frontiers in Political Science

TYPE Original Research
PUBLISHED 29 July 2026
DOI 10.3389/fpos.2026.1888614

OPEN ACCESS

EDITED BY
Alexandru Bodislav,
Bucharest Academy of Economic Studies, Romania

REVIEWED BY
Helder Ferreira Do Vale,
Federal University of Bahia (UFBA), Brazil
Ahmed Shaker Alalaq,
University of Kufa, Iraq

*CORRESPONDENCE
David F. J. Campbell
david.campbell@univie.ac.at

RECEIVED 22 May 2026
REVISED 01 July 2026
ACCEPTED 06 July 2026
PUBLISHED 29 July 2026

CITATION
Sadegh-Zadeh S-A, Kello L, Véliz C, Gordon N and Campbell DFJ (2026) Is a new world order taking shape in the age of AI? Mapping the global distribution of AI readiness and its implications for diplomacy and global order.
Front. Polit. Sci. 8:1888614.
doi: 10.3389/fpos.2026.1888614

Is a new world order taking shape in the age of AI? Mapping the global distribution of AI readiness and its implications for diplomacy and global order

Seyed-Ali Sadegh-Zadeh[1], Lucas Kello[2], Carissa Véliz[3], Neil Gordon[4] and David F. J. Campbell[5]*

[1]Department of Computing, School of Digital, Technology, Innovation and Business, University of Staffordshire, Stoke-on-Trent, United Kingdom, [2]Department of Politics and International Relations, University of Oxford, Oxford, United Kingdom, [3]Institute for Ethics in AI, University of Oxford, Oxford, United Kingdom, [4]Department of Computer Science, Faculty of Science and Engineering, University of Hull, Kingston upon Hull, United Kingdom, [5]Department of Political Science, University of Vienna, Vienna, Austria

Artificial intelligence (AI) is increasingly seen as consequential for power, polarity, and diplomacy, yet direct evidence for its effects on international outcomes remains limited. Using the 2023 Government AI Readiness Index (GARI; 193 countries) and four complementary sources (IMF AIPI, selected Stanford AI Index signals, UN EGDI, and broadband access), we first establish that these indices converge on a common latent construct of national AI capacity, with high inter-index concordance (Cronbach's α = 0.93; first principal component ≈ 76% of variance). We then map how this construct is distributed globally and how it associates with proxies for governance and diplomatic capacity. Cross-sectional results show strong concordance across indices and positive associations between readiness and governance/diplomacy proxies. As directional checks, we estimate short panel fixed-effects models and a staggered event study around countries' first national AI-strategy adoption; effects are small and statistically uncertain, so we interpret all patterns as associations rather than causal impacts. Consistent with this evidence, we frame AI readiness not as a demonstrated driver of present diplomatic influence but as an indicator of latent capacity: it identifies a geographically dispersed yet highly capable group of states that hold the institutional, infrastructural, and organisational foundations most likely to enable influence within future AI-governance regimes. Readiness is broadly distributed across regions rather than concentrated in a single hegemon, consistent with a multi-regional configuration of capability atop a concentrated compute and cloud stack. Policy implications follow: governments should treat AI-readiness (skills, infrastructure, data governance, innovation capacity) as a foreign-policy priority and engage in inclusive, interoperable global governance (e.g., OECD, UNESCO, GPAI, UN) to avoid widening divides. Diplomats should integrate AI tools ethically and transparently while contributing to norm-setting. We report overlap-controlled sensitivities (given partial construct overlap) and align vintages where needed (EGDI biennial). Avoiding AI exceptionalism, we situate AI within a lineage of general-purpose technologies while noting its breadth, dual-use character, network effects, and rapid diffusion.

KEYWORDS

AI readiness, latent capacity, AI governance, data governance, digital government, diplomacy, global governance, multipolarity

Introduction

The diffusion of artificial intelligence (AI) – machine learning systems for perception, prediction, and generation deployed in organisations – might transform international relations (Narayanan and Kapoor, 2024) and may even end the state-centric global order in a post-international order, changing the concept of polarity as we have grown accustomed to it. States that swiftly adopt and govern AI may enjoy strategic advantages in advanced analytics, automation, and innovation, whereas those that lag can lose leverage on the world stage. AI capabilities, from cutting-edge algorithms and computing infrastructure to large data repositories and skilled talent, are becoming key elements of national power (Ding, 2024) (less acknowledged are the risks that AI-heavy countries might be incurring, but we will leave that to the side for the time being).

One of the contributions of this paper is conceptual. We define AI capabilities to mean the inputs and enablers that a state controls: 1-compute and connectivity (hardware, cloud, broadband), 2- data assets & data governance, 3-models, software, and tooling, 4-human capital and organisational know-how, and 5-institutions that can integrate AI into workflows. Because AI is a tool, its strategic salience depends on how it is used, in public service delivery, intelligence and security, defence, economic planning, diplomacy and strategic communication, rather than in the mere possession of the technology. Indeed, the history of technological revolution in security affairs – from the rise of the modern submarine and mechanized warfare to cyberspace – shows that a new powerful tool badly governed can be more harmful to national interests and international order than not possessing it at all (Kello, 2022; Kello, 2017).

These inputs, and the readiness indices that proxy them, should be kept analytically distinct from the outcomes they may enable. We treat AI readiness, state capacity, governance effectiveness, diplomatic influence, and international power as related but separate constructs. AI readiness denotes the stock of enabling inputs (skills, infrastructure, data, and institutions) available for AI adoption. State capacity refers to the broader administrative ability to formulate and implement policy. Governance effectiveness concerns the quality, accessibility, transparency, and accountability of public-service delivery, which we proxy with EGDI and e-participation. Diplomatic influence is the ability to shape outcomes, coalitions, and norms in international fora, and international power is the aggregate capacity to affect others' behaviour across military, economic, and technological domains. These dimensions are correlated but not interchangeable: high AI readiness does not automatically confer diplomatic influence, and strong digital governance does not by itself generate geopolitical leverage. The links among them are conditional and, as our temporal analyses show (Supplementary Appendices A7, A8), far from immediate. Maintaining this separation is essential to interpreting our results, which speak most directly to readiness and governance capacity and only indirectly, and associatively, to diplomatic influence and power.

Leading scholars of security and international relations now identify AI leadership as a central dimension of 21st-century power competition. In particular, the United States and China are often cast as rivals in a global AI race that encompasses military, economic, and diplomatic domains (Williams, 2025). Crucially, however, a broad range of other states are also rapidly building AI capacity. Many advanced European and Asian economies, as well as small states, are making strategic investments in AI governance, research, and infrastructure. These trends are consistent with movement toward a multilayered AI landscape, with multiple states building meaningful capabilities rather than a strict U.S.–China duopoly. The diffusion of AI power reflects a broader power transition in the international system: the shift from U.S. hegemony toward multipolarity (Zarif and Karimi, 2023). The emergence of multiple regional powers increasingly challenges the dominance the United States has maintained during the past three decades. Yet China is not the sole contender. Others, too, are rising, potentially challenging not just America's ambitions but also China's. In such a fragmented order, multiple states, blocs, non-state actors or networks of varying actors would wield significant AI capabilities, each contributing to global decision-making and innovation. While our evidence points to multi-regional readiness, the underlying hardware and data stack is concentrated. Advanced logic chips, specialised tooling, and hyperscale cloud remain controlled by a small number of firms and/or jurisdictions. Many countries counted among "AI-capable" states therefore depend on imported compute, cloud access, or foreign IP, exposing them to chokepoints and weaponised interdependence.

Our analysis should thus be read as dispersion of enabling conditions and governance capacity, not a claim of universal domestic control over data centres or semiconductor manufacturing. In this paper, AI readiness is not assumed to be good in itself. We evaluate it against public-interest benchmarks: 1-legality and rights (e.g., privacy, equality of opportunity), 2-accountability and transparency, 3-fostering democratic participation (e-participation and avenues for contestation), and 4-effectiveness and access in service delivery. When we refer to "success in the AI era," we mean realising public value under these constraints. Our empirical proxies (e.g., EGDI and e-participation) reflect service quality, access and participatory features. Will the AI revolution reinforce existing great-power disparities or allow new actors to rise? How will AI change the practice of diplomacy? And what new institutions or norms are needed to govern this technology? To consider these issues, this study addresses three research questions:

> *RQ1:* How is AI readiness distributed globally?
>
> *RQ2:* How does AI readiness associate with proxies of diplomatic capacity and governance?
>
> *RQ3:* What do these patterns imply for emerging global AI governance and institutional realignment?

We address these questions using the 2023 Government AI Readiness Index (GARI; 193 countries) and four complementary sources: the IMF AI Preparedness Index (AIPI), selected signals from the Stanford AI Index, the UN E-Government Development Index (EGDI), and World Bank broadband data. We treat GARI as a country-level readiness proxy (0–100) spanning government capacity,

technological innovation base, and data infrastructure. We then examine (1) the global dispersion of readiness, (2) its associations with diplomatic and governance proxies, and (3) implications for global AI governance. A multi-index design reduces dependence on any single measure, and alignment rules (e.g., carrying forward EGDI biennial data) ensure year-matching.

