

2026

Who Governs the Intelligence Age?

AI, geopolitical
power and the
politics of who
gets to shape the
future

MUNICH SECURITY CONFERENCE 2026
REPORT



Who Governs the Intelligence Age?

AI, geopolitical power and the politics of who gets to shape the future

A strategic policy collaboration based on the Ahead roundtable held alongside the Munich Security Conference: “Artificial Intelligence as a Geopolitical Force: Rethinking Diplomacy and Governance in an Era of Disorder”

February 13, 2026 | European Parliament Liaison Office, Munich
Convened by Ahead, hosted by the European Parliament Liaison Office in Munich
In collaboration with Bitsight Technologies



FOREWORD: THE ROOM IS PART OF THE POLICY

Artificial intelligence is often discussed as though power begins with the model and ends with the rulebook. It does not. Power also sits in the institutions that define the problem, determine who speaks with authority and decide which ideas travel from conversation into policy, capital and implementation.

This paper emerges from the Ahead roundtable, “Artificial Intelligence as a Geopolitical Force: Rethinking Diplomacy and Governance in an Era of Disorder,” held on February 13, 2026, at the European Parliament Liaison Office in Munich alongside the Munich Security Conference. Convened by Ahead and hosted by Udo Bux, the exchange brought political, business, academic, technical, multilateral and next-generation perspectives into direct conversation. The collaboration with Esya Centre carries that exchange beyond the room, combining Ahead’s ability to assemble consequential influence with Esya’s empirical and policy expertise. The purpose is not to record an event, but to clarify the larger contest over who will govern the intelligence age.

A central premise runs through the analysis: representation is not a closing paragraph about inclusion, but a discipline of institutional judgment. The quality of global decisions depends on who is present early enough to shape the terms of debate. Bringing together Sergey Lagodinsky, Stephen Boyer, Dr. Markus Siewert, Dr. Wambui Kahara, Freddy Kita, Prof. Jacob N. Shapiro, Laura Galante, Giorgia Martínezgarnelo y Calvo, Brunello Rosa, Julia Dudenko, Dr. Vera Schmitt, Dr. Chenghao Sun, Arantza Caballero González-Rojas, Lucia de Luca and Girija Shivakumar placed legislative authority, cyber expertise, academic implementation, African statecraft, entrepreneurship, youth leadership and global convening power in direct exchange. That combination was not incidental; it supplied the discussion’s intellectual range and strategic depth.

What changes when those usually treated as audiences become authors of the agenda?

The answer is visible in the range of issues that surfaced together: Europe’s dependence on foreign compute; the danger of regulating technologies before implementation capacity exists; agentic security and Model Context Protocol vulnerabilities; the future of African digital infrastructure; youth participation; capital-market fragmentation; and the need to turn trust into a commercial advantage. A narrower room might have produced a cleaner consensus. It would almost certainly have produced a poorer diagnosis.



EXECUTIVE SUMMARY

Artificial intelligence is becoming part of the operating system of geopolitical power. It affects military intelligence, economic productivity, critical infrastructure, public information, scientific discovery and the capacity of states to make choices without external permission. The contest is not merely over which country builds the most capable model. It is over who controls compute, capital, cloud, chips, standards, talent and the institutions able to direct technology toward public purpose, regulate its risks and translate technical capability into durable strategic advantage.

The scale of the asymmetry
\$285.9bn in US private AI investment in 2025, more than 23 times China's recorded private total. Europe is responding with an InvestAI ambition of €200bn, including €20bn for up to five AI gigafactories. These figures are not simply investment statistics. They are measures of future strategic room to maneuver. [3][4]

The Munich Security Conference was therefore integral to the argument, not merely its backdrop. It is one of the few places where military risk, diplomacy, markets, technology and alliance politics are read as a single strategic picture. In 2026, its program explicitly addressed the risks and rewards of the AI race. In July 2026, the Financial Times reported that Europe's digital leadership now describes AI as a potential "geopolitical weapon," citing the vulnerability created by dependence on foreign models and infrastructure. [1][5]

The exchange at the European Parliament Liaison Office in Munich, convened by Ahead and hosted by Udo Bux, moved the argument from strategic recognition to democratic execution. Can Europe build the infrastructure and capital markets required to act while preserving the rule of law? Can it turn the "Brussels effect" from the extraterritorial reach of regulation into what the Member of the European Parliament (MEP) Sergey Lagodinsky called a "Brussels business effect," where trustworthy AI is profitable as well as compliant?

