POLICY ANALYSIS

7 / July 2026

THE GREAT AI GAME: THE RACE THAT ISN'T ONE.

Johannes Späth

THIS POLICY ANALYSIS WAS PRODUCED AS PART OF THE COOPERATION BETWEEN OIIP AND THE MINISTRY OF DEFENCE.

Österreichisches Institut für Internationale Politik
Austrian Institute for International Affairs

AUTHOR

Johannes Späth, MA
is a Research Associate at the oiip. His work focuses on the intersection of geopolitics, the transatlantic relationship, and the security implications of emerging technologies.

At the institute, his research centers on U.S. foreign policy and the geopolitical influence of major tech actors. He further examines institutional shifts and governance trends, with a particular interest in the dynamics of autocratization and political economy in the Middle East and North Africa.

IMPRESSUM:

Österreichisches Institut für Internationale Politik – oiip,
Austrian Institute for International Affairs
A-1090 Vienna, Währinger Straße 3/12,
www.oiip.ac.at, info@oiip.ac.at

DOI: https://doi.org/10.83003/pa_7_26

KEYWORDS:
Geopolitics of AI, AI stack, Global AI race, compute governance, semiconductor chokepoints, managed dependency, US-China technology competition, military AI

ZUSAMMENFASSUNG

Diese Analyse beleuchtet die Infrastruktur, auf der künstliche Intelligenz basiert, sowie den geopolitischen Wettbewerb, der sich um sie herum entwickelt. Das zentrale Argument lautet, dass es sich bei diesem Konkurrenzstreben nicht um ein einziges Wettrennen handelt, bei dem alle Staaten mit unterschiedlicher Geschwindigkeit dieselbe Leiter erklimmen. Vielmehr handelt es sich um zwei unterschiedliche Wettkämpfe mit unterschiedlichen Regeln, unterschiedlichen Einsätzen und daher auch unterschiedlichen Strategien zum Sieg.

KEY FINDINGS

1. KI IST BEREITS HEUTE VON STRATEGISCHER BEDEUTUNG UND NICHT NUR EINE ZUKUNFTSPERSPEKTIVE.

Fähigkeiten wie autonomes Zielerfassen, KI-gestützte Logistikoptimierung und hyperautomatisierte Fertigung werden bereits operativ eingesetzt. Folglich treten traditionelle Machtindikatoren wie Bevölkerungszahl oder Militärgröße zunehmend in den Hintergrund zugunsten der strukturellen Grundlagen, die diese modernen Systeme antreiben: Rechenkapazität und Datenzugang.

2. EINE VOLLUMFÄNGLICHE KI-SOUVERÄNITÄT IST FÜR PRAKTISCH JEDEN STAAT STRUKTURELL UNREALISTISCH.

Der KI-"Stack" ist auf jeder Ebene mit Engpässen und Quasi-Monopolen übersät. Kein Staat außerhalb der USA-China-Dyade kann realistisch gesehen einen unabhängigen Zugang zu allen Ebenen sichern, unabhängig vom eingesetzten Kapital. Die strategische Entscheidung fällt daher nicht zwischen Abhängigkeit und Souveränität, sondern zwischen verschiedener Formen kontrollierter Abhängigkeit.

3. ES GIBT NICHT EIN, SONDERN ZWEI VERSCHIEDENE WETTRENNEN.

Für die Vereinigten Staaten und China gilt bei der KI tatsächlich das Prinzip „First-Past-the-Post". Die Überlegenheit bei Pioniermodellen, die Dominanz bei der Rechenleistung und die Konzentration von Humankapital sind genuin entscheidende Faktoren im Kampf um die globale Vorherrschaft. Für alle anderen Staaten ist die Situation völlig anders, da für sie strategisch ausreichende KI-Fähigkeiten auch ohne die Entwicklung eigener Pioniermodelle verfügbar sind.

4. DIE „GUMMIBAND-DYNAMIK" IST, ZUMINDEST VORERST, DER GROSSE AUSGLEICHSFAKTOR.

Das beständige Muster, wonach Open-Source-Modelle innerhalb weniger Monate die Lücke zu Pioniermodellen schließen, ist der Hauptmechanismus, der es auch Staaten aus der zweiten Reihe ermöglicht, Zugang zu modernster KI zu erhalten. Die Zukunft dieser Dynamik hängt vom Ausgang eines sich verschärfenden Wettbewerbs zwischen Unternehmen ab, die ihre führenden Modelle schützen wollen, und Akteuren, die versuchen, diese zu geringeren Kosten nachzubilden.

EXECUTIVE SUMMARY

This policy analysis maps the infrastructure underpinning Artificial Intelligence and the geopolitical competition forming around it. Its central argument is that this competition is not a single race in which all states are climbing the same ladder at different speeds. Rather, it is two distinct contests, with different rules, different stakes, and therefore different strategies for winning.

KEY FINDINGS

1. AI IS ALREADY STRATEGICALLY CONSEQUENTIAL, NOT JUST A FUTURE HORIZON.

Capabilities like autonomous targeting, AI-assisted logistics optimization, and hyper-automated manufacturing are already operationally deployed. Consequently, traditional power metrics, such as population or military size, are increasingly giving way to the structural foundations that power modern systems: compute capacity and data access.

2. FULL-STACK AI SOVEREIGNTY[1] IS STRUCTURALLY IMPLAUSIBLE FOR VIRTUALLY EVERY STATE.

The AI stack is riddled with chokepoints and near monopolies at each layer. No state outside the US-China dyad can realistically secure independent access to all layers, regardless of capital deployed. The strategic choice is therefore not between dependency and sovereignty, but between different forms of managed dependency.

3. THE COMPETITION IS NOT ONE RACE BUT TWO.

For the United States and China, AI is genuinely first-past-the-post. Frontier model superiority, compute dominance, and talent concentration are key determinants in the struggle for global hegemony. For everyone else, the game is categorically different, as strategically sufficient AI capabilities are accessible without domestic frontier model development.

4. THE RUBBER-BAND DYNAMIC IS THE GREAT EQUALIZER, FOR NOW.

The consistent pattern of open-source models closing the gap to frontier performance within months is the primary mechanism keeping bleeding edge AI accessible to second-tier states. The future of the race depends on the outcome of an intensifying contest between firms seeking to protect frontier capabilities and actors attempting to replicate them at lower cost.

1 "AI sovereignty" in this paper is defined as autonomous control over every layer of the AI value chain. Unlike traditional strategic commodities like oil, where supply dependencies can be diversified across multiple global markets, several key layers of the AI stack are controlled by singular (near-) monopolies or duopolies with no viable near-term substitutes. Because absolute self-sufficiency across this stack is structurally impossible for nearly all states, the strategic reality is not a choice between "sovereignty" and "dependence". Rather, states face a spectrum of managed dependency: a structural necessity wherein national strategy is defined not by escaping foreign reliance, but by actively selecting, and where possible, diversifying, and negotiating the terms of that reliance.

