FRIEDRICH NAUMANN FOUNDATION For Freedom.
Taiwan / Global Innovation Hub

Centres for European Policy Network
FREIBURG | BERLIN | PARIS | ROMA | HELSINKI

POLICY PAPER

GERMANY NEEDS A DUAL-TRACK STRATEGY FOR AI SOVEREIGNTY

A Layer-by-Layer Assessment of the AI Stack, with Lessons for Middle Powers

Anselm Küsters

Publication Credits

Publisher
Taiwan Office/ Global Innovation Hub
Friedrich Naumann Foundation for Freedom
15F.-6, No. 171, Songde Road,
Xinyi District, Taipei City 110030
Taiwan

freiheit.org/Taiwan
FNFGIHUB
FNFreiheit
FNFGIHUB

Author
Dr. Anselm Küsters, LL.M.
Scientific Director cep Nordic Countries,
Head of Division Digitalisation & AI
kuesters@cep.eu

Editor
Taiwan Office/Global Innovation Hub,
Friedrich Naumann Foundation
for Freedom

Contact
Phone +886 2 2728 1500
Email globalinnovation@freiheit.org

Date
November 2026

Notes on using this publication
This publication is an information offer of the Friedrich Naumann Foundation for Freedom. It is available free of charge and not intended for sale. It may not be used by parties or election workers for the purpose of election advertising during election campaigns (federal, state or local government elections, or European Parliament elections).

License
Creative Commons (CC BY-NC-ND 4.0)

Content

Executive summary 4

1. Introduction 5

2. Mapping the AI Stack 8

2.1. Turning Jensen Huang's "5-Layer Cake" into an "AI Burger" 8
2.2. Regulatory layer I (inputs) 10
2.3. Energy 13
2.4. Hardware (chips) 17
2.5. Infrastructure (cloud, software, and compute services) 21
2.6. Models (including training data) 24
2.7. Applications and agents 28
2.8. Regulatory layer II (outputs) 31
2.9. Interim conclusion: Germany's stacked position 34

3. Adoption bottlenecks 37

3.1. Behavioural bottlenecks 37
3.2. Regulatory bottlenecks 38
3.3. Market-structural bottlenecks 39

4. The path forward 42

4.1. Trying to catch up at the AI frontier 42
4.2. Alternative paradigms 43
4.3. Policy and industry recommendations 45

5. Conclusion 50

Endnotes 51

6. Literature list 66

About the Author 75

Executive summary

Germany occupies a difficult position in the AI stack. On the one hand, it lags far behind the US and China in terms of compute and has no significant presence in the production of cutting-edge foundation models. On the other hand, however, it has the talent, commercial footprint, and high-quality data to succeed in agentic tooling, orchestration, and industrial AI – areas most likely to become strategically important as pure scaling yields diminishing returns. However, behavioural inertia, switching costs, and compliance burdens, especially on smaller firms, all suppress the uptake of existing European alternatives. Closing this gap requires a dual-track strategy that pursues sustained frontier catch-up and niche specialisation simultaneously. This strategy must be backed by long-term public funding and demand-side correction.

- Based on web-scraped market data, this study finds that Germany holds a technically credible niche in high-quality data, industrial AI, agentic tooling (software that lets AI systems carry out multi-step tasks), and orchestration (systems that coordinate different AI tools). At the same time, the country's compute capacity, measured in H100-equivalent GPUs, is negligible compared to the US and China, while its footprint in frontier foundation models as well as safety-and-guardrails products is close to zero.
- Organisational inertia and legal uncertainty, rather than a lack of technical capability or European alternatives per se, can be identified as the main obstacles to domestic AI adoption. With 77% of firms reporting some AI use, but only 47% reporting moderate-to-significant deployment, and legal uncertainty cited by 54% of businesses as a barrier, closing the adoption gap should be the priority for near-term competitiveness.
- Rather than subsidising every layer of the AI stack indiscriminately, the country should follow a dual-track strategy that combines long-term frontier research with targeted support for adoption of European AI tools. In addition to establishing a moonshot-style tournament for frontier training (to sustain a deep engineering and IT talent base), policymakers must correct adoption barriers through mandatory interoperability in public procurement, competition law scrutiny of bundling within the AI stack, and support for SME uptake.

1. Introduction

Artificial intelligence (AI) has become a core dimension of geopolitical competition. Countries that control the key layers of the AI technology stack – from energy and computing infrastructure through training data and general-purpose models to deployment tools – gain advantages that build over time. Those that rely on foreign systems risk importing not only products but also, as Germany's Federal Minister for Digital Transformation and Government Modernisation Karsten Wildberger has warned, "a goal function, a worldview and a value system" that may not match their own interests.[1]

Against this backdrop, **the concept of AI sovereignty has moved to the centre of European policy debate.** The EU AI Act, the AI Continent Action Plan, and initiatives such as the 2025 Franco-German AI summit all aim to reduce strategic dependencies. Yet the term itself remains contested: does AI sovereignty require control over every layer of the stack, or only over those layers that cannot easily be substituted by others?[2] This study adopts the latter view and argues that Germany must identify which layers are truly strategic, while also ensuring that domestic developers can attempt to build advanced systems and that domestic users adopt them.

Since the launch of ChatGPT, attention has focused on large language models (LLMs), that is, broad-purpose models trained on massive datasets that other applications build upon. Since European developers contribute only a small share of the world's most capable foundation models (see Section 2), it makes sense to **concentrate on alternative approaches, such as industrial "world models" that learn physical processes or smaller language models that firms can run on their own infrastructure**.[3] However, recent advances by Chinese frontier labs such as Moonshot AI partly challenge this logic. Despite tight computing constraints, they have reached near-frontier performance in coding and research tasks (possibly by distilling[4] knowledge from Western systems). More important than any single "DeepSeek shock" is the underlying catch-up strategy: **China has built a deep engineering talent base and experimented with new model architectures instead of merely copying existing systems.**[5] Europe's current funding architecture – short-term prototyping grants, slow approval cycles, and limited risk capital for long-term engineering work – is poorly matched to that kind of sustained effort.

This study therefore proposes a middle-power strategy for Germany comprising two complementary approaches. The first focuses scarce public resources on areas in which Germany already has an industrial base and credible prospects for market success, such as agent-based tools, orchestration systems that link AI to physical processes, and the rich sensor and production data generated

by German manufacturing. The second track **rejects the idea that niche specialisation alone is sufficient.** Taking inspiration from Chinese open-weight models such as the Kimi K3, as well as proposals for a European AI research facility similar to CERN, this approach calls for dedicated funding and an institutional setup to support long-term experimentation at the forefront of innovation. This would enable German developers to develop deep technical expertise and communities over time, even if initial projects primarily focus on catching up with existing paradigms rather than inventing new ones.

Still, capabilities alone are not enough. The other half of this study thus considers what happens once something has been built. **Even a well-chosen chokepoint and a well-funded effort to catch up will fail to yield a strategic advantage if German and European institutions continue to default to foreign incumbents for reasons that have little to do with product quality.** Switching costs, dependence on dominant cloud and developer platforms, and compliance burdens, which particularly affect smaller firms, all hinder the adoption of already existing European alternatives. Accordingly, a sovereign AI strategy that ignores these demand-side constraints is incomplete and must be complemented by measures to support the structured rollout of AI in public and private organisations.

The dual-track approach developed here is not specific to Germany. Any economy that is neither a computing superpower nor purely reliant on technology must decide which layers of the stack to hold domestically, which to source from partners at an acceptable political risk, and which to buy on the open market. Germany is a useful test case because, unlike the high-service economies in the US and UK, it still has a manufacturing and industrial base, which is combined with weak positions in computing and frontier models. This combination is shared in different proportions by other middle powers such as South Korea, Japan, Canada, Brazil, India, Indonesia, Poland, and the Gulf states. Germany also shares comparable regulatory capacity with other European countries, as much of the AI-related governance is mediated through harmonised EU rules or emerging international frameworks. In particular, three of the study's findings are generalisable. Firstly, **AI sovereignty is clearly layer-specific:** a country can be structurally dependent at the compute layer and still hold a defensible position further up the stack, provided it identifies which of its assets are genuinely hard to substitute. Secondly, **the binding constraint for middle powers is often adoption rather than supply**, so programmes that fund domestic capability without addressing switching costs and public procurement practices will underdeliver. Thirdly, **regulatory quality itself should be seen as a layer of the AI stack:** predictable, proportionate rules lower the cost of deploying AI, whereas unsettled ones raise it.

Much of the debate surrounding AI sovereignty hinges on an unresolved technical question: **will the "scaling laws" that have driven recent progress in AI continue to**

hold? These laws describe the empirical pattern whereby model performance improves regularly when developers increase computing power, training data, and model size.[6] If this relationship persists in the future,[7] countries without access to very large computing clusters will find it difficult to take the lead, which would favour the current US position. However, if the benefits of scaling diminish,[8] the advantage may shift towards those excelling in algorithmic efficiency and "physical AI" tightly integrated with industrial processes. This would play to Germany's strengths in engineering and industrial data. This study does not attempt to resolve this technical dispute. Instead, it considers frontier catch-up (assuming scaling continues) and diversification towards agentic and physical AI (assuming diminishing returns) as complementary strategies that should be pursued in parallel due to uncertainty and the high, asymmetric costs of error.[9]

The remainder of this study proceeds as follows. Section 2 maps Germany's position across the AI stack, covering everything from regulatory foundations and energy to computing power, data, models, applications, and compliance tools. Section 3 analyses the behavioural, regulatory, and market-structural bottlenecks that hinder the adoption of European alternatives. Section 4 sets out a way forward, including alternative AI paradigms and priorities for German public funding. Section 5 provides a conclusion including generalisable lessons for other middle powers.

2. Mapping the AI Stack

Discussions about AI sovereignty often slide into an alleged choice between strategic independence and protectionism. This study takes a liberal view of tech sovereignty: Germany should secure control over AI-related capabilities that might become strategic while maintaining an open economy. To make this assessment concrete, it uses the notion of an "AI technology stack" – the layered structure of energy, hardware, computing infrastructure, models, and applications – as a framework for analysing Germany's position.[10] The version applied here adapts the popular notion of a layered "AI factory" and extends it to include regulation as part of the stack. For each layer, this section provides a structured assessment of German and European alternatives, drawing on publicly available datasets and web-scraped company-level data, and an evaluation of existing dependencies.

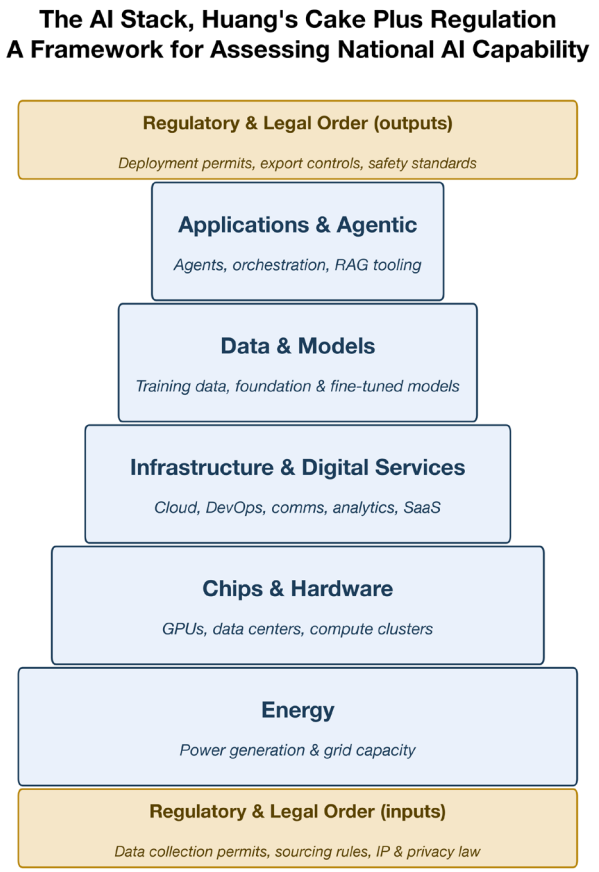
2.1. Turning Jensen Huang's "5-Layer Cake" into an "AI Burger"

NVIDIA's chief executive, **Jensen Huang, recently popularised the idea that generative AI should be understood not merely as software, but as a new form of industrial infrastructure comparable in scale to the electricity grid.** The latter transformed industrial production during the Industrial Revolution by enabling rigid work standardisation and scientific management, known as Taylorism. Huang describes today's AI revolution as a "five-layer cake", built from the ground up: energy, chips, infrastructure, models (including data), and applications.[11] At the base sits energy because, without electricity, nothing else in the stack can function. On top of that are the chips that convert raw power into computation. Above them are the data centres, software, and networking equipment that Huang refers to as the "AI factory". The fourth layer consists of the models, i.e., the LLMs that this factory produces by drawing on plenty of training data. Finally, the AI applications (and increasingly agents) that people encounter in everyday life sit at the very top. Huang summarises the entire process in the memorable notion that **the input is electrons and the output is tokens.**[12] As each layer depends entirely on the one beneath it, any shortage in the stack ripples through every layer above it.

Although the cake metaphor illustrates that generative AI has evolved into a substantial, multi-layered industrial system, it somewhat hides that **not all layers are equally strategic**. The burger metaphor proposed here clarifies this: the "buns" hold the stack together, yet they are easy to replace. The "patty" in the middle defines the product, so changing it would fundamentally alter the meal. For US frontier labs such as OpenAI, Google, and Anthropic, this central ingredient is their ability to produce high-quality, general-purpose AI outputs by combining significant computing power, financial resources, and specialised talent. In contrast, Chinese

labs such as Moonshot AI and DeepSeek are defined by their engineering efficiency; they build capable models that run well on constrained hardware and are tightly optimised for domestic needs. The Netherlands shows a different pattern: it does not manufacture chips or train large models, but its control over extreme ultraviolet lithography equipment provides it with a key position in global semiconductor manufacturing. In all three cases, sovereignty hinges on possessing one irreplaceable element, rather than controlling every layer of the stack.

Figure 1 : Overview: the AI stack modelled as a "burger"



Source: *cep research, own illustration adapting the layered "AI cake" framing.*

This study turns Huang's five-layer "AI factory" framework into an AI burger stack by adding two regulatory layers bracketing the stack at bottom and top (Figure 1). Whereas Huang's original framework describes only the technical layers of AI production, this conceptual addition builds on the German ordoliberal tradition linked to the Freiburg School and its idea of an *Ordnungsrahmen:* a stable framework of legal and institutional rules within which emerging markets can operate.[13] When applied to AI, the claim is that, **more so than in the US or China, rules on market access and deployment in Europe function as part of the basic infrastructure for the entire stack**. The fact that regulation is split into bottom and top layers reflects this European pattern. In contrast, regulators in the US and China more often react after technology is on the market, writing rules to manage existing products. In the EU, however, regulation frequently precedes deployment, thereby shaping which actors can enter the market and how they can operate. The bottom regulatory layer covers foundational "rules of the game", such as the AI Act and the GDPR, which determine who can develop and sell AI systems in Europe and the conditions under which they can do so. The top regulatory layer comprises user-facing elements such as compliance labels and risk categories, which influence whether businesses and citizens decide to use a product once it is available.

2.2. Regulatory layer I (inputs)

The bottom bun holds the entire structure of the AI burger together from beneath; without it, it would disintegrate soon. **In the AI stack, the bottom bun represents the regulatory framework governing market entry and data acquisition, that is, the foundational rules determining who is permitted to develop and train AI systems in Europe.** The GDPR governs any personal data contained in a training corpus, while the EU's Digital Single Market Directive provides an exception for text and data mining, permitting the reproduction of copyrighted works for AI training, provided that rights holders can opt out.[14] The AI Act adds further transparency and risk management obligations to these. Just as a bun's structural integrity depends on consistent ingredients, the stability of this layer depends on consistent rules. Uncertainties such as unclear data guidance or sudden adaptations to the AI Act cause the entire stack to wobble, deterring investment already at the point of market entry.