We do not presume that AI is categorically different from earlier technologies of state power. From the printing press and the telegraph to electrification, nuclear technologies and cyber capabilities, technological change has repeatedly reconfigured administrative capacity as well as tools of interstate coercion and diplomacy. We understand AI as part of this historical continuity. At the same time, several features plausibly amplify its consequences: its general-purpose nature (tooling many sectors rather than a single domain), dual-use applications, network externalities, and unusually fast diffusion. We examine where these attributes correlate with contemporary governance and diplomatic proxies without assuming novelty *a priori*.

We remain cautious about causal claims regarding whether high AI readiness drives diplomatic power or simply correlates with it. Our core analysis is correlational (conditional associations), supplemented by two event-study designs: a short panel country fixed-effects regression (2019–2023) and a staggered difference-in-differences around the adoption of each country's first national AI strategy. These designs test if increases in AI readiness tend to precede improvements in governance indicators (e.g., UN EGDI, broadband access), using country and year fixed effects with clustered errors. We interpret all quantitative results conservatively as consistent associations [following Angrist and Pischke (2009) and Hernan and Robins (2020)], supported by robustness checks (pre-trend tests and placebo checks) in Supplementary Appendices A7, A8.

Our contributions are threefold. First, drawing on five complementary indices, we map a distinctly multipolar geography of AI readiness, identifying a widely distributed cluster of high-capability states rather than a single hegemon. Second, we document robust cross-sectional associations between AI readiness and diplomatic/governance outcomes; short panel and event-study diagnostics indicate plausible directionality but no uncontroversial causal effects, consistent with a deliberately cautious interpretation. Third, we derive policy implications for diplomacy and global AI governance, emphasising capacity-building for lagging states and inclusive norm-setting to ensure that the benefits of AI are shared broadly and equitably. We provide a triangulated, uncertainty-first map of AI-readiness multipolarity and its correlates. Our findings can inform cooperation priorities.

We are deliberate about what this evidence can and cannot support. Our strongest and most robust findings concern the measurement and global distribution of AI readiness; our temporal analyses (Supplementary Appendices A7, A8) yield only small and statistically uncertain short-run effects on governance and diplomatic proxies. We therefore do not claim that rising readiness causes immediate gains in diplomatic influence. Instead, we interpret AI readiness as an indicator of latent capacity, through a three-step argument: (i) existing AI-readiness indices converge on a common underlying construct (Cronbach's $\alpha = 0.93$; first principal component $\approx$ 76% of variance); (ii) this construct identifies a geographically dispersed but highly capable group of states; and (iii) these states hold the institutional, infrastructural, and organisational capacities most likely to enable influence within future AI-governance regimes. This framing preserves the paper's substantive relevance for diplomacy and global order while remaining fully consistent with our evidence and with our uncertainty-conscious analytical stance. Where we discuss diplomatic and institutional change, we do so as plausible, capacity-conditioned trajectories rather than as established short-run effects.

The paper is organised as follows. The data and methods section describes our data and methods, including construction of a composite AI signal from Stanford's index. The results section presents our empirical findings. First, we map the global AI readiness landscape. GARI scores reveal that AI capacity is widely distributed across continents, suggesting emerging multipolarity. We then analyse how this distribution intersects with global order and diplomacy. We discuss how a multipolar AI readiness pattern could democratise innovation and stabilise power, but also introduce risks (new competition, inequality, increased surveillance). We discuss how AI may be reshaping the practice of diplomacy, enabling data-driven strategy, advanced communication, and public outreach, while raising ethical challenges if not properly managed (e.g., transparency, bias).

The results section then examines institutional realignment: how major powers and blocs are articulating differing AI governance models, and how new multilateral initiatives (from OECD guidelines to UN advisory bodies) are taking shape to address AI's global impact.

The discussion provides an assessment of the results and their implications. We distil actionable policy recommendations for governments and international organisations, emphasizing that enhancing AI readiness must be treated as both a domestic and a foreign policy priority, which includes investing in digital infrastructure and AI education, closing the global AI divide through cooperation, and actively shaping international AI norms (via bodies like the OECD, UNESCO, and GPAI).

The conclusion notes that AI readiness is becoming an increasingly important dimension of state capacity and international influence. Countries with foresight and capacity stand to gain influence, but all nations, together with development-cooperation actors (multilateral development banks, UN agencies, and regional funds), must adapt to manage AI's risks and opportunities collaboratively.

Data and methods

Data sources

Our primary measure of national AI readiness is the 2023 Government AI Readiness Index (GARI) (Oxford Insights/IDRC), which covers 193 governments on a 0–100 scale. This index evaluates a country's AI preparedness across key pillars: (1) government capability, (2) technology sector strength, (3) human capital, (4) data availability, and (5) infrastructure. To reduce reliance on any single indicator, we triangulate GARI with four complementary indices:

1 IMF AI Preparedness Index (AIPI): a macro-level index (0–100, latest year) covering 174 + countries, which aggregates indicators of digital infrastructure, human capital, and policy environment for AI.
2 Stanford AI Index (selected signals): we construct a composite from country-level signals in the 2023 Stanford AI Index report, including private AI investment (USD), share of global AI publications, an AI talent proxy (skills availability), and a count of national AI policy initiatives. Each component is

standardized (z-score per year, with higher = better), and we average them to form a composite score for each country.

3 UN E-Government Development Index (EGDI): an index of public-sector digital capacity (0–1 scale) covering 193 UN member states. It comprises sub-indices for online services, telecommunications infrastructure, and human capital. Since EGDI is released biennially (most recently in 2022), we carry the latest value forward by one year (e.g., using 2022 for 2023) for alignment with other 2023 data, while preserving the original values for analysis.

4 World Bank WDI: Fixed Broadband Subscriptions, a basic digital access metric (subscriptions per 100 people, latest year available, ~2021–2024), included as an additional proxy for digital infrastructure.

Each source captures different aspects of "AI readiness" or related capacity. Supplementary Appendix A1 summarises these data sources and coverage. Supplementary Appendix A2 details how each indicator maps to theoretical readiness components (human capital, infrastructure, innovation, and governance). In brief, GARI and AIPI include governance/policy dimensions alongside technology, Stanford's signals emphasize R&D outputs and talent, EGDI focuses on government digital service capacity, and WDI broadband provides a narrow infrastructure gauge. Together, they offer a broad view of the factors that may position a country for success in the AI era. In this study, digital governance effectiveness means delivering public services that are timely, accessible, and reliable while enabling participation, transparency, and accountability. Operationally, we use as proxies EGDI (and its OSI/TII/HCI subcomponents) and e-participation, acknowledging that these capture service capacity, access and participatory affordances rather than the full spectrum of rights.

We harmonised country naming and coverage across datasets and merged them for analysis. We restricted most analyses to countries with data available in all key indices to ensure consistency. This method yielded a core cross-sectional sample of approximately 149 countries (Supplementary Appendix A3 provides a detailed table of data availability by country). The sample spans all world regions and income levels, though a few very small states and conflict-torn countries are excluded due to missing data (e.g., some indices omit certain microstates or countries like North Korea for which reliable data are unavailable). Crucially, the data cover the majority of players in international relations. Where an index was missing for a country, that observation is omitted from analyses requiring that index (we use listwise deletion within each model). Supplementary Appendix A3 illustrates the coverage for each index (GARI, AIPI, Stanford, EGDI, and Broadband), showing that most countries (including all major states) have complete data for 2023.

Panel construction and analysis design

In addition to cross-sectional analysis, we construct a short panel dataset (2019–2023) to explore temporal patterns. We align indices by year: GARI is annual (2017–2023), but AIPI and Stanford signals correspond to 2023 (with some historical values available for a subset), and EGDI is biennial. For panel regressions, we use lagged GARI as the key independent variable (to represent prior AI readiness) and examine subsequent changes in outcomes. Specifically, we align EGDI and other outcome measures to the year after GARI (e.g., using GARI_2022 to predict EGDI_2023) so that we can test if higher AI readiness precedes improvements in those outcomes.

We employ country fixed-effects (FE) regressions on the panel, which control for all time-invariant country characteristics, and year fixed effects to absorb any global trends. This stringent approach tests within-country changes. Do countries that improve their AI readiness see above-average gains in outcomes in the subsequent year? Given the short time frame and possible data lags, we interpret these results cautiously as directional checks rather than definitive causal identification. We cluster standard errors by country to account for serial correlation. After aligning and merging all available data, the raw panel spans up to 899 country-year observations across 193 countries over 2019–2023 (Supplementary Table A7.1). Because the lagged-GARI specification drops the first available year and applies listwise deletion on the aligned outcome measures, the estimation sample for the main model is roughly 495 country-year observations from about 177 countries (see Supplementary Appendix A7 for coverage details).

We also implement a staggered difference-in-differences (DiD) event study to examine one specific policy milestone: the adoption of a country's first national AI strategy. We identified the year each country first implemented a formal AI strategy (if at all) from public sources. Using those as 'treatment' events, we estimate event-time coefficients for up to 3 years before and after adoption, with untreated countries serving as controls (Sun and Abraham, 2021; Callaway and Sant'Anna, 2021). This approach, detailed in Supplementary Appendix A8, includes country and year fixed effects and clustered errors. It tests whether, on average, diplomatic or governance outcomes improve post-strategy adoption (relative to pre-trends). Again, given the short post-adoption windows and potential confounders, these results are treated as suggestive robustness checks rather than proof of impact.