1. INTRODUCTION

Artificial intelligence (AI) should no longer be viewed as an emerging technology. It is already reshaping military doctrines and the economics of industrial production across the globe, thereby profoundly altering traditional equations of state power. The states and firms that control the infrastructure underlying these capabilities are deploying a myriad of policies to maintain their strategic advantage and deny competitors access. Advanced AI systems depend on a set of highly concentrated and increasingly strategic inputs: computational power, data, specialised semiconductor manufacturing, energy infrastructure, cloud capacity and technical talent. Access to these resources is unevenly distributed and, in many cases, controlled by a small number of firms and countries. As a result, competition forms increasingly around the control of the underlying ecosystem that makes such applications possible. This ecosystem is commonly referred to as the AI stack. It encompasses the full chain of resources and capabilities required to develop, train, deploy and operate advanced AI systems. As AI becomes a foundational source of economic and military power, securing access to the AI stack is emerging as a central objective of state strategy. Competition over these resources is already reshaping trade policy, industrial policy, export controls, investment flows and alliance structures. Understanding the emerging geopolitical order therefore requires understanding the structure of the AI stack itself.

This paper maps that infrastructure, the so-called "AI stack", and the geopolitical competition forming around it. Its central argument is that this competition is not a single race in which all states are climbing the same ladder at different speeds. Rather, it is two distinct contests, with different rules, different stakes, and therefore different strategies for winning. This misconception makes most states play the wrong game by optimising for frontier competition they cannot win, rather than the managed dependency they can learn to navigate.

AI IS ALREADY MORE STRATEGICALLY CONSEQUENTIAL THAN YOU THINK

AI as a key enabler and might multiplier in warfighting and military projection is not a distant future but already reality on the ground. Modern warfighting in Ukraine and Iran shows how AI is already successfully employed in intelligence fusion, targeting, (semi-) autonomous kinetic engagement and logistical support. Throughout these various security applications, AI employment minimizes the need for human engagement, giving militaries that employ them a significant, concrete advantage in speed and cost that makes automation a near necessity on the modern battlefield.

Data sifting and fusion

Modern targeting cycles have decoupled from human cognitive limits. AI driven software employed by the Ukrainian army sifts autonomously through vast quantities of satellite imagery to detect and flag Russian positions (Ritters 2026). Simultaneously, data fusion tools such as Palantir's Gotham aggregate various data silos, ranging from satellites imagery and drone footage and ground-based sensors to open source intelligence like Russian soldiers' social media posts, in near real time into a comprehensive battlefield assessment that would take human analysts days at a time (Bergengruen 2024). This architecture of machine-speed targeting has been employed globally. In Gaza, the Israel Defence Forces (IDF) utilize an AI platform called "The Gospel" (Habsora) to rapidly identify military infrastructure and buildings, alongside a system known as "Lavender" that has processed vast amounts of surveillance data to flag tens of thousands of individuals as suspected militants for potential airstrikes (Abraham 2024). With these systems in place, the IDF could generate around 100 targets a day, compared to 50 targets per year prior to its introduction (Davies et al. 2023). Similarly, the United States prepared their opening salvo against Iran in February 2026 with AI. Powered by Anthropic's Claude chatbot and Palantir's data engine, the Pentagon's Maven Smart System (MSS) helped military planners select over 1,000 Iranian targets within 24 hours (Copp et al. 2026). This shift to machine speed doubled the strike volume of the 2003 Iraq War's "shock and awe" opening salvo (Irvine 2003). In testing, an MSS targeting cell of just 20 people matched the speed and efficiency of the 2,000-person cell used in Operation Iraqi Freedom, which was previously considered the most efficient in U.S. history (Probasco 2024). Ultimately, MSS enabled "Operation Epic Fury" to strike 13,000 targets in 38 days, averaging more than 340 strikes per day (Harper 2026). The AI system was so extensively employed that peak daily usage reached roughly 20 billion tokens, which is equivalent to the daily output of around 80,000 heavy civilian AI users (Freedberg Jr. 2026). By compressing days of analytical data fusion into mere minutes, these systems dramatically accelerate military planning but also create a further bottleneck, upping the pressure to transition the next phase of warfare, targeting and engagement, from human deliberation to machine-speed execution. We can already observe that shift.

Automated engagement and the human in the loop illusion

Moving up the "kill chain", the process from intelligence gathering to a completed strike, machine speed quickly shifts from an advantage to a baseline requirement for tactical survival. As the former Ukrainian Minister of Defence Mykhailo Fedorov notes, technological dominance requires being faster than the enemy at every single stage (Fedorov 2026). Advanced Ukrainian drones are equipped with on-board cameras and AI computing power, eliminating the need for an uninterrupted line of communication with the pilot, which makes classic first-person view (FPV) drones highly vulnerable to jamming efforts. AI-powered drones can autonomously lock onto a possibly moving target and engage it directly without the need for human interference (Armyinform 2024); thereby also delivering much more precise strikes than a human drone pilot usually could, such as reliably hitting specific vulnerabilities in an armoured vehicle or a tank (Ritters 2026). Naturally, this raises the success rate of a given strike significantly. Ukraine's defence technology market now encompasses more than 200 companies involved in the production of AI-powered drones, with the stated goal of eventually equipping all drones deployed to the front with machine vision and AI capabilities (Myronyshena 2026). Ukraine insists that its drones engage only after human strike confirmation. While this may currently still be true, the need for human confirmation creates the next bottleneck in terms of speed, which will likely need to be sacrificed at the altar of tactical survival in the near future. Moreover, human vetoes are technically impossible in heavily jammed environments where communication cuts out. Even with a live connection, a drone traveling at terminal velocity leaves an operator with only fractions of a second to react. The sheer tempo of modern operations, discussed in the previous chapter, already makes meaningful human review mathematically implausible at scale. This logic is already visible beyond the drone: in January 2026, the Droid TW-7.62 ground robot autonomously detected and tracked targets well enough to take three Russian soldiers prisoner. In April 2026, the Ukrainian army seized an enemy position exclusively via unmanned platforms. A first in military history (Melkozerova 2026).

