications. The rocket that put a satellite in orbit was a close cousin of the missile that could reach an adversary's cities. National prestige was the public framing. Strategic deterrence was the underlying logic. The sovereign AI race works similarly: the civilian and commercial applications are real, but the strategic significance driving state investment is primarily military and geopolitical. The language of national survival gets applied because the people using it genuinely mean it.

What this framing produces, in practical terms, is the same kind of state mobilization the Cold War generated. Governments treating AI dominance as an existential priority don't approach it through ordinary research funding and market incentives. They deploy extractive capacity to secure critical mineral inputs before rivals do. They use coercive instruments — export controls, investment restrictions, technology transfer prohibitions — to deny adversaries access to key capabilities. They invest in informational capacity to shape the global norms and governance frameworks that will determine what AI development is permitted to look like. Algorithms and compute infrastructure have been written into national security doctrine not as metaphors for importance but as literal strategic assets whose loss or compromise would constitute a meaningful blow to national power. The Cold War precedent suggests that once a competition gets framed that way, it tends to escalate further rather than stabilize.

AI as an Unprecedented General-Purpose Technology

Electricity and the internet transformed economies by enabling scale and improving how physical and informational systems connected. Both were genuinely disruptive. Neither touched cognition directly. AI does something categorically different: it replicates and extends the reasoning, learning, and generative capacities that were previously exclusive to human minds. For states, that's not just a productivity gain — it's a simultaneous amplifier across every instrument of national power. Intelligence collection, military planning, administrative decision-making, information operations — all of them become faster, more scalable, and less dependent on human bandwidth when

AI is embedded in them. That's what makes AI a force multiplier in a sense that electricity and the internet weren't. The historical comparisons help locate the stakes but none of them fully contain what's happening. The industrial revolution gave early movers decisive advantages in projecting physical power and sustaining empire. The nuclear era created a specific deterrence architecture — stable in its way, but built around the logic of mutual assured destruction rather than competitive advantage. AI combines the broad, economy-wide advantages of industrialization with strategic stakes that

approach the existential, while adding autonomous cognitive capability that neither historical precedent involved. It's not that the analogies are wrong. It's that each one captures part of the picture and misses something the others catch.

The deeper transformation is structural. The Westphalian international order organized power around territorial sovereignty and physical military force. What's replacing it — not quickly, but unmistakably — is an order organized around data accumulation, computational capacity, and control over the digital infrastructure that increasingly mediates how economies function, how governments operate, and how militaries fight. States that build genuine sovereign AI capability will write the rules of that order. States that remain consumers of technology produced elsewhere will navigate a world shaped by decisions they had no part in making — which, as the dependency dynamics this paper documents suggest, is a position that tends to become more constrained over time rather than less.

B. The Emergence of a New Global Hierarchy: Algorithmic Hegemony

The analytical picture that emerges from this paper is not optimistic about AI's effects on global equality. The technology is concentrating power — in specific states, specific corporations, and specific infrastructure nodes — faster than any countervailing force is distributing it. The gap between countries that lead in AI and countries that consume what others build is a structural feature of how this technology develops, not a temporary lag that catches up on its own.

The international order this is reshaping was built on a specific set of foundations: territorial sovereignty, fixed borders, and military force as the primary currency of state power. Those foundations haven't disappeared, but they're no longer sufficient to describe where consequential power actually sits. A state can have secure borders and a capable military and still find its population governed through foreign algorithms, its data processed under foreign legal jurisdiction, and its critical infrastructure dependent on vendors answerable to foreign governments. Territorial integrity, in that scenario, is largely formal. The substance of self-governance has migrated elsewhere.

What's emerging in place of the old framework isn't a clean replacement — it's an overlay. Algorithmic hegemony operates through control over data flows, computational infrastructure, and the digital systems that increasingly mediate how economies function and how governments make decisions. The states that dominate those systems exercise influence over other states that no treaty recognizes and no conventional military response addresses. Sovereignty in this environment is less about what happens at your borders and more about whether the systems shaping life inside those borders answer to your laws or someone else's. Most of the world, as this paper has argued, is currently on the wrong side of that question.

The Bipolar Concentration of Power The AI hierarchy has a clear top tier, and it currently has two occupants. The US

and China control the capital flows, hardware ecosystems, and research infrastructure that determine who sets global technical standards — and both governments have concluded that the competition between them isn't primarily technological. It's geopolitical, and it's being prosecuted accordingly.

The structural approaches differ. The US mobilizes through its private sector and venture capital networks, with the federal government providing direction and backstop while Silicon Valley provides scale and speed. The \$500 billion Stargate Project is the most visible expression of this model — a public-private infrastructure bet designed to entrench American AI dominance before rivals can close the gap. China's approach is more directly administered: the New Generation AI Development Plan 2030 and Made in China 2025 treat AI as an instrument of state power, integrating it with military modernization, domestic surveillance infrastructure, and international technology exports through frameworks like the Digital Silk Road. Different architectures, convergent ambitions.

China's rare earth position gives its standing in this hierarchy a material dimension that's easy to underestimate in discussions focused on software and models. Controlling approximately 60% of global REE extraction and 85% of refining capacity means that the physical inputs for every advanced semiconductor and AI hardware system pass through Chinese-controlled supply chains at some point. States building AI infrastructure on imported chips are, structurally, dependent on that processing capacity — regardless of which country's technology brand is on the finished product. That dependency doesn't disappear when relations are cordial. It becomes a coercive instrument when they aren't, as the 2010 Japan episode demonstrated.

