



TAKSHASHILA
INSTITUTION

Narratives and Chokepoints in AI Geopolitics

Bharath Reddy, Pranay Kotasthane

Takshashila Discussion Document 2026-31
Version 1.0, October 2026

The essay explores the narratives that impact AI policy and the geopolitics of the AI stack

Recommended Citation:

Bharath Reddy, Pranay Kotasthane, "Narratives and Chokepoints in AI Geopolitics", Takshashila Discussion Document 2026-31, Version 1.0, October 2026, The Takshashila Institution

©The Takshashila Institution, 2025

Contents

1 Abstract	2
2 Narratives Shaping AI Strategies And Governance	3
2.1 AI as a Conventional Arms Race	3
2.2 AI as Nuclear Weapons	4
2.3 AI as a Geopolitical Innovation Race	4
2.4 AI as a General Purpose Technology	5
2.5 AI as a Global Commons	5
3 Geopolitics Across The AI Technology Stack	6
3.1 Compute	7
3.2 Data	8
3.3 Models	9
4 Policy Implications	10

1 Abstract

Artificial intelligence has become an area of geopolitical contestation, and the policy instruments states reach for are shaped by the narratives by which we understand the technology. This chapter identifies five competing framings — AI as a conventional arms race, as a nuclear weapon, as a geopolitical innovation race, as a general-purpose technology, and as a global commons — and shows that the narrative that dominates a given conversation often reveals political and strategic objectives rather than mere technical realities. It then maps these framings onto the instruments deployed across the AI stack: export controls and monitoring mandates at the compute layer, localisation requirements at the data layer, and weight restrictions and openness debates at the model layer. It concludes with pathways for middle powers like India to mitigate near-term risks while building long-term innovation capacity.

This essay argues that the technology and policy ecosystem around artificial intelligence has become an arena of geopolitical contestation. It begins by exploring the powerful and competing narratives that shape national strategies and policy around AI. An arms race, a nuclear weapon, a geopolitical innovation race, a general-purpose technology, or a global commons: each framing serves some interests and constrains others. Which narrative dominates which conversation often reveals more about AI geopolitics than the technical realities the narratives claim to describe.

Technology is an important determinant of economic and military power, and the narratives that emphasise the links between technology and national power have become more assertive. Nuclear weapons make large-scale conventional wars unlikely, and globalisation has made economic decoupling prohibitively costly, if not entirely unfeasible. Although the events over the past few years challenge both of these claims, the past few decades have seen relatively free trade and limited conflicts. In this landscape, technology is becoming both a means and an end of geopolitical contestation.

Artificial intelligence is transforming both the sources and uses of national power – both economic and military. It introduces new elements of power, changes the importance of existing sources of power, and alters the goals that states should pursue. For instance, access to compute, organisational adaptation, and effective data pipelines become increasingly important to leverage AI. In contrast, population size decreases in importance as a source of power while highly skilled talent grows in importance. The policy measures states adopt to navigate this shift are not chosen in a vacuum; they are justified, contested, and made legible through narratives.

The essay proceeds in three parts. First, we lay out the five dominant narratives in AI geopolitics and the audiences each

This essay was originally written for Strategic Priorities for India – Essays on AI Policy, a compilation edited by Urvashi Aneja and Venkatesh Hariharan with support from the Digital Futures Lab and Samāgata Foundation. The full collection of essays is available [here](#).

serves. Second, we map these narratives to specific instruments across the AI technology stack — compute, data, and models — to show how narrative choices generate concrete policy. Third, we draw out recommendations that serves India's national interest.

2 Narratives Shaping AI Strategies And Governance

Narratives help us make sense of new developments through stories that we can relate to. They play an essential role in clarifying the uncertainty around emerging technologies and the impact they will have on society.

An example is the narrative that AI's impact on the workforce will be similar to past innovations like the internal combustion engine, which rendered some jobs obsolete but increased productivity and created many more jobs in the automobile manufacturing industry. This is one of many competing narratives on the impact that AI is likely to have on the workforce. It makes the case for embracing the technology regardless of the short-term disruption that it is likely to cause, as it would lead to increased productivity and employment in the future.