All quantitative analyses use outcome measures that proxy countries' diplomatic and governance capacities in the international arena. In practice, we focus on outcomes for which global data are available and that plausibly relate to a country's influence or governance performance. The primary outcome highlighted is the UN E-Government Development Index (EGDI). Although EGDI is a domestic digital-government metric, it is tightly associated with broader governance quality and external engagement proxies. UN DESA's 2024 Survey reports strong correlations between EGDI and WGI-adjacent outcomes (e.g., corruption perceptions), as well as FDI inflows; higher-EGDI countries also show stronger public-sector performance and income, even with notable outliers. These patterns suggest that digitally capable states are better positioned to participate in, and influence, multilateral digital-governance processes (UN, OECD, MDBs) (Linhartova, 2022). In line with our frame, we read higher EGDI (and e-participation) as the capacity to deliver services and enable participation, not as an endorsement of any surveillance-heavy or rights-eroding use of AI. However, it should be noted that there tends to be a correlation between the use of machine learning and surveillance, given that the former depends on processing large quantities of data, much of which is often personal data. Where democratic risks are salient (e.g., opaque automated decision-making), we discuss guardrails and oversight in the Discussion and Policy sections.

We also examined a political stability index (WGI's Political Stability and Absence of Violence, PV.EST) and an e-participation component of EGDI (which captures citizen engagement in digital governance), as these relate to governance robustness and inclusive policy, factors that could influence diplomatic capacity. In addition,

for the event study, we consider EGDI as a representative outcome, reasoning that improvements in domestic digital governance could eventually extend to greater diplomatic clout. All continuous variables are inspected for outliers; in robustness checks, we winsorize extreme values at the 1% level, with no qualitative change in results (Supplementary Appendices A4, A5).

Because indices use different scales (GARI 0–100, EGDI 0–1, etc.), we often report standardized coefficients for comparability. Where raw coefficients are presented, we scale them in interpretable units (e.g., per 10-point increase in GARI, or per 0.1 increase in an index on 0–1 scale). This allows a rough comparison of effect sizes across measures. For panel FE models using log-transformed GARI (to treat a percentage increase in readiness), coefficients can be read as semi-elasticities: a coefficient β on ln(GARI) implies that a 1% increase in GARI is associated with a $\beta \times 0.01$ change in the outcome (approximately), and a 10% increase corresponds to $\beta \times \ln(1.10) \approx 0.0953\beta$. We will interpret such results in plain terms alongside the statistical estimates.

Measurement validation

We perform several checks to validate that our indices capture a common underlying concept of "AI readiness" or capacity and that our findings are not artefacts of measurement choice. Firstly, we compute pairwise correlations and rank concordance among the five indices. These are reported in Supplementary Appendix A4.

In summary, we find Pearson correlation coefficients between GARI and the other measures ranging from $r = 0.74$ to $r = 0.89$ (all $p < 0.001$), indicating strong linear associations. Spearman rank correlations are similarly high (0.72 to 0.87), showing that country rankings are broadly consistent across indices. For example, GARI correlates $r \approx 0.86$ with IMF's AIPI and $r \approx 0.89$ with EGDI. The Kendall's W coefficient of concordance across all five indices is ≈ 0.81, denoting high agreement in the joint ranking of countries. Furthermore, treating the five indices as items in a scale, we find a Cronbach's $\alpha \approx 0.93$, indicating excellent internal consistency (the implied average inter-item correlation is very high). A principal components analysis shows that the first principal component explains ~76% of the total variance, while the remaining components each contribute relatively little. In plain terms, all these indices appear to be tracking the same latent 'AI capacity' factor: whenever one index rates a country highly, the others tend to as well (and vice versa). This justifies our use of GARI as a representative measure, and it also motivated us to construct a combined index (L_{AIR}) as the first principal factor of GARI and AIPI (we use this latent AI readiness factor in some regressions to enhance signal-to-noise). High concordance also means that our results should be robust to using alternative readiness measures, a point we will demonstrate in the robustness section.

Secondly, we verify that our key findings are robust to alternative index choices. Instead of GARI, we re-run the main cross-sectional association tests using AIPI, the Stanford AI composite, EGDI, and broadband in turn as the 'readiness' predictor. The outcomes remain the same general type (diplomatic/governance proxies). Supplementary Appendix A5 presents a comparative summary of these results. In all cases, a higher AI readiness score (however measured) is associated with significantly higher values of the outcome. The effect sizes (standardized) range from about 0.45 to 0.66, all significant at $p < 0.001$. This pattern aligns with expectations: indices that emphasize governance (GARI, EGDI) show the strongest associations (around 0.62–0.66 SD differences between leader countries and others), while a pure infrastructure proxy (broadband access) still shows a positive but somewhat smaller effect (~0.45 SD). In short, our conclusions do not hinge on one specific AI index.

Thirdly, we conduct outlier and influence diagnostics to ensure that results are not driven by a few atypical countries. We calculate variance inflation factors (VIFs) for multicollinearity and Cook's distance for influential observations in our regressions. These are reported in Supplementary Appendices A1, A4 (for multicollinearity). The largest VIF in the baseline multivariate model is under 3, indicating no concerning multicollinearity among our regressors (which include region dummies in some specifications). Cook's distance identified a couple of high-leverage points (e.g., very low-readiness fragile states) but excluding them did not meaningfully change coefficients (Supplementary Appendix A1 shows a summary of influence diagnostics). We also ran rank-based (Spearman and Kendall) correlations as a robustness check on linearity, which yielded the same qualitative relationships (Supplementary Appendix A4).

Finally, because our key explanatory variables are indices that could contain measurement error or arbitrary scaling, we applied a simulation-extrapolation (SIMEX) procedure and Oster's δ for robustness to measurement error and omitted variables. In Supplementary Appendices A2, A7, we present these more technical checks. The SIMEX analysis (simulating increased noise in the AI readiness variable) suggested that any reasonable level of classical measurement error would if anything attenuate the observed association, not reverse it, meaning our significant findings are unlikely to be false positives due to noise. Oster's method (varying the assumed R^2max and comparing coefficient shifts with and without controls) indicated that unobserved confounding would have to be exceptionally strong (with $\delta > 1$) to explain away the readiness–outcome relationship; in fact, controlling for known determinants like human capital and infrastructure already explains most of the covariance (consistent with AI readiness working through those factors). In sum, these validation exercises strengthen confidence that the results reflect a real and generalizable pattern rather than an artifact of any single data source or model specification.

These results are decision aids, not eligibility thresholds. They must not be used to justify sanctions, exclusion from standards bodies, or conditional access to finance. We publish uncertainty bands, disclose construct overlap, and avoid league tables to reduce performativity risk.

Results

As throughout, we interpret empirical estimates as associations rather than causal effects. Panel FE and event-study results are treated as directional checks and remain small and statistically uncertain, consistent with our cautious approach.

Global AI readiness landscape

Per GARI 2023, the global AI readiness landscape is highly uneven. The United States sits at the high end of the distribution with an 84.8/100 score, reflecting strengths across all AI pillars (a world-class tech sector, robust digital infrastructure, and strong governance). Not far behind is Singapore (82.0), an exemplar of how a strategic small state can 'punch above its weight' by investing in digital

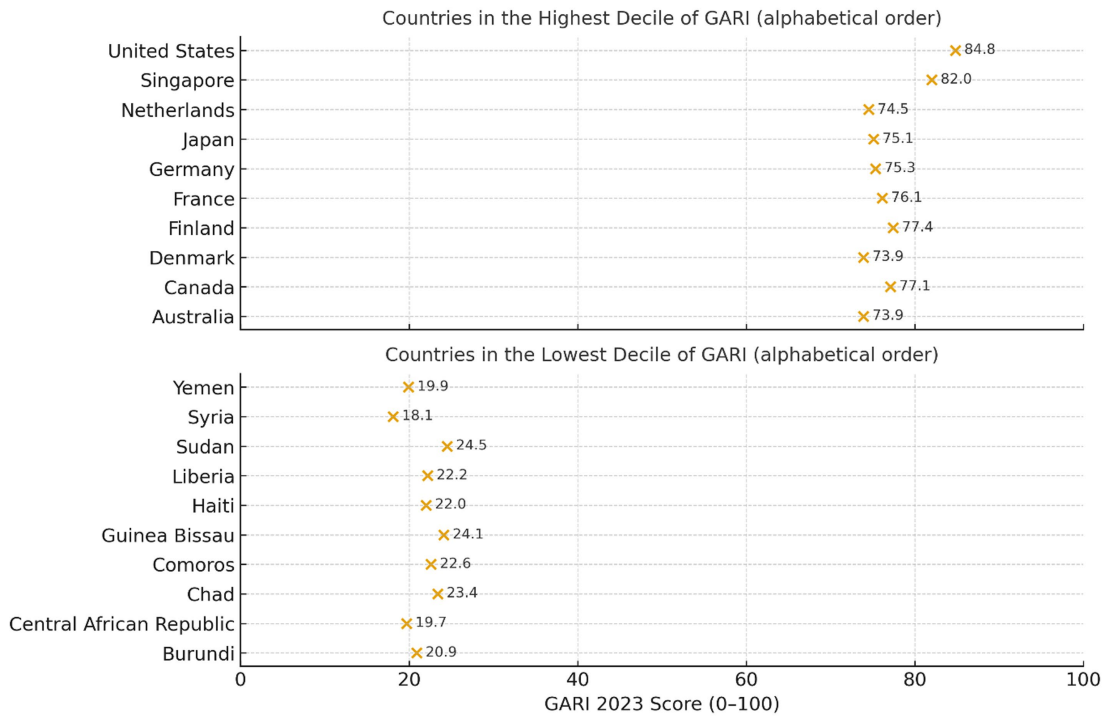

FIGURE 1
AI readiness across the highest and lowest cohorts of the 2023 Government AI Readiness Index. Values are descriptive and are not eligibility thresholds or prescriptive rankings. The figure displays the highest- and lowest-scoring cohorts of the analytic sample rather than the full sample, with countries ordered alphabetically within each cohort to discourage rank interpretation; full coverage for all countries is reported in Supplementary Appendix A3.