The EU is making a serious attempt to establish a third position through regulatory influence — using the AI Act, GDPR, and its substantial market size to export governance norms that multinationals operating in Europe must comply with globally. That's genuine power, and it shapes the terms on which AI gets deployed across a significant portion of the world economy. But normative influence over existing systems is categorically different from controlling the infrastructure those systems run on. The foundational layer of the AI era — the hardware, the compute, the capital — is being built in Washington and Beijing. Brussels is working to govern what others have built, which is a meaningful role, but not the same as building it.

Network Effects, Economic Stratification, and Digital Colonialism AI is concentrating wealth fast. Faster, probably, than any technology before it — because it doesn't just automate physical tasks. It amplifies thinking. The gains compound for whoever gets there first, and the gap widens from there.

PwC projected AI would add around \$15.7 trillion to the global economy by 2030, with North America and China capturing most of it. That's not a coincidence. It follows from who owns the infrastructure, the training data, and the engineers.

For countries that don't have those things, the options are bad. Build your own — which takes capital and years. Or buy

what's available, which usually means licensing platforms built abroad, shaped by foreign priorities, running on servers you don't control. Most of the Global South ends up with option two, not because they chose it but because option one isn't realistic.

This is what people mean by "digital colonialism." Residents generate data, foreign algorithms process it, profits leave. The country ends up running a system it didn't design and can't meaningfully audit. Switching costs are high enough that even negotiating is difficult. You're just a user.

That positioning isn't accidental. It's structural.

The Militarization of AI and the Obsolescence of Deterrence The military piece is where this gets harder to look away from.

AI is now inside weapons systems — not just logistics software or facial recognition at borders. Autonomous weapons, drone swarms, targeting algorithms that process battlefield data faster than any human chain of command. The countries developing these aren't just gaining an edge. They're changing what conflict looks like in a way that makes older deterrence frameworks basically irrelevant. If your doctrine assumes rough parity, and parity no longer exists, the doctrine doesn't hold.

Cyber is the same story. Machine learning makes attacks faster, harder to attribute, and more precise. The same infrastructure runs large-scale disinformation — deepfakes, synthetic media, coordinated influence operations that are genuinely difficult to detect in real time. That's not hypothetical anymore. It's been the operational reality of several conflicts in the past few years.

The dynamic this creates resembles a Cold War arms race, except the barrier to entry isn't nuclear material — it's technical capacity. Which means more countries are caught in it, with fewer options. Either develop your own AI military capability, which most countries can't afford, or attach yourself to a power that has one.

That second choice is sold as a partnership. It functions more like dependency. You get security guarantees; you give up a degree of sovereignty over your own systems that's hard to quantify until something goes wrong.

Multinational Tech Corporations as Quasi-State Actors The companies building this infrastructure aren't just businesses anymore and haven't been for a while.

Google, Microsoft, OpenAI, Huawei, Tencent, Baidu. They run the platforms that governments depend on, the cloud systems hospitals use, the data pipelines underneath modern finance and media. That's not a commercial relationship in any ordinary sense. When critical public infrastructure runs on private servers, the company operating those servers has leverage that no amount of market competition fully neutralizes.

The accountability gap is the real issue. Banks, utilities, broadcasters — powerful institutions, but they operate inside regulatory frameworks with actual enforcement mechanisms. These companies mostly don't. They set the technical standards other industries have to follow. They shape what information reaches people and what doesn't. They decide who

gets platform access, under what conditions, and for how long. Very little of that passes through any democratic process.

"Quasi-state" is the term people reach for, and it captures something real. But states have formal obligations to citizens — constitutional ones, legal ones, ones you can sue over. These companies have terms of service and shareholder meetings. The gap between those two things is where a lot of the actual power lives.

In the end, algorithmic predominance represents a significant shift in the global order. The ability to project algorithmic power has replaced military might and territorial span as the primary determinants of global power, enabling dominant states and a small number of multinational tech companies to set the normative, security, and economic rules for the rest of the world.

Here is the expanded, detailed version of Section 11.3, integrating the structural drivers, technical realities, and geopolitical consequences drawn directly from your source materials.

C. The Fate of Emerging Economies: Permanent Dependence or Strategic Vulnerability?

The question underneath all of this is simple, even if the answer isn't: does this arrangement change, or does it just continue?

The structural pressures suggest it continues. Countries without domestic AI capacity or meaningful data governance don't have many real options. They license foreign platforms, generate data that flows out, and watch decisions get made in boardrooms and legislatures they have no access to. That gap doesn't naturally close — it widens, because the countries already ahead keep reinvesting while everyone else is still working on the basics.

"Digital colonialism" and "data colonialism" are the terms that have stuck in the research literature. They're useful because they connect what's happening now to something older. The mechanism is different — server farms instead of trade monopolies, behavioral data instead of raw cotton — but the basic shape is recognizable. Value concentrates at the center. The periphery supplies inputs and receives finished products it had no hand in designing.

The harder thing to sit with is how invisible the infrastructure is. Historical extraction was physical. Ships, ports, ledgers — it left a material record. This runs in the background, embedded in platforms people use daily without thinking about where the data goes or who profits from it. That invisibility isn't incidental. It makes the whole arrangement easier to maintain and harder to contest.

Structural Bottlenecks and Infrastructure Dependence
The dependency isn't just political. It's physical, and that matters because physical constraints are harder to legislate away.

Building AI from scratch requires money, engineers, stable electricity, and access to the hardware underneath it all. Most countries in the Global South are short on several of these simultaneously. The electricity problem alone is severe — hundreds of millions of people in Sub-Saharan Africa don't

have consistent power. You cannot run data centers or train models without it. There's a basic computing floor that AI development requires, and large parts of the world are still beneath it.