With multiple competing concerns around AI, such as the impact on society, climate, the economy and balance of power between nations, narratives have an essential role in influencing policymaking. They help frame priorities, shape public opinion, mobilise resources, and set the stage for policy-making. Various narratives such as stories of prosperity, decline, hope, fear, and ambiguity are frequently used in policymaking¹. There is a battle of competing narratives around AI geopolitics. We identify four prominent narratives below.

2.1 AI as a Conventional Arms Race

This narrative treats AI research and development as an arms race. This narrative serves the US security establishment, frontier AI labs seeking national-security legitimation, and a broader China-hawk coalition. It treats AI research and development as a zero-sum race in which falling behind is unacceptable, and has gained traction in academic and policy discussions and the popular imagination. This narrative rests on contested expectations that the exponential progress of AI systems seen in the past few years will continue, and we will reach superintelligence in the near future^{2 3}. Critics have argued these assumptions reflect Silicon Valley hype cycles more than technical inevitability⁴.

There have been several policy measures aimed at maintaining a lead over adversaries which seemed unimaginable just a few years ago. These include export controls on advanced semiconductors, high bandwidth memory and advanced packaging tools⁵. Further

up the value chain, there is alignment by allies such as the Netherlands and Japan on export restrictions on lithography machines⁶. There are several other legislative proposals that include export control whistleblower programmes, restrictions on semiconductor investments in China, and certification requirements for advanced semiconductor exports⁷. The narrative also leads to prioritising speed over other concerns such as safety or ethics in order to maintain technological leadership.

2.2 AI as Nuclear Weapons

A more extreme form of the arms race narrative includes proposals to treat AI as similar to "catastrophic dual-use nuclear, chemical, and biological technologies"⁸. This one too serves the interests of the US security establishment and frontier AI labs seeking national-security legitimisation.

Policy measures in this framing aim to pursue deterrence and limit proliferation at all costs in addition to having a competitive advantage over rivals. Any state pursuing unilateral AI dominance is likely to be met with deterrence measures through kinetic or cyber means. Hendrycks et al introduce the concept of mutual assured AI malfunction (MAIM) resembling mutual assured destruction (MAD) dynamics seen in nuclear deterrence⁹. Non-proliferation measures take the form of stringent controls on compute and models to ensure they do not get into the wrong hands. Some of the more extreme proposals under consideration such as hardware based monitoring of AI chips and location tracking mandates¹⁰ are closer to the non-proliferation logic than the arms race logic.

2.3 AI as a Geopolitical Innovation Race

In contrast to an arms race, an innovation race is a positive-sum game in which multiple actors compete and collaborate across national boundaries. Schmid et al. propose that a "geopolitical innovation race" more accurately describes the current situation where great power politics and technology development are deeply entangled¹¹. They contend this narrative is driven by a mix of motivations that include national security, economic leadership and national status. It involves a complex network of companies, states, and research institutions competing and collaborating on technology research and development, instead of a state-military nexus that would drive efforts in an arms race. It can have both zero-sum and positive-sum payoffs for the actors involved.

Policy measures in this narrative include those that maintain an edge in innovation while maintaining global collaboration. It might include some measures that limit a rival's progress on technology such as export controls on inputs such as semiconductors. It might also justify the subsidies and incentives towards building data centres seen in many regions including the EU, US, China

and India aiming to make domestic companies more competitive on the global stage.

2.4 AI as a General Purpose Technology

This narrative views AI not as a potentially superintelligent technology in the foreseeable future but as a transformative general-purpose technology similar to electricity or the Internet. Widespread adoption would be on the time scale of decades and will be limited by speed breakers such as bridging the capability-reliability gap, learning curves, organisational reforms, compliance with laws and norms and plain inertia¹². This narrative demands a far greater emphasis on safety and ethics over speed of development.

States should place greater emphasis on building the capacity to “diffuse”, or widely adopt, innovations than on employing protectionist measures to build national innovation capabilities¹³. Diffusion depends not just on the ability to innovate but on institutions that promote technology transfer, human capital, and openness to trade.