infrastructure and open data. Other upper decile countries span North America, Europe, and Asia; for example, the United Kingdom (~78.6), Finland (77.4), Canada (77.1), France (76.1), South Korea (75.6), Germany (75.3), Japan (75.1), and the Netherlands (74.5). The geographic spread of these leaders suggests AI readiness leadership is multi-regional rather than monocentric.

Readiness is multi-regional with wide dispersion. Figure 1 shows GARI 2023 scores for the highest- and lowest-scoring cohorts, displayed alphabetically within each cohort rather than by rank, emphasising the size of the gap over the ordering within it. The lower-decile cohort is concentrated among fragile and low-income states; the upper decile spans multiple regions. For example, bottom-of-the-ranking states suffer from weak electricity and internet infrastructure, minimal open data, and unstable institutions, conditions that make public-sector AI capacity virtually non-existent. By comparison, even the tenth-ranked country (Netherlands, ≈74.5) outscores the lowest-scoring states in our sample (e.g., Syria, ≈18.1) by more than 55 points. This stark gap underscores the risk of an 'AI readiness divide' wherein the most technologically advanced governments race ahead in leveraging AI, while the least ready countries fall further behind. Absent targeted capacity-building support, AI could thus exacerbate global inequalities, creating a new dimension of the digital divide. The dispersion we show reflects readiness indices (skills, governance, adoption, and access), not domestic control of compute and chip fabrication. In practice, model-grade compute (GPUs/TPUs), lithography/tool chains, and hyperscale data infrastructure are concentrated in a few countries and firms. Thus, several Global North leaders and most Global South improvers leverage external supply chains and the cloud. This is consistent with a multipolar user base operating atop a concentrated production base, which also explains why shocks in chip/tooling supply reverberate globally.

It is noteworthy that AI readiness only loosely tracks traditional power metrics like GDP or military strength. While there is overlap (e.g., the U.S. and major Western economies rank highly), some large economies do not make the very top tier, and some small countries outperform much larger ones. China, for instance, is often seen as an AI superpower and indeed ranks in the upper tier (GARI ≈ 70, around the top fifteen countries globally). However, it does not crack the very top group in GARI-2023, trailing several Western European and Asian economies. This is because GARI's holistic methodology penalises shortcomings in areas like governance transparency or human capital. China's strengths in research output and data are offset by relatively modest scores on indicators related to open government and ethical AI.

Conversely, several small states emerge as global leaders in certain aspects of AI readiness. For example, Estonia and Uruguay, neither of which is a great power in the traditional sense, rank among the leaders in civic AI readiness, thanks to their innovative digital government services and open data policies. Such cases demonstrate that targeted digital strategies can enable smaller countries to amplify their influence in specific domains like e-government or digital trade, even if they lack the scale of a superpower. In short, AI capability appears

multi-regional and not monopolised by any single bloc: we observe a distributed cluster of advanced states in North America, Europe, and Asia, rather than a simple U.S.–China duopoly. Evidence from AI-based PTA network mapping indicates a multi-regional configuration with European centrality, consistent with a dispersed capability landscape (Sadegh-Zadeh, 2025). This emerging landscape hints at a more multipolar configuration of power based on technological capacity. We also note that today's upper-decile cohort skews toward the Global North, with notable Global South exceptions in niche governance or service delivery; this pattern aligns with the current geography of compute and capital.

Association with diplomatic and governance outcomes

How does AI readiness relate to a nation's diplomatic clout or governance performance? We treat the patterns reported below as cross-sectional associations that characterise latent capacity rather than as estimates of causal effect; the directional checks in Supplementary Appendices A7, A8 caution explicitly against a causal reading. We explored this via cross-sectional regressions linking the readiness indices to various outcome proxies. In our baseline OLS model, we find that a one standard deviation higher AI readiness (as measured by GARI or the latent factor L_{AIR}) is associated with roughly 0.6–0.9 standard deviations higher scores on governance outcomes like EGDI (digital government effectiveness). For instance, in a bivariate regression of EGDI on L_{AIR} (AI readiness factor), the standardized coefficient is about 0.82 ($R^2 \approx 0.82$), indicating that countries with high AI readiness almost invariably have high e-government performance. In raw terms, a 10-point increase in GARI (out of 100) corresponds to about a 0.06–0.08-point increase in EGDI (on a 0–1 scale) on average, holding regional effects constant, a substantively large difference given that the global EGDI range is typically 0.3 to 0.95.

These associations remain positive and significant even after controlling for potential confounders like GDP per capita, population, and regional fixed effects. They also hold across alternative AI indices. Whether we use the IMF's AIPI, the Stanford AI composite, or broadband access as the predictor, countries in the top decile of AI readiness consistently score higher on governance/diplomacy outcomes than those outside the top decile (Supplementary Appendix A5). All estimated effects are positive of comparable magnitude ($\beta \sim 0.5$–0.7), and statistically significant ($p < 0.001$).

In plain language, this tells the same story from different angles: countries classed as AI 'leaders' tend to perform substantially better on governance proxies than other countries. For example, being an AI leader (upper decile) is associated with roughly half to two-thirds of a standard deviation higher EGDI score, all else equal. The probability of seeing such large gaps by chance (if no true effect existed) is extremely low given $p < 0.001$ across the board.

It is important to note that correlation does not prove causation. High AI readiness could both result from and contribute to strong institutions. However, the strength and consistency of these associations support the notion that building AI capacity goes hand-in-hand with building state capacity. Even controlling for general development level, the relationship persists, suggesting it is not merely a proxy for wealth. That said, we are looking only at a particular moment in time. It is possible that strong institutions led to the development of AI but that higher use of AI will lead to the weakening of institutions in the long run. It may not be a coincidence that we are seeing an erosion of democracy (e.g., according to the Economist's Democracy Index) at the same time as the rise of digital tech.

One of the authors of this paper has argued elsewhere that there is a fundamental tension between surveillance and democracy, and how AI currently depends on surveillance (Véliz, 2020). Without privacy, some of the most important pillars of democracy, such as anonymous peaceful protest, investigative journalism, lawyer-client confidentiality, and the secret ballot—collapse. Code also tends to be notoriously authoritarian and opaque, allowing users little leeway for negotiation or appeal. The case of China is a prime example of how it's possible to be highly AI ready and authoritarian, suggesting there is no necessary connection between AI readiness and strong institutions (in a liberal democratic sense of strength). Indeed, AI could in principle act as the serpent that eats itself, being a product of but undermining democratic institutions.

In our analysis, we checked whether any single region or outlier drives the result; it appears not. Removing, for instance, the United States or China, or regional dummies, still leaves a significant positive association. Nor is it driven solely by authoritarian vs. democratic governance styles; both the UAE and Estonia are AI leaders, for example, despite different political systems, implying the key commonality is investment in AI and digital infrastructure.

From a different perspective, we can ask, how much of the variance in diplomatic outcomes can we 'explain' with AI readiness? Readiness is strongly associated with EGDI; high concordance is expected given construct overlap, so we interpret this as descriptive alignment, not causal explanation (Supplementary Appendix A4 shows $R^2 \approx 0.82$ in the baseline model). This suggests that knowing a country's AI readiness level gives a very good prediction of its digital governance effectiveness. Of course, EGDI is itself partly a component of what makes up GARI (there is conceptual overlap in infrastructure and human capital elements), so high correlation there is expected. More interestingly, AI readiness also correlates strongly with independent measures like WGI Political Stability (PV.EST) ($r \approx 0.65$) and with foreign policy reach metrics (for those we have limited data, e.g., the number of diplomatic missions abroad tends to be higher in AI-ready states). The broader interpretation is that AI-ready countries have on average, more effective governments and greater global engagement capacity. They likely possess better bureaucratic institutions, more skilled human capital, and more resources to project influence, all factors that contribute to success in diplomacy.