Hardware is a separate wall. The semiconductors that make AI run are produced by a small number of companies, mostly in Taiwan, South Korea, and the US. The rare earth elements those chips depend on flow through supply chains that emerging economies don't control and largely can't enter. If you can't manufacture your own chips and can't reliably source what's needed to make them, you buy from whoever can — on their terms, not yours.

That's the practical reality behind the purchasing decisions that look, from the outside, like policy choices. When a government needs smart city infrastructure, biometric border systems, or predictive policing tools, the supplier list is genuinely short. American companies, Chinese companies, a few European ones. The country on the receiving end of that transaction has limited negotiating power, limited visibility into what it's actually buying, and limited ability to audit it afterward. It takes what's available.

The Algorithmic Dependence Model: The opacity is built in, not incidental — and that distinction is worth sitting with. When a government imports a closed AI system, it gets the interface, not the engine. The code is proprietary. The algorithmic logic stays with the vendor. Local authorities can see what the system produces but have no way to examine how, what assumptions shaped it, or what it was originally trained on. If the model flags certain populations at higher rates, or was built with data that reflects someone else's priorities, there's no mechanism to detect that from the outside. You're just running it.

Vendors typically retain remote access to update and patch the software. Which means the system a government bought and approved last year may not be the system running today, and the government may not know what changed or why. The dependency isn't just about purchasing — it's structural, written into how the technology works.

Then there's the data. These systems collect behavioral information from residents — movement patterns, communications metadata, financial activity — and route it to servers abroad. Once it leaves, it's subject to foreign law, not the host country's. Under the US CLOUD Act, American authorities can compel access to data held by US companies regardless of where it was gathered. China's Cybersecurity Law creates equivalent obligations for Chinese firms. The country whose citizens generated that data has no standing in either framework and no practical recourse. It has, in any meaningful sense, lost control of it.

Elite Capture and Passive Adoption Dependency this deep doesn't sustain itself. Someone keeps choosing it.

The "urban modernization" and "public safety" framing does a lot of work. It makes the adoption of foreign surveillance infrastructure sound neutral — technocratic, obviously sensible, barely worth debating. But the procurement decisions underneath that framing are often neither. In fragile democra-

cies, bilateral agreements with foreign tech vendors regularly bypass normal processes. No public tender, no parliamentary review, no independent audit. The deal happens between a company and whoever currently holds power.

The logic isn't complicated. Surveillance tools are useful to incumbents. A platform that can track political opposition, flag dissent, and monitor communications is valuable to a government trying to stay in power — and doubly valuable because accountability is diffuse. If something goes wrong, or someone asks uncomfortable questions, the algorithm becomes a convenient buffer. We just use what the system tells us.

What gets given up in this exchange rarely gets named directly. Governments that sign these agreements without requiring source code access, local data storage, or human rights safeguards aren't just making a questionable procurement call. They're handing effective legal authority over their own populations to a foreign company. The domestic legal system still exists in form. But the actual rules — how data is collected, what triggers a flag, what happens to the information afterward — were written by someone else, in another country, for different purposes. That gap between formal sovereignty and operational reality is where a lot of the damage lives.

Strategic Vulnerability and the Feedback Loop of The endpoint of this is a country that cannot trust its own infrastructure — and may not know it until something goes wrong. The kill-switch scenario is the most concrete version of that risk. If a nation's defense systems and policing tools run on vendor-controlled software, the vendor can degrade or disable them. Remotely, quietly, with little warning. That's not paranoia — it's a logical consequence of the architecture. A country in a geopolitical dispute with the US or China could find its own security systems manipulated or switched off by the same foreign entities that sold them. At that point, sovereignty is a word that applies on paper and nowhere else.

The political erosion is slower and harder to point to, but it compounds. AI surveillance systems that bypass domestic courts don't just create legal grey areas — they hollow out the institutions that are supposed to provide accountability. When citizens can see that the state's tools operate outside any framework they can challenge or even observe, trust breaks down. Not all at once, but steadily. The gap between official claims and operational reality becomes something everyone senses and nobody can formally prove.

What makes this difficult to reverse is that the incentives all run the wrong way. Dependent governments have little practical motivation to build alternatives. Vendors have every reason to deepen the dependency. Elites who negotiated these agreements often benefit from the opacity. And the populations most affected — surveilled, profiled, politically managed by systems they can't see — have the least power to change any of it.

Without serious structural intervention, this doesn't correct itself. The tools sold as safety infrastructure become the mechanism of a more permanent subordination, and the Global South's place in the AI order gets harder to renegotiate with every system installed and every data agreement signed.




Fig. 9. AI Surveillance

D. Breaking the Dependency Cycle Through Strategic Policy

The structural picture is bad. It's not fixed.

Countries that have made real progress on AI dependency share one trait: they stopped treating adoption as a procurement decision and started treating it as a sovereignty question. That reframe changes what gets negotiated. It changes what gets refused. And it changes which ministries are in the room when the contracts get drafted.

Procurement reform is the fastest lever. Governments that require open-source architectures, mandate independent algorithmic audits, and enforce data localization laws change the terms vendors have to accept if they want market access. Brazil's LGPD isn't perfect, but it forced foreign companies to adapt to Brazilian legal requirements rather than bypass them. The underlying principle is transferable: a large enough market can attach conditions, and vendors will comply because the alternative is losing the contract.