AI-assisted logistics
AI is also making inroads in further military fields such as logistics and supply. Take for example the US army's AMC predictive Analytics Suite (APAS) that processes vast amounts of data, such as maintenance records, usage patterns, environmental conditions, and operational tempo, thereby providing insights that would be impossible to derive manually (Tyler 2023). It enables commanders to anticipate when critical systems will likely fail and preposition spare parts and maintenance capabilities accordingly. The system has a predictive quality, offering a holistic view of readiness by tracking equipment failures, part availability, delivery timelines, and estimated repair times to actively mitigate operational risks (Hill 2024). These real-time analytics once again deliver a key speed and planning advantage compared to adversaries acting without them. The bogging down of the Russian motorcade in the outskirts of Kiev due to fuel shortage at the beginning of Russia's full-scale invasion may have well been decisive for Ukraine's survival. This is only one of endless historic examples of the critical importance of logistics in warfare.

The shift to machine speed warfare underscores a fundamental transformation in the architecture of global security. As the Chief Digital and Artificial Intelligence Officer (CDAO) of the Pentagon, Cameron Stanley, observed regarding this technological transition: "We've handed our warfighters a Ferrari, and my only sleepless nights come from… making sure we never, ever run out of the high-octane fuel that they need, which is compute" (Harper 2026). Ultimately, 21st-century hard power will not be measured merely by the volume of physical platforms, but by the structural foundations that power them: compute and data.

Reshaping economic state power
This military edge is the national security expression of a broader structural shift: AI fundamentally alters traditional economic capacity. Generative AI alone is projected to add up to $4.4 trillion annually across global use cases, matching or exceeding the entire 2025 GDP of the United Kingdom (Chui et al. 2023, IMF 2025). We are already observing clear production leaps in hyper-automated manufacturing. Xiaomi's flagship smart factory produces a smartphone every second with virtually no human labour (Bochis 2026). Car manufacturers' BYD's Xi'an plant operates at approximately 97% autonomy, deploying AI-driven robotics, automated guided vehicles and intelligent warehousing systems. The cost advantage this generates is concrete: internal analysis of BYD's manufacturing processes concludes that

a BYD car comparable to the Tesla Model 3 costs 15% less to produce than at Tesla's own Shanghai Gigafactory (Darley 2025). The same dynamic is visible in solar panel production with leading automation-heavy facilities achieving production costs below $0.20 per watt; traditional manufacturing costs are more than double (Mose Solar 2026). Naturally, these dynamics apply in a similar way to almost any manufacturing sector.

For most of modern history, economic output, industrial capacity and military strength were ultimately constrained by access to labour, capital and natural resources. AI alters this equation. As increasingly capable systems assume tasks previously performed by humans, the relative importance of sheer population size diminishes while the importance of technological infrastructure grows. States that can deploy advanced AI across military and civilian domains will be able to generate greater economic output, scientific innovation and military effectiveness from the same underlying resource base.

THE AI STACK

Practitioners refer to the AI "stack", as the employment of the technology is dependent on multiple layers being stacked upon one another. Crucially, states need access or control to all layers of the stack to be on the forefront of AI technology.

The physical layer

All frontier AI models are trained and run on advanced computer chips. While major chips companies such as Nvidia or AMD design ever more advanced chips, they outsource their production. Therefore, an estimated 92% of the most advanced chips are manufactured by a single company: Taiwan Semiconductor Manufacturing Company (TSMC) (Varas et al. 2021, Jones & Krulikowski 2024). The remaining 8% come primarily from South Korea, mainly Samsung. This enormous geographic concentration of a key AI input constitutes a major geopolitical chokepoint and supply chain risk around the Taiwan Strait and the Strait of Malacca. Narrow maritime lanes through which vessels must pass through on their export routes to the rest of Eurasia and Africa. To mitigate this risk, target countries offer hefty subsidies to entice TSMC to build factories within their territory: The companies' Arizona expansions were negotiated directly with the U.S. Department of Commerce including up to US$6.6 billion in direct funding under the CHIPS and Science Act (TSMC 2024). Similarly, Germany subsidizes a TSMC factory in Dresden with EUR 7.5 billion, which is scheduled to open in 2027 (Der Spiegel 2024). The subsidy is partly justified as advancing EU semiconductor sovereignty under the European Chips Act, aiming to reduce the continents dependency on foreign supplies. For Taiwan, the export of manufacturing capacity is a dangerous game. Every fab that opens abroad fractionally erodes the "Silicon Shield", which is the deterrent logic by which TSMC's irreplaceable role in the global economy makes a Chinese military takeover extremely costly. Consequently, Taiwan's most advanced process nodes remain on the island, and likely will for as long as the shield logic holds.

This strategically necessary concentration, however, is precisely what nudges partners seeking more technological sovereignty to cultivate homegrown alternatives. In April

2026, Japan's industry ministry announced an additional $4 billion in state funding for Rapidus, bringing total public investment in the semiconductor startup to over $16 billion, as it targets mass production of 2-nanometre chips within the next fiscal year (Mak 2026). In the United States, Intel announced that it will join Elon Musk's "Terafab" AI chip project, which aims to manufacture processors in Austin, Texas, in order supply the tech billionaire's robotics and data centre ambitions (Sophia 2026). However, there is yet another level beneath the computer chips production bottleneck: The Dutch firm ASML, Europe's largest technology company, supplies almost all of the world's lithography machines used to print microscopic circuitry onto silicon wafers, a central step in chips production (Wral News 2025). The de facto monopoly derives from the complexity and intricacy of these machines, combining decades-long research, proprietary complex software and so on, making it extremely difficult for a potential competitor to replicate. Crucially, these machines also rely on critical U.S. components, allowing Washington to pressure the Netherlands into not issuing export licenses to China by threatening to cut their supplies in turn (Sterling 2025). As Beijing aimed to circumvent export restrictions by stockpiling older, less advanced, lithography machines, right before these would once again be included in the export regime, the US reacted by adjusting its de minimis rule to 0%. This means that Washington gains a "long-arm" authority over any product from a non-American company even if the product contains no obvious US linkage (Van der Lugt 2024). Unsurprisingly, this clear violation of trade sovereignty is a friction point between the US and the EU, especially as China constitutes ASML's biggest target market as of 2025 (Wral News 2025). Beijing meanwhile plays the long game and continues to find workarounds in the face of American protectionism: by combining multiple older machines, Chinese companies achieve a similar market-edge less reliably and at higher cost, but the state can compensate for that (Mirraslavska 2026). Moreover, in what has been dubbed China's Manhattan project, a team of former engineers from ASML reverse-engineered the company's extreme ultraviolet lithography machines in a high-security Shenzhen laboratory (Potkin 2025, Stonor 2026). While constituting a major breakthrough, technical challenges remain. It's unclear whether the machine is fully functioning as of June 2026. Still, these developments have alarm bells ringing in Washington sending its export restrictions into overdrive, now also including spare parts, tools and single components (Bradford et al. 2025). In an effort to strengthen techno-sovereignty, China introduced a new law in late 2025 requiring chipmakers to use at least 50% domestically made equipment for adding new capacity (Reuters 2025). Also, the US is actively trying to get their own lithography companies off the ground (Mirraslavska 2026). Additionally, lawmakers have proposed the U.S. MATCH Act in May 2026, requiring US allies to follow Washington's export controls as part of efforts to curb China's ability to make advanced semiconductors (Sterling 2026).