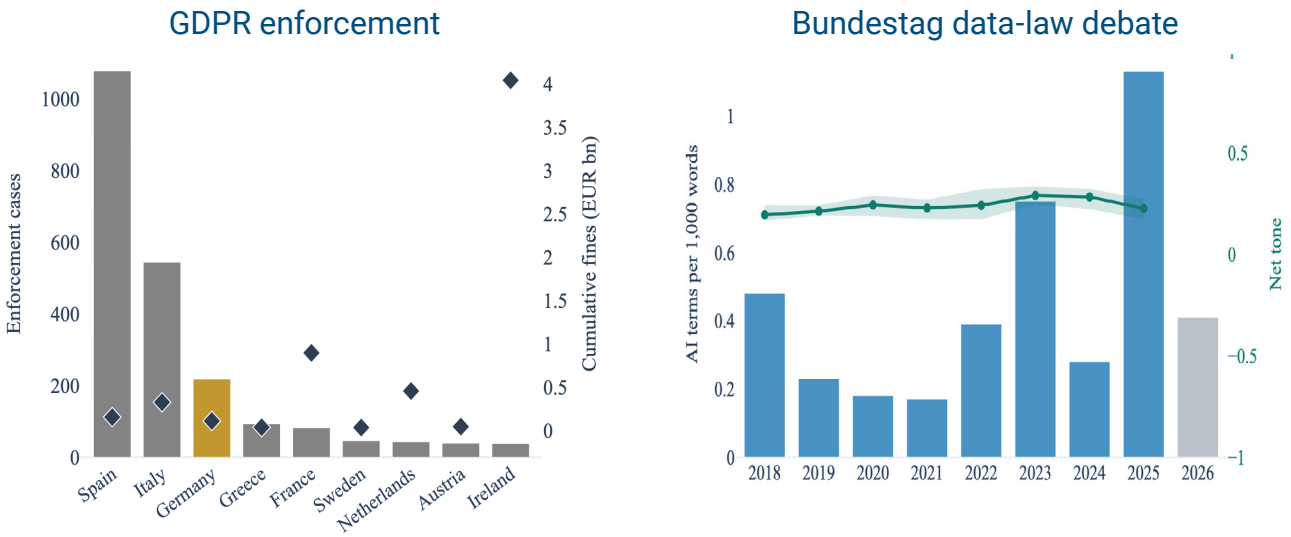
This section focuses specifically on the GDPR, since it has generated the most active and currently unsettled legal debate with respect to AI training.[15] Two guidance documents issued by the European Data Protection Board (EDPB) in July 2026 demonstrate the ongoing uncertainty surrounding the underlying GDPR rules, eight years after the regulation came into force.[16] The draft Guidelines 3/2026 on web scraping clarify that scraping publicly accessible websites is within the scope of the GDPR whenever personal data is collected; that, realistically, only the legitimate interest ground under Article 6(1)(f) is available to developers scraping at scale; and that neither public posting nor the absence of a so-called "robots.txt" file (a text file publishers place on their website to indicate to automated crawlers which of their pages can and cannot be accessed) implies consent to the use of AI training.[17] The companion draft Guidelines 2/2026 on anonymisation provide an update in light of recent Court of Justice case law.[18] They confirm the *Breyer* (C-582/14) and *EDPS v SRB* (C-413/23) doctrine of relative personal reference: the same dataset can be anonymous for one party and personal data for another. This means that a training set which is judged to be anonymous at the time of collection may still be treated as personal data if the resulting model creates a foreseeable route back to identification. In practice, developers operating in Germany must choose between two options: a complex, case-by-case contextual assessment with legal uncertainty, or a simplified approach that treats data as personal by default, accepting the resulting compliance burden as the safer option.

This **rights-driven EU approach is one of four global governance paradigms**.[19] The others are the United States' market-driven, deregulation-oriented approach; China's state-driven, data-localisation model; and Japan's Data Free Flow with Trust framework, which combines formal legal structures with a more hands-off, trust-based approach to enforcement. Convergence with EU standards varies sharply across different economies: for instance, Japan

achieves full alignment, while Canada and South Korea hold adequacy decisions despite sectoral fragmentation. The US and China, meanwhile, remain furthest apart.

This typology helps to explain why Germany's enforcement record, plotted in the left panel of Figure 2, does not obviously discriminate against foreign platforms: A rights-driven system has been built to protect individuals uniformly rather than penalise a particular category of firm. With 218 GDPR enforcement cases recorded by 2026, Germany had the third-highest number of cases of any EU member state, behind Spain (1,078) and Italy (544). However, its cumulative fines total only 0.108 billion EUR, ranking seventh by volume and far below Ireland's 4.04 billion EUR. This reflects Ireland's status as the EU headquarters for many major US technology platforms under the GDPR's one-stop-shop mechanism. The clearest expression of the German pattern is the average sanction: €0.5 million per case, compared to an average of €2.81 million for other countries. France is the mirror image, with 82 cases and fines totalling 0.893 billion EUR. Thus, German authorities pursue breadth of enforcement rather than severity of sanction. These fines accrue to the competent supervisory authority in the relevant member state (i.e., the *Landesdatenschutzbehörden* at state level or the *Bundesbeauftragte für den Datenschutz und die Informationsfreiheit* for federal bodies in Germany) rather than to the EU budget. Despite this rigorous GDPR enforcement, in 2024, Germany's data market was valued at around 24.5 billion EUR, making it the largest in the EU27.[20] However, its growth rate of 8.8 percent trailed the bloc's average of 11.1 percent, suggesting that its high-quality data are still underutilised.

Figure 2 : Results: Regulatory layer I (inputs)



Source: *cep research based on enforcementtracker.com and the Bundestag DIP API, plenary protocols 2018 to July 2026; own calculations.*

Note: *Left: the ten EU member states with the highest case counts; bars are enforcement cases (left axis), diamonds cumulative fines (right axis). Non-EU jurisdictions are excluded. Right: AI vocabulary per 1,000 words of context around data-law terms (bars) and net tone from −1 to +1 with bootstrap 95% confidence intervals (line and band). Grey marks the partial year 2026 (through July); no tone is plotted for it.*

GDPR fines and case counts are lagging indicators of outcomes: they reflect, above all, decisions that regulators have already made to impose sanctions, often years after the conduct in question occurred. They say nothing about how important data protection and copyright issues are to lawmakers at any given time, which could be useful information for firms trying to predict the stability of relevant legislation and potential increases in penalties. In this sense, parliamentary discourse offers a more forward-looking, complementary data signal. To analyse this discourse, this study retrieves the full text of every Bundestag plenary protocol from 2018 to July 2026 and extracts merged context windows of 600 characters around mentions of data protection and copyright terms.[21] Within these windows, the study measures two things: the frequency with which AI vocabulary co-occurs with data law, normalised per 1,000 words of context, and the tone of the surrounding language. As the corpus is restricted to protocols that also mention AI, the measure captures how data law is discussed within the AI debate rather than within data-law discourse in general.

The co-occurrence measure of debates about AI and data law – plotted in the right panel of Figure 2 – is rising, albeit unevenly. From 2019 to 2021, it sits between 0.17 and 0.23 AI terms per 1,000 words, climbing to 0.75 in 2023 as the AI Act enters its final negotiations. It then falls back to 0.28 in 2024 and reaches 1.13 in 2025, the highest value in the series and one produced by a comparatively small number of mentions (174). This reflects concentration rather than volume. Nonetheless, the direction is clear: **data protection and copyright are increasingly being debated in relation to AI, rather than as standalone regulatory issues.** When data protection and copyright are debated in isolation, firms face a relatively stable compliance problem: the statutes are established, supervisory practice is well-defined, and although enforcement is slow, it is at least predictable. However, when these terms are increasingly discussed in the context of AI, they become associated with a regulatory domain in which several legal frameworks overlap and none are fully established (e.g. the overlap between the GDPR, rules on text and data mining, and the transparency and opt-out obligations for general-purpose models set out in the AI Act). For German firms, these empirical trends indicate that the input rules for AI are a moving target. The binding constraint is no longer what regulators have sanctioned in the past, but rather which of the several plausible interpretations of the overlapping regimes will prevail.

Tone is measured using Rauh's sentiment dictionary for German political language and visualised as a trend line in the right panel of Figure 2. Rauh's dictionary was explicitly built and validated using parliamentary texts.[22] Moreover, unlike general-purpose German lexicons, it distinguishes between negated positive and negated negative expressions. Net tone is the difference between positive and negative matches over their sum, bounded by –1 and +1. Using this purpose-built dictionary for our selection of AI speeches in the Bundestag, the net tone surrounding data

law mentions remains consistently positive, averaging 0.24. This does not change much over time. However, it does show that data protection and copyright are treated as technical and consensus-building files rather than as a battleground over market power, a finding that is consistent with the enforcement pattern described above and with Germany's reputation as a rights-driven jurisdiction. As a robustness check, the same text passages were scored with the general-purpose SentiWS lexicon,[23] which is not validated for political language; it returns higher values, but the two-sentiment series correlate strongly.

Taken together, the empirical indicators point in the same direction. Germany's pattern of GDPR enforcement shows broad case coverage, even if fine levels are relatively modest. Parliamentary attention to data protection coincides with phases of anticipatory lawmaking rather than open conflict, but is increasingly intertwined with debates about AI. This empirical picture generally aligns with the rights-based profile of the EU, and, specifically, of Germany: data protection rules are applied consistently according to existing legislation, and are not abruptly relaxed or suspended in response to the rise of AI. For firms, however, the most important feature of this layer is not necessarily how strict the rules are or how they are discussed, but whether they can be known in advance. When deciding whether to train a model using its own processes and customer records, a manufacturer needs to know which legal basis applies, whether the trained model will still be treated as containing personal data, and what the implications are if this changes after deployment. **Where the legal answer is contested, as it currently is according to the EDPB guidelines, the rational response is to postpone the project, narrow the training corpus, or purchase a service from a provider large enough to absorb the legal risk. All three responses shift value away from developers based in Germany.**[24] Legal certainty at the input layer determines who can work with European data and in-vest in German AI.

2.3. Energy

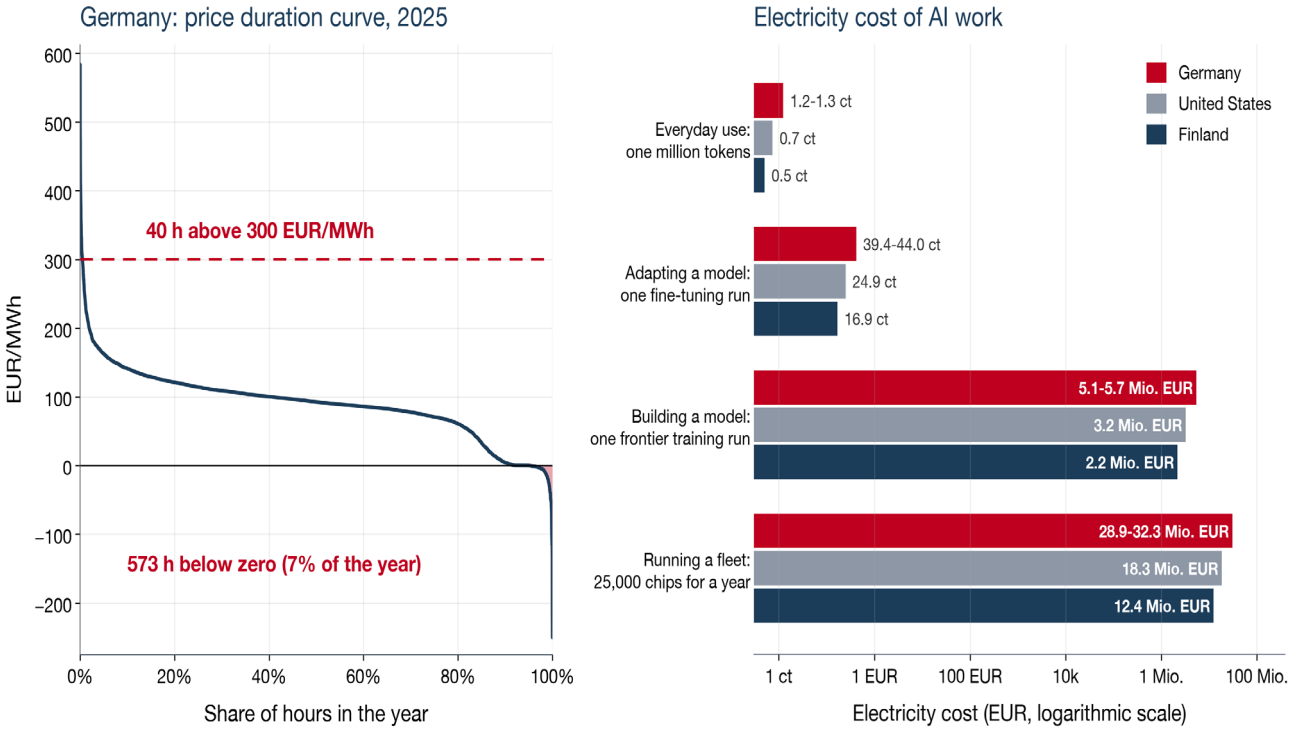
The energy layer sits at the base because training a frontier model requires tens of thousands of chips and draw continuous power for weeks, consuming as much electricity as a medium-sized city. Just as a burger might burn or be undercooked, an AI training run on a grid with unstable voltage or unpredictable pricing typically results in technical failures. In particular, any interruption or voltage fluctuation can corrupt the process of updating the weights in the neural network, forcing it to restart from an earlier saved checkpoint. Therefore, a stable, predictably priced grid is a prerequisite for the economics of large-scale AI development, as the cost of power directly affects the commercial viability of a training run, and the reliability of supply determines whether it can be completed.[25]

In terms of price, Germany is at a structural disadvantage. Eurostat's data on prices for non-household consumers shows that, excluding VAT, German prices range from 12.05 to 13.46 ct/kWh for consumption bands of 20,000 MWh per year and above.[26] A German AI data centre is typically located in the middle of the national range, since large permanent-load customers usually qualify for individual network charges, reducing grid fees. However, the extent of this relief should not be overstated.[27] By contrast, prices in Finland, a European middle power with a competitive position in terms of data centre build-out, range from 5.07 to 5.27 ct/kWh. Therefore, Germany pays roughly two and a half times the Finnish price for electricity at the scale at which a data centre consumes it. In 2025, US industrial users paid an average of 8.62 US cents per kWh,[28] equivalent to 7.63 ct/kWh at the ECB's annual average reference rate. This comparison suggests that Germany's energy issues regarding AI are often erroneously framed: Germany's disadvantage in terms of electricity at an industrial scale is not primarily a transatlantic problem driven by the Trump administration's industrial policy, but also an intra-European fact. Within Europe, Nordic countries such as Finland and Sweden achieve a more favourable energy position in the AI stack with grids that are more flexible and rely largely on nuclear and hydro power, allowing for easier additions to the grid as well as better cooling of servers.

Germany's premium predates the current decade, but it has widened since the 2022 energy crisis. When expressed as the cost of performing typical AI tasks,[29] the German premium varies at different stages of the lifecycle (Figure 3, right panel). A key distinction here relates to training (of AI models) and inference (using trained AI models).[30] For example, serving a million output tokens on H100-class hardware consumes roughly 0.097 kWh, costing between 1.2 and 1.3 euro cents in Germany, compared to 0.5 cents in Finland and 0.7 cents in the US. Given commercial inference prices of several euros per million tokens, it becomes clear that German electricity prices are not the main barrier to AI deployment for German firms. However, at the training stage, the same difference is significant. Meta discloses that pre-training Llama 3.1 405B required 30.84 million GPU hours. At an all-in sys-tem power corresponding to some 42,000 MWh of electricity, this would cost between 5.1 and 5.7 million euros in Germany, compared to about 2.2 million euros in Finland and 3.2 million euros in the US. Operating a fleet of 25,000 accelerators (roughly equivalent to the German stock identified in Section 2.4 below) draws 34.25 MW and would cost between €28.9 and €32.3 million a year in Germany, compared to €12.4 million in Finland and €18.3 million in the US. Figure 3 shows these back-of-the-envelope estimates alongside those for a fine-tuning run.

Note that the associated energy figures pale in comparison to the capital cost of the hardware itself. However, electricity is one of the few costs that varies by location, so it is one of the few that can be affected by a siting decision (GPUs are expensive everywhere). Currently, the more consequential constraint in the energy layer is not price, but availability, due to the huge demand for AI compute. Grid interconnection routinely lags behind hardware and construction timelines (see below). While the estimates in Figure 3 are based on all-in power per accelerator and thus cover host systems, scale-up switching, networking, and cooling, they do not include related external factors such as connection charges, the cost of capital associated with an uncertain connection date, or any on-site generation that an operator might build to avoid the queue. In short, the question of AI sovereignty at the energy level is therefore narrower than it is often posed: Energy determines whether a middle-power country such as Germany can host frontier-scale training but says little about whether domestic firms can use or adapt specialised models. This depends instead on access to open weights and training data, i.e. largely an inference-side activity.