Domestic capacity building is slower and harder to sustain politically because the returns don't show up in one term. But the goal doesn't have to be building a domestic OpenAI. It's building enough local expertise that governments aren't completely reliant on foreign vendors to understand what they've purchased — and enough of a domestic ecosystem that alternatives exist when a vendor's terms become unacceptable. Regional coordination changes the negotiating math for smaller economies. No single country in Sub-Saharan Africa can face down Google or Huawei alone. Collectively, the picture is different. The African Union's Continental AI Strategy and ASEAN's digital frameworks are early attempts to aggregate that leverage. Neither has fully delivered, but the logic is correct and the alternative — every country negotiating individually — consistently produces worse outcomes.

India is the case worth studying for middle powers. It hasn't picked a side in the US-China competition so much

as positioned itself to matter to both — engaging with the Quad’s technology working groups and the iCET initiative while maintaining enough independence to extract real concessions from each direction. That’s not neutrality. It’s leverage, deliberately constructed and maintained.

None of this dismantles the existing hierarchy quickly. The interests sustaining these dependency cycles are well-resourced and have long time horizons. But the determinism that sometimes creeps into “digital colonialism” framing overstates the case. Countries have agency here. The more honest question is whether they’ll use it — or keep signing agreements that quietly sign it away.

IX. POLICY RECOMMENDATIONS: A BLUEPRINT FOR TECHNOLOGICAL AUTONOMY

Everything so far has been diagnosis. This part is about what to actually do.

The countries best placed to escape dependency aren’t necessarily the richest or most technically advanced. They’re the ones treating AI governance as a strategic priority — where these decisions happen at the level of heads of state and national security councils, not buried in procurement offices or delegated entirely to IT ministries. That positioning matters because the decisions that create dependency are political ones, made by people with power. Reversing them requires the same. What follows draws on what middle powers and regional blocs have actually tried — what moved the needle, what got captured by the same elite interests this paper has described, and what failed quietly and got relabeled. It isn’t a universal blueprint. Contexts differ too much for that, and anyone selling a one-size solution to the Global South’s AI problem is probably selling something else too.

The goal isn’t to match US or Chinese compute power. It’s to build enough domestic capacity, legal architecture, and regional coordination that governments stop being passive recipients of other people’s technical choices. To get into the room where the terms are set, rather than reading about them afterward in a licensing agreement.

That’s achievable. It hasn’t happened yet mostly because it’s been nobody’s urgent priority — and because the current arrangement, for all its asymmetries, is quite convenient for the people currently in charge of it.

A. *Enacting AI Sovereignty and Sovereign Procurement Strategies*

Procurement is where dependency gets formalized. It’s also where it can be interrupted.

Every time a government signs a contract for a closed AI system without demanding source code access, audit rights, or local data hosting, it’s making a choice — usually a quiet one, buried in technical specifications that receive almost no public scrutiny. Reversing that requires treating these decisions differently. Not as commercial transactions delegated to procurement offices, but as sovereignty decisions that belong at the level of national policy, with the legal architecture to match.

Source code access needs to stop being optional. Governments deploying AI in policing, border control, or public infrastructure need to be able to see how those systems actually work — not trust that they work, but verify it independently. That means procurement contracts legally requiring full access to source code, training data, model documentation, and APIs. It means domestic authorities having the right to audit, modify, or decommission systems without vendor permission or foreign update cycles they can’t control.

There’s a secondary benefit worth naming. Engineers and technical staff who work directly with real systems learn faster and build more transferable skills than those kept at arm’s length from closed platforms. Transparency requirements don’t just protect against manipulation — they accelerate the domestic capacity building that reduces dependency over time.

Data localization is the other half. Cloud-based AI systems routinely route data to overseas servers, where it becomes subject to foreign law regardless of where it was collected. The US CLOUD Act and China’s Cybersecurity Law both create compelled access obligations for companies under their jurisdiction. Governments without strict data localization requirements are effectively sharing their citizens’ biometric and behavioral data with foreign legal frameworks whether they intended to or not.

The models exist. The EU’s GDPR, Brazil’s LGPD, and India’s data localization rules aren’t flawless, but they establish the legal architecture that forces foreign vendors to adapt to domestic norms rather than bypass them. The principle isn’t complicated: sensitive public data on domestic servers, under domestic law.

None of this works without enforcement. Independent technology review bodies — with real authority to examine contracts, conduct algorithmic impact assessments before deployment, and block systems that fail human rights or constitutional standards — are what separates policy on paper from policy in practice. Saõ Paulo’s AI procurement rules, which require vendors to disclose training data and system inputs, show this is workable at the city level. Kenya’s Communications Authority is building similar standards nationally. Neither is a complete solution, but both demonstrate something important: you don’t have to wait for a perfect national framework to start making vendors accountable. You can begin with what you have and build from there.

B. *Building Public Compute and Digital Infrastructure*

Software gets most of the attention in AI policy debates. The hardware problem is just as serious and considerably harder to fix.

Training, hosting, and maintaining AI systems requires physical infrastructure — data centers, GPU clusters, reliable electricity, and access to the semiconductors underneath all of it. Countries that can’t provide that domestically are dependent on whoever can. Right now that means American or Chinese cloud providers, operating under legal frameworks the host country didn’t write and can’t easily contest. That’s not a neutral arrangement. A cloud provider subject to its home

government's legal demands is a potential point of leverage, and governments without alternatives have limited options when that leverage gets applied.

Compute needs to start being treated the way electricity is treated — as infrastructure too critical to leave entirely in foreign hands. The analogy isn't perfect, but it captures something real. Modern governance, public safety systems, healthcare records, financial infrastructure — all of it increasingly runs on compute. Treating that as a commodity to be procured from the cheapest available foreign supplier is the same category of mistake as outsourcing water treatment to a company subject to another country's laws.