Ahead ensured that the exchange did not collapse into another transatlantic conversation conducted by a familiar circle about the rest of the world. Equal representation of women and men, combined with participation spanning European institutions, cybersecurity, universities, African leadership, multilateral governance, entrepreneurship and youth networks, altered the substance of the exchange. Its method rests on a demanding proposition: influence should be assembled for consequence, not optics. The right combination of authority, technical depth and lived experience can make a discussion more predictive, more legitimate and more useful.



The question is not only who writes the rules, but who is present when the purpose of the technology is decided. Representation is not decorative. It changes what risks are seen, what opportunities count and whose future is being designed.

GIRIJA SHIVAKUMAR | CO-FOUNDER AND CHIEF GROWTH OFFICER, AHEAD

Seven conclusions follow. First, sovereignty is the capacity to choose, not a fantasy of isolation. Second, Europe requires applied capability, commercial scale and strategic investment, not only regulatory distinction. Third, its principal weakness is the gap between rules and implementation. Fourth, in an era of disorder, adaptability and preparedness matter more than fixed forecasts or speculative assumptions about artificial general intelligence. Fifth, AI security is moving from model risk to system exposure. Sixth, African states must be co-authors rather than downstream adopters of global frameworks. Seventh, the composition of decision-making rooms is itself a governance variable.

1. WHY MUNICH SECURITY CONFERENCE MATTERS NOW

When does a technology policy become a security doctrine?

It happens when dependence can be converted into pressure. The contemporary security order is increasingly shaped by export controls, sanctions, cloud contracts, payment systems, satellite access, undersea cables, semiconductor equipment and data flows. AI sits across all of them. It can strengthen intelligence and logistics, automate cyber operations, alter the information environment and increase the productivity of both firms and states.

MSC matters because it brings together actors who usually encounter these questions in separate institutional silos: heads of government, ministers, legislators, defense and intelligence leaders, technology companies, investors, academics and civil society. Its value lies not only in prestige, but in strategic compression. Strategic consequences that might take months to move between institutions can be confronted in the same place, while alliances and assumptions are still fluid.



The 2026 context sharpened that function. Europe is navigating war on its continent, renewed pressure on defense spending, a more transactional United States, Chinese technological reach and the return of economic coercion. AI is becoming a test of whether Europe can remain a rule-maker without becoming a permanent technology-taker. The FT's warning about potential foreign "kill switches" is vivid because it translates a technical dependency into a political reality: access can be withdrawn, conditioned or priced by actors outside European democratic control. [5]

If access to intelligence can be switched off, is it still merely a commercial service?

The Wall Street Journal's reporting on the rise of lower-cost Chinese open-weight models adds another complication. Strategic dependence does not always arrive through expensive monopoly. It can also arrive through attractive, affordable tools that spread rapidly because they solve a commercial problem. The policy challenge is therefore not to close markets reflexively, but to distinguish resilience from protectionism and security from regulatory capture. [6]

This is why Ahead's role at Munich was substantive rather than ceremonial. Ahead works at the intersection of access, influence and representation across consequential global platforms. Its collaboration with Esya created a bridge between the density of a high-level exchange and the discipline required to turn it into a policy proposition. Ahead assembled forms of authority that rarely interrogate one another directly, then worked with Esya to translate the resulting insights into a framework capable of traveling beyond Munich.



2. THE POLITICS OF PRESENCE

Can an influential room still be intellectually provincial?

Yes. A gathering can contain eminent names yet remain narrow in the questions it permits. Prestige is not the same as range. The composition of influence must be designed with the same care as the agenda. Bringing in women, next-generation leaders and perspectives from African states is not an act of representational benevolence. It changes the strategic field.