Underlying all these dynamics is a further structural asymmetry that Beijing holds at the base of the supply chain: rare earth elements, the critical minerals used in semiconductors and the broader electronics ecosystem. China controls not only the majority of global rare earth deposits (ca. 70%) but dominates the refining and processing capacity (ca. 90%) needed to turn raw ore into usable inputs (Ang 2026). Meaning that even deposits mined elsewhere largely flow through Chinese facilities. Both the United States and the European Union have moved to address this vulnerability through a combination of domestic mining incentives, strategic stockpiling, and bilateral agreements with resource-rich partners such as Australia, Canada, Greenland and various African nations, aiming to develop alternative refining capacity and reduce the chokehold China maintains at this foundational layer of the AI supply chain (Natural Resources Canada 2024, Australian Government 2026, EU Commission 2026, Okebiorun 2026, Reuters 2026).

The physical layer of the AI stack is characterized by various potential chokepoints on different levels, each primarily controlled by a different actor. This leaves the US, EU, China and Taiwan in a mutual web of dependency. But its an asymmetric one: Washington's leverage over ASML's exports rests on the Netherlands' own security and component dependence on the US, a constraint China's rare-earth chokepoint doesn't share.

The infrastructural layer

Moving up the AI stack, data centres sit at the intersection of the physical and the infrastructural layer. These physical facilities, packed with high-end computer chips, massive cooling units, and energy storage, are needed to train and run AI models. The larger the data centre, the more compute it provides, and the more sophisticated and powerful the AI model that can be trained or deployed. While training an AI model carries a large upfront compute cost, running a finished model is comparatively modest. However, as models grow more sophisticated and are tasked with more complex operations, inference costs rise accordingly. Meaning that deploying advanced agents or handling demanding workloads requires substantially more compute than simply querying a basic model. Whereas inference constituted only around 25% of data centre capacity demand in 2025, with the remainder driven by model training, this relationship is expected to flip by 2030, with inference workloads overtaking training as the primary driver of data centre demand (JLL Research 2026).

How much data centre capacity a state needs therefore depends entirely on what it seeks to do with AI. Training frontier models from scratch demands an altogether different scale of infrastructure than hosting and running complex agent systems, which in turn far exceeds what is required to simply deploy a ready-made open-source model for limited governance applications. The necessary investment is, in short, a function of ambition. Clarity on this question should therefore precede any national AI infrastructure strategy, as the answer shapes everything from procurement and energy planning to the international partnerships a state will need to pursue.

As of now, AI-ready compute capacity is highly unevenly distributed globally. Measured in GPU cluster performance, the United States accounts for roughly three-quarters of it, with China in second place at 15% and the EU a distant third at around 5% (Pilz et al. 2025a). Global capacity is expected to double between 2026 and

2030 (JLL Research 2026), without fundamentally shifting this balance, as expansion will be driven overwhelmingly by US hyperscalers. This will likely further accelerate the ongoing privatization of compute: the private sector's share of global AI computing capacity reached 80% in 2025, up from 40% in 2019 (Pilz et al. 2025b). The much-publicized Stargate project, a $500 billion initiative bringing together US Big Tech and the UAE's MGX sovereign investment fund, aims to build a supercomputer infrastructure spanning multiple continents, targeting a total capacity of 10 gigawatt (GW) by 2028 (Stargate How 2026). Amazon, Google, and Meta are all planning, building, and in some cases already opening data centres across the United States that will eventually each compile multiple gigawatts of capacity (Sigalos 2025).

China is also aggressively expanding AI data centre infrastructure, with compute capacity expected to nearly triple between 2025 and 2030, rising from 10.5 to 28.9 GW (Deng et al. 2026). Here too the private sector plays a central role: China's four largest internet companies, Alibaba, Tencent, ByteDance, and Baidu, have committed a combined $84 billion in AI infrastructure investment by 2027, a 60% increase from 2025 levels (Blackwood 2026).

Gulf states have also emerged as significant players. During President Trump's visit to the region in May 2025, Saudi Arabia, the UAE, and Qatar collectively committed to around $2 trillion in AI-related investments (Dent 2025). Abu Dhabi has announced the UAE–US AI Campus, the largest such AI infrastructure project outside the United States, eventually targeting 5 GW of associated data centre capacity and set to host several US hyperscalers (Embassy of UAE in Washington 2025). Saudi Arabia, meanwhile, is targeting 1.8 GW of AI-ready data centres by 2030 through $18 billion in total hyperscale investment (Arab News 2026). As a result of these commitments, the MENA region's data centre capacity is expected to triple by 2030 (Shiwani, Abbasi & Levack 2025).

Similarly, the European Union aims to also at least triple its own data centre capacity over the next five to seven years anchored by the Cloud and AI Development Act (CADA) and a public-private push to build AI gigafactories, accompanied by €20 billion in committed investment. The rollout has, however, faced repeated delays, with the bidding process currently scheduled to open in July 2026 (Kyriasoglou, Volpicelli & Doenecke 2026). With US cloud providers already dominating the European market and being vastly outspent by Big Tech in terms of AI data center construction, the most the EU can realistically hope for in the near term is a managed dependence on US hyperscalers.

What this landscape makes clear is that, even accounting for the significant investments underway across the Gulf, Europe, and Asia, the structural dominance of the United States, and to a lesser extent China, in AI compute is unlikely to be meaningfully challenged in the foreseeable future. The scale advantages, the private capital, and the incumbency of US hyperscalers are simply too large to close within a single investment cycle. For most states, the realistic choice is therefore not between dependency and sovereignty, but between different forms of dependency: on US cloud infrastructure, on Chinese state-backed alternatives, or on some combination of both. The strategic question for most states is therefore not how to escape this dependency, but how to structure and mitigate it; i.e. which layer to secure, which to concede, and to whom.

Crucially, the vast majority of the world's AI compute capacity is not in state hands at all, but controlled by a small number of private corporations, predominantly American ones. The geopolitics of AI infrastructure is therefore not simply a competition between states, but increasingly a question of the relationship between states and these corporations. Governments that wish to shape their AI futures will need to negotiate not only with rival powers, but with the firms on whose infrastructure their digital economies increasingly depend. In this sense, the increasing demand for compute capacity should also be seen as a growing structural source of private power.