2.5 AI as a Global Commons

This narrative is often promoted by international organisations and civil society to advocate for AI as a tool to benefit all humanity¹⁴. There are several ethical arguments that are used to back such a claim. These include the possibility that a large amount of human labour and content go into training AI systems which might eventually impact livelihoods at scale. Similarly, massively power-hungry data centres are likely to have externalities on the environment, the costs of which will be borne disproportionately by the more vulnerable sections of society while benefits accrue to a few.

While this narrative does not align with how incentives around AI development and deployment work, nation-states, particularly middle powers or low and middle-income countries, use it as a diplomatic strategy. It offers a way out of the current landscape where the US and China dominate technology research, development and norm-setting. For example, this aligns with India’s diplomatic position in forums like the G20 or GPAI (Global Partnership on AI), where it has championed AI for all.

Narratives are not just framing devices. They help build consensus, mobilise resources and lay the groundwork for policymaking. Different stakeholders seek to shape the narratives around AI governance so they can shape the rules around its governance. A state that adopts the arms race narrative reaches for compute denial and model export controls; one operating under the nuclear weapons logic adds non-proliferation and hardware monitoring; the innovation race produces subsidies

and ecosystem building; the GPT narrative emphasises diffusion infrastructure, regulatory capacity, and human capital; the global commons supports multilateral governance and redistribution mechanisms. The map below shows which narratives currently dominate each layer of the AI stack — and which layers are most contested. We now move to examining how geopolitics determines access across the key layers of the AI technology stack.

3 Geopolitics Across The AI Technology Stack

The narrative a state adopts determines which layers of the stack it treats as strategic and which instruments it adopts. For instance, under the Biden administration, the “small yard, high fence” approach was adopted by the U.S. as a means of protecting sensitive technologies with a high fence of trade and investment controls. However, regardless of who is in office, the yard has continued to expand in scale even if enforcement has not always been effective. AI systems are built on key building blocks such as data, compute, models and applications. Across these building blocks are bottlenecks that are being leveraged for geopolitical ends.

Considering the entanglement of technology with national power, states are likely to pursue strategies to secure their relative advantage in the AI domain over their rivals. Strategic objectives could take the form of denial, outpacing adversaries and removing major bottlenecks. Table 1 outlines the key instruments of AI geopolitics and their often-overlooked repercussions.

Strategic Objective	Instruments Used	Underrated Repercussions
Denial	Export controls and sanctions	Difficult to sustain; incentives for adversaries to accelerate domestic innovation or pursue backroom deals
	Hardware-based monitoring and location tracking of AI chips	Raises ownership, privacy, and surveillance concerns in partner countries
	End-use restrictions	Encourages allies and adversaries to develop domestic capacity
	Screening investments, mergers and acquisitions	Slows down cross-border innovation and capital flow
	Restrictions on the movement of talent	Can slow down technological progress
Outpace Adversary	Industrial espionage for model weights or proprietary datasets	Invites stricter cybersecurity and IP protection regimes
	Build strategic partnerships	Sustained coordination is politically and economically difficult
	Indigenisation and industrial policy	Difficult to sustain
	Increase dependence and control	Helps manage the adversary's pace to an extent
	Economic coercion	Self damage
Increase Supply Chain Resilience	Promoting open source and open weight models	Diffusion, trust and adoption may increase
	Build partnerships to create resilient AI supply chains	Sustained coordination is politically and economically difficult
Influence Minds and Actions	Algorithmic influence via closed AI models	Subtle shaping of narratives; limited accountability mechanisms
	AI-driven information operations or data poisoning	Limited long-term impact on national power; risk of escalation
	Export of 'trustworthy AI' standards and safety frameworks	Soft power tool; shapes norms but may limit domestic innovation

Table 1: Instruments of AI Geopolitics

3.1 Compute

Compute governance is influenced by the arms-race, nuclear-weapons and geopolitical innovation race frames. Computing power is crucial for the development and deployment of AI systems. Compared to other inputs, “compute is a particularly effective point of intervention: it is detectable, excludable, and quantifiable, and is produced via an extremely concentrated supply chain”¹⁵, making it an attractive leverage point for governing AI.