Dr. Wambui Kahara's insistence that machine decision-making must retain clear human limits introduced a rights and accountability test. Honorary Minister Freddy Kita's emphasis on national AI strategies, financing and 2030 infrastructure moved the debate from European rulemaking to the realities of state-building. Giorgia Martínezgarnelo y Calvo's focus on youth participation challenged the assumption that future generations should be consulted only after systems are deployed. Dr. Vera Schmitt's argument for "applied law" tested whether European ambition could survive contact with technical systems. Stephen Boyer's warning about MCP vulnerabilities brought the discussion down to authentication, exposure and operational control.

These were not parallel contributions. They revealed how strategic choices interact. A regulation designed without operational cybersecurity may be legally elegant and technically hollow. An African AI strategy without finance may be ethically ambitious and politically inert. A technology agenda without youth participation may optimize near-term adoption while misunderstanding the generation that will live longest with the consequences.



AI policy must be context-sensitive. A one-size-fits-all approach does not work across countries and regions. Preparedness and adaptability matter more than fixed projections.

LUCIA DE LUCA | CO-FOUNDER AND CEO, AHEAD

The balanced participation engineered by Ahead demonstrated a larger proposition: diversity is most valuable before consensus forms. Equal representation of women and men mattered because it was embedded in a discussion about hard power, capital, cybersecurity and institutional design, not isolated in a side session about inclusion. That is the difference between visibility and influence.



3. EUROPE: SOVEREIGNTY WITHOUT STRATEGIC THEATER

What is sovereignty worth if it cannot be financed, deployed and scaled?

Europe has regulatory reach, advanced research, industrial expertise and a market of continental scale. It also has fragmented capital markets, uneven access to compute and deep dependence on non-European cloud and model providers. The risk is not that Europe fails to produce a European version of every American or Chinese company. It is that it loses the capacity to make independent choices in critical systems.

MEP Sergey Lagodinsky's phrase "strategic parallelism" is more useful than the language of full autonomy. Europe should preserve essential ties with the United States and China while reducing points of coercive dependence. That means identifying where redundancy is necessary, where shared infrastructure is efficient and where public procurement can create a durable European market.



"Europe should pursue strategic parallelism: preserve essential ties with the United States and China while building the technological sovereignty to act on its own terms."

SERGEY LAGODINSKY | MEMBER OF THE EUROPEAN PARLIAMENT

The numbers demonstrate the scale of the challenge. Stanford estimates US private AI investment at \$285.9bn in 2025. The Commission's InvestAI initiative seeks to mobilize €200bn, including €20bn for up to five gigafactories, and aims to triple EU data-center capacity over five to seven years. [3][4] Europe is not short of announcements. The question is whether capital, energy, procurement and market integration will move at the same speed as political intent.

The test is not whether Europe can predict every technological turn. It cannot. The more credible objective is institutional adaptability: building regulatory, industrial and investment systems able to learn, respond and reallocate capacity as evidence changes. In a period of geopolitical disorder, preparedness is a harder and more useful discipline than foresight presented as certainty.



Could Europe win the rulebook and lose the market?

Europe's innovation imperative

The European Union has positioned itself as a global regulatory leader through the AI Act and related digital-governance frameworks. That approach reflects a distinctly European conviction that technology must remain aligned with rights, accountability and privacy through a risk-based model. Europe's strength lies in benchmark-setting and rule-making. The roundtable was equally clear, however, that normative leadership does not automatically become geopolitical weight.

The central weakness is not a shortage of principles but the gap between regulation and implementation. Europe must strengthen innovation capacity, reduce market fragmentation, deepen access to capital in line with the Letta and Draghi recommendations and support firms as they scale across the single market. Without this, the EU could remain normatively influential yet strategically dependent in a field still driven by competition between the United States and China.

Europe should therefore treat AI policy as part of a wider strategic agenda linking regulation, industrial policy, secure cloud capacity, advanced digital infrastructure, energy and capital-market reform. It can lead in setting global benchmarks, but only with close attention to technical maturity, different development paths and the need for flexible adaptation as technologies and risks evolve.