The model and data layer

The relevant frontier AI models themselves are also technically all in private hands. A frontier model is a highly advanced, general-purpose AI system representing the cutting edge of current capabilities. There are currently a handful of models from different AI companies that compete for top performance across several benchmarks (Zafar 2026). Of these, all but one, belong to US or Chinese tech companies. The sole exception being a model of France's Mistral AI. In terms of raw intelligence, these frontier models are clustered very tightly. Because their general capabilities are roughly identical, the competition is shifting towards domain-specific performance, reliability and latency as well as, crucially, cost or efficiency (AI Index 2026).

As an AI model is only as good as the data it was trained on, finding ways to access higher quality, and preferably exclusive data, is a possible avenue of maintaining the strategic edge. As most of the publicly available data has already been exhausted for model training, the most lucrative data sources now lie in private, proprietary contexts: reinforcement learning from human feedback (RLHF) generated through product use, expert annotation by domain specialists, and institutional data from healthcare, law, and finance. Access to these sources predominantly remains concentrated in private, US and Chinese tech firms (Varela Sandoval & Wilkinson 2026). Driven by this domestic exhaustion of Western and East Asian internet data, a secondary race has emerged to extract value from the last remaining linguistic and cultural data frontiers, particularly within the Global South. Big Tech firms are increasingly looking to emerging markets in Africa, Latin America, and South Asia to harvest local data, driving concerns over a new form of digital extractivism where regional data is exported to train frontier models in developed nations without local value retention (Signé 2026).

In terms of security applications, also the vast amounts of data generated at the frontlines of 21st century conflicts, constitute a highly strategic resource. In this context, the Ukrainian Ministry of Defence has launched the Brave1 Dataroom, which grants over 100 allied defence companies access to more than 2 million hours of combat drone footage to feed directly into AI training (Mykhailenko 2026). As with every commodity along the AI stack, also cross-border data flows that seemed routine only a few years ago, now increasingly face stricter oversight and restrictions, through state policies aiming to shield this strategic resource. Governments from the EU to China are implementing directives to keep sensitive and valuable data within their borders (Esposito 2025). Ultimately, this landscape entrenches a rigid global hierarchy. As advanced states increasingly securitize and ring-fence their proprietary data assets, the barrier to entry rises, making it nearly impossible for digitally nascent nations to overcome their structural data disadvantage. This means that training cutting edge frontier models domestically is structurally implausible for most states.

However, states actually don't need domestically produced frontier models to be independently competitive in AI. This is due to the availability of highly competitive open-source models that can be acquired and relatively inexpensively repurposed by states lacking their own proprietary models. Depending on the ambition and target application, this still usually requires some fine-tuning training of the open-source model, to make it fit for local context and the desired task. Still, this procedure constitutes a relatively cost-efficient alternative to training frontier models from scratch. Even though the capability gap between closed frontier models and open-source ones has widened between 2024 and 2025, the gap is still largely negligible for most applications a state would desire. Historically, the relationship between closed and open-source models has followed a rubber-band logic: proprietary labs conquer a new capability frontier through massive capital investment, stretching the gap temporarily, before the open-source community inevitably snaps back, catching up as frontier performance reaches a new temporary plateau (AI Index 2026). Moreover, the time it takes to close the gap has narrowed: open models now catch up to closed frontier ones within 3-6 months rather than the 6+ months observed in early 2024 (Nagel & Yue 2025).

A key driver of this snap-back effect, is a technique called distillation, whereby a sophisticated "teacher" model trains a "student" model on reasoning and output capability without the latter having to learn it from scratch. This practice allows to near replicate a model at a fraction of the cost a genuine frontier model development would require. While distillation is a legitimate technique employed for many AI use cases, US Big Tech accuses primarily Chinese tech companies of free-riding on American innovation and capital expenditure by distilling their models into open-source ones. To better shield themselves against "distillation attacks" OpenAI, Anthropic and Google entered a rare collaboration in which they share information on adversarial distillation among each other (Ghaffary & Eastland 2026).

By replicating comparable capabilities at a fraction of the cost, distillation creates a massive asymmetric advantage for late-movers and geopolitical competitors. It erodes the strategic edge the US derives from compute, data and innovation dominance. As the geopolitical race for AI dominance accelerates, the incentive to capture cutting-edge capabilities through low-cost methods will extend beyond China. Despite the introduction of counter measures, reliable prevention of illicit distillation at scale will remain unlikely (Singh 2026).

The talent layer

The ultimate layer of the artificial intelligence stack is ironically human intelligence. Globally, there is only an extremely limited number of researchers, some estimates put the total under 1,000, that are capable of building and advancing AI frontier models (Jackson 2025). This scarcity reflects not merely the novelty of the field, but the unusually demanding combination of skills required: systems engineering at massive scale, deep machine learning theory, and the empirical intuition that only comes from running large training experiments over years. This profile takes roughly a decade to produce, which is why no amount of capital can rapidly expand the pool.

This scarcity has led Big Tech companies to engage in salary bidding wars, with Meta, Google, OpenAI and others offering the most promising AI researchers multi-year contracts in the hundred-million-dollar compensation range, comparable to top NBA stars (Isaac, Tan and Metz 2025). Naturally, this has led to a high concentration of top-tier talent at these companies. As in other layers of the stack, the US currently holds the lead in terms of absolute numbers of AI researchers. However, India as well as the EU and UK combined don't lag far behind on this metric (Pal, Schneider & Lazzaroni 2026). The UK in particular punches above its weight, with DeepMind in London representing one of the densest concentrations of frontier AI research outside the United States. Reliable figures for the Chinese context aren't publicly accessible, but informed estimates put it roughly on par with the US in absolute numbers of AI researchers (Maes & Sawaya 2023).

However, raw headcount figures mask a critical dynamic: the EU and UK's strong absolute numbers are substantially undermined by structural brain drain. The salary difference between American labs and European institutions is so extreme that European-trained researchers systematically migrate to the US, meaning a significant share of European-counted talent effectively works for American firms. China faces a different but related problem: despite strong domestic numbers, Chinese AI researchers educated abroad have historically remained in the US at high rates. Post-2022 geopolitical tensions and tightening US visa conditions have accelerated return migration, though the number still remains quite low (Sha 2026).

Across the world, states have introduced AI talent retention and attraction policies in response. China is reportedly restricting overseas travel for individuals involved in advanced AI work (Nusca 2026), the UAE has introduced a new AI Specialist Visa (Department of Economic Development 2026), and the recently concluded EU-India free trade agreement explicitly prioritises ease of tech talent mobility. The EU Commission also aims to attract US talent with the €500 million 'Choose Europe' initiative (Grey 2026).

