AI multilateralism: How artificial intelligence can be the infrastructure for new forms of international cooperation

Anonymous Author(s)
Affiliation
Address
email

Abstract

Twentieth-century international cooperation rested on rules-based multilateralism. That system is fraying. We propose extitAI multilateralism: build shared, co-owned AI infrastructure—compute, data, models, evaluation, applications—and govern it together. A recent cross-border MVP shows this is feasible. We outline near-term pilots in climate-aligned trade, nature valuation, and biosecurity, and a pragmatic 36-month path for a small coalition of middle powers and Global South partners to deliver them.

1 Introduction

On September 2, 2025, a coalition of AI labs and cloud computing providers located across the globe did something refreshingly practical: they pooled compute to ship a public-use AI product. You can try it here. It has the same feel as ChatGPT; but a different DNA. Led by Swiss AI (with Swisscom and CSCS), AI Singapore, Amazon Web Services, and Australia's NCI, the five-week MVP shows many actors can build a shared AI stack—compute, data, models, applications—to tackle challenges that no single actor can address alone.

This achievement is important in a world where the foundations of international cooperation are in flux. Twentieth-century international cooperation came to be defined by multilateralism: treaty-based bodies such as the U.N., the Bretton Woods institutions, and the World Trade Organization where three or more states coordinated behavior under generalized rules. While far from perfect, these venues privileged negotiation, reciprocity, and the pursuit of "win–win" outcomes. In the twenty-first century, that order is fraying. Great-power rivalry, veto politics, and structural inequities have turned many forums from problem-solving arenas into sites of paralysis or coercion. Development finance and humanitarian aid are retreating. Meanwhile the cast of actors has widened: cities, firms, philanthropies, and standards bodies now shape cross-border outcomes.

AI itself lays bare the challenges. A handful of tech superpowers in Silicon Valley and Shenzhen are racing to dominate the infrastructure and operating systems of tomorrow's global economy, leaving most other actors with a stark choice between dependency and irrelevance. Without a viable "AI third way," middle power nations in particular are increasingly exposed to risks, including geopolitical coercion and vendor lock-in, foregoing opportunities for capture and redistribution of AI's benefits, and inability to proactively manage AI's emerging environmental and social impacts.

Because so many actors face a similar set of risks, there is an opening for a different kind of cooperation: building together in order to govern together. If twentieth-century multilateralism privileged negotiation and norm-setting, an AI-era variant would privilege building common infrastructure for shared challenges—and then subjecting that infrastructure to joint governance. In concrete terms,

Submitted to the NeurIPS 2025 Workshop on Algorithmic Collective Action. Do not distribute.

this means a co-owned AI stack and co-develop applications. All with transparent rules of access, audit, and accountability. Call it "AI multilateralism."

1.1 Building as the new negotiation

An "Airbus for AI" is a useful, if imperfect, analogy for what is possible. In the 1970s, European governments pooled capital, talent, and industrial policy to create an aircraft manufacturer capable of competing with America's giants. A coalition of middle powers could do the same in AI: pool compute from national research labs and local cloud providers to produce models, data pipelines, and products that rival hyperscaler ecosystems on performance while wiring the common good into their governance. A club model would give participating countries voice and upside while dynamic equity stakes based on active contributions (of energy, compute, and talent) and usage (of applications).

Proving the value of the approach will come through building applications that address global challenges that no single actor can solve alone. The first could be a green energy trade transparency platform. By integrating satellite data, grid telemetry, facility-level disclosures, and customs records, a shared platform could generate auditable estimates of embedded emissions in traded goods such as steel, cement, hydrogen, and critical minerals. Lower due-diligence costs, harmonized carbon-accounting methodologies and subsidy rules, and better financing terms would follow for producers that can prove low-carbon credentials. Because the algorithms and evaluation standards are co-owned, no single actor can weaponize the data.

A second application is a global nature-valuation engine. Federating existing state-of-nature assets likeICARUS and SEED and linking to locally-generated knowledge graphs like RESTOR. Scale up of Interspecies Money could compensate farmers, fishers, and Indigenous communities for measurable ecological outcomes while keeping sensitive data local. Shared standards would curb greenwashing without forcing one-size-fits-all disclosure.

A third is biosecurity early warning. Using privacy-preserving techniques to federate national biosecurity systems would allow continuous scanning of pathogen signals—from genomics and wastewater to wildlife sensors and syndromics—and would support joint dashboards and response playbooks for earlier, coordinated action. Pathogens cross borders; detection and response should do so too.

Executing capabilities at this scale requires more than hardware and code; it requires an applied science of collaboration. Economists shaped twentieth-century policy by offering rigorous tools for allocating scarcity. Today's problems demand equally rigorous methods for designing institutions, incentives, and applications for shared action. A "CERN for Collective Intelligence" could develop, test, and disseminate these methods for AI use cases, just as European particle physics leapt forward when laboratories were pooled.

1.2 Where to begin

The initial MVP hints at how to start. An initial group of five to seven middle powers, supported by philanthropic capital, could fund a three-year, 50–150 million runway—modest by AI-infrastructure standards. A small "Airbus AI" team would negotiate pooled-compute agreements across public and private providers; open-source a federated inference stack optimized for public workloads; launch the three pilots above with clear service-level targets and governance charters; and seed an applied research unit—a CERN-for-CI—for human–AI problem solving and institutional design.

Existing strengths can be bundled rather than invented from scratch. The European Union could pool leading national AI labs while leading safety, evaluation, and data governance. Countries like Australia and Canada can leverage abundant green power to expand public provisioning of data centers, with compute carve-outs for public-interest and sustainability workloads. Japan and Korea can specialize in edge compute and lead model pipelines for manufacturing and logistics. Switzerland or Singapore can provide neutral administration and legal interoperability and steward public data trusts and audit programs. A Global South anchor—such as Indonesia, Mexico, or Kenya—can lead biodiversity and health data trusts, local deployments, and community benefit-sharing mechanisms.

AI multilateralism will fail under familiar pressures: if pooled compute and data are treated as bargaining chips rather than shared utilities; if governance bodies are captured by a few states or vendors; if procurement, IP, and liability rules choke joint builds; if the energy footprint cannot be

decarbonized; and if Global South partners are cast as implementers rather than co-owners. The minimum conditions for success are equally plain: binding, multi-year pooled-compute and data-trust agreements; open interfaces and model licenses with safety guardrails; an independent, well-resourced evaluation commons with authority to halt deployments; green-power procurement at scale; and entry-and-exit rules that deter both free-riding and lock-in. Absent these, the project should not proceed.

Three years from now middle powers could still be lamenting the demise of rules-based order while watching AI giants ossify geopolitical fault lines. Or they could be pursuing a new path: one that transforms multilateralism from a 20th century relic to a 21st century competitive advantage.