India's AI Mission is the most concrete model currently available for the Global South. A \$1.25 billion commitment to build sovereign GPU capacity — targeting at least 10,000 GPUs — alongside plans to develop accessible pools of civic and environmental data for domestic researchers and startups. The goal isn't to out-compete Nvidia or Google. It's to ensure Indian researchers and companies can build and run models without routing everything through foreign infrastructure first. That's a narrower ambition, but it's a real one.

Public-private partnerships are unavoidable — the capital requirements for this kind of infrastructure exceed what most emerging economy governments can fund alone. But the terms are everything. PPPs structured badly just recreate dependency with a domestic logo on it. Structured well, they use public investment and regulatory leverage to pull private capital into infrastructure that retains local talent, generates local intellectual property, and stays under national law. The difference between those two outcomes lives entirely in who drafts the contract and what they're willing to insist on.

Then there's the minerals problem, which is where this gets openly geopolitical. AI hardware depends on rare earth elements. China controls roughly 60% of global extraction and about 85% of refining capacity. Countries in the Global South that sit on these deposits are currently exporting raw materials cheaply and importing finished chips expensively — the same extractive pattern this paper has described throughout, just applied to physical goods rather than data.

The Minerals Security Partnership offers a way to change that calculus. Countries with REE deposits can use plurilateral frameworks to attract investment in domestic refining and processing rather than just extraction. A country that supplies processed materials to AI hardware supply chains has more leverage than one that ships out unrefined ore. That leverage is real and largely untapped. Using it seriously is part of what building sovereign infrastructure actually requires.

C. Cultivating Open-Source AI Ecosystems

Local innovation is stifled and "vendor lock-in" is essentially enforced when proprietary, "black-boxed" AI systems from foreign multinational businesses are relied upon. Recipient states are unable to audit judgments, rectify biases, or change platforms to meet local demands since these imported systems conceal underlying algorithmic logic and source code. Rather, foreign suppliers are given operational authority and

are solely responsible for managing system fixes and remote upgrades. Emerging economies must use `\textbf{open-source AI mandates}` as a strategic instrument to democratize technology, develop domestic engineering talent, and ensure that AI systems represent local language and cultural realities in order to break this structural reliance.

- **Implement Open-Source Mandates in Government:** Legally mandating that all non-classified algorithmic tools used by the state be developed on open-source designs would drastically alter sovereign procurement practices. Contracts for procurement should require both local and international providers to give complete access to model documentation, source code, and Application Programming Interfaces (APIs). By requiring open architectures, governments enable local institutions to independently check systems for biases, modify software to comply with domestic constitutional safeguards, and do away with the need for distant updates managed by foreign entities. Adoption of open-source software also speeds up indigenous technical capability and breaks the loop of black-box reliance by enabling local engineering staff to train on and modify practical tools.
- **Fund Localized LLMs and Indigenous R&D:** National research funding and public-private partnerships must significantly fund indigenous R&D in order to counteract "digital colonialism," in which foreign algorithms educated on Western data determine local reality. In particular, nations have to give top priority to the creation of sovereign Large Language Models (LLMs) that are trained on regional languages, political environments, and cultural norms.
 - **Singapore** serves as a prime example: the city-state established an SGD 70 million national multimodal LLM initiative specifically intended to reflect the various cultures, languages, and contexts of Southeast Asia, solidifying its position as a regional technological hub, after realizing the strategic necessity of sovereign capabilities.
 - Similarly, **Brazil**, through its *Programa Brasileiro de Inteligência Artificial (PBI)*, has strategically targeted the development of a Portuguese-language LLM to preserve its cultural sovereignty and elevate its geopolitical discursive power across the Lusophone world. By funding localized AI models, states can mitigate the biases inherent in imported technologies while securing cognitive and cultural autonomy.
- **Create Innovation Sandboxes:** Overly rigid regulations can crush nascent domestic startups, while a complete lack of oversight invites exploitation. To foster a thriving domestic tech ecosystem, governments must establish dynamic "regulatory sandboxes". **These are controlled, state-supervised environments where local universities, tech enterprises, and startups can safely test and iterate new AI applications in real-world scenarios**

(such as smart agriculture, public health, or fintech) without being paralyzed by heavy initial bureaucratic burdens. The European Union has successfully pioneered the use of digital sandboxes and open data initiatives to nurture collaborative research while protecting citizen rights. Implementing similar innovation sandboxes in the Global South allows domestic developers to experiment with open-source tools to solve community-specific problems, ultimately bridging the gap between academic research and commercial deployment.

D. Leveraging South–South Collaborations and Regional Governance

Individual nations in the Global South often lack the economic leverage to negotiate favorable terms with multinational tech giants. To counter this asymmetry, states must aggregate their data markets and coordinate regulatory frameworks through regional blocs, presenting a unified front to foreign vendors.

- **Operationalize Regional AI Frameworks:** Governments should actively implement regional standards, such as the African Union’s Continental AI Strategy or the ASEAN Digital Data Governance Framework. These frameworks pool resources, establish shared ethical baselines, and create cross-border data governance protocols that increase collective bargaining power.
- **Establish Cross-Border AI Research Centers:** Form pan-regional AI labs to share technical standards, open-source toolkits, and computing infrastructure. This prevents repetitive spending and accelerates capacity-building across neighboring states.
- **Adopt “Swing State” Plurilateral Diplomacy:** Emerging economies should avoid exclusive alignment with either the U.S. or China. By engaging in plurilateral technology coalitions (like the U.S.–India iCET or the Global Partnership on AI), nations can secure technology transfers, co-develop hardware, and shape global norms without sacrificing strategic autonomy.