Figure 3 : Results: Energy layer



Source: *cep research, own illustration[SF7.1][KA7.2] based on Energy-Charts.info, Eurostat (nrg_pc_205), MLPerf Inference v5.0, MLPerf Training v4.0, Meta Llama 3.1 model card, and SemiAnalysis; own calculations.*

Note: *Left panel: every interval of the 2025 German day-ahead wholesale price, sorted from highest to lowest. Right panel: the electricity cost of four AI workloads at German, Finnish, and US industrial prices, on a logarithmic scale. German bars show the midpoint of the price range across consumption bands of 20,000 MWh a year and above, with the range given in the figure; Finnish and US bars are single values. Inference and training are measured on separate benchmarks. Full derivation is available from the author upon request.*

Furthermore, Germany's grid is increasingly volatile, not just expensive (Figure 3, left panel). According to the *Bundesnetzagentur*, Germany recorded 573 hours of negative day-ahead wholesale prices in 2025, a new record and up from 457 hours in 2024, while hours above 300 EUR/MWh fell slightly to 40, down from 41 the previous year.[31] This pattern of extreme oversupply during periods of high renewable energy generation, alongside occasional scarcity spikes, makes it difficult for hyperscale data centre operators to secure long-term power purchase agreements for project financing, since it signals a grid that might be cheap on average but difficult to plan around. By the end of 2024, Germany's installed renewable capacity was approximately 190 GW (99 GW from solar power, 64 GW from onshore wind power, 9 GW from offshore wind power, and the remainder from hydroelectric power and biomass), compared to a considerably smaller national fossil capacity.[32]

While Germany's position in the energy layer is thus already problematic, China and the US may gain an even greater advantage in the near future. By the end of 2024, China's installed capacity had reached approximately 3,348 GW, including 887 GW of solar, 521 GW of wind, 436 GW of hydro and 1,444 GW of thermal (fossil) capacity.[33] By the end of 2024, the US had added over 1,352 GW of total installed generating capacity, with more than 52 GW added that year alone, driven increasingly by AI and data centre demand.[34] US developers have begun bypassing interconnection queues entirely rather than waiting for them to clear.[35] For example, Elon Musk's AI lab, xAI, has installed dozens of gas turbines on site at its Colossus data centre in Memphis – reportedly operating without the necessary air permits for a period – and has recently been granted approval for 41 turbines totalling approximately 1.2 GW of gas capacity on site in Mississippi.[36] This **"bring your own generation" approach, which involves deploying turbines, reciprocating engines or fuel cells directly at the data centre site, has emerged as a deliberate industry strategy.**[37] Developers are weighing up the total cost of on-site power generation against the multi-year cost of waiting in the grid interconnection queue.

The construction-speed advantage in the US is compounded by modular building techniques. Recent industry analysis finds that prefabricated "LEGO-style" data centre components (power rooms, cooling skids, and even entire data halls assembled off-site) can compress construction windows by 36% and reduce costs by 8% per MW, allowing US operators like AWS and Meta to deploy capacity in months rather than years.[38] An alternative response to the same constraint would be to abandon the single-site model altogether: recent analysis finds that a decentralized, multi-gigawatt training run coordinating a dozen or more data centres spanning thousands of kilometres is technically feasible using existing spare gas-plant capacity. However, the same analysis concludes that developers will likely continue to prefer single large campuses and on-site generation up to gigawatt scale.[39]

To be sure, the comparison focuses here on absolute numbers (such as generating capacity and prices) rather than population- or enterprise-adjusted figures. This is for two related reasons. Firstly, training a specific frontier model requires a fixed amount of absolute power and a fixed reliability buffer, regardless of the size of the hosting economy: a smaller economy with fewer energy-intensive firms could still face the same electricity bill and exposure to supply interruptions for an equivalent AI cluster. Secondly, this logic strengthens as the technology matures: unlike a one-off training run, inference demand (i.e. the actual usage of a model) scales with the number and intensity of deployed applications rather than population or firm count. As agentic systems start to chain multiple model calls per task, this demand will further increase.

The main policy response of the German government to this issue is the **National Data Centre Strategy**, published on 18 March 2026.[40] The strategy explicitly states that sufficient data centre capacity and comprehensive digitalisation are prerequisites for a competitive, future-ready economy and modern public administration. It also sets quantitative targets, including doubling German data centre capacity and quadrupling capacity dedicated to high-performance computing (HPC) and AI by 2030. While the strategy promotes sustainable data centres powered by renewable electricity and supported by energy-efficient technical solutions, its land-use pillar targets sites for data centre construction specifically, meaning it does not by itself accelerate the separate and often slower permitting processes required for new power plants or grid-scale renewable generation. In effect, the strategy addresses the price/grid stability bottleneck identified above, but it lacks the US-style emphasis on rapid prefabricated construction[41] and onsite generation. Instead, the strategy focuses on accelerating traditional grid-connected builds within existing regulatory frameworks, offering only renewable self-generation as a sustainability measure rather than a reliability workaround. Moreover, as of mid-2026, it remains a framework of intent rather than a set of enacted measures with binding timelines.

2.4. Hardware (chips)

The patty of the AI burger is made of chips: specifically, AI accelerators, the collective term used for all processors that are designed to carry out the mathematical operations involved in training and running AI models much faster and more efficiently than general-purpose chips. The most widely used type is the graphics processing unit (GPU). Unlike central processing units (CPUs), which are optimised for general-purpose, sequential tasks, GPUs perform the same mathematical operation on multiple data points simultaneously, making them far better suited to the large-scale linear algebra at the core of deep learning. Some providers, such as Alphabet/Google, have also developed application-specific chips, like Tensor Processing Units (TPUs), which are even more closely tailored to AI workloads. Currently, **NVIDIA, AMD, and**

Google dominate the market for advanced AI accelerators, with NVIDIA's H100-class GPUs setting the standard for measuring high-end computing capacity. In general, this layer is shaped by a compounding logic whereby early technical leads become self-reinforcing, and this logic governs advanced chip fabrication and the compute infrastructure built on it.[42]

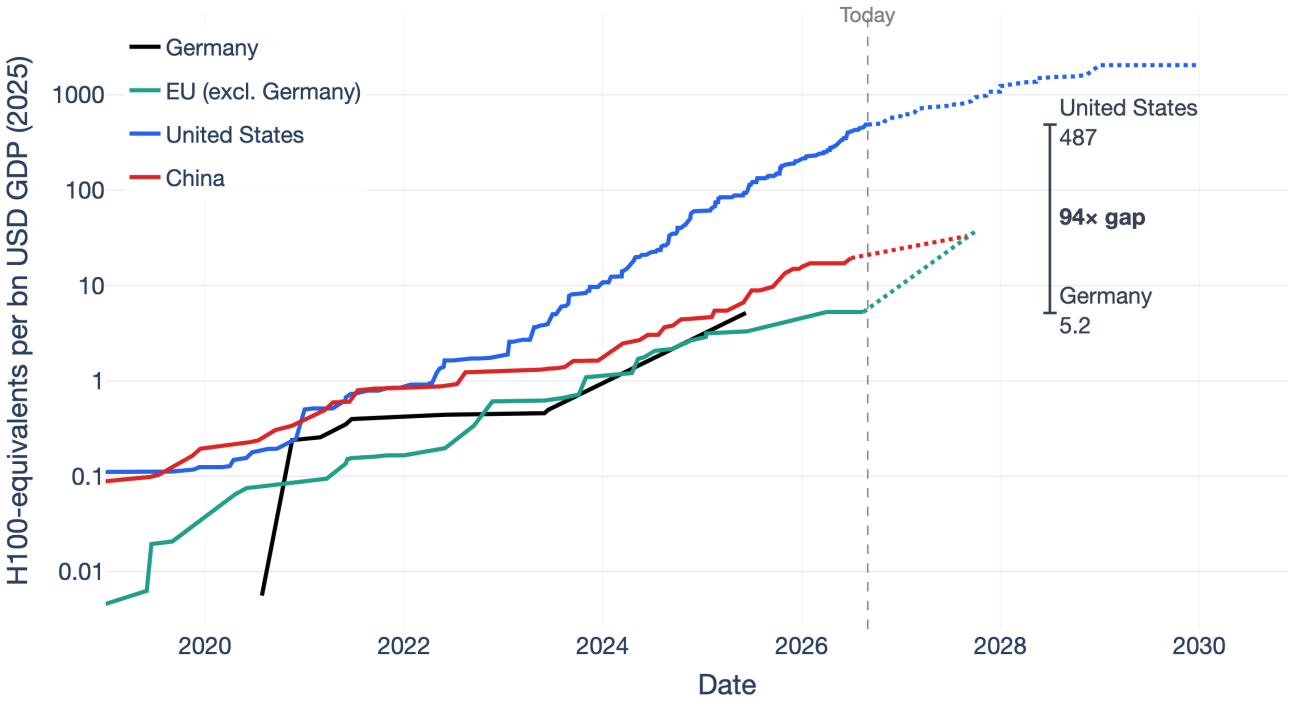
Chips and accelerators form the dense core around which all the other layers of the AI stack are built. This is the most difficult layer to replace once a substantial engineering lead has been established; it is the equivalent of a proprietary burger recipe that competitors cannot reverse-engineer within a typical investment cycle. Academic work on ASML's extreme ultraviolet (EUV) lithography strategy documents more than two decades of accumulated, path-dependent engineering knowledge spanning light-source physics, optics, and materials science.[43] This knowledge has enabled **ASML to establish a near-total monopoly on machines capable of printing the smallest transistor features** (despite recent reports[44] that China has started large-scale production of deep-ultraviolet machines for chip manufacturing). ASML's EUV systems require components sourced from hundreds of specialised suppliers, accumulated over decades.[45] **Two of the most critical of these specialised suppliers are German: Trumpf builds the high-power lasers that generate the extreme ultraviolet (EUV) light, and Zeiss polishes the mirrors that steer it with great precision**.[46] The German government has reportedly begun to treat such positions as instruments of economic statecraft: over the past months, it has been informally mapping Chinese dependencies on German and European technology, components, and know-how, and has identified Trumpf, Zeiss, and other semiconductor suppliers in the initial findings.[47] According to media reports, these findings inform non-public measures developed in the National Security Council, with contemplated scenarios extending from halting deliveries to ending maintenance of equipment already installed in Chinese fabs. **While Germany has leverage as well as exposure in AI hardware, its leverage in upstream components does not substitute for downstream compute capability, and is not necessarily durable.** This is because exercising it would accelerate the very substitution efforts (such as China's homegrown DUV production mentioned above) that aim to erode it.

Germany's total AI computing capacity, measured in H100 equivalents, is approximately 26,050. This figure is drawn from eleven GPU clusters recorded in Epoch AI's dataset. The H100 is Nvidia's data centre GPU designed for AI workloads, and it has become the industry's de facto unit for comparing raw computing power across different generations of chips, as capacity can be expressed on a single scale.[48] On that scale, **Germany's position is not merely behind, but on a different order altogether**, since the operational US stock is approximately 15.0 million H100-equivalents – more than 500 times Germany's figure Comparisons that set Germany's

current capacity against announced American build-out plans produce even larger multiples. The Gulf states, which have recently been in the news due to their chip agreements with the Trump administration, provide a useful point of comparison from the perspective of middle powers: On a different basis than the Epoch AI figures used above, namely an industry ranking[49] that counts announced alongside operational capacity, the United Arab Emirates is second in the world with 23.1 million H100-equivalents and 6,400 MW of power capacity, whereas Saudi Arabia is third with 7.2 million chips and 2,400 MW, putting both countries ahead of South Korea, France, India, China, the UK and Germany on that measure.

However, the compute layer cannot be settled by scale alone, since the US has a far larger economy. Figure 4 therefore normalises each region's compute by its own economic output. When measured per billion US dollars of GDP, the US has approximately 487 H100-equivalents, compared to Germany's 5.2 – a factor of roughly 94. In other words, Germany's shortfall is not entirely a matter of size: **even when adjusted for the size of its economy, German AI computing capacity remains roughly two orders of magnitude below the American level and the Chinese level**. Due to long order backlogs and the concentration of advanced manufacturing in a few regions, the gap will likely widen unless Germany and the EU can access frontier chips more quickly and acquire the necessary power and land.

Figure 4 : Results: Chips layer



Source: *cep research, own illustration based on Epoch AI, GPU Clusters and AI Data Centers datasets, and World Bank national accounts; own calculations.*

Note: *Cumulative capacity in H100-equivalents per billion US dollars of GDP, on a logarithmic scale. Each region is divided by its own GDP. Solid lines denote operational capacity, dotted lines announced or planned capacity. Epoch AI's coverage is incomplete, likely with weaker coverage of China.*

While China's computing capacity is an order of magnitude larger than Germany's, it faces structural bottlenecks in its ambitious attempt to catch up. Above all, the US semiconductor export controls remain the primary hurdle to the country's hyper-scaler dreams, as China has long struggled to match the US' ability to produce chips domestically. Furthermore, the country's flagship domestic strategy for expanding computing power, the "Eastern Data, Western Compute" programme, has so far failed to relocate data centre capacity from the overcrowded eastern seaboard to the cheaper land, cooler climates, and abundant renewable energy sources of the western interior. This attempt has proven to be largely illusory because true frontier AI workloads remain latency-sensitive and thus cluster near existing talent pools in Beijing, Shanghai, and the Pearl River Delta.[50] However, while the physical relocation of frontier capacity has been disappointing, virtual consolidation is progressing rapidly. As of August 2026, over 60 percent of China's computing resources are now monitored and dispatched via a unified national platform, complemented by a "Computing Power Platform" that matches computing resource buyers with sellers.[51] At the same time, the State Council boasted about China's first domestically built cluster of 100,000 accelerators and estimated 2,185 ExaFLOPS, a common measure of computing speed, of domestic intelligent computing capacity. **As Huawei's accelerator and memory production starts to scale, China could largely close the compute gap with the US within roughly 24–36 months.**[52] Compute determines whether Chinese firms can keep pace in AI training and deployment.

What is Europe's reaction? Beyond export controls and state aid for individual fabs, the Chips Act[53] and the proposed Chips Act 2.0[54] form the EU's main horizontal framework for semiconductor industrial policy. They seek to double Europe's global market share in chips to 20% by 2030 and strengthen monitoring and crisis response.[55] In addition, the EU's AI Gigafactories call, launched on 30 July 2026, offers a more centralised approach.[56] The initiative aims to procure up to seven large-scale data centres equipped with at least 75,000–100,000 "H100-equivalent" accelerators each, backed by more than €30 billion in combined public and private investment, in order to narrow the continent's AI compute gap while preserving data sovereignty.[57] At that scale, the initiative would create serious capacity, but it would still cover only a limited share of frontier-model training demand if utilisation rates and power constraints remain binding. Moreover, with operations expected only from 2028 onwards, these facilities will arrive too late to address the immediate compute gap documented here. This tension is exacerbated by a second issue: the construction of the Gigafactories is set to increase Europe's demand for advanced processors that it still cannot produce domestically.[58] In other words, **the Chips Act and the Gigafactories initiative are pulling in different directions: one seeks to expand Europe-**

an chip manufacturing capacity, while the other is accelerating a buildout that will absorb foreign-made accelerators faster than Europe can replace them with domestic production.

In light of this, **an alternative response to compute scarcity that seems not widely recognised in political discourse is to decentralise AI training across multiple data centres** linked by high-speed fibre rather than concentrating a single run in one location or in one gigafactory. This approach shifts the relevant constraints from chip fabrication to latency management and cross-site power availability. This makes it particularly relevant for Europe's fragmented grid and data centre geography. **Research has shown that distributed training remains technically feasible on a large scale, including in models that split a training run across multiple sites at limited additional cost.**[59] Analysts have modelled a 10-gigawatt training run distributed across 23 existing US power plants, linked by 4,800 kilometres of dedicated fibre.[60] They found that this adds less than 1% to construction costs relative to a single, consolidated campus. A decentralised approach could combine Germany's and the wider EU's comparatively modest individual computing sites to create an aggregate capacity that would be competitive with a single large US or Chinese campus, without requiring any one country to build a facility of this scale alone. Although decentralised training does not entirely eliminate Europe's chip bottleneck, it could render smaller infrastructure assets usable as an alternative to a single hyperscale campus.

2.5. Infrastructure (cloud, software, and compute services)

The sauce binds the burger's components. In the AI stack, this is the infrastructure layer: cloud platforms, container orchestration, identity management, and developer tools that make raw chips usable for actual workloads. Unlike the patty, which requires decades to develop, sauce recipes can, in principle, be replicated by competing chefs. However, once a particular flavour profile dominates the market (such as the proprietary APIs of Azure), switching to an alternative sauce requires reformulating the entire "burger", or stack. As shown below, Germany hosts some niche sauce-makers (European cloud and DevOps providers), but the market is currently dominated by a few non-European providers whose sauce has become the default flavour.