Lagodinsky's proposed shift from a "Brussels regulation effect" to a "Brussels business effect" is therefore central. Compliance should not operate only as a cost or a threat of fines. It should create market value through trusted procurement, lower insurance risk, interoperable standards and credible certification. The goal is not deregulation. It is regulation that produces investable advantage.

Ahead's contribution lies in connecting policymakers with leaders who experience these frictions differently: legislators concerned with legitimacy, entrepreneurs confronting scale, technical experts confronting system risk and representatives of states for whom sovereignty cannot be separated from development. A global platform becomes useful when proximity to influence is converted into a better policy diagnosis, rather than celebrated as access in itself.

What happens when claims of concentration are tested against firm behavior?

Esya Centre's recent empirical work offers a useful corrective to a debate often dominated by assumptions imported from the largest American technology markets. Its 2026 assessment of India's AI value chain surveyed 227 firms across infrastructure, foundation models, data and applications. It found an ecosystem that remained dynamic and contestable: firms were multi-homing across cloud providers, combining open-source and proprietary models and experimenting with decentralized deployment.



The finding matters beyond India. It suggests that concentration at the frontier-model layer does not automatically translate into a closed market at every layer of the value chain. Policy should therefore distinguish carefully between genuine bottlenecks and competitive adaptation before prescribing remedies. [12]

This is where the Ahead–Esya collaboration becomes strategically valuable. Ahead brings into the same exchange the actors who experience power from different positions: those who legislate, build, finance, secure and live with technological systems. Esya tests those propositions against evidence. Ahead widens the field of intelligence; Esya disciplines it through data. Together, they make it harder for policy to confuse visibility with dominance, ambition with capacity or inclusion with influence.

4. APPLIED AI AND THE DISCIPLINE OF USEFULNESS

Is the obsession with frontier models obscuring where public value will actually be created?

Lagodinsky argued that Europe should concentrate on applied, sector-specific AI in health, industry and transport rather than treating competition in general-purpose large language models as the only test of leadership. This is not an argument for technological modesty. It is an argument for strategic concentration.

The same logic informed his skepticism about relying only on gigafactories. Large-scale infrastructure may be essential for strategic capability, but smaller, targeted AI factories can be more accessible to start-ups, less energy-intensive and better aligned with sectoral applications. Europe needs a tiered ecosystem: frontier capacity where scale is indispensable and distributed capacity where adoption, experimentation and commercialization occur.

Stephen Boyer's call to focus public communication on real productivity rather than speculative artificial general intelligence also matters. Several participants returned to the same principle: governments and institutions should prioritize preparedness for systems already being deployed rather than organize policy around uncertain short-term assumptions about AGI. OECD evidence shows meaningful gains in some controlled tasks, but effects vary by occupation, skill level, workflow and organizational readiness. The immediate task is to communicate demonstrable productivity gains and real-world applications while building the safeguards, skills and accountability required for responsible deployment. The serious story is not that AI automatically produces growth. It is that institutions capable of redesigning work, data and accountability may gain rapidly, while others purchase tools without changing outcomes. [11]



Ahead's reach beyond technical specialists is a strategic advantage. AI policy will fail if it is intelligible only to those already inside the field. Ahead's network of business leaders, philanthropists, public figures, next-generation leaders and global institutions creates a constituency able to translate technical choices into organizational, investment and public consequences. That translation is not dilution; it is a condition of implementation.

5. FROM LEGISLATION TO APPLIED LAW

What is a rule worth if no institution can inspect the system it governs?

Dr. Vera Schmitt identified the fault line between “normative ambition and technical reality.” The EU has invested heavily in principles and legislation. The next contest is operational: regulators need technical staff, testing environments, evidence-preservation protocols, access to systems and the authority to act at digital speed.



For AI-enabled threats, the focus must move from legislation to applied law: technically enforceable rules, prevention at scale and traceability that makes accountability real.