To ensure we are not conflating AI readiness with general development, we ran a placebo test. We checked the association between future AI readiness and current outcomes (which should show no effect if our interpretation is correct). Indeed, when using 2023 GARI to 'predict' 2022 EGDI (one-year reverse), the relationship weakens and becomes statistically insignificant in a panel context (Supplementary Appendix A7), consistent with the idea that it is improvements in readiness that coincide with improvements in outcomes rather than the other way around. We also introduced lags and found that past values of readiness did not strongly predict future deviations in outcomes in fixed-effects models (the coefficients were near zero; see Supplementary Table A7.2 in Supplementary Appendix A7). These null findings in the panel (discussed below) suggest that while high-readiness countries have high outcome levels, a short-term increase in readiness does not yet show an immediate payoff, implying the relationship may play out over longer periods or be due to underlying third factors. In summary, AI readiness correlates with, but, in the short term, does not cause, higher diplomatic and governance

performance, which aligns with our theoretical expectation that building AI capacity is part of a broader state modernization process, but that for strong democracies to stay strong, they must govern AI and not allow AI to govern over them.

Implications for diplomacy and global governance

Our findings carry several implications for diplomacy and the evolving global order. First, the emergence of a multipolar, or even post polar, AI readiness landscape could democratise innovation and diffuse power more widely. Unlike in the 20th century, where cutting-edge technological capacity (e.g., nuclear technology) was confined to few superpowers, AI capabilities are proliferating across a diverse group of states and non-state actors, from G7 economies to smaller nations like Finland or Israel that excel in niche areas. This decentralization suggests an increasingly multilateral and networked international environment in which a larger set of actors can contribute to and shape AI-driven developments. In addition, AI capabilities are spreading well beyond the state domain. While states remain primary wielders of power, AI powerbrokers also include large multinational firms such as OpenAI, Google, and NVIDIA that can operate beyond the institutions of the state and sometimes against its wishes or interests. This marks a departure from earlier technological revolutions such as the tank or nuclear weapons, where inventive power was largely confined to governments, though it extends the trend of growing influence by private actors beginning with the expansion of cyberspace in the last two decades (Kello, 2013). In short, AI is driving not only a transition of power in the international system but also its diffusion beyond the state.

For policymakers, this trend is a double-edged sword. On the one hand, broader diffusion of AI know-how can prevent any single actor from dominating; it creates opportunities for middle powers, coalitions and networks of state and non-state actors to have a say in setting norms and standards. Many countries can pursue niche leadership; for example, a small state might specialise in AI ethics or data governance and become influential in those global discussions. Indeed, we already see initiatives like the Global Partnership on AI (GPAI) that include dozens of countries collaborating beyond the US-China binary. On the other hand, multipolarity without clear hierarchy can breed instability. With no single hegemon, or even state monopoly on AI advancement, there is less clarity on enforcement of norms or deterrence of misuse. Unlike the bipolar Cold War era—when superpower confrontation was often displaced into violent proxy conflicts across the Global South—the emerging AI- multipolar and networked landscape features more fluid alignments and less predictable competition (see Westad, 2005).

This convertibility of latent capacity into governance influence is, at present, observable descriptively rather than established causally. Participation and leadership in the principal AI-governance fora are concentrated among high-readiness states: the OECD AI Principles—the first intergovernmental AI standard and the template for the G20 principles—originated among, and are adhered to chiefly by, high-capability economies, and the Global Partnership on AI was launched by fifteen founding members drawn from the upper readiness tiers, with its expert centres hosted by Canada, France, and Japan, before its 2024 integration into the OECD expanded it to roughly forty-six states (OECD, 2024). The deliberate, ongoing effort to recruit low- and middle-income members is itself a marker of their current under-representation. This pattern is consistent with our latent-capacity interpretation: the dispersed-but-capable states identified by the readiness indices are also those most present in the institutions where AI norms are being written. We stress that this is an association, not evidence that readiness causes governance influence; a systematic test—modelling membership, leadership, sponsorship, and standards-body participation on readiness and its sub-components (the first and second pathways above)—is the natural next study, and one our framing is designed to enable.

How, more concretely, might such 'soft' capabilities translate into influence over the emerging AI architecture? Three mechanisms are plausible, and each is in principle testable. First, standard-setting authority: states with mature regulatory and ethical frameworks can export their templates, so that domestic rules become de facto international reference points—the trajectory already visible in the external reach of the European Union's data and AI regulation. Second, epistemic and institutional credibility: states with deep human capital and trusted institutions supply the experts, secretariats, and technical assessments on which multilateral bodies depend, converting expertise into agenda-setting and drafting power. Third, coalition-building and norm entrepreneurship: credible actors can convene like-minded states, sponsor resolutions, and help define what counts as legitimate practice. On this view, influence in AI governance need not track raw compute; it may accrue disproportionately to states that become authoritative in defining legitimate standards through expertise, institutional quality, ethical governance, and sustained participation in the OECD, UNESCO, GPAI, and UN processes. Our descriptive evidence is consistent with the first step of this logic—high-readiness states are already the most present in these fora—but whether soft capacity actually converts into rule-making influence, and under what conditions, is an empirical question we do not resolve here and to which we return in the conclusion.

Secondly, AI is deeply entangled in global supply chains and knowledge networks, making interdependence a defining feature of this multipolarity. No country is self-sufficient in AI. Advanced semiconductor chips come from cross-border value chains (e.g., design in the US, fabrication in Taiwan and Korea), and AI research is an inherently international enterprise with open science and talent mobility. This interdependence means countries are both empowered by and vulnerable to each other. For example, a nation strong in chip fabrication or data resources can hold leverage over others that need those inputs. This dynamic is evident in the current US–China tech rivalry.

The United States' curbs on China's access to advanced semiconductors illustrate weaponised interdependence in Farrell and Newman's sense: exploiting control over critical network hubs and chokepoints to coerce or constrain a rival. Specifically, U.S. jurisdiction over key nodes in the chip design–tooling–fabrication stack (EDA software, lithography inputs, and U.S.-origin technology rules) enables the extraterritorial reach seen in the October 2022 and October 2023 measures. In turn, China's seizing control of critical rare earth minerals necessary for building digital technologies and limiting their exports (Bradsher, 2025) can potentially nullify the multipolarity described here.

Conversely, this mutual interdependence also incentivises coalition-building. The U.S.–EU Trade and Technology Council has coordinated transatlantic approaches to AI risk management and standards; democratic coalitions such as Five Eyes and GPAI facilitate information-sharing and the diffusion of "trustworthy AI" practices. At the same time, critics note that Washington leverages

these same networks to align export controls and standards in ways that preserve a U.S. technological edge, illustrated by allied coordination on advanced-chip and tool restrictions, raising wider questions about how security logics shape ostensibly cooperative AI governance. Diplomats must therefore navigate a landscape where economic security and tech policy are integral to foreign relations. Traditional instruments, treaties, sanctions, export controls, and the revival of tariffs, are being redeployed in technology domains, where they face some limitations, while new forums are emerging to manage these issues.

Third, AI is transforming diplomatic practice. Internally, foreign ministries and international organisations are already using AI for horizon scanning, rapid briefings, and stakeholder communication; externally, delegations deploy AI to track negotiating positions, map coalition partners, and monitor norm diffusion across venues. These changes matter for our question because they mediate how 'AI readiness' translates into voice, coalition-building capacity, and rule-shaping influence in multilateral settings. Advanced analytics might help diplomats process big data; for instance, algorithms can analyse social media or satellite data to provide early warnings of crises. AI-driven language translation (e.g., real-time translation systems) is breaking down language barriers in international dialogue, potentially enabling more inclusive global discussions (Barrault et al., 2023; Costa-Jussà et al., 2022). During complex negotiations or crises, AI decision-support systems might assist diplomats in wargaming scenarios or identifying win-win solutions (Pokhriyal and Koebe, 2023). Additionally, states are increasingly using AI to shape public diplomacy, employing algorithms to target and tailor their messaging abroad, sometimes verging into propaganda or disinformation campaigns, and raising ethical concerns.

These innovations might make diplomacy more effective, and informed (or misinformed) by data. However, they also introduce ethical and security challenges. Issues of transparency and bias in AI outputs are critical; diplomatic decisions based on opaque AI recommendations could be problematic, and algorithmic bias might skew situational awareness (Bellaby, 2024). AI is mostly being built by a few companies for the purposes of profit, thereby making them inherently not neutral. There is also the risk of an AI-enabled information arms race, where states deploy bots and deepfakes to influence public opinion in other countries, undermining trust. Thus, while AI tools might enhance diplomats' capabilities, they must do so ethically and judiciously, developing guidelines for responsible use, and never allowing human beings to lose their place at the helm of the ship of our societies. International cooperation is needed to establish norms against malicious AI-driven information warfare, analogous to agreements on cyber warfare.