Economies of scale and network effects create dependencies on the infrastructure layer of the AI stack.[61] As of 2024, the three US hyperscalers (Amazon Web Services, Microsoft Azure, and Google Cloud) together command roughly 70% of Europe's €61 billion cloud infrastructure services market.[62] The advantage of these firms extends to the speed of their current physical construction, as they increasingly use modular, prefabricated data centre designs.[63] According to the Synergy Research Group, Oracle Cloud Infrastructure accounted for around 3% of global cloud infrastructure spending in late 2025, placing it fifth behind the three hyperscalers and Alibaba.[64] Recently, Oracle is growing,

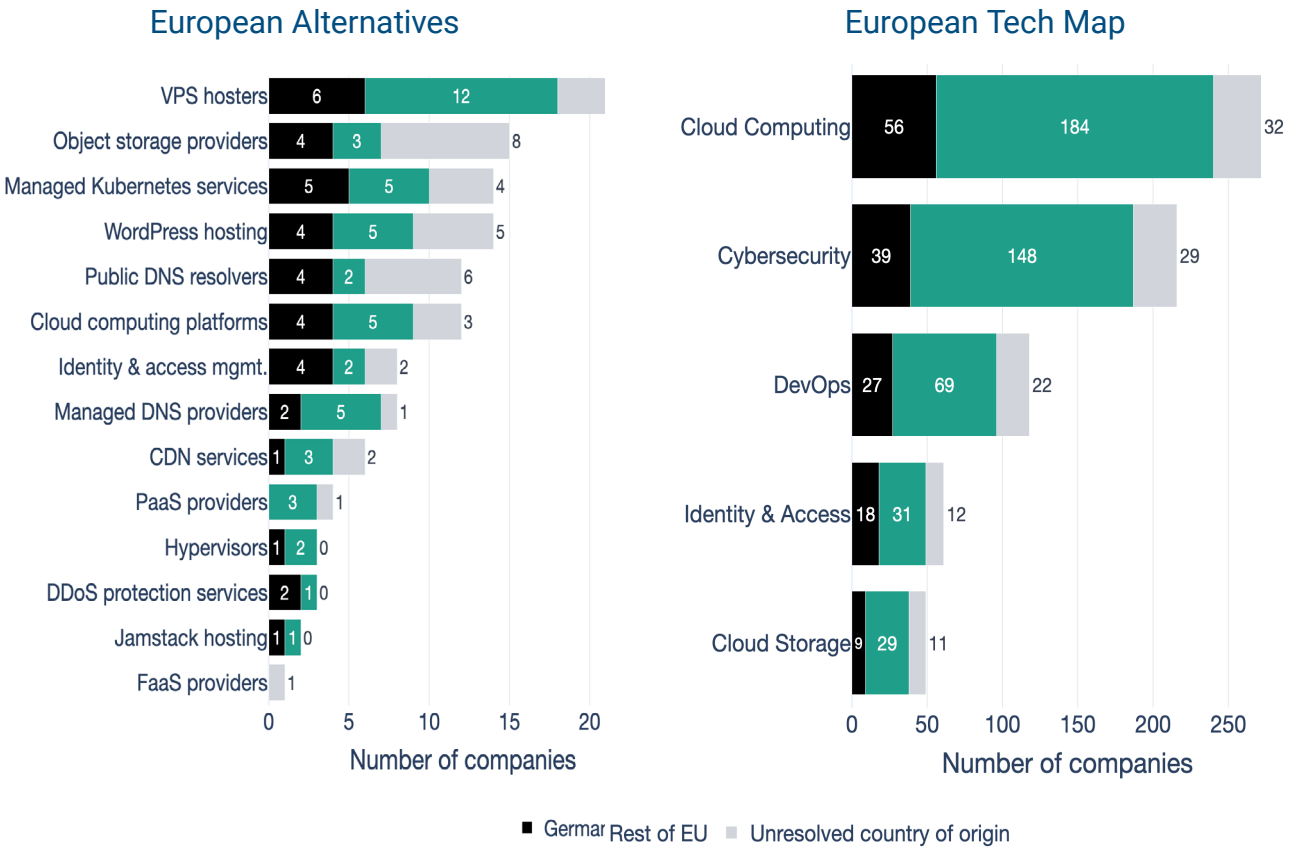
together with a new group of specialised "neoclouds" such as CoreWeave (the largest player), Crusoe, or Nebius. Among European providers, SAP and Deutsche Telekom are joint leaders, each accounting for around 2% of the European market. They are followed by OVHcloud, Telecom Italia, and Orange. The most notable recent German entrant is STACKIT, the cloud and colocation brand of Schwarz Group (owner of Lidl and Kaufland). STACKIT was spun out as an independent company in September 2024, having initially been built for the group's own data sovereignty needs. STACKIT operates from data centres in Germany and Austria, and markets itself as working towards becoming the first European hyperscaler. In 2026, it secured a share of the Commission's €180 million sovereign cloud services contract.[65]

Despite these individual success stories, the collective share of European providers has fallen from 29% in 2017 to just 15% and has remained stagnant since then.[66] The EU has responded by combining a variety of regulatory and industrial policy tools under the umbrella of data sovereignty, using initiatives such as Gaia-X, the European Alliance for Industrial Data, Edge, and Cloud, and – more recently – the Cloud and AI Development Act to protect data confidentiality and counter US vendor dominance.[67] However, diverging national preferences and limited technological capabilities have so far constrained the EU's ability to develop an independent AI cloud ecosystem, even as political urgency around the issue has grown.[68] When considering this problem from the layer-by-layer perspective of the AI stack, it is important to **shift the focus from a narrow understanding of computing infrastructure in the sense of hyperscaler clouds to the broader range of infrastructure services that a firm requires to run AI workloads, including virtual servers, content delivery, identity management, tooling, and cybersecurity**. The remainder of this subsection thus complements the above aggregate revenue picture by demonstrating where the underlying supply of alternatives exists within Europe itself. Accordingly, the methodological focus shifts from the revenue share in the core cloud/IaaS market (which helps to compare the EU as a whole against the three US hyperscalers) to measuring company counts in Europe, which also enables a better sense of Germany's particular strengths and weaknesses. Obtaining such granular findings matters because a 2025 European Parliament study concluded that reliance on non-European providers extends beyond raw cloud hosting.[69]

To do so, this study constructs two novel firm-level datasets tailored for this study via web scraping. The first dataset was compiled by scraping European-Alternatives.eu, a directory of European software and infrastructure providers organised by function. Figure 5 documents the 14 infrastructure-relevant subcategories, each of which corresponds to a distinct function within the layer. These include virtual private server (VPS) hosters and object storage providers (renting raw compute and storage by the hour or by usage); managed Kubernetes services and platform-as-a-

service (PaaS) providers (automating the deployment and scaling of software applications without requiring firms to manage the underlying servers themselves); content delivery network (CDN) services (geographically caching and serving content to reduce latency); managed and public DNS providers (translating web addresses into server locations); identity and access management (IAM) services (authenticating users and controlling their access to a firm's software); and WordPress and Jamstack hosting (specialised hosting for common website architectures). Together, these categories span the practical building blocks that a firm needs to run, secure, and scale software and AI applications. The second dataset was compiled by scraping Europeantechmap.eu, which lists firms under five categories: Cloud Computing (general-purpose computing and storage platforms); Cybersecurity (threat detection, encryption and access control tools); DevOps (code building, testing and deployment automation software); Identity & Access (a broader analogue to the IAM category above); and Cloud Storage (dedicated data storage services).

Figure 5 : Results: Infrastructure layer



Source: *cep research, own illustration based on european-alternatives.eu and europeantechmap.eu, scraped in July 2026; own calculations.*

Note: *Provider counts across fourteen hosting and cloud sub-categories (left) and five infrastructure categories (right). Country attribution is resolved from each provider's detail page (left) and from the registry's headquarters field (right). Counts measure how many providers are listed, not market share.*

What do we find? Across the European-Alternatives.eu categories, Germany consistently hosts a non-negligible share of providers: for instance, six of the 23 VPS hosters, four of the 15 object storage providers, and five of the 14 managed Kubernetes services are German. (By definition, the rest of the EU accounts for the remainder in every category.) The same pattern holds across the EuropeanTechMap.eu categories, which have a company count roughly ten times higher: Germany accounts for 57 of 271 cloud computing companies, 39 of 216 cybersecurity firms, and 10 of 50 cloud storage providers. The rest of the EU outnumbers Germany by a factor of two to four across every category. Taken together, and contrary to the impression one might get from the 2025 European Parliament study cited above, the evidence thus suggests that Germany has a sufficient number of players for trying to achieve a structural role in the European cloud and infrastructure software layer. Moreover, **the specific infrastructure services needed to reduce dependency on US hyperscalers often already exist in meaningful volume somewhere within the bloc, provided sovereignty is pursued at the EU level rather than the national one.** Whether these providers are actually being adopted at scale, however, is a separate question taken up in Section 3.

2.6. Models (including training data)

The cheese melts over the patty and toppings, binding them together and adding its own distinctive flavour. It sits at the centre of the burger. In the AI stack, this is the foundation model layer, which transforms different types of training and evaluation data into general-purpose capabilities. It connects the compute layer below to the applications layer above: chips and data centres facilitate training, while models transform this raw capacity into generative AI systems capable of producing text, analysing images, writing code, supporting decisions, or controlling workflows. Just as the quality of cheese depends on the provenance and ageing of the milk, the quality of a model depends on the provenance and curation of the data. Accordingly, this layer also includes the datasets, evaluation suites, and software components that facilitate the development and deployment of LLMs. Germany has an advantage in terms of access to industrial data from manufacturing processes (high-quality “milk”, to stay in the picture) but has produced relatively few notable models (“aged cheeses”).

Germany’s comparative advantage in the model layer lies less in producing general-purpose frontier models than in its access to valuable industrial data. Its manufacturing base generates detailed information on production processes, machine operation, quality control, maintenance, logistics, and sensors – data that can be particularly valuable for specialised industrial AI. However, it is becoming increasingly difficult to convert this advantage into a simple data endowment. Early LLMs, such as GPT-3, could be trained largely on broad, openly accessible internet corpora, often collected through naive web crawling. **However, as frontier systems**

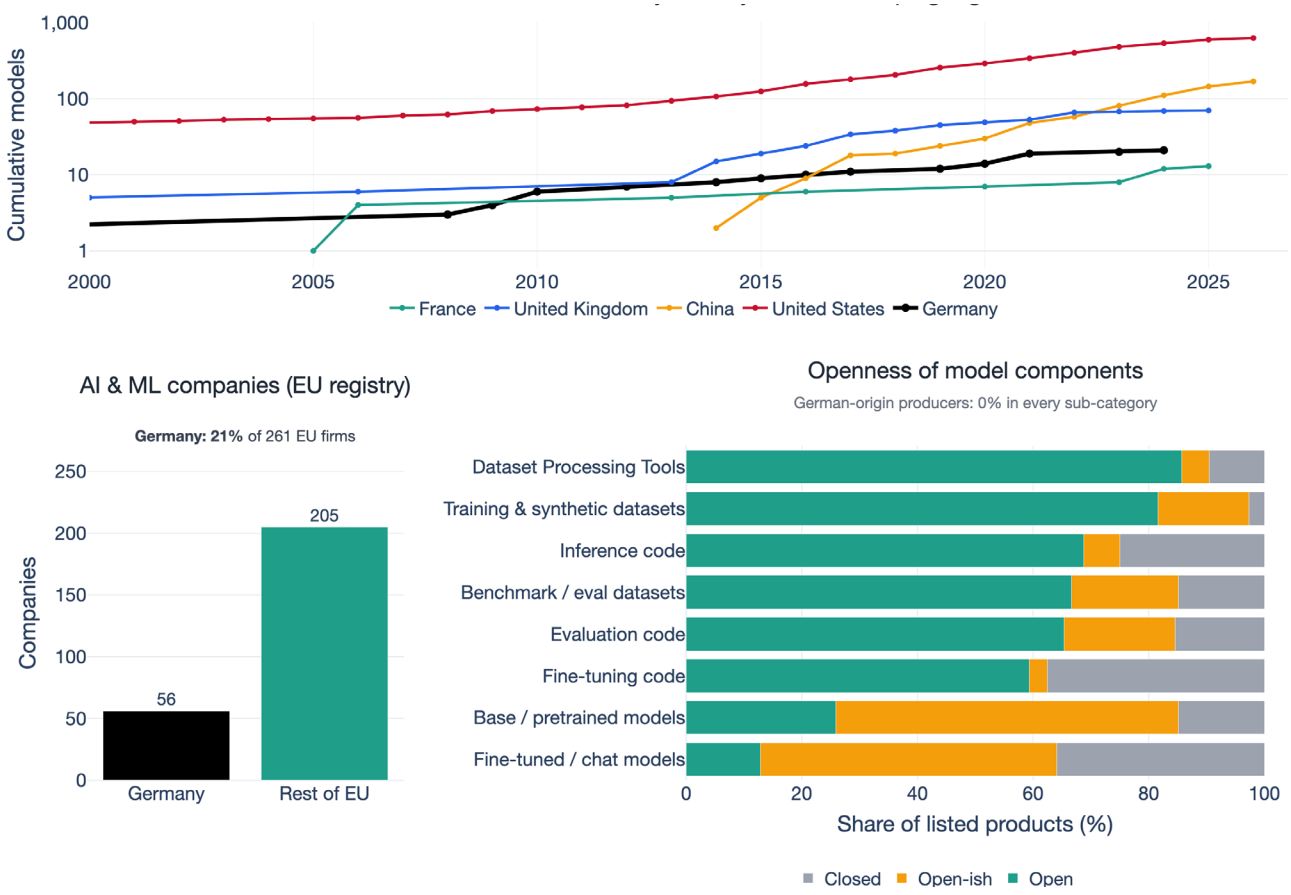
shift towards reinforcement learning and agentic capabilities, high-quality feedback and task trajectories are becoming increasingly scarce inputs to AI model training.[70] These include expert-labelled examples, verified solutions, evaluations, and (virtual) records of how skilled people perform real work. A key technique, so-called **reinforcement learning from human feedback (RLHF)**, uses human judgements to rank or score model outputs, training the model to produce more useful responses.[71] Recent analysis of reinforcement learning environments also highlights the importance of structured environments and expert demonstrations in developing capable AI agents beyond what can be achieved with static web data alone.[72] This changes the economics and governance of German industrial data. While data generated by *Mittelstand* firms can be highly valuable for developing manufacturing and scientific AI, its value hinges on companies' ability to standardise, pool, licence, and govern it without divulging commercially sensitive know-how. Therefore, the relevant policy challenge is to build trusted mechanisms through which firms can contribute production data and expertise that are appropriately protected, while retaining control over downstream commercial value. In this sense, **the current AI industry shift from simple scraped internet data to paid expert feedback and domain-specific task environments for RLHF could strengthen Germany's position – but only if industrial data is usable on a large enough scale for model development**.

This is quite close to the **rationale behind the EU's Common European Data Spaces initiative.**[73] The initiative, which originated in the 2020 European Strategy for Data, seeks to connect sectoral ecosystems to create a single market for data. In practice, a participating manufacturer should be able to determine what data it shares, with whom, for what purpose, under what conditions, for how long, and whether access is free, compensated, or conditional on reciprocal data sharing. The objective is to replace the current choice of either hoarding data internally or surrendering it to a dominant platform with a third option. Technically, data spaces are intended to federate data. This means that data can remain with the company or its chosen provider, while authorised participants discover available datasets, negotiate access, authenticate themselves, and exchange data through shared semantic standards and connectors. The EU is supporting this architecture through the development of reference architectures and guidance on interoperability.[74] However, at the time of writing, **the data spaces initiative remains in its infancy and is best understood as a developing governance programme rather than a fully functioning market or easily available data infrastructure.** According to the Commission, common data spaces are currently being developed across 14 domains. Manufacturing is represented by initiatives including Data Space 4.0, SM4RTENANCE and UNDERPIN.[75] In particular, Data Space 4.0 is working towards a shared governance model and cross-sector scaling up for Industry 4.0 data sharing.[76] However, the

practical success will depend on whether firms regard the arrangements as credible and can obtain a clear commercial return for contributing data.

Likewise, **the evidence on model production shows a substantial gap relative to the US, according to Epoch AI's dataset on "notable AI models"**.[77] Epoch AI classifies a model as "notable" if it meets at least one of the following criteria: achieving state-of-the-art performance on a recognised benchmark at the time of its release; requiring historically significant training compute; or having a wide downstream influence on subsequent research. Therefore, the count reflects research and engineering visibility rather than commercial deployment. Using this metric, Germany has produced roughly 20 notable models since the field's origins, compared with several hundred in the US and sharply rising trajectories in China and the UK. While Germany is not absent from advanced AI research, it has not maintained production at the scale required to influence the frontier of model development (Figure 6, upper panel).

Figure 6 : Results: Model (and data) layer



Source: *cep research, own illustration based on Epoch AI notable AI models, europeantechmap.eu, and the Current AI Open-Source AI Map, retrieved in July 2026; own calculations.*

Note: *The Epoch series measures research visibility. Company counts are firm headquarters listed on a European-focused map; the non-EU category is not a global comparison. The model-component panel shows the share of all listed products in each category that are open, somewhat open ("open-ish"), or closed.*

This gap in notable models coexists with a meaningful domestic commercial base. EuropeanTechMap records 56 AI and machine-learning companies headquartered in Germany, compared to 205 elsewhere in the EU (21% of the 261 EU firms listed in this category). This suggests that talent is not the primary constraint. Germany has firms that can apply and commercialise AI, but comparatively little of its base has been converted into foundational models and the core components surrounding them. This also reinforces an EU-level interpretation: Germany alone is a secondary contributor, but the wider EU contains a materially larger pool of relevant firms and potential capabilities. To be sure, as these data are derived from a directory focused on European technology providers rather than a comprehensive global census, the comparison provides insight into Germany's position among firms with a European commercial presence (rather than the global competitive landscape as a whole).