By levying a hefty $100,000 fee on new H-1B visas, the Trump administration appeared to impede its own ability to attract foreign tech workers. While across the broader tech sector new H-1B visa applications did plummet significantly, the opposite proved true for the capital-rich, leading AI labs, which were ready to absorb the higher cost amid rising application numbers (Munis 2026).

Finally, the extreme concentration of indispensable expertise in so few individuals introduces a strategic vulnerability distinct from every other layer of the stack. Unlike compute or data, talent cannot be stockpiled or export-controlled in any conventional sense.

A TWO-TIERED RACE

The preceding analysis has shown geopolitical manoeuvres by states and companies to carve out strategic advantages and curtail competitors' ambitions, across every layer of the AI stack. This may suggest a single, unified competition on the global stage, in which relative position on that ladder determines strategic standing. That framing, while intuitive, is misleading. The AI stack is not one race with many contestants. It is two fundamentally different competitions, running in parallel, with different rules, different finishing lines, and therefore different winning strategies.

For the United States and China, AI is genuinely first-past-the-post. Frontier model superiority, compute dominance, and talent concentration are key determinants in the struggle for global hegemony. Breaking into that loop from the outside is categorically different from sustaining it from within. For these two actors, and realistically only for these two, being second is a meaningful and compounding strategic loss.

For everyone else, the logic is categorically different and considerably more forgiving. As the preceding layers have shown, full-stack AI sovereignty is a structural near-impossibility for all but the most resource-rich states. To varying degrees, every state outside the US-China dyad, will likely remain constrained by one or more chokepoints across the AI stack. Rare earth processing runs overwhelmingly through Chinese facilities. The lithography machines that print the world's most advanced chips come almost exclusively from a single Dutch firm subject to American extraterritorial pressure. The vast majority of the most advanced semiconductors are manufactured by a single company on a contested island. Compute capacity is 80% privately held, dominated by a handful of American hyperscalers. The very limited number of researchers capable of building frontier models are concentrated in a handful of labs, primarily in the United States. At every layer, a different actor holds the chokepoint, and the web of mutual dependencies that results is a map of shared vulnerability on the one hand, and a hierarchy of leverage on the other. The room to manoeuvre within that dependency is not constant across the stack. Where a layer is oligopolistic (compute and talent) states retain real bargaining power over the terms of reliance. Where it is monopolistic the only choice left is to become dependent or not to play the game at all.

And yet, the strategic implications for most states are not nearly as dire as this architecture implies. The reason is the rubber-band dynamic described at the model layer. Distillation allows a sophisticated model's reasoning capabilities to be replicated at a fraction of the original development cost. The incentive to replicate this approach will only intensify as the geopolitical stakes rise, and reliable prevention of illicit distillation at scale will remain unlikely. However, the rubber-band dynamic is an empirical observation about the recent past, not a structural law. This logic may dissolve if we were to reach a point where distillation is no longer technically feasible, which would leave second-tier states that have built their

AI strategies around open-source access stranded. The one layer that does not yield to this more accessible logic is talent: the very limited number of AI researchers that is capable of pushing the frontier cannot be distilled, stockpiled, or replicated at low cost. This makes talent the most durable source of structural AI advantage.

What this means in practice is that the capabilities sufficient for most state ambitions, such as battlefield targeting, logistics optimisation, autonomous systems, economic automation, domain-specific governance tools, are accessible without sovereign frontier model development. The capabilities reshaping security doctrine and economic competitiveness today are, for the most part, reachable by states willing to invest in the institutional capacity to deploy AI effectively at the relevant layers of the stack, without needing to own or control all of them. Still, the more capable AI systems become, the bigger the rewards they will reap and the higher their inference costs (likely) will climb. This dynamic will make privileged access to compute even more strategically vital than it is today, where a modest infrastructure footprint is still sufficient if a state primarily wants to employ models rather than train them.

The great AI game, properly understood, is not a single ladder that all states are climbing at different speeds. It is a contest with two distinct tiers. In the first, Washington and Beijing compete for frontier dominance in a genuinely first-past-the-post dynamic where being second compounds into strategic disadvantage over time. In the second, the rest of the world competes not for full sovereign control of the stack but for the terms of their dependency on it, and for the institutional capacity to translate accessible AI capabilities into concrete military, economic, and governance outcomes. The capabilities that are already reshaping how states fight, produce, and project power do not, for the most part, require sitting at the cutting edge of every layer. They require knowing which edge matters, for which purpose, and at what cost.