Finally, our results speak to global AI governance and institutional realignment. Similar to the history of Internet governance, the wide dispersion of AI capacity can mean that no single country dictates global governance of AI; it will require multilateral cooperation and inclusive norm-setting, though we expect states – especially large ones – to disagree about those norms, sometimes fundamentally. Although the United States and China are both going through phases in which they seem ill disposed towards international cooperation, encouragingly, we see nascent efforts in the opposite direction. For example, the OECD's AI Principles (adopted by dozens of countries) and UNESCO's Recommendation on AI Ethics are attempts to set international standards. That even smaller and middle-income countries have a stake (and capability) in AI suggests forums like the United Nations need to incorporate voices beyond the usual great-power circle when crafting AI regulations or guidelines.

It may be necessary to create new institutional mechanisms, perhaps an international AI Governance Forum, where countries can regularly meet to discuss AI's societal impacts, share best practices, and coordinate policies (Olugbade, 2025), while also integrating the input of relevant private actors. We might also see realignment of existing institutions. For instance, the UN Security Council has begun informal discussions on AI's implications for global security, and one could imagine an advisory body of AI experts to guide the Council. Different blocs are promoting competing models of AI governance. An open, human-centric model championed by the EU and like-minded democracies versus a state-controlled, security-centric model favoured by authoritarian regimes. These differences echo broader ideological divides, but there is also common ground in avoiding catastrophic misuse of AI (e.g., autonomous weapons production). How institutions evolve, whether through new treaties on AI in warfare or agreements on data-sharing for global good (climate, health), will shape the new world order.

Our data highlight that many countries and non-state actors have the capacity to be rule-makers, not just rule-takers, in AI governance. But it's not something that can be taken for granted or that will happen by default, especially given the dominance of a few countries in both software and hardware production. Diplomats from a broad spectrum of states should be proactive in these emerging negotiations to ensure the resulting frameworks are equitable and reflect diverse values.

Discussion

Interpreting our findings through a historical lens, we treat AI as one more technology through which states organise information, capabilities, and influence. The associations we document are consistent with long-run patterns whereby general-purpose technologies reshape administrative capacity and external projection. Where AI appears distinctive in our analysis is not in breaking history, but in the combination of breadth (multi-sector reach), dual-use, and the scale and velocity of effects that can compound across policy areas. We therefore avoid essentialising AI as sui generis and emphasise institutional choices and complementarities (skills, infrastructure, and governance) as the operative levers. Our findings support a view of readiness dispersion across multiple regions without implying dispersion of foundational inputs. The chip and computing stack and hyperscale data infrastructure remain concentrated, which conditions how 'AI capabilities' translate into influence. Countries can exercise voice through governance capacity, talent, and institutional adoption even while relying on external compute, a pattern consistent with interdependence rather than autarky. This apparent tension between dispersed readiness and concentrated infrastructure is, we argue, the organising fact of AI multipolarity rather than a contradiction within it. Following Farrell and Newman (2019), jurisdictional control over central network hubs, advanced logic fabrication, electronic-design-automation tools, and hyperscale cloud, generates "panopticon" and "chokepoint" effects that allow a small number of states to monitor and throttle the flows on which many otherwise-capable states depend. Readiness therefore measures the capacity to use and

govern AI, which is broadly distributed, whereas the capacity to produce its foundational inputs remains narrowly held; the two are analytically separable, and our dispersion results pertain only to the former. This distinction reframes 'multipolarity' as a multipolar user base operating atop an oligopolistic production base: most states are adopters and rule-shapers whose strategic autonomy is bounded by their position within these networks. It also explains why measures that capture skills, governance, and adoption can show wide dispersion even as material leverage stays concentrated, and why shocks to chip or tooling supply reverberate globally. For lagging and middle-power states, the practical implication is that readiness yields durable influence only when paired with strategies that reduce single-point dependence, through diversified and trusted compute access, coalition-building, and the retention of lower-tech and analogue fallbacks discussed below.

Two elements that can aid in achieving a more balanced geopolitical horizon are, first, democratic alliances that can grant more leverage to countries that would otherwise be too small to make a difference on their own. Second, using AI in a way that does not create complete dependence on it, such as combining it with other kinds of more readily available digital technology (e.g., spreadsheets) and, especially, retaining the power to run on analogue. The more individuals and institutions (companies and governments alike) retain analogue systems as either parallel systems or backup systems, the less dependent they will be on any one technology company, product, or country.

We recognise the concern that pervasive government use of AI can harm democratic structures through surveillance, discriminatory profiling, opacity, and automation bias. Accordingly, our use of "success" presumes democratic guardrails: lawful basis and proportionality, privacy and data protection by design, algorithmic transparency and auditability, human-in-the-loop for consequential decisions, impact assessments and public consultation (especially for vulnerable groups), contestability and effective redress, and independent oversight. Absent such conditions, higher "readiness" is not normatively desirable. Our empirical results should therefore be read as capability under constraints, not as carte blanche for expansive adoption. Whether AI readiness strengthens or erodes democratic governance is, accordingly, conditional rather than automatic. The same capabilities, large-scale data integration, biometric identification, predictive analytics, and automated decision-making, can either improve service delivery and citizen participation or enable mass surveillance, discriminatory profiling, and the suppression of dissent, depending on the institutional environment in which they are deployed. Readiness is democratically enhancing where it is embedded in enforceable rights protections, independent oversight, auditability, effective contestation and redress, and a free press, and democratically corrosive where these safeguards are weak or absent. China illustrates the latter pole: high readiness coupled with pervasive state surveillance demonstrates that technical capability and liberal-democratic strength are dissociable. Because our governance proxies (EGDI, e-participation) capture service capacity and participatory affordances rather than these rights conditions, a high score should not be read as evidence of democratic quality; capability and democratic safeguards must be assessed separately, and our policy recommendations are addressed to the case in which the latter are present.

Our analysis provides empirical support for several key insights relevant to policymakers and international relations practitioners. First, AI readiness is emerging as a foundational pillar of national power, alongside traditional measures such as economic size and military might. Countries that cultivate digital infrastructure, human capital, data capabilities, and AI innovation are positioning themselves to wield greater influence in global affairs. This is not merely theoretical; we see it in how closely AI readiness aligns with effective governance and international engagement. For diplomats and policymakers, this means that boosting national AI capacity is not just an economic or domestic goal but a strategic foreign policy imperative, as long as strong democratic governance guardrails are put in place. A country with strong AI and digital foundations will have more credible soft power (e.g., being seen as technologically advanced and governance-efficient) and potentially more hard power in future conflicts that might involve cyber and autonomous systems, as long as it retains a strong rule of law and autonomy in the form of not being dependent on any one technology.

Second, our finding of broad-based multipolarity in AI capacity suggests that the international system may avoid an overwhelming AI hegemony by one actor, but it also means states must coordinate to manage the technology's risks in an inclusive fashion. Without coordination, there is a risk of regulatory fragmentation (different rules in different jurisdictions) or a race to the bottom in AI ethics. Global governance mechanisms thus need to catch up quickly. The data show that dozens of countries have something to contribute. This underscores the importance of inclusive global governance. International organisations (UN, ITU, UNESCO, etc.) must accommodate the perspectives of emerging AI players, not only the usual powers. Global challenges like AI-driven disinformation, autonomous weapons, or algorithmic bias in international supply chains require collective action. Our work reinforces calls (Roberts et al., 2024) for strengthening global AI governance structures, whether through new treaties or expanding mandates of existing bodies. There is momentum in this direction (e.g., the recently proposed Global AI Safety Institute), and our findings add urgency; coordination failures could deepen the divide we identified, leaving weaker states even further behind.

Third, we acknowledge the limitations of our short-term causal analysis. We did not find evidence that increases in AI readiness immediately cause improvements in diplomatic outcomes (Supplementary Appendices A7, A8). The panel FE and event-study estimates were statistically indistinguishable from zero for the few years observed. However, this does not mean AI investments are futile for diplomacy; rather, it suggests that the benefits may take time to materialise or may be conditional on certain thresholds. It could be that governance improvements lag behind capacity-building, or that a conducive policy environment is needed to translate readiness into influence. It also reminds us that correlation is not causation. Countries strong in AI are generally strong across the board, so one must be careful not to attribute all their success to AI prowess alone. And we must be careful in continuing to monitor trends, lest an increase in the use of AI leads to a deterioration of democratic institutions. For practitioners, the lesson is to pursue AI advancement as part of a holistic development and diplomatic strategy, not a silver bullet. We advise treating our quantitative results as consistent with a

reinforcing cycle (AI readiness ↔ state capacity) rather than a one-way street.

Policy recommendations

In light of the above, we outline several actionable recommendations for governments and international organisations:

Make AI readiness a foreign policy priority

Governments should treat AI readiness as a core foreign policy priority. Building national strength in AI is not just a matter of industrial competitiveness or research excellence; it is also a growing basis of international power. Nations must explicitly incorporate AI capacity-building into their national security and diplomatic strategies. To compete effectively, states must invest not just in digital infrastructure and compute capacity, but also in the intellectual resources and institutional machinery that underpins diplomatic efforts. This means ministries of foreign affairs working closely with science and technology ministries to present a coherent agenda in international discussions and to ensure that the nation's external agenda aligns with scientific trends and domestic capacity.