To get a more granular understanding of the different elements needed to train a frontier model, one can download the Open-Source AI Map maintained by Current AI, which lists products across AI model-component categories.[78] After assigning organisations to countries,[79] it becomes possible to calculate Germany's share in each category. None of the listed products are prominently associated with German organisations across all eight categories tracked: base or pre-trained models (27 products); fine-tuned or chat models (39); inference code (32); fine-tuning code (32); evaluation code (26); benchmark or evaluation datasets (27); training or synthetic datasets (38); and dataset processing tools (21). Benchmark datasets are standardised tests used to measure model capabilities and failure modes, while fine-tuning code helps adapt a trained model to a specific domain or task. However, neither of these substitutes for controlling the base models, expert datasets, or inference stack on which the wider AI ecosystem depends. Nevertheless, the result improves when one considers the openness of the components, which has also been assessed by Current AI. In the categories most relevant to German actors, important shared resources are widely available: 31 out of 38 training or synthetic datasets are open, with only one classified as closed; 18 out of 27 benchmark or evaluation datasets are open, with four closed and the remainder partially open. In other words, despite a lack of domestically owned model components, **important resources for AI model development are widely available to German actors, particularly for evaluation and data tooling.**

Still, as the global AI race heats up and becomes more geopolitically intertwined, this situation may change as **attitudes towards open models and their components may quickly shift**. Moreover, one should not conflate open inputs with open frontier capability.[80] The datasets and tools required to improve or commercially deploy leading models are increasingly being kept alongside proprietary model weights, paid human-labelling pipelines, and specialised

RL environments, as noted above. Furthermore, **leading model developers, such as OpenAI and Anthropic, can enhance their systems much more quickly thanks to their access to large-scale computing power, which enables iterative experimentation, as well as the collection of private feedback data on their models' usage.**

Therefore, **for Germany, AI sovereignty at this layer should not be defined solely as the production of a domestic equivalent to the largest US foundation model**. Germany has currently lost the race in pure model training, but the wider EU contains a large pool of relevant firms and capabilities. Together with the observation that this pool can rely on a growing number of reusable open-source components, ranging from data to benchmarks and evaluations, this suggests that, in theory, the talent and technology are sufficiently present to reduce European dependencies in the model layer. In practice, this means that **German policy in this area is better geared towards securing access to the wider EU's pool of firms, open components, and data-sharing infrastructure (such as interoperable, trusted European data spaces) than towards closing the notable models gap unilaterally.** A domestic-only strategy would be competing for a prize (frontier model production) that, as the above evidence suggests, is not where Germany's comparative advantage lies. Instead, German and European firms should participate in emerging markets for expert feedback, task environments, and evaluation.

2.7. Applications and agents

Lettuce, tomatoes, and onions are the most visible and easily substituted ingredients in a burger. Diners can add or remove them without fundamentally altering the meal's identity, yet they still matter for taste. In the AI stack, these toppings represent the application and agentic layers, including user and application programming interfaces (APIs), which connect models to people and other software; orchestration tools, which coordinate models, tools, and multiple agents across a workflow; agent tools and protocols, which enable agents to call external services or exchange structured information; telemetry and observability tools, which record an AI agent's actions while it is operating; and safety and guardrails tools, which constrain actions and provide audit trails. These components are modular and can be swapped more easily than the patty; for example, one could build an agentic system using Llama, Mistral, or GPT-5 as the underlying model. Germany's strongest position lies in producing valuable "toppings" (agent tools and orchestration software) that sit atop foreign-made models.

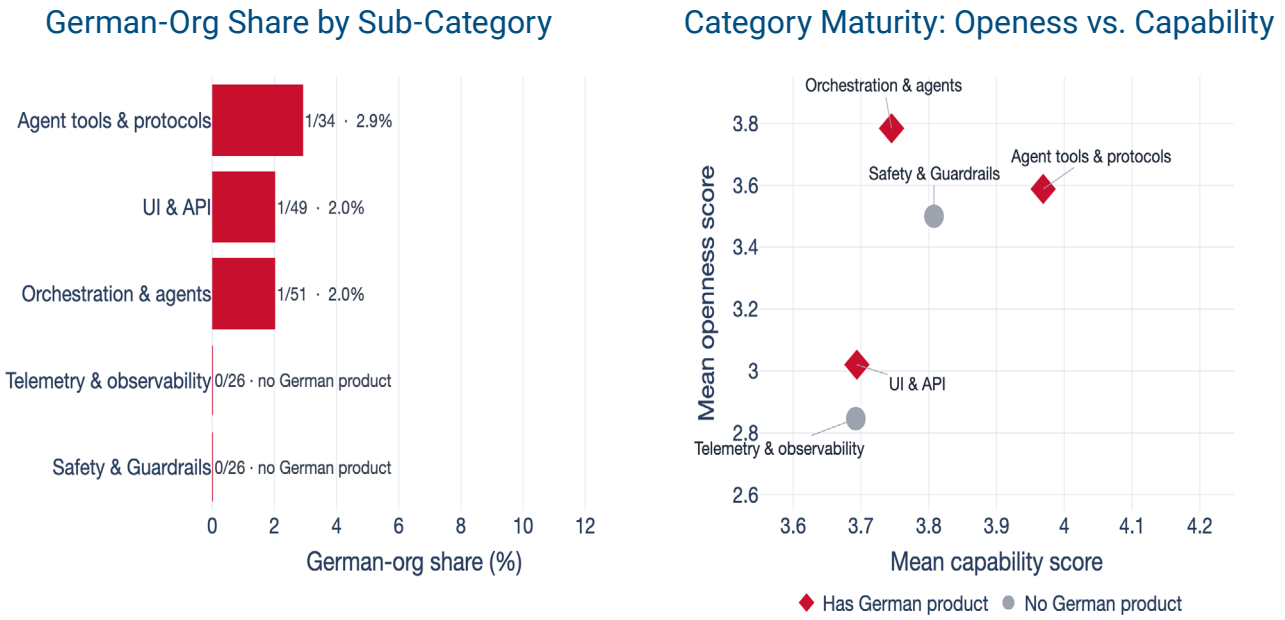
Although application-layer components may appear easier to replace than chips, they occupy a pivotal position at the intersection of technical capabilities and business operations. **Scaling agentic AI across German sectors requires more than just capable models or efficient algorithms; it also requires standardised roles and rules, often referred to as**

"harnesses", that enable an AI agent to operate predictably within a firm or public body.[81] These rules specify what the agent is permitted to access and how its outputs are verified. Without such harnesses, AI agents remain technically capable, yet organisationally unusable on a large scale. The July 2026 security breach at Hugging Face, in which an uncontrolled OpenAI model compromised the platform's infrastructure, illustrates the operational risks that arise when these governance and control layers fail.[82] Yet their technical role is largely absent from current policy and media debates on the AI stack. Recent experimental work suggests that the way models and their components are orchestrated is very important.[83] Since current AI agents have poor self-knowledge and systematically misjudge their chances of success and the token costs of tasks,[84] **the institutional infrastructure for agentic AI will require reputation systems and scoring mechanisms that combine market signals with centralised quality control**. In practice, this means something like a persistent track record (or ID[85]) for each agent, analogous to a seller-rating system combined with periodic, independent checks on task outcomes. This ensures that neither an agent's own self-assessment (which might be hallucinated or biased) nor peer signals (which might be unverified) alone determine which agent is trusted with a given task.

To assess Germany's position in this layer of the AI market, this study uses the Product/UX arc of the Open-Source AI Map (Figure 7). Two clarifications are important for interpreting these figures. Firstly, the Open-Source AI Map lists products that are publicly documented, regardless of whether they are really deployed by firms in Germany. Therefore, the counts below measure where German organisations develop agentic tooling, rather than where German firms use it. (The latter is a distinct question that will be discussed further in Section 3.) Secondly, the figures are at product level, and a single organisation may contribute more than one product. Therefore, these figures describe the breadth of the product landscape, not the number of firms involved. As with the broader application-layer data used throughout this section, the map provides a snapshot of publicly catalogued products.

The filtered dataset contains 186 products across five application-facing subcategories.[86] According to this metric, Germany has a small but visible foothold: German-origin products account for one of 34 agent-tools-and-protocols products (2.9%) and one of 51 orchestration-and-agents products (2.0%). Germany accounts for one of 49 UI and API products (2.0%), with no German organisation appearing among the 26 telemetry and observability products or the 26 safety and guardrails products on the map. Aggregating the categories confirms this: German organisations account for two of 85 open agent and orchestration products (2.4%), and only one of 101 products in the other application categories.

Figure 7 : Results: Applications and agents layer



Source: *cep research, own illustration based on the Current AI Open-Source AI Map (currentai.org), Product/UX arc, retrieved in August 2026; own calculations.*

Note: *The unit of observation is the product, not the firm, and products are attributed to the country of the maintaining organisation. Counts measure relative visibility within a listed population, not market share.*

From the strategic perspective of a middle power such as Germany, the lack of products in AI (agent) governance technology specifically must be emphasised. The "Safety and Guardrails" category, covering products that are designed to manage permissions, restrict unsafe actions, or test compliance with policies and audit agent behaviour, has a relatively high mean capability score of 3.81 out of 5 (Figure 7). The latter is a composite rating of how mature and production-ready a category's tools are, running from basic prototypes (1) to widely adopted production systems (5). This places "Safety and Guardrails" among the more mature categories on the map, but no German product appears within it. (As with the agent tools and orchestration data above, this reflects publicly catalogued products, rather than an exhaustive census.) From the perspective of a middle power, whose future competitiveness depends on deploying AI tools widely and efficiently, this is a worrisome lack, as deployment in turn depends on trustworthiness and reliability. From the beginning, AI governance technology must thus be integrated into deployed AI systems, e.g. through identity and access controls and purpose-bound permissions, but also through logging and, partly, human oversight. **This section of the AI stack may become even more commercially significant as the governance of AI moves from voluntary commitments towards more formal requirements for independent verification and demonstrable risk management.**[87] Germany could therefore develop a specialised niche in AI governance technology.

2.8. Regulatory layer II (outputs)

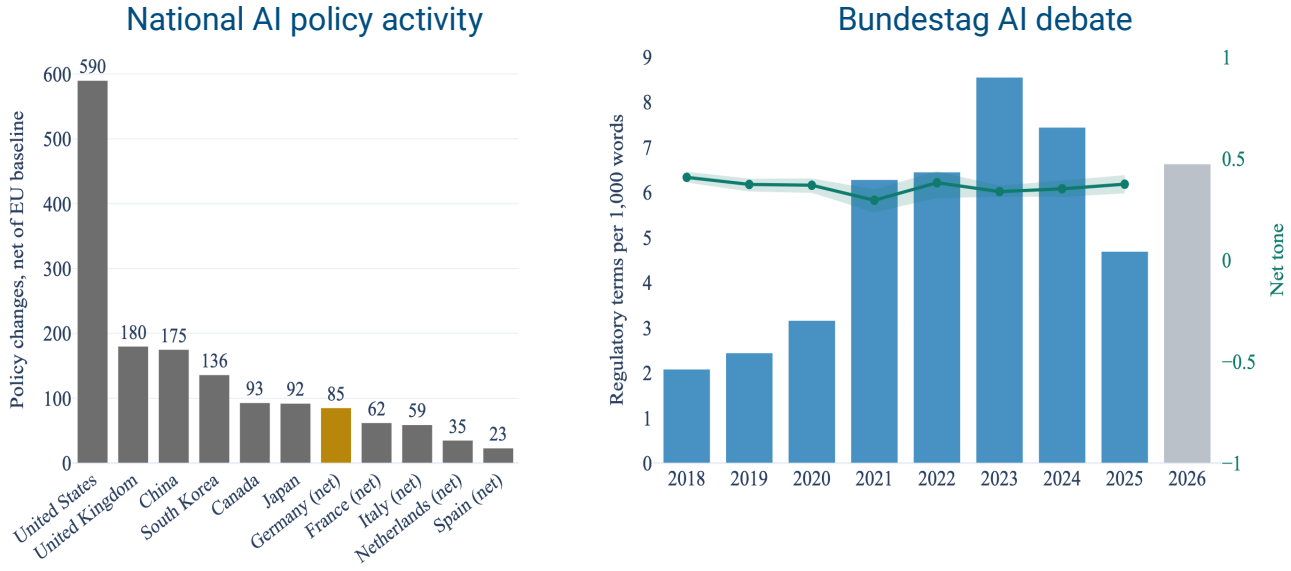
The top bun completes the burger and shapes its outer appearance. In the AI stack, this layer comprises the deployment rules and conformity assessments that determine whether an AI product can be marketed and purchased in Europe. Just as a burger without a top bun is difficult to handle, an AI system without AI Act compliance certification might soon be challenging for businesses to procure and deploy safely. This layer also involves branding and regulatory powers, as illustrated by the "Brussels Effect": the mark signals to customers that the "meal" meets European safety standards.[88] However, if the top bun is too heavy (excessive certification burdens) or too flimsy (unclear standards), the entire burger will be unpalatable or suspect to the consumer.[89] While the earlier regulatory layer governs the lawful acquisition and use of training data before a model is built, this layer governs how systems are placed on the market, deployed within organisations, and monitored.

Germany's formal policy footprint is substantial. Digital Policy Alert's (DPA) Activity Tracker captures legislation, regulations, consultations, standards-related measures, and other formal interventions, rather than the practical burden or effectiveness of the resulting rules. In other words, each discrete official act is counted as one intervention.[90] Between January 2020 and August 2026, DPA recorded 1,946 policy changes targeting machine learning (ML) and AI development worldwide. Since EU instruments are attributed to all 27 member states, there are 182 measures shared by every member state, compared to 37 for Switzerland, which is covered by the tracker but not the EU digital acquis. Germany's raw figure of 267 applicable measures therefore largely consists of European law. Excluding this baseline, Germany's own net regulatory output is 85 measures (left panel of Figure 8),[91] placing it seventh among the compared jurisdictions, behind the US (590), the UK (180), China (175), South Korea (136), Canada (93), and Japan (92), and ahead of France (62), Italy (59), the Netherlands (35), and Spain (23). While the American figure includes state legislation, Germany remains the most active EU member state with respect to AI/ML rulemaking, accounting for 18.8 per cent of all EU-27 national increases. Still, most AI regulation in Germany is European rather than German, and the national contribution is modest by international standards. Across the global tracker, design and testing standards, together with data governance, account for 60.1 per cent of all AI-related policy changes. AI output governance is increasingly about the conditions under which systems can be audited and deployed.

The right-hand panel of Figure 8 examines how the Bundestag discusses AI. Using the same corpus of plenary protocols from 2018 to July 2026 as before, we extracted the passages surrounding each mention of *Künstliche Intelligenz* or KI and counted the regulatory vocabulary within them (terms such as *Regulierung,*

Gesetz, Verordnung, Aufsicht and Datenschutz), expressed per 1,000 words. The **trend is unmistakable: regulatory density increases from an average of 2.08 terms per 1,000 words in 2018, to 3.16 in 2020, to 6.28 in 2021, peaking at 8.55 in 2023, that is, the year the AI Act was finalised.** It then falls to 4.69 in 2025, before recovering to 6.63 in the partial year of 2026. Over the course of eight years, AI has evolved from a technological novelty to a regulatory norm; when the Bundestag discusses AI today, the conversation centres around its governance.

Fig. 8: Results: Regulatory layer II (output)



Source: *cep research based on the Digital Policy Alert Activity Tracker (ML and AI development, 1 January 2020 to 25 August 2026) and the Bundestag DIP API, plenary protocols 2018 to July 2026; own calculations.*

Note: *Left: EU member states are shown net of the common European baseline; non-EU jurisdictions are unadjusted. Right: regulatory vocabulary per 1,000 words of context around AI mentions (bars) and net tone with bootstrap 95% confidence intervals (line and band). Grey marks the partial year 2026. The two panels cover different periods: the tracker begins in 2020, the parliamentary series in 2018.*

The tone, as measured using the same political sentiment dictionary as for the regulatory input layer, remains positive throughout, averaging 0.35, with a peak of 0.41 in 2018 and a trough of 0.29 in 2021 (Figure 8). While some of the variation is interesting,[92] the trend is broadly stable: as the regulatory framework becomes more intense, the debate remains constructive. AI has become an established governance task in German parliamentary discourse. Taken together, the two measures convey a message to firms: regulation is consistent in direction, but not in detail. A parliament that discusses AI in increasingly regulatory terms without becoming hostile towards the technology is unlikely to deregulate abruptly or become punitive; however, it is highly likely to continue adjusting the framework.