DR. VERA SCHMITT | TU BERLIN / DFKI

Dr. Markus Siewert's proposal for policy clinics and fellowships of practice offers a concrete institutional response. Policymakers and technical experts should stress-test rules together, learn from implementation and build shared language. Universities can serve as trusted conveners because they connect research, training, public purpose and comparative learning.

The partnership is especially credible at this point in the argument. Ahead brings the political and relational infrastructure required to put different forms of influence into meaningful exchange. Esya brings the analytical discipline to test claims against evidence and convert dialogue into recommendations. Together, they address a familiar weakness of global forums: insights too often remain trapped in the room that produced them. This paper is designed to make them usable.



Esya's survey of 300 Indian firms on the Digital Personal Data Protection Act reaches a related conclusion from the perspective of implementation. It found that restrictions affecting the use of publicly available personal data and the absence of familiar lawful bases could raise compliance costs and constrain routine activity for AI start-ups and smaller businesses. The broader lesson is not that rights protections should be diluted. It is that regulation must be tested against the operational realities of the firms expected to comply with it, especially those without the legal and technical resources of global incumbents. [13]

Could a rule intended to curb concentrated power inadvertently strengthen the firms already best equipped to absorb its costs?

6. SECURITY HAS MOVED BEYOND THE MODEL

Who is accountable when an AI agent can act, authenticate and expose systems on its own?

Boyer warned that the proliferation of AI agents and Model Context Protocol connections is expanding the attack surface. The issue is not only whether a model produces an unsafe answer. It is whether an agent can access internal tools, retrieve sensitive data, invoke external services or make changes without robust authentication and monitoring.

His recommendation was deliberately practical: use existing policy tools and shine a spotlight on AI-specific exposure rather than inventing an entirely new framework for every risk. Organizations should be required to know which agents are deployed, what permissions they hold, how credentials are managed, what logs exist and which human remains accountable.

This operational realism also sharpens Ahead's broader proposition. Influence without responsibility is theater. Access to global rooms must be matched by the willingness to ask what institutions can actually do on Monday morning. A high-level exchange has value only if it travels into authentication standards, procurement rules, training budgets and accountable ownership.



7. AFRICA CANNOT BE AN IMPLEMENTATION FOOTNOTE

Will the intelligence age reproduce the extractive architecture of earlier technologies?

The answer depends on whether African countries are treated as markets for imported systems or as authors of the governing architecture. Honorary Minister Freddy Kita pointed to the African Union's 2024 Continental AI Strategy and called for national strategies, dedicated financing, youth participation, data protection and ethical standards embedded from the outset. Dr. Wambui Kahara stressed that regulation must not close the door to innovation in emerging economies and that machine judgment must remain bounded in high-risk systems.



African AI governance must be development-oriented, ethical and locally owned, with infrastructure, financing, youth participation and cybersecurity built in from the start.

FREDDY KITA | HONORARY MINISTER, DEMOCRATIC REPUBLIC OF THE CONGO

This is not a regional sidebar to a European paper. Africa's choices will shape global labor markets, data governance, public-service delivery and the legitimacy of international standards. The continent's young population makes education and future employment central to AI policy. A governance framework that protects rights but leaves countries without compute, skills or secure infrastructure will formalize dependency rather than reduce it.

Ahead's global outlook is especially relevant here. It rejects the idea that influence should flow only from established centers to invited audiences. By bringing African political and entrepreneurial perspectives into a high-level European and US security discussion, Ahead treated those voices as strategic authorities rather than beneficiaries of inclusion. That distinction is essential to any credible claim of global governance.

8. YOUTH, TRUST AND THE SOCIAL PURPOSE OF INTELLIGENCE



What are we optimizing for: attention, productivity or human capability?

Giorgiana Martínezgarnelo y Calvo argued that young leaders and STEM talent must participate directly in policy design. Julia Dudenko pushed the debate further: AI systems should create value for new generations rather than reproduce addictive algorithms, and scientists should be made visible as role models.

These interventions expose a deeper question. Governance is not only about preventing catastrophic misuse. It is also about deciding which incentives shape ordinary deployment. Systems optimized for engagement, surveillance or short-term cost reduction may be legally compliant and socially corrosive. The policy conversation must therefore include educational institutions, families, creators and young people, not merely firms and regulators.