Equally important is the cultivation of human capital within the diplomatic and foreign policy establishment. Diplomatic staff should be trained in tech diplomacy, and tech experts should be included in delegations to forums like the UN or WTO when AI issues are on the table. Dedicated training in AI diplomacy should be routine: officials must combine policy expertise with a working understanding of algorithms, compute economics, and platform governance. One way to achieve this is for foreign ministries to second officers to standards bodies or major technology companies – as the United States did in 2011 when it created the world's first dedicated office for cyber issues, under the leadership of Christopher Painter, within its Department of State (in this way, Painter became the world's first de facto cyber ambassador), or as Denmark did in 2020 when it designated Anne Marie Engtoft as "Tech Ambassador" to work closely with companies, including in the AI sector. A permanent AI policy unit within foreign ministries could coordinate positions and capacity across departments, providing rapid analytical support to delegations while ensuring that national interests are represented consistently in venues at the United Nations, OECD, UNESCO, and WTO.

Another important basis for national power in the AI realm is the design of a coherent national AI strategy. Among its main tasks is to clarify the scope of strategic priorities, identify mechanisms for their pursuit, and align foreign policy activity with the pursuit of security, economic, and social objectives.

The evolution of cyber diplomacy offers a cautionary parallel: it demonstrates the cost of institutional inertia and slow strategic adaptation. After the 2007 distributed denial of service attacks against Estonia's financial and governmental systems – a watershed incident that exposed the technical vulnerability of advanced digital societies to political hacking – many governments were slow to adapt conceptually and bureaucratically. Although the strategic importance of cyber threats was widely recognized, few states had yet published a national strategy to address them in a coherent programme of activity – even as their societies immersed themselves in digital systems. Britain and Australia were early movers, publishing their strategies in 2009, but others lagged behind (Germany and France in 2011). Even the first mover, the United States, which issued a national cyber strategy in 2003, lacked an integrative institutional plan linking technological risk to foreign and security policy. Muddled thinking marred strategy building. Germany's experience is instructive. The country's first national cyber strategy (from 2011) defined cyberspace narrowly as "IT systems linked at data level"; it explicitly excluded "isolated virtual systems." The omission was not merely semantic. By treating cyberspace as coterminous with the Internet, Germany's cyber strategy constrained its own scope of protection: it neglected critical infrastructures – such as the industrial controller at the Isar 2 nuclear power station or the BASF's high-hazard chemical facility in Ludwigshafen – whose compromise could have major strategic effects. The absence of a coherent cyber strategy undermined Germany's foreign policy and cybersecurity efforts. The country did not appoint a cyber ambassador until 2020 and binding obligations for critical infrastructure operators arrived only in 2015 and 2016 – later than other European countries like Britain.

The German cyber strategy episode provides a lesson for the AI era: an inchoate home strategy with a clear map of national interests can undermine foreign policy efforts. In the absence of precise conceptual frameworks and trained officials who understand them, governments risk misinterpreting the very technologies they seek to secure and govern. A credible AI strategy should be able to double as a foreign policy playbook setting external objectives: securing access to trusted clouds and compute, protecting critical data flows, shaping safety baselines in trade negotiations, leading in standards setting forums, and de-risking global supply chains. Alignment between the foreign policy establishment and domestic stakeholders should be routine rather than episodic: embassies, for example, should run standing consultations with national champions, cloud providers, universities, and civil society groups. Technology-sector attachés, or the more ambitious roving 'tech ambassador', working closely with large firms can give delegations the technical footing and the access to bargain from knowledge. Progress should be measured according to concrete foreign policy outcomes: positions submitted, votes won in standards committees, mutual recognition agreements, joint exercises held, reciprocal arrangements on compute access, etc. When the AI aspects of foreign policy are organised in this way – explicit interests, assigned instruments, engaged stakeholders, measurable outcomes – domestic readiness can turn into external influence.

Invest in digital infrastructure and human capital at home

Our findings highlight infrastructure and skills as core components of readiness. Nations, especially those lagging, should invest in expanding broadband access, reliable electricity, and computational infrastructure, and ensure these investments reach all residents, including rural, remote, and low-income communities (e.g., Indigenous and tribal lands), where internal digital divides persist, alongside education and training to develop AI talent and digitally literate bureaucracies. These investments not only spur economic growth but also enhance governance and a country's attractiveness as a partner. Development agencies and IFIs (International Financial Institutions) should support such projects in low-income states to prevent widening the global AI divide.

Close the global AI divide through international cooperation

Advanced countries and international organisations must support capacity-building in less-prepared states. This could include inclusive technology transfer initiatives, mentorship programs (pairing AI researchers and government digital teams from high-readiness and low-readiness countries), and funding mechanisms akin to climate finance for AI infrastructure. A concrete step could be establishing a Global AI Readiness Fund under the UN or World Bank to provide resources and expertise to countries with GARI scores in the bottom quartile. Without deliberate effort, the gap we identified will persist or widen, with destabilising consequences.

Promote inclusive norm-setting and governance frameworks

Diplomats from AI-leading nations should resist the temptation to form exclusive clubs; instead, they should champion inclusive processes (e.g., involve developing countries in drafting AI norms). All states have a stake in ensuring AI is used responsibly and does not threaten international stability. We recommend expanding the mandate of existing multilateral groups (like the OECD AI Policy Observatory or the GPAI) to interface more with developing world concerns and perhaps creating a UN high-level panel or special envoy on AI to harmonise efforts across agencies. New treaties or agreements may be needed in areas such as autonomous weapons (building on the Geneva discussions) and cross-border data governance.

At the same time, major barriers confront the normative enterprise of AI on the international scene. States diverge because their regime types embrace different first-order values and policy priorities. China's party led model subordinates AI to social stability and national security: model developers must register algorithms, submit security reviews, and ensure that they align with "socialist core values" amid tight content controls. By contrast, the European Union's "rights-first" model applies ex ante precautionary regulation. It bans "unacceptable risk" uses (e.g., social scoring by public authorities) and grounds enforcement in fundamental rights protection under its AI Act. The American market led model is largely sectoral and executive action driven: rather than resorting to state direction or binding rights obligations, it relies on voluntary standards and guidance (e.g., NIST's AI Risk Management Framework). Inevitably, these differences in regime type and political values will yield incompatible preferences on questions such as audits, transparency, compute controls, and citizen redress – thereby complicating universal norm setting beyond mere technical standards.

Here, too, the record of cyber norm building is instructive. After several rounds of intergovernmental discussion at the UN Group of Governmental Experts, states affirmed in 2015 – to much fanfare – that the UN Charter principles and international law apply in cyberspace. Yet such declaratory progress did not measurably stem interstate offensive conduct. Election interference, wholesale data theft, infrastructure destruction, and other harmful activity continued unabated. The implication for AI norm setting is clear: any declaratory consensus that emerges will still face the impediment, in practice, of divergent preferences on transparency, auditability, and redress – complicating the search for common enforced practice.

Temper the hype

Ensuring AI readiness does not mean surrendering to the hype in AI inspired by tech companies. AI is not the solution to every problem; most social and political challenges need social and political responses that have little to do with technology. AI is not even an adequate tool in many contexts, especially those in which it is an overkill and more simple, cheap, and reliable technologies would be more appropriate, and those in which fairness is more important than efficiency (Cordasco and Véliz, 2026). AI readiness must be advanced at the same time as strengthening alternative technologies, especially analogue ones.

Embed ethics and transparency in diplomatic use of AI

To reduce the risks posed by automated accounts and AI-driven analytics, states and international organisations should promote the adoption of non-use and transparency pledges (e.g., commitments to refrain from state-sponsored mass disinformation). They should also clearly label state-affiliated automated accounts, and commit to joint attribution and public call-outs of state and non-state actors that orchestrate manipulative campaigns. A practical channel is the G7 Rapid Response Mechanism (RRM), which already coordinates information-sharing and responses to foreign information manipulation; governments could formalise an AI-specific protocol within the RRM for rapid technical attribution, coordinated messaging, and proportionate countermeasures. Similarly, international cooperation on AI ethics and safety research (to prevent unintended harmful behaviour by AI systems) should be bolstered, recognising this as a transnational public good akin to pandemic preparedness.

By pursuing these steps, states can better manage the transition to an AI-influenced world order. The overarching theme is adaptation and cooperation. Just as nations had to adapt to the nuclear revolution through new doctrines and treaties, so too must they adapt to the AI revolution by updating institutions and forming new norms. The diplomatic community has a window of opportunity now, while AI governance is still malleable, to steer it in a direction that enhances global stability and shared prosperity.

Conclusion

AI readiness is becoming an increasingly important dimension of state capacity in the 21st century. Our study has shown that countries leading in AI capacity are often leaders in governance and poised to shape international norms, while those falling behind in AI risk losing influence and developmental ground. We are likely witnessing the birth of a new multipolar order where power is distributed not only by military or economic heft but also by technological and algorithmic capacity. Importantly, this AI-era multipolarity is more geographically and politically diverse than earlier epochs; it is not a simple bipolar or unipolar structure but a complex network of many capable states.