Ethical AI Positioning and Multi-Stakeholder Governance

To prevent AI from becoming a tool of digital authoritarianism or elite capture, policy must aggressively mandate ethical deployment. States must align their AI strategies with international human rights standards while establishing robust domestic accountability mechanisms.

- **Institutionalize Multi-Stakeholder Oversight:** AI governance cannot be left solely to security ministries or tech corporations. Implement a multi-stakeholder governance model that actively integrates civil society, local governments, technical communities, and human rights watchdogs into AI procurement and oversight processes.
- **Mandate Algorithmic Explainability and Redress:** Citizens must have the legal right to challenge automated decisions (e.g., in predictive policing or welfare distribution). Governments must enact algorithmic accountability laws that hold both domestic agencies and foreign

vendors jointly liable for harms caused by biased or malfunctioning AI deployments.

- **Adopt and Localize Global Ethical Standards:** Formally adopt international frameworks such as the UNESCO Recommendation on the Ethics of AI or the ISO/IEC 42001 AI management system standards. By embedding these ethical principles into domestic technical design, states establish a legally binding baseline for transparency, fairness, and human oversight.

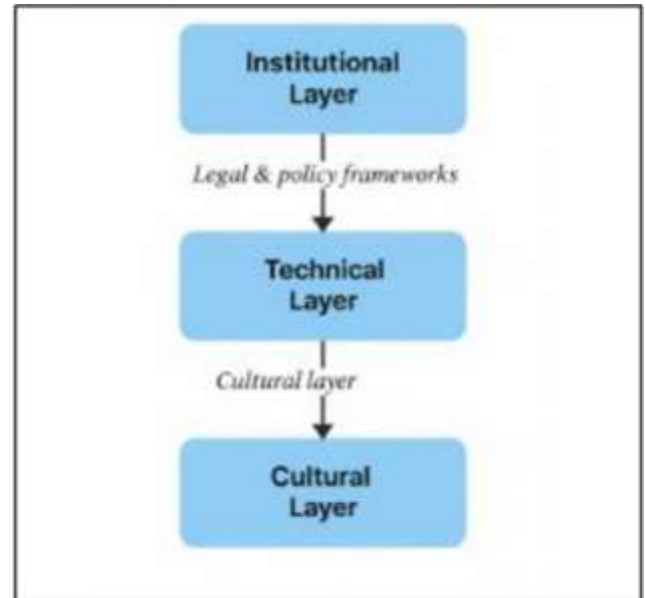


Fig. 10. Enter Caption

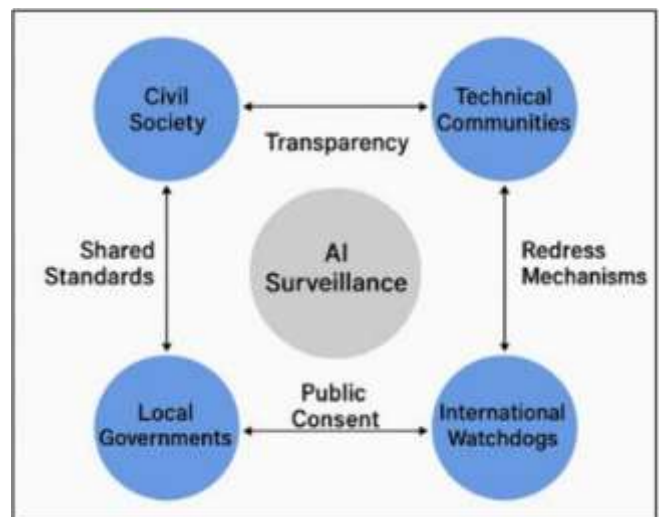


Fig. 11. Enter Caption

X. CONCLUSION

The rapid advancement of artificial intelligence has fundamentally restructured the 21st-century international order.

No longer confined to the realms of computer science or commercial enterprise, AI has evolved into a multidimensional geopolitical domain that actively reshapes how states project authority, secure their economies, and negotiate global norms,.

A. Summary of Key Findings and Answers to Research Questions

This study set out to investigate how foreign AI technologies reshape domestic autonomy, particularly for emerging economies in the Global South, and how these nations can reclaim their sovereignty in an asymmetrical digital world. The findings provide a clear answer: the importation of opaque, turnkey AI infrastructures actively erodes national autonomy through a systematic process defined in this paper as the **Algorithmic Dependence Model**.

The research demonstrates that global power is increasingly concentrated in a bipolar rivalry between the United States (championing a market-driven, defense-oriented innovation model) and China (promoting a state-directed, integrationist model). Because developing nations often lack the massive capital, domestic computing infrastructure, and critical material supply chains (like Rare Earth Elements) to compete, they are forced to import closed-loop foreign AI systems. This structural deficit leads directly to "source code blackboxing," offshore data hosting, and remote vendor lock-in. Ultimately, this dependency creates a new form of "digital colonialism," wherein the Global South is relegated to exporting raw behavioral data while importing biased algorithmic decisions, severely weakening democratic accountability and domestic legal oversight.

However, the study also concludes that this dependency is not inevitable. To answer the question of how states can break this cycle, the research highlights that emerging economies must adopt aggressive "technological statecraft." By implementing sovereign procurement strategies—such as mandating open-source architectures, enforcing strict data localization laws, and pooling resources through regional governance frameworks—middle powers like India can act as "swing states" to safeguard their algorithmic sovereignty and multipolarize the global AI ecosystem.