Ahead's work with women and next-generation leaders is central to this challenge. It creates pathways from future-facing constituencies into current decision-making. The point is not to ask young people what they think after policy is drafted. It is to give them enough proximity, preparation and authority to influence what is drafted in the first place.



9. TEN PROPOSITIONS FOR ACTION

1

Define strategic dependencies

Map foreign exposure across compute, cloud, models, chips, energy and critical public systems. Sovereignty begins with knowing which choices can be constrained from outside.

2

Treat AI policy as part of Europe's wider strategic capacity

Connect risk-based regulation to industrial policy, secure cloud capacity, advanced digital infrastructure, energy, capital-market reform and a tiered compute ecosystem combining frontier-scale facilities with smaller, sector-specific AI factories accessible to start-ups, universities and public-interest users.

3

Create the Brussels business effect

Use procurement, certification, liability and insurance to make trustworthy AI commercially valuable, not merely legally mandatory, and close Europe's main weakness: the gap between regulation and implementation.

4

Fund implementation, not only legislation

Create technical units, regulatory sandboxes, policy clinics and fellowships of practice able to inspect systems and enforce rules, building normative influence while avoiding strategic dependency.

5

Apply existing security duties to agents

Require inventories, authentication, permission controls, audit logs and accountable human ownership for agentic systems and MCP connections.

6

Prioritize preparedness and applied capability

Focus on preparedness rather than speculative assumptions about how AI will evolve, which are difficult to predict and unlikely to materialize in the short term.

Prioritize real-world applications in health, industry, transport, energy and public administration, where Europe has sectoral depth, while maintaining access to frontier capability. Communicate measurable productivity gains, prepare institutions for systems already being deployed and avoid allowing speculative AGI scenarios to displace immediate governance responsibilities.



7

Protect pathways from research to scale

Strengthen capital markets and strategic investment so European and African scale-ups are not systematically acquired before they become durable institutions.

8

Finance African implementation

Connect national and continental strategies to secure infrastructure, skills, data protection, procurement and long-term financing.

9

Treat participation as an influence metric

Measure whether women, youth and under-represented regions shape agendas, recommendations and follow-through, not simply whether they attend.

10

Build durable networks across rooms

Use collaborations such as Ahead and Esya to carry insights from high-level convenings into policy research, business practice and public communication.



CONCLUSION: WHO GETS TO AUTHOR THE FUTURE?

The intelligence age will not be governed by law alone, nor won by capital alone. It will be shaped by the interaction of infrastructure, markets, technical capability, political legitimacy and the people granted authority to define the problem.

The exchange at the European Parliament Liaison Office in Munich, held alongside the Munich Security Conference, demonstrated why Ahead's approach belongs within the substance of AI governance. When legislators, cyber experts, African leaders, academics, entrepreneurs, multilateral voices and young leaders confront the same problem, the result is not simply a more representative conversation. It is a more accurate account of power. Ahead's contribution is to make that breadth consequential: to assemble influence across sectors and geographies before the agenda closes, and to ensure that voices too often invited as audiences can operate as authors.

Esya's contribution is equally substantive. Its empirical work on 227 firms across India's AI value chain challenges simplistic claims that the ecosystem is already irreversibly concentrated, while its survey of 300 firms shows how even well-intended data rules can impose disproportionate costs on smaller innovators. Taken together, these findings point to a larger principle: technology policy must be evidence-led, layer-specific and alert to second-order effects. A rule that looks neutral on paper may entrench incumbency in practice. A market that appears concentrated at one layer may remain contestable at another.

For Europe, the outcome is exacting. Regulatory leadership remains an asset, but it will endure only if matched by innovation, implementation and strategic investment. Europe's influence will depend less on its ability to foresee a single technological future than on its capacity to adapt across several: to update rules as technical maturity changes, finance firms through scale, secure critical infrastructure and translate abstract ambition into practical readiness.





Can global AI governance become more sophisticated than the binaries that currently contain it: innovation or regulation, sovereignty or openness, scale or inclusion?