For example, the Dutch company, ASML, has a monopoly on building EUV (extreme ultraviolet) machines required to fabricate advanced semiconductor chips. NVIDIA dominates the AI chips market and has become the most valuable company in the world, surpassing many established technology giants. This huge market share is due to its early investments in building GPUs and its widely used proprietary software interface, CUDA, which allows developers to utilise the hardware's capabilities effectively¹⁶. Further downstream, Microsoft Azure, Google, and Amazon Web Services together account for nearly two-thirds of the global cloud infrastructure market¹⁷. These providers are the touchpoint for most users deploying AI at scale. They

leverage economies of scale and vertical integration to maintain a competitive advantage over rivals. These are just a few of the bottlenecks in the extremely concentrated supply chain.

Some of the policy measures in the US deny compute access to adversaries. The now-rescinded AI diffusion framework, for instance, cleaved the world into three tiers based on levels of trust that determine access and security requirements for importing advanced AI chips and certain AI model weights¹⁸. However, this has been replaced by measures that strengthen chip-related export controls on China¹⁹. Hardware-based mechanisms are also reportedly under consideration to monitor and restrict the usage of AI chips²⁰. The US has also introduced legislation mandating built-in location tracking for AI chips to prevent their diversion to China, Russia, and other countries of concern²¹. Such measures are problematic, as they raise ownership, privacy and surveillance concerns.

Strategic blocs also help to build collective resilience. The Pax Silica is a U.S. led strategic partnership initiative to secure the global supply chain for AI spanning critical minerals, energy, advanced manufacturing, semiconductors, AI infrastructure and logistics²². Key partners include the United States, Japan, South Korea, the Netherlands, the United Kingdom, Australia, Singapore, Israel, United Arab Emirates, Qatar, Sweden and India.

Many countries are also building out massive domestic data centres to incentivise domestic innovation and address sovereignty concerns. These include the EU's AI factories²³, India's AI Compute Mission²⁴ and China's state support for data centres²⁵.

3.2 Data

Narratives that determine data governance are more mixed. Sovereignty and localisation reflects the innovation race while open dataset type initiatives gesture toward a commons. Data is an integral input to training AI. Large language models are trained on almost the entire corpus of written human content. As high-quality publicly available human-generated content for training AI is nearly exhausted²⁶, proprietary datasets with high information content and multimodal data sources such as images or video become invaluable data sources. Technology platforms collect vast amounts of data in exchange for the free services they provide users. It is used to feed their ad-based revenue models. In addition, data within walled gardens such as X and Reddit could become invaluable sources of high-quality user data.

Data localisation requirements limit cross-border data flows due to privacy, security and sovereignty concerns. The EU's General Data Protection Regulation, China's Personal Information Protection Law, and India's Digital Personal Data Protection Act and other sector-specific regulations all require varying degrees

of data localisation. This leads to increased compliance burdens due to varied regional regulations to house their data locally.

Governments are also creating datasets with positive externalities or security implications, such as the strategic and science datasets planned under the US AI Action Plan²⁷ or the language datasets under Bhashini.

3.3 Models

Narratives that determine model governance sit between innovation race (efforts to promote domestic champions) and arms race (restrictions on use cases and model weights). Most of the frontier AI models are designed by companies based in the US or China. There are different gradients of openness of models, with the most maximally open ones being those where all model artefacts, including training data and code, are released under open licences. While these maximally open models are mostly used for open science, the less open models, with only model weights, documentation and tools released under open licenses, also allow downstream innovation. They are especially useful for applications where users would not want to share sensitive data with a model vendor. This points to a distinction that the debate around open models tends to collapse: the dependency a model creates is a function of its deployment topology, not its provenance. A proprietary model reached through a foreign API creates a continuous dependency — every query is processed within a foreign jurisdiction, and access can be withdrawn by an export control decision or a change in terms of service, at which point the workflow breaks. An open-weight model running on domestically located hardware creates a one-time dependency. The weights are a static file, and the releasing lab has no API to revoke or visibility into the deployment. The open-weight approach outperforms on data residency and access continuity, including when the weights originate in a rival jurisdiction. Both China²⁸ and the US²⁹ have clearly articulated in their national AI plans the need to encourage their open models to be adopted by the world, indicating their strategic importance.