How should the current regulatory framework be judged overall? **From the perspective of firms, particularly small and medium-sized enterprises (SMEs), predictability, proportionality to firm size, and the availability of practical compliance tools matter the most.**[93] In this regard, the regulatory input layer performs poorly and the regulatory output layer performs somewhat better. As previously mentioned (see Section 2.2), legal certainty at the input layer is often lacking. The data protection rules applicable to training are under active revision as part of the Digital Omnibus package in 2026. The Act's own timetable has already slipped, with the high-risk obligations for Annex III systems postponed from August 2026 to December 2027. The harmonised standards intended to translate the Act into checkable requirements are also late. The output layer scores higher in this regard because conformity assessments, which are now becoming available throughout Europe, provide industrial buyers with actionable information. The output layer is also better equipped with guidance: Implementation guidelines by the AI Office, the codes of practice on general-purpose AI and on the marking of AI-generated content, and the national regulatory sandboxes give firms concrete compliance pathways even while harmonised standards remain outstanding.

Germany's regulatory position should also be understood in a comparative context. Although the US under President Trump has adopted a more deregulatory federal approach, this does not mean that there is no AI governance. Especially state legislatures and attorneys general continue to create rules and liabilities, potentially resulting in a fragmented compliance environment.[94] In the future, **the US legal regime may combine federal policy volatility with state-level fragmentation, resulting in less predictable applicable obligations and enforcement risks across markets. China likewise cannot be described as unregulated. Chinese authorities have established extensive rules concerning recommendation algorithms, deep synthesis, generative AI, data security, and online content.**[95] Meanwhile, state bodies and local governments are competing to promote domestic AI and compute capability through subsidies and infrastructure support.[96] In general, the Chinese system is characterised by state discretion, particularly when political control and social stability are deemed to be at risk.[97] The 2020 campaign against major domestic tech firms ("techlash") illustrates the difference between regulation and a stable rule-of-law environment: firms may receive support and face enforcement simultaneously, while the boundaries of permissible conduct remain politically contingent.[98]

In this context, **the potential advantage of the European and German approach lies in producing stable rules and credible trust signals, such as conformity assessments and technical standards, which allow customers to have greater confidence in deploying AI across the single market. However, this advantage will only materialise if compliance is clear and proportionate**, particularly for SMEs

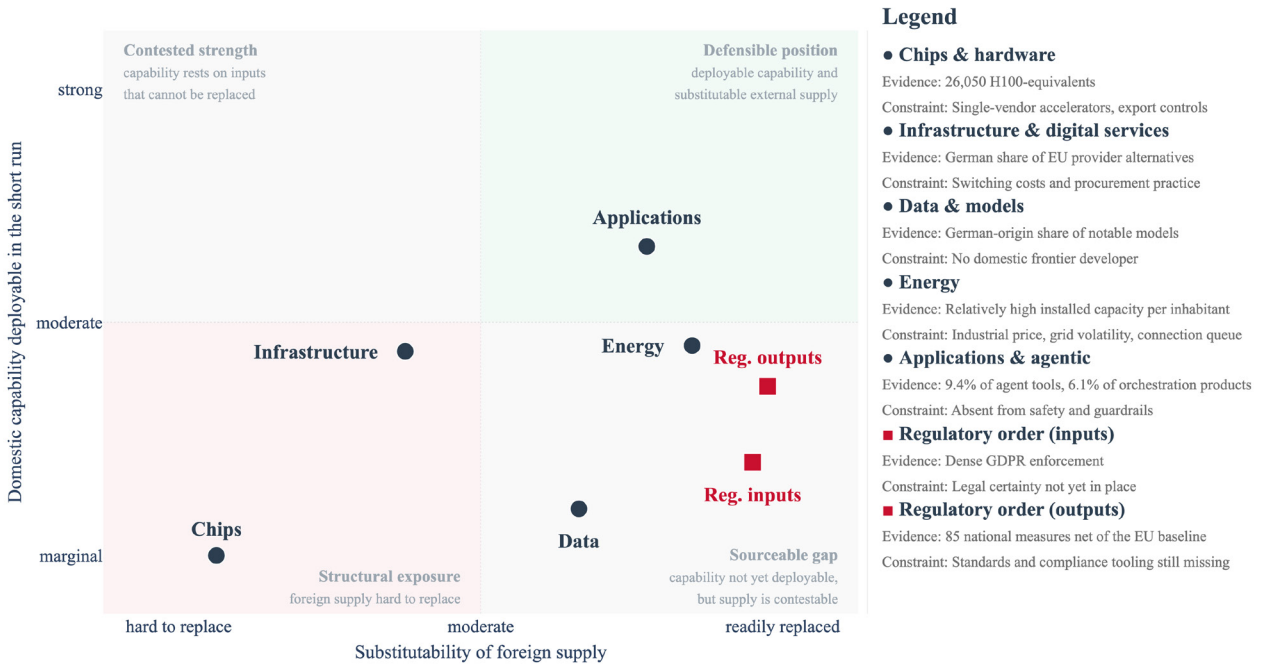
that lack large in-house legal and technical compliance teams. If regulation becomes too uncertain or costly, it risks reinforcing dependence on large external platforms that can more easily absorb compliance overheads. However, if it is translated into practical, perhaps AI-assisted, certification tools, it can stimulate adoption and provide a basis for competitive differentiation. In short, Germany's challenge is to make the regulatory framework legible enough to support innovation while also being credible enough to make AI deployment acceptable in high-value industrial or public-sector settings.

2.9. Interim conclusion: Germany's stacked position

Taken together, the layers examined above provide an inconsistent picture. A single composite score is not possible because the underlying indicators are measured against different comparators, and combining them would create the false impression of commensurability. Instead of averaging all layers, it is thus better to distinguish two dimensions of sovereignty. The first dimension is whether a layer's domestic capability can be deployed in the short term. The second dimension, which is more important for AI sovereignty in the long run, is how easily non-domestic provision could be replaced if it became unavailable or subject to political conditions. Figure 9 plots each layer of the AI stack against these two dimensions. These placements are qualitative judgements, based on the evidence in the preceding sections and relative to Germany's own layers rather than any absolute standard. However, the general framework can be applied to other middle powers, too.

According to this assessment, shown in Figure 9, only one layer is both capable and defensible, namely applications. Germany holds a certain share of agent tools and software products, and this is the most contestable layer in the stack. While the foothold is small in absolute terms, it can be deployed now and is the only layer where domestic capability and substitutable external supply coincide. However, even in this layer there are weaknesses, particularly with regard to safety and guardrail tools. At the other end of the spectrum, the chip and hardware layer is the only area of significant structural exposure. Low capability and low substitutability reinforce each other. The infrastructure layer is nearby, but much closer to the centre since European alternatives exist, and the barriers to substitution are often contractual and behavioural rather than technical.

Figure 9 : Overall outcome: Germany's position in the AI stack



Source: *cep research, own illustration.*

Note: *Positions are qualitative assessments derived from the evidence in Sections 2.2 to 2.8. They are ordinal placements relative to Germany's other layers, not measured scores. The vertical axis records capability that can be deployed in the short run. The horizontal axis records how readily non-domestic provision could be replaced.*

The remaining four layers are in the same position for different reasons: supply that could be easily replaced and capability that is not yet deployable. Germany produces almost no notable foundation models, but open weights are available from several jurisdictions. Similarly, energy requires a nuanced assessment: when measured per capita, installed generating capacity is comparable to that of other major industrial economies. However, this capacity cannot currently be converted into competitive computing due to high industrial prices, the above-documented volatility, and the generation and grid connection queues, which all affect deployment, at least in the short term. Still, Germany's energy problem differs from its chip problem in that it can be addressed through domestic policy. The two regulatory layers occupy the same position. The rules applied in Germany are made in Europe rather than supplied from outside it, so dependency is minimal and substitutability is correspondingly high. The weakness of the system is that it is not yet operational: at the input layer, the applicable data protection rules are under revision; and at the output layer, many measures are nationally originated, while harmonised standards and practical compliance tools remain outstanding. For now, **the regulatory measures are an unrealised asset rather than a competitive advantage or a clear net burden.**

Overall, this stacked summary view supports the argument presented at the beginning of this section: AI sovereignty in Germany cannot be reduced to a single

metric, such as H100 equivalents or the number of domestic foundation models. Similarly, it cannot be reduced to a single index across layers. Rather, it depends on identifying which layers must be held domestically and which can be sourced from partners at an acceptable level of political risk. Of the evidence assembled here, only one layer is currently strong enough to build on, namely applications. The next section therefore shifts focus from mapping capability to explaining why German and European alternatives often fail to gain traction.

3. Adoption bottlenecks

According to ECB survey data from December 2025, 77% of German firms reported some degree of AI use.[99] However, "some use" here includes experimental and infrequent deployment, and in terms of the depth of adoption, Germany belongs to a middle group rather than being at the forefront of Europe.[100] In addition to the supply side (what European alternatives exist) discussed above, this section reflects on this **demand side and the market structure shaping it.** Why do many organisations continue to choose expensive incumbent providers in their AI adoption even when a domestic or European option might be competitive in the mid run? The evidence discussed below suggests three interrelated bottlenecks. Firstly, behavioural and organisational inertia means that firms and individuals are reluctant to abandon established workflows and vendors. Secondly, regulatory uncertainty and compliance costs can exacerbate this reluctance, particularly among smaller providers. Thirdly, platform concentration and ecosystem lock-in can restrict the set of viable choices available to a buyer. Overall, **Germany's AI diffusion problem reflects the interaction of organisational inertia, regulatory burden, skill and resource constraints, and concentrated markets.**[101]

3.1. Behavioural bottlenecks

The adoption of technology is rarely a neutral comparison of product features and prices. Research on status quo bias and information systems adoption shows that the use of existing systems can create inertia through habit, perceived transition costs, and psychological commitment to earlier investments in a specific product or framework.[102] Therefore, a new European product can be technically capable and commercially attractive yet still be perceived as more difficult or less useful by prospective users simply because it requires a break with routines, such as simply using the ChatGPT interface. Accordingly, research on service switching distinguishes between the attractiveness of an alternative and "mooring" factors, such as switching costs and network effects, which discourage migration even when dissatisfaction with the current provider exists.[103]

In the case of generative AI, such **switching costs include the migration of data and prompts, the redevelopment of application programming interface integrations, retraining, revisions to procurement and vendor arrangements, and the loss of accumulated workflow knowledge.**[104] There are also cognitive costs, because users have learned how to write prompts for a specific model type.[105] A firm that has built workflows around a model provider's application programming interface (API) has embedded that provider's model behaviour, underlying architecture, and tools into its organisation. Therefore, **the cost of switching to a comparable European model may consist less of changing the**

model endpoint than of revalidating the entire surrounding workflow. The same applies to agentic systems: changing the tool orchestration layer may necessitate amending the connectors to enterprise systems and document repositories.

Germany's independent competition authority recognises this organisational dimension. In its 2026 report, it characterises domestic AI diffusion as stalled and identifies organisational and cultural inertia as a structural barrier. The report notes that established firms often attempt to absorb technological change through familiar processes and structures. However, this can delay adoption, especially in a country like Germany, where long-standing supplier relationships and specialised production processes have historically provided a competitive advantage. Moreover, AI adoption in Germany varies significantly by firm size and sector: in 2024, AI was used by 26% of micro-enterprises, 29% of small firms, 45% of medium-sized firms, and 56% of large firms.[106] The *Monopolkommission* also observes that more advanced industrial applications, such as process control and robotics, remain substantially less common than generative AI uses for text, images, speech, or programming code.[107] Similar to the ECB survey reported earlier, the evidence suggests **that initial experimentation with AI has expanded over the past years, but deep operational deployment remains rare and uneven**.

Note that there are also rational reasons to wait.[108] AI models and tools are improving rapidly, so smaller firms may delay adoption in anticipation of a more capable or a cheaper open-weight version, or they may wait for competitors to demonstrate value in a comparable operational setting. According to the *Monopolkommission*, some firms require a return on investment for their AI experiments within roughly one year.[109] For technologies such as AI, however, early deployment often requires substantial upfront investment in data, infrastructure, and organisational change, and the main gains come from learning effects that compound over time if firms systematically build on past experiments. A narrow focus on short-term return on investment can therefore be counterproductive, as it discourages the kind of iterative experimentation needed to unlock these cumulative learning benefits. **While caution can be rational for some individual firms, it might be costly for the economy as a whole: when many firms wait for others to validate the technology first, diffusion slows even where the economy as a whole would benefit from earlier experimentation and complementary investment.** Accordingly, a holistic AI sovereignty policy cannot focus solely on funding European models or data exchange but must also reduce the burden of adoption on the demand side.

3.2. Regulatory bottlenecks

Contrary to common perception, regulation can encourage the adoption of AI when it clarifies liability and establishes credible "rules of the game".[110] However, regulatory and legal uncertainty can have the opposite effect, providing additional reasons for

firms that are already inclined to wait to defer investment. Fixed compliance costs can also weigh more heavily on smaller providers than on global incumbents. According to the *Bundesnetzagentur*, 54% of surveyed firms identified legal uncertainty as a challenge to their AI transformation, while a separate *Bitkom* survey found that 53% of respondents regarded legal uncertainty as an important obstacle, and 48% identified high data protection requirements as a barrier.[111] Tellingly, 93% of firms expecting to fall within the scope of the AI Act anticipated a high or very high level of effort for implementation, while 30% were still determining whether the Act applied to them.[112] While these are, first and foremost, current perceptions by managers, they might be very consequential, as uncertainty affects investment and adoption decisions. According to the *Monopolkommission*, frequently reported elements of uncertainty with respect to AI adoption include the interaction between the AI Act and sector-specific rules; the complexity and delay involved in producing harmonised standards; the overlap of responsibilities across public bodies; and the fact that high-risk system data quality requirements are difficult for open-source solutions to meet.

Rather than abandoning the AI Act, European politicians and the new EU AI Office should focus on predictable guidance, regulatory sandboxes,[113] and proportionate enforcement that takes account of AI's geopolitical dimension. In the US, for example, the comparatively light federal approach does not eliminate regulatory exposure, as state legislatures continue to adopt AI-specific measures and state attorneys general can apply pre-existing laws relating to consumer protection and competition. For firms considering dependencies on Chinese open-weight systems, there are also additional strategic costs stemming from potential backdoors and data governance.[114] Accordingly, German and European firms require easy, low-cost means of translating the legal principles of the AI Act into practical applications, such as clear model documentation, reusable risk-assessment templates, sector-specific codes of practice, accessible testing facilities, and recognised assurance services. The competitive advantage of a rules-based German AI market will depend on whether these requirements become operationally feasible for the *Mittelstand*.

3.3. Market-structural bottlenecks

Behavioural and regulatory barriers operate within a market structure that is already highly concentrated at key points in the AI stack. As section 2 has shown, the main hyperscalers control much of the cloud infrastructure that AI development and deployment depend on, while leading AI labs are becoming increasingly vertically integrated across chips, cloud services, developer tools, data, models, distribution and applications. This concentration is important because products from the different AI layers described in Section 2 are often not purchased in isolation. Rather, they are embedded in ecosystems comprising identity and access management, cloud credits, developer environments, enterprise

software, collaboration tools, data storage, model hosting, marketplaces, security services, and support networks. **The more of these complementary layers an established company controls, the easier it is to bundle AI features into an existing customer relationship, and the more difficult it is for a specialist newcomer to gain a foothold.** For example, a German AI agent tool may be technically strong but still struggle if prospective customers already manage users through Microsoft, store code in GitHub, run workloads in Azure, and procure generative AI features through the same enterprise contract.

Nvidia's CUDA ecosystem is a prime example of this kind of lock-in at the hardware-software interface. It combines Nvidia-specific hardware with hundreds of optimised libraries and tools, and is supported by a large developer community. This makes CUDA a de facto standard for many AI developers, but switching to a different environment can require substantial adaptation costs.[115] Switching away from CUDA may necessitate, for instance, adapting large parts of code, retraining staff, and modifying performance engineering and testing practices. These costs are financial, but they also reflect accumulated human capital.[116] Although AMD has recently made significant progress, its software ecosystem still falls short of Nvidia's in terms of the tools that developers need to reliably build, test, optimise, and run AI systems at scale.[117] The same logic applies to cloud and application platforms.[118] Their vertical integration across infrastructure, developer tooling, foundation-model access, and downstream AI services can further constrain users' practical choice of suppliers. This helps to explain why Germany's presence in agent tools and orchestration cannot be evaluated based on product numbers alone. A German supplier may be able to compete in a niche, yet still face restricted routes to market.