B. Final Insight: AI is Not Just Technology, But Power

The most profound realization drawn from this geopolitical analysis is that **AI is not just technology, but power**. In the contemporary era, algorithmic power represents an entirely new modality of sovereignty. Historically, international dominance was measured by territorial span, conventional military might, or the control of hydrocarbons. Today, authority is increasingly exerted through data accumulation, computational capacity, and the control of transnational digital infrastructures.

Algorithms are no longer neutral tools; they are active institutional actors that determine economic access, dictate military targeting in autonomous weapons systems, and shape political discourse. Control over the source code and the silicon that powers it effectively equals control over the

future. Multinational technology corporations that build these systems now exercise quasi-state functions, operating beyond traditional democratic borders. Consequently, the race for sovereign AI is not simply a contest for market share—it is an existential struggle over who will dictate the moral, legal, and security architecture of the human future,.

C. Future Research Scope

While this study provides a comprehensive framework for understanding algorithmic dependency, the rapid evolution of AI necessitates continued academic inquiry. Future research should expand upon the following critical areas:

1) · **Longitudinal Impacts of AI on Democracy: Evaluating Institutional and Civic Erosion:** Empirical, longitudinal studies are urgently needed to assess the evolving impact of foreign AI surveillance technologies on democratic structures in the Global South. While current literature highlights the immediate risks of importing turnkey surveillance solutions—such as facial recognition and predictive policing platforms—there is a critical gap in tracking the long-term, systemic erosion of domestic autonomy and democratic norms. Future research must longitudinally map how phenomena like "policy co-optation" and "elite capture" gradually subordinate domestic legal frameworks to the opaque operational defaults of foreign vendors. By systematically tracking these adoptions over a multi-year horizon, scholars can empirically document how the outsourcing of critical security and administrative functions leads to the progressive hollowing out of institutional independence, bypassing traditional judicial review and parliamentary oversight.

Furthermore, research must systematically evaluate the prolonged impacts of AI on **electoral integrity**. The integration of AI into public administration and political intelligence often equips incumbent regimes with advanced tools to monitor political rivals, preempt peaceful protests, and suppress opposition. When combined with the broader proliferation of AI-driven disinformation campaigns, hyper-realistic deepfakes, and targeted propaganda, these technologies pose existential threats to democratic processes. Longitudinal studies are required to track how state-sponsored cyber interference and algorithmically amplified extremism manipulate public discourse and influence voter behavior across multiple election cycles.

Finally, empirical research must investigate the sociological effects of algorithmic dependence on **public trust and civic engagement**. As foreign technological control becomes deeply embedded in local politics, it actively drives a "Feedback Loop of Democratic Fragility," wherein unchecked surveillance leads to weakened accountability and severe civic trust erosion. Future studies should empirically test this feedback loop over time, examining how sustained exposure to opaque, biased, and non-auditable AI decision-making impacts societal cohesion and the perceived legitimacy of state institutions. Additionally, tracking the long-term effectiveness of civil society resistance movements—such as digital rights coalitions demanding algorithmic transparency and moratoriums

on foreign AI—will provide vital insights into how emerging economies can successfully defend human rights and reclaim their algorithmic sovereignty.

2) **Algorithmic Decolonization and Indigenous AI: Countering Digital Hegemony:** Future technical and sociological research must urgently prioritize “**algorithmic decolonization**”, a paradigm aimed at dismantling the extractive practices of digital colonialism and reclaiming technological autonomy in the Global South. Currently, developing nations are structurally forced to import AI systems from Western and Chinese vendors. These foreign models are frequently trained on non-representative datasets from entirely unrelated geopolitical contexts, meaning they inherently encode foreign assumptions regarding “risk,” “normalcy,” and societal behavior. Consequently, these imported systems introduce severe algorithmic biases—often exacerbating discrimination across race, gender, and socioeconomic status—and fundamentally misalign with the local legal and cultural norms of the host nation.

To counter this algorithmic colonization, research must focus on the development of **Indigenous AI**. This involves deliberately designing “**lightweight, explainable, and interoperable AI tools suited for resource-constrained environments**”. Rather than relying on black-boxed, foreign-controlled Large Language Models (LLMs), emerging economies must cultivate AI systems explicitly rooted in local languages, historical experiences, and indigenous cultural values.

This movement is already taking shape through several key practical and strategic pathways:

Localized Language Processing and Datasets: Universities and startups in nations like Brazil, Kenya, and Indonesia are initiating AI research clusters specifically focused on ethical design and the creation of “local datasets”. A primary goal of these initiatives is to fill critical technological gaps by developing natural language processing capabilities for indigenous dialects, ensuring marginalized communities are not excluded from the digital economy.

Sovereign LLMs for Cultural Representation: Technologically capable nations are actively working to consolidate their linguistic communities to resist the homogenization caused by dominant global AI platforms. For example, Brazil is pushing for a sovereign Portuguese-language LLM to enhance its discursive power and cultural sovereignty, while the United Arab Emirates is developing Arabic models.

Regional AI Ecosystems: Singapore has launched a flagship national AI initiative explicitly aimed at creating “multimodal and localized LLMs that reflect the context and values of Southeast Asia’s diverse cultures and languages”. By tailoring AI technologies to regional realities, Singapore aims to foster deeper collaboration and build a unified technological ecosystem that accurately represents Southeast Asian diversity. Ultimately, advancing algorithmic decolonization is not merely a technical challenge; it is a profound sociological and geopolitical necessity. By investing in indigenous AI, nations in the Global South can **disrupt the extractive paradigms of surveillance capitalism, mitigate the harmful biases of for-**

eign algorithms, and successfully reclaim their algorithmic sovereignty and cultural representation in the digital age.