REFERENCES

- Abraham, Y. (2024). 'Lavender': The AI machine directing Israel's bombing spree in Gaza. +972 Magazine. https://www.972mag.com/lavender-ai-israeli-army-gaza/
- Ang, A. (2026). Beijing's dominance in rare earth processing leaves others scrambling to close the gap. Fortune. https://fortune.com/2026/03/11/china-us-rare-earth-processing-critical-minerals/
- Arab News. (2026). Saudi data center capacity expands sixfold since Vision 2030 launch. https://www.arabnews.com/node/2639053/business-economy
- ArmyInform. (2024). AI and machine vision in drone training. https://armyinform.com.ua/2024/06/22/sou-zastosovuyut-shtuchnyj-intelekt-i-mashynnyj-zir-dlya-navchannya-svoyih-droniv/
- Australian Government. (2026). Australia and United States advance cooperation on critical minerals and rare earths supply chains. Department of Industry, Science and Resources. https://www.industry.gov.au/news/australia-and-united-states-advance-cooperation-critical-minerals-and-rare-earths-supply-chains
- Bergengruen, V. (2024). How tech giants turned Ukraine into an AI war lab. Time Magazine. https://time.com/6691662/ai-ukraine-war-palantir/
- Blackwood, S. (2026). Data center build-out in China. Jon Peddie Research. https://www.jonpeddie.com/news/data-center-build-out-in-china/
- Bochis, N. (2026). Xiaomi's dark robot factory can make a phone every second without people. BGR. https://www.bgr.com/2087200/xiaomi-dark-robot-smart-phone-factory/
- Bradford, A., Li, E., & Waxman, M. C. (2025). How domestic institutions shape the global tech war. Harvard National Security Journal, 16, 75.
- Chui, M., Manyika, J., Miremadi, M., Henke, N., Chung, R., Nel, P., & Malhotra, S. (2023). The economic potential of generative AI: The next productivity frontier. McKinsey Global Institute. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier
- Copp, T., Dwoskin, E., & Duncan, I. (2026). Anthropic's AI tool Claude central to U.S. campaign in Iran, amid a bitter feud. The Washington Post. https://www.washingtonpost.com/technology/2026/03/04/anthropic-ai-iran-campaign/
- Darley, J. (2025). How BYD is manufacturing EVs at an unprecedented scale. Manufacturing Digital. https://manufacturingdigital.com/news/how-chinas-byd-is-using-ai-to-scale-global-ev-manufacturing
- Davies, H., McKernan, B., & Sabbagh, D. (2023). The Gospel: How Israel uses AI to select bombing targets in Gaza. The Guardian. https://www.theguardian.com/world/2023/dec/01/the-gospel-how-israel-uses-ai-to-select-bombing-targets
- Deng, S., Moe, T., Haase, C., & Haugom, E. (2026). China's data center capacity set to top 60 GW by 2030, driving a doubling of power demand. Rystad Energy. https://www.rystadenergy.com/news/chinas-data-center-capacity-doubling-of-power
- Dent, E. (2025). Unpacking Trump's 2025 Gulf investment tour. The Washington Institute for Near East Policy. https://www.washingtoninstitute.org/policy-analysis/unpacking-trumps-2025-gulf-investment-tour
- Department of Economic Development, Abu Dhabi. (2026). Abu Dhabi golden visa for talents in digital technology. https://www.added.gov.ae/en/live/long-term-residency/abu-dhabi-golden-visa/for-brilliant-minds/for-talents-in-digital-technology
- Der Spiegel. (2024). EU-Kommission genehmigt Milliardensubvention für TSMC-Fabrik. https://www.spiegel.de/wirtschaft/unternehmen/tsmc-in-dresden-eu-kommission-genehmigt-milliardensubvention-fuer-fabrik-a-44cc5646-2dc3-4c6b-8f80-c914df9363c0
- Embassy of the United Arab Emirates in Washington, D.C. (2025). UAE, US presidents attend unveiling of phase 1 of new 5GW AI campus in Abu Dhabi. https://www.uae-embassy.org/news/uae-us-presidents-attend-unveiling-phase-1-new-5gw-ai-campus-abu-dhabi
- Esposito, M. (2026). AI geopolitics and data centres in the age of technological rivalry. World Economic Forum. https://www.weforum.org/stories/2025/07/ai-geopolitics-data-centres-technological-rivalry/
- European Commission. (2026). EU and US launch strategic partnership on critical minerals. https://ec.europa.eu/commission/presscorner/detail/en/ip_26_862
- Freedberg, S., Jr. (2026). 'Insatiable appetite' for AI: Maven usage surged for strikes on Iran, Pentagon AI chief says. Breaking Defense. https://breakingdefense.com/2026/05/insatiable-appetite-for-ai-maven-usage-surged-for-strikes-on-iran-pentagon-ai-chief-says/
- Ghaffary, S., & Eastland, M. (2026). OpenAI, Anthropic and Google unite to combat model copying in China. Australian Financial Review. https://www.afr.com/world/north-america/openai-anthropic-google-unite-to-combat-model-copying-in-china-20260407-p5zlr5
- Grey, C. (2026). Europe's battle for American talent. EURACTIV. https://www.euractiv.com/opinion/europes-battle-for-american-talent/
- Harper, J. (2026). DOD planning to address compute 'bottleneck' that could hinder AI proliferation. Defense Scoop. https://defensescoop.com/2026/05/07/dod-planning-to-address-compute-bottleneck-ai-proliferation/

- Hill, C. (2024). Analytics at the edge. United States Army Acquisition Support Center. https://asc.army.mil/web/news-analytics-at-the-edge/
- International Monetary Fund. (2025). GDP, current prices [Data set]. https://www.imf.org/external/datamapper/NGDPD@WEO/OEMDC/ADVEC/WEOWORLD/GBR
- Irvine, J. (2003). Iraq War: Shock and awe assault on Baghdad begins [Video]. ITN Frontline. https://www.youtube.com/watch?v=m8KimNtB9HI
- Isaac, M., Tan, E., & Metz, C. (2025). A.I. researchers are negotiating $250 million pay packages—just like NBA stars. The New York Times. https://www.nytimes.com/2025/07/31/technology/ai-researchers-nba-stars.html
- Jackson, S. (2025). The AI talent wars are white hot. Business Insider. https://www.businessinsider.com/ai-talent-wars-recruiting-hiring-white-hot-2025-6
- JLL Research. (2026). 2026 global data center market outlook. JLL. https://www.jll.com/en-us/insights/market-outlook/data-center-outlook
- Jones, L., & Krulikowski, S. (2024). Taiwan—The Silicon Island (Executive Briefing on Trade). U.S. International Trade Commission. https://www.usitc.gov/publications/332/executive_briefings/ebot_silicon_island_taiwan_semiconductor.pdf
- Kyriasoglou, C., Volpicelli, G., & Doenecke, P. (2026). EU's AI data center plans stumble due to delays, funding. Bloomberg. https://www.bloomberg.com/news/articles/2026-06-02/eu-s-ai-data-center-plans-stumble-due-to-delays-funding-issues
- Maes, W., & Sawaya, A. (2023). How businesses can close China's AI talent gap. McKinsey & Company. https://www.mckinsey.com/capabilities/quantumblack/our-insights/how-businesses-can-close-chinas-ai-talent-gap
- Mak, R. (2026). TSMC confronts new contenders for chipmaking crown. Reuters. https://www.reuters.com/commentary/breakingviews/tsmc-confronts-new-contenders-chipmaking-crown-2026-04-15/
- Melkozerova, V. (2026). Robots captured Russian army positions for first time in history, Zelenskyy says. Politico. https://www.politico.eu/article/volodymyr-zelenskyy-robotic-systems-russia-army-positions-ukraine/
- Mirraslavska, A. (2026). The chip monopoly is at the centre of the new cold war. Cutting Takes. https://cutting-takes.com/the-chip-monopoly/
- Mose Solar Manufacturing. (2025). How robots are revolutionizing solar panel production lines. https://www.moserbaersolar.com/solar-pv-manufacturing-and-development/how-robots-are-revolutionizing-solar-panel-production-lines/
- Munis, J. (2026). OpenAI and Nvidia CEOs didn't flinch at Trump's $100,000 H-1B visa fee, and now they're paying up as their application numbers soar. Fortune. https://fortune.com/2026/06/10/sam-altman-jensen-huang-h1b-visa-fee-application-numbers-trump-lawsuit/
- Mykhailenko, D. (2026). Ukraine opens AI dataroom with real battlefield data for defense companies. United24 Media. https://united24media.com/defense-tech/ukraine-opens-ai-dataroom-with-real-battlefield-data-for-defense-companies-19741
- Myronyshena, T. (2026). Ukraine builds AI-driven defense ecosystem as over 200 companies develop drone technologies. The Kyiv Independent. https://kyivindependent.com/ukraine-builds-ai-driven-defense-ecosystem-as-over-200-companies-develop-drone-technologies/
- Nagle, F., & Yue, D. (2025). The latent role of open models in the AI economy (Georgia Tech Scheller College of Business Research Paper No. 5767103). SSRN. https://ssrn.com/abstract=5767103
- Natural Resources Canada. (2024). Government of Canada and the United States co-invest to strengthen critical mineral value chains. https://www.canada.ca/en/natural-resources-canada/news/2024/05/government-of-canada-and-the-united-states-co-invest-to-strengthen-critical-mineral-value-chains.html
- Nusca, A. (2026). China restricts travel for top AI talent. Fortune. https://fortune.com/2026/05/27/china-restricts-travel-for-top-ai-talent/
- Okebiorun, O. (2026). Africa faces new competition for critical minerals investment as U.S. secures Greenland rare earth deal in bid to cut China reliance. Business Insider Africa. https://africa.businessinsider.com/local/markets/africa-faces-new-competition-for-critical-minerals-investment-as-us-secures-greenland/mjplplw
- Pal, S., Schneider, C., & Lazzaroni, R. (2026). Talent in, talent out. Interface. https://www.interface-eu.org/publications/talent-in-talent-out
- Pilz, K., Erdil, E., Besiroglu, T., & Sevilla, J. (2025a). The US hosts the majority of GPU cluster performance, followed by China [Data insight]. Epoch AI. https://epoch.ai/data-insights/ai-supercomputers-performance-share-by-country
- Pilz, K., Erdil, E., Besiroglu, T., & Sevilla, J. (2025b). Private-sector companies own a dominant share of GPU clusters [Data insight]. Epoch AI. https://epoch.ai/data-insights/ai-supercomputers-performance-share-by-sector
- Potkin, F. (2025). Exclusive: How China built its 'Manhattan Project' to rival the West in AI chips. Reuters. https://www.reuters.com/world/china/how-china-built-its-manhattan-project-rival-west-ai-chips-2025-12-17/
- Probasco, E. (2024). Building the tech coalition: How Project Maven and the U.S. 18th Airborne Corps operationalized software and artificial intelligence for the Department of Defense. Center for Security and Emerging Technology, Georgetown University. https://cset.georgetown.edu/publication/building-the-tech-coalition/
- Reuters. (2025). Exclusive: China mandates 50% domestic equipment rule for chipmakers, sources say. https://www.reuters.com/world/china/