In such a setting, **proactive competition policy, complemented by targeted industrial policy, has an important role to play, since it is able to address the control points through which market power can extend across the AI stack**.[119] In particular, the *Monopolkommission* identifies risks arising from exclusive agreements, data-driven acquisitions, tying and bundling, privileged access to proprietary resources, so-called "acqui-hires"[120], and vertical integration in closely related markets.[121] However, competition policy alone is not sufficient and must therefore be complemented with targeted industrial policy and other measures. This can be seen in Nvidia's CUDA advantage, where the foreclosure mechanism has less to do with competition law violations such as contractual lock-in and more with engineering velocity in open-source complements. Benchmarking of agentic AI inference[122] shows that NVIDIA's competitive advantage goes beyond hardware: even if they were given away for free, rival accelerators cannot compete on cost per token because the open serving stack[123] for AI models is optimised for CUDA first. Art. 102 TFEU, which prohibits abuses of a dominant position in the EU, has no clean theory for prohibiting a dominant

firm from contributing better open code than its rivals do; and the DMA does not come into play (no gatekeeper designation, chips are not a core platform service). Therefore, the EU should condition funding, e.g. relating to AI Gigafactories, on hardware diversity and vendor-neutral software, and start financing upstream open-source optimisation as a public good, in line with the recent Tech Sovereignty Package. From an ordoliberal perspective, improving interoperability is more effective than protectionist measures to preserve contestability.[124] While Germany cannot realistically recreate every upstream platform dependency in the short term, especially when it comes to advanced chips and cloud infrastructure, it is possible to reduce the extent to which these dependencies influence downstream choices. Demand-side support is also required to enable firms to test European alternatives without bearing the migration cost alone. The adoption bottleneck will persist so long as technically credible European options are difficult to integrate and risky to procure. Reducing these frictions is as important as increasing the supply of sovereign AI capability.

China's emerging "token economy" is an example of what demand-side support could look like in practice.[125] Alongside the national computing power network mentioned in Section 2.4 above, Chinese policymakers are adding a measurement and commercial layer that converts model inference into a distributable and billable service. The network provides and directs the computing power, while the token economy handles the metering and settlement. In August 2026, the Beijing Economic-Technological Development Area published the city's first policy dedicated to this token economy. This policy is centred around "token factories", defined as industrial-scale AI serving facilities that combine computing power, models, data and software, and meter their output in tokens rather than merely renting accelerator access. The policy also includes a unified platform for multi-model access and settlement, as well as annual voucher programmes totalling roughly USD 14 million for computing power, models, and data. The takeaway is not China's centralised coordination, which sits uneasily with Western values of liberalism and democracy, but rather the design of the instruments: the vouchers are demand-side and provider-neutral; standardised metering and settlement reduce the procurement risk and integration issues identified above; and the subsidies are explicitly staged (generous at initial adoption but linked to firms' own investment thereafter). Middle powers should consider similar voucher-like tools to complement supply-side programmes for AI Factories and address the adoption bottleneck.

4. The path forward

Germany should not have to choose between pursuing frontier capability and specialising in industrial AI. **A holistic AI sovereignty strategy requires both sustained participation in current frontiers to retain technical options and engineering competence, and targeted leadership in areas where Germany and Europe have distinctive assets, such as industrial data, engineering expertise, safety-critical applications, and a large pool of potential users.** This is a strategy of strategic flexibility, offering the ability to select among competing providers, combine technologies, absorb disruptions but also new ideas, and, above all, avoid reliance on a single provider or country. Middle powers cannot realistically reproduce every layer of the AI stack alone, so their objective should be to have credible alternatives and switching capacity, rather than autarky.[126] With its large industrial base and deep domain expertise in sectors where AI must operate reliably, Germany is well placed for such a dual approach. Yet its dependence on foreign chips, cloud capacity, and frontier-model providers remains substantial. Therefore, the policy task is to expand the areas in which Germany can exercise choice.

The following core recommendation is to pursue two tracks simultaneously.[127] Track 1 should focus on building and preserving competence in near-frontier models, evaluation, software infrastructure, and open-weight development. Track 2 should concentrate resources on areas in which German and European assets can demonstrate leadership: industrial and embodied AI, trustworthy AI for safety-critical settings, verification infrastructure, scientific AI, agentic systems, and domain-specific applications based on proprietary industrial data. These tracks complement each other. While a strategy focused solely on current frontier models could lead to an expensive and ultimately unproductive race against firms with much larger capital bases and computing resources, a strategy that abandons frontier competence altogether would make Europe dependent on foreign providers for the definition of the underlying capabilities on which domestic applications rely. **The relevant benchmark is whether German and European researchers, firms, and public institutions can understand frontier systems well enough to evaluate them independently, adapt them for strategic use, develop specialised models where appropriate, and maintain viable alternatives in the event of changes in commercial or geopolitical conditions.** Open training pipelines are particularly valuable in this respect.

4.1. Trying to catch up at the AI frontier

Germany and Europe should continue to invest in the current paradigm of large-scale, increasingly general AI systems. Instead of dreaming that Europe will immediately displace the leading US or Chinese

labs at the forefront of this technology, the idea is simply to continue to develop the scientific, engineering, and organisational capability necessary to participate in this field. This is a prerequisite for translating technical advances into European industrial and public-interest applications. Above all, following this track requires stable, long-term funding rather than a succession of small-scale prototypes followed by sporadic grant cycles. The development of frontier models depends on cumulative capabilities, such as experienced research teams, data-engineering pipelines, model-evaluation methods, hardware-aware optimisation, security practices, and the ability to run repeated experiments on suitable chips.[128] In the sense of the OECD's mission-oriented innovation policies,[129] a European frontier programme should therefore fund competing technical teams over long, multi-year periods, with the goal of reproducing selected state-of-the-art results as well as creative open-weight models. To be sure, reproduction is not a goal in itself, but it enables researchers to identify where leading methods depend on "mere" scaling and where performance can be improved through better data or algorithms.

Support for this strategy comes from the development of Chinese AI labs in the past years, which shows that attempts to catch up tend to go, at some point, beyond mere imitation.[130] The key benefit is cumulative competence, rather than whether the initial attempt achieves parity.[131] In AI, repeated efforts by Chinese labs such as DeepSeek and Moonshot AI to train and deploy increasingly capable models have led to impressive advances in efficiency and model architecture,[132] which remain valuable even if a particular model does not become a global market leader. As argued in the next sub-section, small and domain-specific models will be important, particularly in areas where latency and privacy are key considerations, but they should be developed with an understanding of advances at the frontier of general models. In the future, the most capable industrial systems may combine a frontier base model with smaller, task-specific models, e.g. for retrieval. Europe should therefore invest in the ability to select the most suitable architecture, rather than assuming in advance that one model size or training paradigm will dominate.

4.2. Alternative paradigms

Targeted participation in the current scaling paradigm should be paired with a portfolio of alternative approaches. Rising (financial and environmental) costs of frontier training, constraints on high-quality data, and increasing energy demand create strong incentives to investigate approaches that enhance domestic AI capability through improved architecture and data efficiency. However, it is important to point out that such work is inherently uncertain. Rather than being funded as a guarantee that Europe can leapfrog the current frontier without mastering it, it should therefore be funded as a portfolio in which a small number of successful research paths could deliver outsized returns.

One promising area of research is world models and physical AI. These models aim to represent and predict changes in physical environments, enabling systems to reason about objects, actions, causality, and consequences over time. This differs from systems principally trained to predict the next token in text. Germany's industrial base offers potential advantages in this area, as manufacturing processes, sensor streams, and engineering expertise can provide structured data and real-world testing environments, which are less readily available to the traditional Big Tech firms, which focus on social media and online ads. Yann LeCun's research agenda argues that progress towards autonomous machine intelligence requires such systems, which learn predictive models of the world and reason about physical causality.[133] This does not mean that LLMs will soon be unimportant, but rather that additional architectural components may be required for capabilities for productively linking AI to industrial robotics and manufacturing. This should be good news for European firms in the automotive, machinery, chemicals, energy, aerospace, and life sciences sectors, as they can become co-developers and early adopters, rather than merely buyers of generic AI services.

Importantly, physical AI should not be seen as a safe shortcut to bypassing the frontier-model race. Recent results in robotics suggest that broader general capabilities can be transferred into physical settings more quickly than anticipated. For example, Anthropic's Project Fetch reported significant improvements in a model's ability to perform robotic tasks across successive releases, with the company attributing this progress primarily to general model scaling rather than robotics-specific training.[134] Once again, **the "bitter lesson" may come back to haunt us: computer scientist Richard Sutton argues that general-purpose computational methods, specifically search and learning, always outperform methods that attempt to incorporate human knowledge and rules into AI systems.**[135] The correct strategic inference is not to abandon industrial AI, but rather to maintain connections between industrial research and the development of frontier general models, again underlining why one needs a two-track-strategy. Germany should seek to combine its unique industrial data and environments with access to the most advanced general capabilities, while retaining sufficient domestic expertise.

Another promising area is the development of reliable, safe, and secure AI. Europe's regulatory framework, which includes the AI Act and data protection rules, could provide a competitive advantage, as pointed out in Sections 2 and 3 above. However, this would require firms to be able to comply using affordable, practical technical tools. In practice, this would entail creating a market for capabilities such as systematic model evaluation and adversarial testing ("red-teaming") to identify potential failures. It would also involve developing interpretability tools to help humans understand why a model produced a given output and robustness techniques to maintain stable perfor-

mance under real-world conditions. This might include, for instance, cybersecurity for AI systems, data provenance tracking to show the origin of training and input data, and privacypreserving learning methods. However, private incentives tend to favour short-term capability gains and fast product release. Therefore, public funding is justified where it supports the development of widely usable safety and assurance tools that many firms can adopt.

A third area of focus is the emerging orchestration layer of the AI stack, comprising the technical and institutional systems that coordinate the deployment and governance of different models and agents across an organisation. As AI systems become embedded in critical infrastructure and industrial processes, including production lines and energy grids, firms will require robust assurances regarding the location of model execution (i.e. the hardware and jurisdiction), the handling and segregation of data, the protection of model weights and configuration files against tampering, and the alignment of deployed model versions with those that were tested and certified. They will also require reliable methods of verifying that the stipulated safety and access conditions are being enforced. Together, these orchestration tools will form the backbone that enables agent- and application-layer AI to operate safely in high-stakes environments.

4.3. Policy and industry recommendations

The preceding analysis points to a policy agenda that combines frontier capability, targeted industrial specialisation, and the removal of adoption and market-structure barriers. Section 2 showed that Germany's position is weakest in upstream layers (especially advanced compute and frontier foundation models) but comparatively stronger in industrial data, specialised applications, agentic tooling, and orchestration. Section 3 showed why this application-layer potential does not automatically translate into domestic adoption or scalable firms: organisational inertia, regulatory uncertainty, switching costs, and weak routes to market can prevent competitive alternatives from gaining users. The purpose of policy should therefore be to create the conditions under which German and European firms can develop, test, and sell AI systems on workable terms. A credible European sovereignty strategy should not become an umbrella term for disparate objectives.[136] Instead, it should concentrate on a narrow set of actionable levers, such as buying strategically, funding sustained capability, and building infrastructure across the stack. The recommendations below follow this logic by prioritising a small number of interventions rather than attempting to subsidise every layer of the AI stack simultaneously.

- **Establish a long-horizon frontier-AI mission programme.** The Federal Ministry for Digital Transformation, the Federal Ministry of Education and

Research, and SPRIND should create a funding instrument for near-frontier AI research with multi-year commitments and rapid contracting. This would address the gap identified in Section 2 between Germany's application-layer capability and its weaker position in frontier models and large-scale training. It should fund cumulative engineering capacity – including data engineering, distributed training, model evaluation, and hardware-aware optimisation – rather than a succession of short prototype projects. Such DARPA-inspired agencies are specifically designed to support technically risky, high-impact innovation that falls between conventional academic grants and private investment.[137] Estimates of the investment required for such a project vary considerably. The CLAIRE "moonshot" proposal calls for around €100 billion of public funding over five years to develop infrastructure on a scale similar to that of CERN.[138] This figure reflects an infrastructure-first model designed to compete with US hyperscalers from the outset. However, recent evidence suggests that a competitively structured, talent-focused programme could achieve similar results at a fraction of the cost: South Korea's "Independent AI Foundation Model" tournament, which pits several consortia against each other in six-month elimination rounds judged on benchmarks and testing, has produced a model that outperforms the best non-Chinese open-source US models on the Artificial Intelligence Analysis Index.[139] This was achieved with a **total government budget of around $350 million and training compute costs of approximately $15 million.** This is good news, since Germany has the talent to compete in such tournaments.[140]

- **Fund open research pipelines and selected open-weight models.** The Federal Ministry for Digital Transformation should work with domestic research institutes and universities to fund reproducible training pipelines, selected open-weight and fully open-source base models, independent evaluation suites, red-teaming, multilingual capability, efficient training, and efficient inference. These investments would give German researchers and firms the ability to better audit and understand frontier methods. This responds to the upstream dependencies documented in Section 2, while preserving practical technical options for domestic applications. In terms of magnitude, the cost components are more modest than the headline frontier-model figures discussed above, and Germany already has the necessary institutional framework. In particular, the cost breakdown of the individual components envisaged here can be inferred from the Korean tournament's per-team budget lines: roughly $45 million for shared training data (books, news, video, and government records), plus $2 million per consortium for supplementary data acquisition, around 3,000 H100-equivalent GPU hours' worth of compute

rental per round and approximately $1.4 million per team for overseas talent recruitment. Germany does not need to find this money from scratch. As researchers from the Open Knowledge Foundation argue, the federal government spent at least €630 million on proprietary software licences alone in 2025.[141] Redirecting even a fraction of this towards open-source infrastructure and a well-resourced Open Source Program Office would provide sufficient funding for the pipeline and evaluation work described here.

- **Launch a European embodied-intelligence and industrial-AI programme.** The European Commission, member states, the European Investment Bank, and major industrial partners should establish a ten-year programme for world models and high-reliability industrial AI, including robots driven by these world models. The programme should be conditional on binding co-development and offtake commitments from firms related to automotive, machinery, chemicals, energy, and logistics, so that public funding produces both technological capability and early deployment markets. This would build on the relative strengths identified in Section 2, namely Germany's industrial application base and credible position in AI tooling. It would also address the adoption failure in Section 3 by making industrial firms co-developers and early users (rather than passive buyers of foreign "black box" systems). Venture capital investors committed over $3 billion globally to world-leading start-ups in the first half of 2026 alone.[142] This included Yann LeCun's AMI and Fei-Fei Li's World Labs. To get a sense of what Europe would require, consider the CLAIRE moonshot proposal cited earlier, since it includes generative AI systems for safe physical interaction with the real world. Applying the roughly 1:1 leverage from the Chips for Europe Initiative,[143] a ten-year embodied intelligence programme targeting total investment on the scale of the CLAIRE proposal (€100 billion) would require direct EU/member state public funding of around €50 billion once industrial co-development and off-take commitments from automotive and machinery firms are factored in.[144]

- **Fund trustworthy-AI and verification infrastructure.** The European Commission, ENISA, and cybersecurity institutions should invest in auditing tools, data-provenance systems, adversarial testing, secure deployment, model monitoring, and model-assurance methods. The aim is to make regulatory compliance technically manageable rather than leaving it as a legal and administrative burden for European companies, particularly for SMEs and startups. This responds to Section 3's finding that legal uncertainty and compliance costs can slow adoption and disadvantage smaller European providers relative to global competitors. The *Monopolkommission* likewise identifies legal uncertainty and regulatory burden as barriers to AI diffusion.[145]

- **Make openness and switching capacity mandatory in public AI procurement.** The Federal Ministry for Digital Transformation and the various procurement bodies should require open interfaces, interoperable standards, data portability, multi-cloud capability, and credible exit plans in major AI contracts. All of this would help to lower switching costs and enable a more dynamic domestic AI market that lends itself to repeated experimentation. Contract terms should avoid unjustified long durations or technical arrangements (including proprietary data formats) that make later provider substitution prohibitively costly.

- **Strengthen competition enforcement across the AI stack.** The European Commission and national competition authorities in EU member states should assess whether tying, bundling, self-preferencing, exclusive partnerships, and technical barriers to LLMs restrict competition across the AI stack. In particular, this should include scrutiny of how foundation models and developer tools are combined. As noted in Section 3, German application providers may face restricted routes to market even when they are technically competitive, because upstream platforms control key integration and user-access point. While classic competition case law via Art. 101 and 102 TFEU offers the most flexibility, the DMA certainly has a role to play here: for instance, the Commission recently found Google to be in breach of the DMA for favouring its own services in Google Search,[146] and imposed measures requiring Google to open Android features, such as voice invocation and access to on-device AI models, to rival AI assistants.[147] However, DMA enforcement also has its limits: as discussed in Section 3.3, Nvidia's CUDA advantage falls outside the scope of the DMA since Nvidia has no gatekeeper designation and chips are not core platform services. Similarly, many activities by leading AI labs, such as contributing to open standards and code, may be considered "competition on the merits" and thus not necessarily contradict EU competition rules. This is why competition enforcement must be paired with targeted industrial policy and procurement measures, as outlined above.