That is the challenge this collaboration is designed to meet. Ahead expands who carries authority into the room. Esya tests what survives scrutiny once it gets there. The value of the partnership lies not in producing consensus, but in producing better disagreement: grounded in evidence, sharpened by global experience and capable of traveling from influential rooms into policy, investment and public understanding.

The roundtable "Artificial Intelligence as a Geopolitical Force: Rethinking Diplomacy and Governance in an Era of Disorder" was convened by Ahead and hosted by the European Parliament Liaison Office in Munich on Friday, February 13, 2026, during the Munich Security Conference.

Ahead thanks Bitsight Technologies for making the organisation of this roundtable possible.

This report was produced with the support of Ahead's knowledge partner, the Esya Centre.



ABOUT THE CO-AUTHORS OF THIS PAPER

Ahead

Ahead is a global leadership and influence platform that brings women, next-generation leaders and male allies into the rooms, networks and conversations where consequential decisions are made. Through strategic convenings, curated access, thought leadership and cross-sector partnerships, Ahead works to broaden participation in global decision-making and strengthen the routes through which diverse leaders build influence.

At the heart of its approach is the conviction that representation acquires meaning only when it is matched by access, agency and a substantive role in shaping policy, capital and institutions.

Esya

The Esya Centre, a distinguished non-profit organization, is dedicated to fostering empirical research and shaping thought leadership to redefine public policy frameworks for the future. Focused on the intersection of people, innovation and value, the Centre aims to generate transformative ideas that influence policymaking and catalyze sustainable development. By engaging with contemporary challenges and opportunities, the Esya Centre works to reimagine India's public policy discourse through data-driven insights and innovative constructs.

1. **Munich Security Conference.** (2026). Outsmarting ourselves? Risks and rewards of the AI race. Munich Security Conference.

The Munich Security Conference describes itself as the world's leading forum for debating international security policy.

2. **Munich Security Conference.** (2026). Munich Security Index 2026. Munich Security Conference.

The report and related technology program materials provide context on current geopolitical and technological security challenges.

3. **Stanford Institute for Human-Centered Artificial Intelligence.** (2026). AI Index Report 2026. Stanford University.

US private AI investment reached \$285.9bn in 2025, more than 23 times the \$12.4bn recorded in China; private investment figures do not capture all state-directed Chinese spending.

4. **European Commission.** (2025). AI Continent Action Plan. European Commission.

European Commission. (2025). InvestAI Initiative. European Commission.

InvestAI aims to mobilize €200bn in AI investment, including a €20bn facility for up to five AI gigafactories. The Commission also aims to triple EU data-center capacity over five to seven years.

5. **Financial Times.** (2026, July 21). AI is becoming a geopolitical weapon, warns EU digital chief.

The article reports European concerns about dependency on foreign AI models, cloud infrastructure and potential technological "kill switches."

6. **The Wall Street Journal.** (2026, July 21). Top American AI execs sound alarm on Chinese models.

The article discusses the strategic and commercial competition surrounding lower-cost Chinese open-weight models, as well as tensions between security, competition and regulation.

7. **Draghi, M.** (2024). The future of European competitiveness. European Commission.
See also Letta, E. (2024). Much more than a market.

8. **European Union.** (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.

9. **African Union.** (2024). Continental artificial intelligence strategy. African Union.

10. **UNESCO.** (2021). Recommendation on the ethics of artificial intelligence. UNESCO.

11. **OECD.** (2025). Is Generative AI a General Purpose Technology? Implications for Productivity and Policy. OECD Publishing.

12. **Bal, M., Qayoom, A., & Oberoi, A.** (2026, May). An empirical assessment of India's AI value chain: Market structure, competition, and innovation dynamics. Esya Centre.
Primary survey of 227 firms.

13. **Bal, M., Venkatesan, S., & Qayoom, A.** (2026, March). A survey-based assessment of the impact of the Digital Personal Data Protection Act on Indian MSMEs and AI innovation. Esya Centre.
Primary survey of 300 firms.