This strategic importance also decides how long the current openness lasts. Chinese labs release open weights partly out of necessity, as export controls constrain the compute required to serve global demand at scale. Beijing has a consistent record of restricting exports once a material becomes strategically significant — gallium, germanium, antimony and rare-earth processing technologies were freely exported until they were not. That open weights still leave without restriction is a statement about how Beijing perceives the current generation of models rather than a stable condition. The asymmetry is that weights, unlike physical goods, cannot be unexported once released; any future restriction would apply only to subsequent generations.

There are also many restrictions on the use of models. The same rescinded US diffusion framework also imposed controls on

the model weights for certain advanced closed-weight dual-use AI models³⁰. The leading AI labs in the US, such as OpenAI, Anthropic, and Meta, have announced deals to allow the US government to use their models for military purposes³¹. Since the terms of use for these models prohibit their use for security purposes by other nations, this raises fears about how this would impact the balance of power between states³².

There are also espionage and sabotage concerns regarding models. These could be industrial espionage to gain access to proprietary data or model weights, which are not available in the public domain. Chinese AI labs such as DeepSeek, Moonshot, and MiniMax have been accused of distilling frontier models from American AI companies. In a nutshell, this involves taking a stronger AI model's outputs to teach a weaker model. The attacks were sophisticated and used deceptive techniques to mask the identity and intent of the distillers. The US AI lab Anthropic claims that this happened on an industrial scale – "16 million exchanges with Claude through approximately 24,000 fraudulent accounts, in violation of our terms of service and regional access restrictions."³³

4 Policy Implications

While most of the frontier research and development in the compute and model layers of the AI stack is being led by the US or China, middle powers such as India have a few pathways to navigate this landscape and secure their technology resilience and autonomy. These range from short-term measures to ensure access to the technology to longer-term measures to strengthen the innovation ecosystem.

First, building a foundational level of base capabilities offers a pathway to developing knowledge and skills over time in critical areas. For example, India's efforts to build semiconductor fabrication capabilities in 28nm chips aims to create an ecosystem around leveraging existing strengths in semiconductor design capabilities. The compute, data, foundation model and skilling pillars of the India AI mission can be seen as trying to achieve a base level of capabilities and establish a domestic presence in various critical areas. However, with only 38% of the 33,000 GPUs empaneled under the compute mission as of March 2026, adoption has been slow. Implementation will be key in ensuring that these initiatives achieve the desired outcomes.

Second, funding and adopting open-tech ecosystems that can compete with proprietary alternatives in areas with high market concentration can address vulnerabilities. These open-tech ecosystems need not be domestic and can still provide resilience due to their public good nature. For example, developing open-source alternatives to CUDA, NVIDIA's proprietary programming platform, could help break the GPU stranglehold in AI. Ongoing efforts such as OpenAI Triton, an open-source

alternative to CUDA that allows AI models to run on diverse hardware, or Apache TVM, which optimises machine learning workloads across different processors could be candidates. There are other lessons from China's open-source playbook ³⁴ ³⁵ such as actively participating in open standard instruction set architectures such as RISC-V that also allows companies to design chips without licensing fees or export restrictions.

While limited open-weight models such as those released by NVIDIA, OpenAI, Meta and DeepSeek can allow downstream uses, they do not allow full access that allows further research and development. Supporting maximally open models, such as those developed by the non-profit EleutherAI, can have positive spillovers for the research ecosystem. Alternatively, targeted government procurement or pilot adoption to create early users can also help support the ecosystem.