For the international relations community, these findings underscore the urgency of integrating technology foresight into strategic planning. Diplomats, security strategists, and development

experts must collaborate to ensure AI's diffusion does not lead to a zero-sum scramble or normative void. The history of nuclear technology taught us that unmanaged technological revolutions can threaten humanity's survival; AI, while different in nature, demands a comparable level of global attention to governance, ethics, and equitable access.

No country can afford to ignore AI in its foreign policy calculus. Smaller states have a chance to carve out niches and contribute disproportionately on the global stage, but they need investment and international support to do so. Major powers, on the other hand, bear a special responsibility to avoid G2 rivalry leading to technospheres of influence that marginalize others. A balanced multipolar AI world, buttressed by robust international institutions, could be one of empowerment, where AI is harnessed to tackle shared challenges and diffuse power, rather than one of entrenchment of existing divides.

These findings also define a clear agenda for future research. If AI readiness is best understood as latent governance capacity, the priority question is whether, and under what conditions, that capacity translates into actual influence over international AI rule-making. Future work could test this directly—modelling participation, leadership, agenda-setting, coalition-building, and norm entrepreneurship within emerging AI-governance institutions as functions of readiness and, in particular, of its government-and-ethics sub-components. A natural hypothesis is that states scoring highly on governance quality and ethical-regulatory capacity become disproportionately influential in shaping norms through the OECD, UNESCO, GPAI, and UN processes, even where they lack frontier compute. Complementary designs could trace regulatory diffusion (whether domestic frameworks are adopted elsewhere), the conditions under which rights-based rather than authoritarian models gain traction, and the extent to which chokepoint control over compute constrains the autonomy of otherwise capable states. Our framework and measures are intended to make this programme tractable: they identify the states with the latent capacity to lead, leaving the question of realised influence to research designs built for causal inference.

In closing, our analysis provides evidence-based reassurance that the AI revolution, while transformative, does not necessarily spell dominance by a single actor. There is room to shape a future where AI's benefits are broadly shared. Whether we achieve that future hinges on policy choices made today. We urge scholars and practitioners to treat AI readiness as a key variable in theories of international power and as a priority in diplomatic agendas. By doing so, and by cooperating across borders, the global community can guide AI's impact toward enhancing global stability and human welfare, fostering a new world order where technology serves as a tool for peace and progress, not conflict.

Limits and reflexivity

Construct overlap: Readiness and governance proxies share components; some associations are partly mechanical. We therefore report overlap-controlled sensitivities.

No short-run causality: Fixed-effects and staggered DiD checks show small, statistically uncertain short-term effects; read all results as descriptive associations.

Vintage alignment: EGDI (biennial) is carried forward to align with 2023; results include a vintage-hygiene sensitivity.

Performativity risk: Scores/maps can alter behaviour; we adopt a harm-use policy and avoid league tables.

Uncertainty-first visuals: We display quantiles/robustness bands.

Policy significance statement

This paper shows that national AI-readiness, human capital, digital infrastructure, data governance and innovation capacity, is broadly distributed across regions and is positively associated with digital-government and diplomatic proxies. Short fixed-effects and staggered event-study checks find small and statistically uncertain short-term effects, so we interpret the relationships as robust associations rather than causal impacts. For policymakers, two priorities follow: (1) treat AI-readiness as a foreign-policy objective, investing in skills, infrastructure, open data and accountable institutions; and (2) engage in inclusive, interoperable international norm-setting (e.g., OECD, UNESCO, UN, GPAI) to prevent widening divides. Practitioners should adopt AI tools ethically and transparently in diplomatic practice. Readiness today shapes voice and influence tomorrow; cooperation can turn dispersion into shared public value.

Data availability statement

All data used in this study are from publicly available sources: the Government AI Readiness Index (GARI, 2023; Oxford Insights/IDRC), the IMF AI Preparedness Index (AIPI), selected country-level indicators from the Stanford AI Index (private AI investment, share of AI publications, AI talent proxy, and counts of national AI policy actions), the UN E-Government Development Index (EGDI, 2022 carried forward to 2023), and the World Bank WDI Fixed Broadband Subscriptions. The appendices, including the full data-availability table by country, are provided as Supplementary material.

Author contributions

S-AS-Z: Conceptualization, Methodology, Software, Formal analysis, Data curation, Visualization, Investigation, Project administration, Writing – original draft, Writing – review & editing. LK: Conceptualization, Investigation, Writing – original draft, Writing – review & editing. CV: Conceptualization, Investigation, Writing – original draft, Writing – review & editing. NG: Methodology, Validation, Writing – original draft, Writing – review & editing. DC: Conceptualization, Supervision, Writing – original draft, Writing – review & editing.

Funding

The author(s) declared that financial support was not received for this work and/or its publication.

Acknowledgments

We thank colleagues at the Artificial Intelligence in Health Lab at the University of Staffordshire for their assistance with data collection, experimental support, and helpful discussions that improved this study.

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declared that Generative AI was not used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

Supplementary material

The Supplementary material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fpos.2026.1888614/full#supplementary-material

References

Angrist, J. D., and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton university press.

Barrault, L., Chung, Y.-A., Coria Meglioli, M., Dale, D., Dong, N., Duquenne, P.-A, et al. (2023). SeamlessM4T: massively multilingual and multimodal machine translation. *arXiv [Preprint]*. doi: 10.48550/arXiv.2308.11596

Bellaby, R. (2024). The ethical problems of 'intelligence–AI'. *Int. Aff.* 100, 2525–2542. doi: 10.1093/ia/iiae227

Bradsher, K. (2025). *Step by Step, How China Seized Control of Critical Minerals*. New York: The New York Times.

Callaway, B., and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *J. Econ.* 225, 200–230. doi: 10.1016/j.jeconom.2020.12.001

Cordasco, C. L., and Véliz, C. (2026). The accuracy-explainability trade-off, the right to explanation, and the implications for organisations. *J. Bus. Ethics*. doi: 10.1007/s10551-026-06305-9

Costa-jussà, M. R., Cross, J., Çelebi, O., Elbayad, M., Heafield, K., Heffernan, K., et al. (2022). No language left behind: scaling human-centered machine translation. *arXiv [Preprint]*. doi: 10.48550/arXiv.2207.04672

Ding, J. (2024). The rise and fall of technological leadership: general-purpose technology diffusion and economic power transitions. *Int. Stud. Q.* 68:sqae013. doi: 10.1093/isq/sqae013

Farrell, H., and Newman, A. L. (2019). Weaponized interdependence: how global economic networks shape state coercion. *Int. Secur.* 44, 42–79. doi: 10.1162/isec_a_00351

Hernan, M. A., and Robins, J. M. (2020). Causal Inference: What If. Boca Raton, FL: Chapman & Hall/CRC.

Kello, L. (2013). The meaning of the cyber revolution: perils to theory and statecraft. *Int. Secur.* 38, 7–40. doi: 10.1162/ISEC_a_00138

Kello, L. (2017). *The Virtual Weapon and International order*. New Haven, CT: Yale University Press.

Kello, L. (2022). *Striking Back: The End of Peace in Cyberspace — And How to Restore It*. New Haven, CT: Yale University Press.

Linhartova, V. (2022). The role of e-government in the evaluation of the quality of governance in the countries of the European Union. *Hrvatska i komparativna javna uprava* 22, 267–287. doi: 10.31297/hkju.22.2.4

Narayanan, A., and Kapoor, S. (2024). AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference. Princeton, NJ: Princeton University Press.

OECD. (2024). Global Partnership on Artificial Intelligence (GPAI). Paris: OECD.AI Policy Observatory. Available online at: https://oecd.ai/en/about/about-gpai (Accessed July 17, 2026).

Olugbade, O. (2025). In search of a global governance mechanism for artificial intelligence (AI): a collective action perspective. *Glob. Public Policy Gov.* 5, 139–161. doi: 10.1007/s43508-025-00113-z

Pokhriyal, N., and Koebe, T. (2023). AI-assisted diplomatic decision-making during crises—challenges and opportunities. *Front Big Data* 6:1183313. doi: 10.3389/fdata.2023.1183313

Roberts, H., Hine, E., Taddeo, M., and Floridi, L. (2024). Global AI governance: barriers and pathways forward. *Int. Aff.* 100, 1275–1286. doi: 10.1093/ia/iiae073

Sadegh-Zadeh, S.-A. (2025). Dynamics of global trade diplomacy: an artificial intelligence multi-dimensional analysis of preferential trade agreements. *J. Comput. Soc. Sci.* 8:63. doi: 10.1007/s42001-025-00393-9

Sun, L., and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J. Econ.* 225, 175–199. doi: 10.1016/j.jeconom.2020.09.006

Véliz, C. (2020). *Why and How you Should take Back Control of Your Data*. London: Bantam Press.

Westad, O. A. (2005). *The Global Cold War: Third World Interventions and the Making of Our Times*. Cambridge, UK: Cambridge University Press.

Williams, B. K. (2025). Winning the defining contest: the US-China artificial intelligence race. *Washington Q.* 48, 153–171. doi: 10.1080/0163660X.2025.2517501

Zarif, M. J., and Karimi, S. (2023). A world in transition: rethinking the current global order. *China Q. Int. Strateg. Stud.* 9, 429–457. doi: 10.1142/S2377740023500203