3) **AI Diplomacy and Normative Power: Exerting Soft Power in the Digital Age:** Future scholarship must prioritize the emerging field of “**AI diplomacy**,” which examines how states negotiate technological standards, ethical norms, and strategic collaborations at bilateral and multilateral levels. In an international system that is heavily influenced by the bipolar technological rivalry between the United States and China, middle powers and emerging economies face the structural risk of marginalization. However, future research should explore how these nations utilize AI diplomacy at multilateral forums—such as the United Nations (UN), the Organization for Economic Co-operation and Development (OECD), and the Global Partnership on AI (GPAI)—to actively shape global governance and exert soft power.

Normative power theory can be extended to analyze how states exert international influence through the dissemination of ethical standards and regulatory frameworks, rather than relying on traditional coercive or material capabilities. The European Union currently exemplifies this approach, translating its regulatory frameworks into tangible soft power that shapes global market behavior and expectations regarding privacy, accountability, and human rights. Future studies should investigate how emerging economies—acting as vital “swing states”—are increasingly adopting similar strategies to assert their geopolitical agency. For example, nations like India actively participate in plurilateral coalitions to promote inclusive, rule-based international resource governance, effectively bridging the gap between advanced economies and the Global South. Similarly, countries like Brazil and Singapore are leveraging localized AI models and digital strategies to enhance their regional influence and discursive power.

Furthermore, researchers should empirically track how middle powers leverage these international platforms to negotiate the **operationalization of ethical guidelines**, such as the UNESCO Recommendation on the Ethics of AI or the OECD AI Principles. By advocating for algorithmic transparency, equitable technology transfer, and democratic accountability at the multilateral level, emerging economies can use AI diplomacy to counter digital colonialism and secure their own digital sovereignty. Understanding these diplomatic dynamics is essential for anticipating how middle powers will negotiate the moral and institutional architecture of the digital age, and how AI will continue to reshape normative expectations, societal trust, and international cooperation frameworks.

4) **The Intersection of Centralized AI and Web 3.0:** Future geopolitical scholarship must analyze the paradox between the centralizing, state-driven nature of sovereign AI and the socially-driven, decentralized architecture of Web 3.0 (e.g., blockchain), examining how states will govern both paradigms simultaneously to project power. Ultimately, artificial intelligence demands a continuous reassessment of classical international relations theories. As AI systems grow in scope and autonomy, our geopolitical, ethical, and regulatory tools must evolve symmetrically to ensure that technological progress

aligns with global stability, equity, and human security,.

REFERENCES

- [1] A. Roy, "AI, rare earths, and the geopolitical algorithm: Strategic intersections of India, China, and the U.S. in the 21st century tech race," *International Journal of Politics and Media*, vol. 4, no. 1, pp. 36–44, 2025.
- [2] S. G. Chari, "Power, pixels and politics: The geopolitics of emerging technologies in the digital age," *London Journal of Research in Humanities & Social Science*, vol. 25, no. 2, 2025.
- [3] C. Colther, J. P. Doussoulin, and G. Tontini, "Artificial intelligence and global power dynamics: Geopolitical competition, strategic alliances, and the future of AI governance," Preprint, n.d.
- [4] C. Yilmaz, "The geopolitics of artificial intelligence: Power, regulation, and global governance," *Journal of Academic Social Studies - JOSSH*, vol. 2, no. 3, 2026. doi: 10.29329/ufusobed.2026.1410.2.
- [5] S. F. Abiade, "Algorithmic sovereignty and the new security dependencies: How foreign AI surveillance technologies reshape domestic autonomy in the Global South," *World Journal of Advanced Research and Reviews*, vol. 27, no. 2, pp. 162–180, 2025. doi: 10.30574/wjarr.2025.27.2.2845.
- [6] Z. Wang, "Generative AI-making and state-making: Sovereign AI race and the future of digital geopolitics," *Politics and Governance*, vol. 13, Art. no. 10222, 2025. doi: 10.17645/pag.10222.
- [7] M. JamalZai and C. Fei, "The geopolitical impact of artificial intelligence: Exploring AI's role in shaping global power structures and international relations," *The Critical Review of Social Sciences Studies*, vol. 3, no. 4, pp. 2244–2263, 2025.
- [8] R. Tiwari, "The geopolitics of digital sovereignty and artificial intelligence (AI)," *Episteme: An Online Interdisciplinary, Multidisciplinary & Multi-Cultural Journal*, vol. 14, no. 3, pp. 66–77, 2025.
- [9] A.-I. C. Popescu, "The geopolitical impact of the emerging technologies," *Bulletin of "Carol I" National Defence University*, pp. 8–21, 2021. doi: 10.53477/2284-9378-21-38.
- [10] K. Osadcha and N. Shumeiko, "Geopolitics, artificial intelligence and the global economy," 2024.
- [11] G. Allison, *The Geopolitics of Artificial Intelligence*. Academic Press, 2021.
- [12] N. Bostrom, *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, 2014.
- [13] F. H. Cardoso and E. Faletto, *Dependency and Development in Latin America*. University of California Press, 1979.
- [14] K.-F. Lee, *AI Superpowers: China, Silicon Valley, and the New World Order*. Houghton Mifflin Harcourt, 2018.
- [15] M. O'Sullivan, *The Levelling: What's Next After Globalization*. PublicAffairs, 2020.
- [16] K. N. Waltz, *Theory of International Politics*. Addison-Wesley, 1979.