Such funding projects could also be sustained in partnership with states with similar goals, such as the EU, which recognises open-tech as a priority. Such a partnership could complement India's software expertise and scale with world-class research infrastructure in the EU. Creating an Open Source Programme Office that can crowd source identifying projects globally and supporting them can make a difference. Such a model has been prototyped at a small scale by Zerodha's FOSS (free and open source software) fund³⁶.

Third, an open-weight model pathway with domestic inference capacity. Models with near frontier capabilities are currently released as open weight models by leading American and Chinese labs. India needs multiple domestic inference providers like OpenRouter that serve multiple models rather than a single national AI cloud. Regulatory attention should shift from the provenance of model weights to the deployment layer, which determines where data flows, who controls the compute, and what audit trails exist. Given that the openness window may be limited, the capacity to absorb, evaluate and fine-tune whatever is available has to be developed. The GPUs running inference in India remain designed by American firms subject to American export controls, which is a reason to continue engaging Washington through arrangements like Pax Silica. A model's refusals and framings are fixed in its weights — domestic inference removes revocability risk, not values risk, which is why evaluation and red-teaming capability must be developed.

Fourth, middle powers could also explore asymmetric options to develop leverage in critical bottlenecks in the AI value chain. These could be niche areas where a middle power can become indispensable and wield disproportionate power. Examples include the expertise developed by ASML and TSMC in critical parts of the semiconductor value chain. For India, this could involve building a strong talent base on open-source machine learning platforms such as PyTorch.

Fifth, while the above measures help build some short-term

resilience, an innovation policy is needed to strengthen the research and development ecosystem in academia and industry. A resilient future lies “not in a narrow focus on LLMs or GPUs, which are transient artifacts, but the societal and academic foundation required to enable conditions and ecosystems that lead to the creations of breakthroughs like LLMs”³⁷. This needs strengthening both the innovation and diffusion capabilities. The Anusandhan National Research Foundation already has this mandate and its scope can be scaled up. They should be empowered to fund future-looking, high-risk, high-reward research, moving beyond traditional academic metrics.

Lastly, India can leverage its human capital as a diplomatic tool. India is poised to be the world’s AI talent hub and can leverage it to build talent partnerships with other countries such as the EU and Japan.

The recommendations above — base capability building, open-tech investment, asymmetric leverage in niche bottlenecks, innovation ecosystem reform, and human capital diplomacy — should be read as the operational arm of the general purpose technology or innovation race story at the domestic level, while the commons rhetoric continues to do work to build collaboration and broaden access at the global level. We argue this is not contradiction but the deliberate deployment of different narratives at different levels to serve different audiences and strategic interests. The scale of the initiatives under the IndiaAI mission pale in comparison to the investments by the private sector in the US or China. India’s comparative advantages in the short term lie in enabling faster diffusion of the technology. The recommendations offered in this essay aim to enable diffusion and mitigate risks in the near term while building long-term capabilities.