china-mandates-50-domestic-equipment-rule-chipmakers-sources-say-2025-12-30/
- Reuters. (2026). Australia and EU seal trade deal to cut reliance on China's rare earths. https://www.aljazeera.com/news/2026/3/24/australia-and-eu-seal-trade-deal-to-cut-reliance-on-chinas-rare-earths
- Ritters, R. (2026). Ukraine is now the world's AI war lab. Deutsche Welle. https://www.dw.com/en/ukraine-is-now-the-worlds-ai-war-lab/video-77409971
- Sha, L. (2026). The returnees: Inside China's AI talent reversal. The Business Times. https://www.businesstimes.com.sg/international/returnees-inside-chinas-ai-talent-reversal
- Shiwani, A., Abbasi, H., & Levack, A. (2026). Unlocking the data centre opportunity in the Middle East. PwC. https://www.pwc.com/m1/en/media-centre/articles/unlocking-the-data-centre-opportunity-in-the-middle-east.html
- Sigalos, M. (2025). Amazon opens $11 billion AI data center in rural Indiana as rivals race to break ground. CNBC. https://www.cnbc.com/2025/10/29/amazon-opens-11-billion-ai-data-center-project-rainier-in-indiana.html
- Signé, L. (2026). How to bridge the global AI divide. Brookings Institution. https://www.brookings.edu/articles/how-to-bridge-the-global-ai-divide/
- Singh, V. (2026). AI distillation attacks in the US–China contest. International Institute for Strategic Studies. https://www.iiss.org/online-analysis/cyber-power-matrix/2026/05/ai-distillation-attacks-in-the-uschina-contest
- Sophia, D. (2026). Intel joins Musk's Terafab AI chip project to power humanoid, data center goals. Reuters. https://www.reuters.com/business/autos-transportation/intel-join-musks-terafab-mega-ai-chip-project-2026-04-07/
- Stanford University Human-Centered Artificial Intelligence. (2026). AI index report 2026. https://hai.stanford.edu/ai-index/2026-ai-index-report/technical-performance
- Sterling, T. (2025). Dutch government excludes most ASML sales to China from 'dual use' export data. Reuters. https://www.reuters.com/technology/dutch-government-excludes-most-asml-sales-china-dual-use-export-data-2025-01-17/
- Sterling, T. (2026). Dutch government objects to proposed US law restricting ASML's China exports. Reuters. https://www.reuters.com/world/asia-pacific/dutch-government-objects-proposed-us-law-restricting-asmls-china-exports-2026-05-14/
- Stonor, J. (2026). ASML shares slip on report of U.S. concern that China has its tech. The Wall Street Journal. https://www.wsj.com/business/asml-shares-slip-on-report-of-u-s-concern-that-china-has-its-tech-f6400c8f
- TSMC. (2024). TSMC Arizona and U.S. Department of Commerce announce up to US$6.6 billion in proposed CHIPS Act direct funding [Press release]. https://pr.tsmc.com/english/news/3122
- Tyler, S. (2023). Army predictive logistics underpins success in Europe. U.S. Army. https://www.army.mil/article/267532/army_predictive_logistics_underpins_success_in_europe
- Van der Lugt, S. (2024). Why ASML is the EU's most important bargaining chip. Bocconi Institute for Policy Making. https://iep.unibocconi.eu/why-asml-eus-most-important-bargaining-chip
- Varas, A., Varadarajan, R., Palma, R., Goodrich, J., & Yinug, F. (2021). Strengthening the global semiconductor supply chain in an uncertain era. Boston Consulting Group & Semiconductor Industry Association. https://www.bcg.com/publications/2021/strengthening-the-global-semiconductor-supply-chain
- Varela Sandoval, F., & Wilkinson, I. (2026, February). How middle powers can weather US and Chinese AI dominance [Research paper]. Chatham House. https://www.chathamhouse.org/sites/default/files/2026-02/2026-02-13-sovereign-ai-strategies-sandoval-et-al.pdf
- Zafar, S. (2026). LLM leaderboard 2026: Best AI models ranked by benchmark, speed and price. ClickRank AI. https://www.clickrank.ai/llm-leaderboard/

www.oiip.ac.at

Österreichisches Institut
für Internationale Politik
Austrian Institute
for International Affairs