Endnotes

1. Deborah Stone, *Policy Paradox: The Art of Political Decision Making: The Art of Political Decision Making* (W. W. Norton & Company, 2011).
2. Daniel Kokotajlo et al., "AI 2027," April 3, 2025, [Link](#)
3. Leopold Aschenbrenner, "Situational Awareness: The Decade Ahead," June 6, 2024, [Link](#)
4. Arvind Narayanan and Sayash Kapoor, "AI As Normal Technology," Knight First Amendment Institute, April 15, 2025, [Link](#)
5. Karen Sutter, "U.S. Export Controls and China: Advanced Semiconductors" (Congressional Research Service, September 9, 2025), [Link](#)
6. *ibid*
7. *ibid*
8. Dan Hendrycks, Eric Schmidt, and Alexandr Wang, "Superintelligence Strategy: Expert Version," April 14, 2025, [Link](#)
9. *ibid*
10. Girish Sastry et al., "Computing Power and the Governance of Artificial Intelligence," arXiv.org, February 13, 2024, [Link](#)
11. Stefka Schmid et al., "Arms Race or Innovation Race? Geopolitical AI Development," *Geopolitics* 30, no. 4 (January 28, 2025): 1907–36, [Link](#)
12. Arvind Narayanan and Sayash Kapoor, "AI As Normal Technology," Knight First Amendment Institute, April 15, 2025, [Link](#)
13. Jeffrey Ding, *Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition* (Princeton University Press, 2024).
14. United Nations, "Governing AI for Humanity: Final Report," September 2024, [Link](#)
15. Girish Sastry et al., "Computing Power and the Governance of Artificial Intelligence," arXiv.org, February 13, 2024, [Link](#)
16. Satya Shoova Sahu and Ashwin Prasad, "A Primer on AI Chips: The Brains Behind the Bots," November 13, 2024, [Link](#)
17. Adeyomola Kazeem, "AWS Vs Azure Vs Google Cloud: Comparison of the Big Three Cloud Providers," *Cloudwards(blog)*, July 19, 2025, [Link](#)
18. Bureau of Industry and Security, Department of Commerce, "Framework for Artificial Intelligence Diffusion," *Federal Register*, January 15, 2025, [Link](#)
19. Bureau of Industry and Security, Department of Commerce, "Department of Commerce Announces Rescission of Biden-Era Artificial Intelligence Diffusion," May 13, 2025, [Link](#)
20. Girish Sastry et al., "Computing Power and the Governance of Artificial Intelligence," arXiv.org, February 13, 2024, [Link](#)
21. Fanny Potkin, Karen Freifeld, and Jun Yuan Yong, "Exclusive: US embeds trackers in AI chip shipments to catch diversions to China, sources say," *Reuters*, August 14, 2025, [Link](#)
22. U.S. Department of State, "Pax Silica," December 11, 2025, [Link](#)

23. European Union, "AI Factories," Shaping Europe's Digital Future, October 30, 2025, [Link](#)
24. IndiaAI, "IndiaAI Compute Capacity," n.d., [Link](#)
25. Kyle Chan et al., "Full Stack: China's Evolving Industrial Policy for AI," RAND, June 26, 2025, [Link](#)
26. Pablo Villalobos et al., "Will We Run Out of Data? Limits of LLM Scaling Based on Human-generated Data," arXiv.org, October 26, 2022, [Link](#)
27. The White House, Winning the Race, 2025, [Link](#)
28. Ministry of Foreign Affairs, People's Republic of China, "Global AI Governance Action Plan," July 26, 2025, [Link](#)
29. The White House, Winning the Race, 2025, [Link](#)
30. Bureau of Industry and Security, Department of Commerce, "Framework for Artificial Intelligence Diffusion," Federal Register, January 15, 2025, [Link](#)
31. Gerrit De Vynck, "AI companies get comfortable offering their technology to the military," The Washington Post, November 8, 2024, [Link](#)
32. Kevin Klyman, "Stanford CRFM," April 8, 2024, [Link](#)
33. Anthropic, "Detecting and Preventing Distillation Attacks," February 23, 2026, [Link](#)
34. Kyle Chan et al., "Full Stack: China's Evolving Industrial Policy for AI," RAND, June 26, 2025, [Link](#)
35. Axel Thévenet et al., "Open Source Software Country Intelligence Report – China," 2023, [Link](#)
36. Kailash Nadh, "Announcing FLOSS/Fund: \$1M per Year for Free and Open Source Projects," October 15, 2024, [Link](#)
37. Kailash Nadh, "DeepSeek, AI Sovereignty, and India," January 29, 2025, [Link](#)



TAKSHASHILA
INSTITUTION

The Takshashila Institution is an independent centre for research and education in public policy. It is a non-partisan, non-profit organisation that advocates the values of freedom, openness, tolerance, pluralism, and responsible citizenship. It seeks to transform India through better public policies, bridging the governance gap by developing better public servants, civil society leaders, professionals, and informed citizens.

Takshashila creates change by connecting good people, to good ideas and good networks. It produces independent policy research in a number of areas of governance, it grooms civic leaders through its online education programmes and engages in public discourse through its publications and digital media.

©The Takshashila Institution, 2025