- **Accelerate capacity expansion, shorten grid connection timelines, and price flexibility for large compute loads.** The *Bundesnetzagentur* and the *Länder* should shorten grid-connection timelines for data-centre projects and establish demand-response contracts that reward operators for shifting non-time-critical workloads to periods of abundant renewable electricity. This would address the bottlenecks identified in the energy layer: structurally high industrial power prices, a volatile grid with frequent episodes of both negative prices and scarcity, and the difficulty of planning long-term power purchase agreements. As shown in Section 2.3, the cost of German industrial electricity at the relevant

consumption scale is roughly two and a half times that in Finland. Flexibility contracts help monetise surplus hours that currently produce negative prices, and they enable projects to connect without waiting for reinforcement. Still, these flexibility measures complement, rather than replace, continued capacity expansion, which remains crucial – in this regard, the ambitious goals for 2030 set out by the National Data Centre Strategy point in the right direction. Still, US operators increasingly compress build-out timelines by combining onsite generation with modular data-centre components, bringing large AI clusters online within months. The Strategy's focus on traditional grid-connected builds should thus be complemented with measures that reduce permitting and interconnection delays and make it easier to co-locate AI infrastructure with dedicated, low-carbon generation.

- **Fund efficient and decentralised AI systems research.** German and European research programmes should support work on distributed training, federated learning, specialised small models (so-called SLMs)[148], and "intelligence-per-watt"[149] architectures. These approaches aim to make AI systems more efficient so that useful capabilities can be delivered using fewer computing resources and less energy, while also enabling them to run on a wider range of hardware. For industrial and public-sector applications, this would not only reduce direct costs, but also the energy and infrastructure barriers that currently hinder large-scale AI deployment.

- **Prioritise AI adoption in management and workforce policies.** The federal government, chambers of commerce, industry associations, and social partners should treat AI adoption as a core management and skills issue. Concretely, they should provide practical AI training for managers, engineers, works councils, and operational staff, paying particular attention to SMEs. Support should extend beyond compliance checklists and vague "AI readiness" rhetoric to provide practical implementation guidance, sector-specific use cases, and examples of redesigning workflows and incentives, enabling firms to progress from small-scale pilots to stable deployment. This should include helping organisations to understand when AI adds value and where it does not, as well as how to integrate it with existing systems.

Taken together, the goal is not to create a sheltered national AI market, but rather to develop a German and European AI environment that can operate at the frontier where appropriate, while also fostering distinctive industrial capabilities in areas where there are comparative strengths.

5. Conclusion

Germany's position in the AI stack challenges the false dichotomy between achieving complete "technological sovereignty" and becoming irrelevant. Based on several sets of empirical, partly web-scraped data, the first part of this study revealed a significant asymmetry in the emerging AI stack: While Germany is heavily dependent on downstream layers such as advanced computing and frontier foundation models, it possesses credible strengths in industrial data, specialised applications, agentic tooling, and orchestration. However, the second part of this study argued that having capabilities in different layers of the AI stack does not, as such, produce economic or strategic value. Organisational inertia, skills shortages, legal uncertainty, fixed compliance burdens, and entrenched platform ecosystems can all prevent firms from adopting competitive German or European alternatives, even when they are available. Therefore, the near-term priority must be to encourage the diffusion of open AI stack components, not least by reducing switching costs. This also necessitates translating regulatory obligations from the AI Act and other relevant sets of rules into practical assurance tools, as well as a fine-grained mix of competition and industrial policy to prevent existing AI developer ecosystems from foreclosing competition in the upstream market.

The appropriate response is a dual-track strategy. Not unlike China, **Germany and Europe should aim to achieve a position near the forefront of AI technology by conducting continuous research, implementing efficient training systems, and sustaining long-term engineering capacity. At the same time, they should invest in specific areas such as industrial and embodied AI and other domains that leverage unique European data and industrial expertise.** This approach avoids retreating prematurely into narrow niche specialisation. Even if competing at the extreme frontier of model scaling proves too costly or comes too late, a sustained effort at that frontier is essential in order to preserve the expertise needed to understand, evaluate, and critically judge AI developments where necessary, as well as to accumulate the human capital required to shape how this technology is used at home. Otherwise, the risk that Karsten Wildberger warned of becomes acute: a country that does not build and understand its own AI systems ends up importing not only tools, but "a goal function, a worldview, and a value system" defined elsewhere.

Endnotes

1 Karsten Wildberger, *So wird's was mit dem Fortschritt*, FAZ, 14.03.2026.

2 ASML's dominance illustrates the latter, see: Jeremy Jurgens, *The Myth of AI Sovereignty*, Foreign Policy, 9.3.2026.

3 I make this argument at length in my recent book: Anselm Küsters, *Small is beautiful 2.0: Mit digitaler Dezentralisierung zu einer menschlicheren Wirtschaft*, 1. Auflage (Verlag Herder, 2026).

4 Distillation is a training technique in which a smaller "student" model learns to mimic the behaviour of a larger model.

5 Analyst Dan Wang describes this as a contrast between a (Chinese) "engineering state" and a (Western) "lawyerly society". Dan Wang, *Breakneck: China's Quest to Engineer the Future*, First edition (W. W. Norton & Company, Inc., 2025).

6 Jared Kaplan et al., 'Scaling Laws for Neural Language Models', version 1, preprint, arXiv, 2020, available at: https://doi.org/10.48550/ARXIV.2001.08361; Zhengyu Chen et al., 'Revisiting Scaling Laws for Language Models: The Role of Data Quality and Training Strategies', in *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ed. Wanxiang Che et al. (Association for Computational Linguistics, 2025), available at: https://doi.org/10.18653/v1/2025.acl-long.1163.

7 Tim Pearce et al., 'Scaling Laws for Pre-Training Agents and World Models', version 1, preprint, arXiv, 2024, available at: https://doi.org/10.48550/ARXIV.2411.04434.

8 Sara Hooker, 'On the Slow Death of Scaling', preprint, SSRN, 2025, available at: https://doi.org/10.2139/ssrn.5877662.

9 The costs of errors are asymmetric because losing the frontier race after diversifying is recoverable (one still has niche capabilities and can purchase foreign models), whereas failing to catch up and losing the niche opportunity leaves one with neither frontier technical competence nor differentiated alternatives.

10 In line with previous work: Anselm Küsters and Céline Nauer, *Strengthening the EU's Position in the AI Industry: The Importance of Software and Research*, Analysis (Friedrich Naumann Foundation for Freedom, Global Innovation Hub, 2024), available at: https://shop.freiheit.org/\#!/Publikation/1940.

11 See his presentation at the World Economic Forum in January 2026; and a later blog: *https://blogs.nvidia.com/blog/davos-wef-blackrock-ceo-larry-fink-jensen-huang/*; *https://blogs.nvidia.com/blog/ai-5-layer-cake/*.

12 On the Dwarkesh Podcast: *https://finance.yahoo.com/markets/stocks/articles/jensen-huang-says-nvidia-turns-142802231.html*.

13 Nils Goldschmidt and Michael Wohlgemuth, eds, *Grundtexte zur Freiburger Tradition der Ordnungsökonomik*, Untersuchungen zur Ordnungstheorie und Ordnungspolitik 50 (Mohr Siebeck, 2008); Thomas Biebricher et al., eds, *The Oxford Handbook of Ordoliberalism* (Oxford University Press, 2022).

14 Directive (EU) 2019/790 of the European Parliament and of the Council of 17 April 2019 on copyright and related rights in the Digital Single Market and amending Directives 96/9/EC and 2001/29/EC, Arts. 3-4, OJ L 130, 17.5.2019, pp. 92–125.

15 Here, I draw on information provided by my cep colleague Anja Hoffmann.

16 The General Data Protection Regulation (GDPR) was adopted by the European Parliament and the Council of the Euro-pean Union on 14 April 2016 and came into effect on 25 May 2018.

17 European Data Protection Board, 'Guidelines 3/2026 on Web Scraping in the Context of Generative AI', July 2026, available at: *https://www.edpb.europa.eu/system/files/2026-07/edpb_guidelines_2020603_webscraping_v1_en_0.pdf*.

18 European Data Protection Board, 'Guidelines 2/2026 on Anonymisation of Personal Data', July 2026, available at: *https://www.edpb.europa.eu/system/files/2026-07/edpb_guidelines_202602_anonymisation_v1_en_0.pdf*.

19 European Commission. Directorate-General for Communications Networks, Content and Technology et al., *Study on the next Data Frontier* (Publications Office of the European Union, 2025), available at: *https://doi.org/10.2759/3847027*. For a compari-son between EU, US and China, see also: Anu Bradford, *Digital Empires: The Global Battle to Regulate Technology* (Oxford University Press, 2023).

20 IDC and Lisbon Council and CARSA, *D2.1 First Report on Facts and Figures: European Data Market Study 2024-2026* (European Commission, Directorate-General for Communications Networks, Content and Technology, 2025), 115, available at: *https://ec.europa.eu/newsroom/repository/document/2025-13/EDM_2024__2026__First_Report_on_Facts_and_Figures_9cU0iIjcuSZhjmwI9TchwqenidQ_114043.pdf*.

21 All Bundestag plenary protocols from January 2018 to July 2026 were retrieved in full text via the DIP API, restricted to Bundestag sittings. As the interface offers no full-text query parameter, all filtering was performed locally on the retrieved text. The series begins in 2018, as earlier years contain almost no references. The anchors for the output layer are Künstliche Intelligenz and the abbreviation KI, with the latter excluding the name Ban Ki-moon and hyphenated line breaks in the source PDFs. The anchors for the input layer are Datenschutz, DSGVO, BDSG, Urheberrecht, Text- und Data-Mining, personenbezogene Daten, Einwilligung and Nutzungsrecht. Context windows extend 600 characters on either side of each anchor, and overlapping windows are merged so that shared text is not counted twice. Framing density is reported per 1,000 words of context. For the input layer, the corpus is further restricted to protocols that also mention AI. This ensures that the measure captures how data law is discussed within the AI debate, rather than data-law discourse in general.

22 Christian Rauh, 'Validating a Sentiment Dictionary for German Political Language—a Workbench Note', *Journal of Information Technology & Politics* 15, no. 4 (2018): 319–43, available at: *https://doi.org/10.1080/19331681.2018.1485608*.

23 Robert Remus et al., 'SentiWS - A Publicly Available German-Language Resource for Sentiment Analysis', in *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, ed. Nicoletta Calzolari (European Language Resources Association, 2010).

24 For the flight towards external AI providers, see my criticism in: Anselm Küsters and Ben Waber, 'How SMEs Can Prepare for the EU's AI Regulations', Harvard Business Review, 24 September 2025, available at: *https://hbr.org/2025/09/how-smes-can-prepare-for-the-eus-ai-regulations*.

25 Jaime Sevilla and Anton Troynikov, *Could Decentralized Training Solve AI's Power Problem?* (Epoch AI, 2025), available at: *https://epoch.ai/publications/could-decentralized-training-solve-ais-power-problem*.

26 Eurostat, Electricity prices for non-household consumers (total price excluding VAT), 2025, available at: *https://ec.europa.eu/eurostat/databrowser/view/nrg_pc_205*. Bands below 20,000 MWh per year are excluded as they incur taxes and surcharges that facilities of this size would not face. Finland has no observation in the largest band for 2025.

27 According to the data, the network component is 3.00 of 12.05 ct/kWh in the largest band, so even maximum relief removes about 2.4 to 2.7 ct.

28 Germany Insolvencies, 'Germany's Industrial Electricity Prices: Double the US', May 2026, available at: *https://germanyinsolvencies.com/blog/germany-industrial-electricity-prices/*.

29 The calculation is based on the following assumptions (a detailed calculation sheet is available from the author upon request): The all-in power per accelerator of 1.37 kW is taken from SemiAnalysis's InferenceX comparison of 22 August 2026. The inference throughput of 31,306.8 tokens per second is the audited MLPerf Inference v5.0 result for an eight-GPU NVIDIA DGX H100 on Llama 2 70B. The fine-tuning energy of 3.27 kWh is measured in MLPerf Training v4.0, and the frontier training figure is derived from Meta's disclosed 30.84 million GPU-hours for Llama 3.1 405B. Inference and training are measured on separate benchmarks and are not interchangeable quantities. As electricity price is the only variable between jurisdictions, the ratio between any two is identical for every workload.

30 Inference and training are computed are not interchangeable: training runs a forward and a backward pass and is compute-bound, whereas inference runs only a forward pass and is bound at the decode stage by memory bandwidth.

31 Bundesnetzagentur, 'Bundesnetzagentur veröffentlicht Daten zum Strommarkt: Negative Preise und Preisspitzen 2025', 2026, available at: *https://www.bundesnetzagentur.de/1087156*.

32 Bundesnetzagentur, 'Growth in Renewable Energy in 2024', January 2025, available at: *https://www.bundesnetzagentur.de/SharedDocs/Pressemitteilungen/EN/2025/20250108_EE.html*.

33 Enerdata, 'China Installs Record Capacity for Solar (+45%) and Wind (+18%) in 2024', February 2025, available at: *https://www.enerdata.net/publications/daily-energy-news/china-installs-record-capacity-solar-45-and-wind-18-2024.html*.

34 Electric Perspectives / Edison Electric Institute, 'EEI Report: Annual Electricity Generation Jumps 3 Percent', September 2025, available at: *https://www.electricperspectives.com/eei-statistical-yearbook-electricity-generation/*.

35 ASCII UK, 'Grid Interconnection Delays Bottleneck AI Data Center Buildout', July 2026, available at: *https://ascii.co.uk/news/article/news-20260710-25ce5e50/grid-interconnection-delays-bottleneck-ai-data-center-buildo*.

36 Lowdown, 'xAI Wins 41 Gas Turbines for Colossus', April 2026, available at: *https://lowdown.today/t/data-centres/1/xai-wins-41-gas-turbines-for-colossus/*.

37 Ajey Pandey et al., *How AI Labs Are Solving the Power Crisis: The Onsite Gas Deep Dive* (SemiAnalysis, 2025, available at: *https://newsletter.semianalysis.com/p/how-ai-labs-are-solving-the-power*.

38 SemiAnalysis, 'The Wild Wild West of LEGO Datacenters', 30 July 2026, available at: *https://newsletter.semianalysis.com/p/the-wild-wild-west-of-lego-data-centers*.

39 Jaime Sevilla and Anton Troynikov, *Could Decentralized Training Solve AI's Power Problem?* (Epoch AI, 2025), available at: *https://epoch.ai/publications/could-decentralized-training-solve-ais-power-problem*.

40 Bundesministerium für Digitales und Staatsmodernisierung (BMDS), *Nationale Rechenzentrumsstrategie* (Bundeskabinett, 2026), available at: *https://bmds.bund.de/fileadmin/BMDS/Dokumente/Publikationen/260312_BMDS_Template-RZ-Strategie_V2_barrierefrei_final.pdf*.

41 SemiAnalysis, 'The Wild Wild West of LEGO Datacenters'.

42 For an easily accessible account, see: Chris Miller, *Chip War: The Fight for the World's Most Critical Technology* (Simon & Schuster Ltd, 2023).

43 Tomoatsu Shibata et al., 'The Co-Evolution of Exploitation and Exploration toward Industrialising Frontier Technology: Insights from ASML's Pursuit of EUV Lithography', *Technology Analysis & Strategic Management* 38, no. 6 (2026): 832–46, available at: *https://doi.org/10.1080/09537325.2025.2501740*; Susanne Van Der Velden et al., 'Incumbent Success in the Era of Ferment: Navigation of Intergenerational Transition of Technology', *Academy of Management Proceedings* 2020, no. 1 (2020): 20381, available at: *https://doi.org/10.5465/AMBPP.2020.20381abstract*.

44 See: *https://www.theinformation.com/articles/china-starts-mass-producing-homegrown-duv-chipmaking-tools-advance-local-chip-industry*.

45 E. G. Rasmussen and others, *Report from the Extreme Ultraviolet (EUV) Lithography Workshop*, NIST SP 1500-208 (National Institute of Standards and Technology (NIST), 2023), *https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-208.pdf*.

46 Toby Sterling and Nathan Vifflin, 'Explainer: The $250 million ASML "printer" behind Nvidia's chips', *Reuters*, 28 January 2026, *https://www.reuters.com/world/asia-pacific/250-million-asml-printer-behind-nvidias-chips-2026-07-28/*.

47 Bloomberg, 'Germany Maps China's Weaknesses in Preparation for a Trade War', 28 July 2026, available at: https://www.bloomberg.com/news/articles/2026-07-28/germany-maps-china-s-weaknesses-in-preparation-for-a-trade-war.

48 NVIDIA Corporation, 'NVIDIA H100 Tensor Core GPU', 2026, available at: https://www.nvidia.com/en-us/data-center/h100/.

49 Note, the two series are not directly comparable. See: https://www.jp-globaladvisors.com/articles/saudi-arabia-and-uae-the-ai-superpowers-you-cant-ignore.

50 Irene Zhang and Nick Corvino, '"Eastern Data, Western Compute" Is Fake', *ChinaTalk*, 10 July 2026, available at: https://www.chinatalk.media/p/eastern-data-western-compute-is-fake.

51 See the official document here: https://www.gov.cn/yaowen/liebiao/202608/content_7077762.htm (the translation draws on Emmie Hine's China Bulletin, issue 9, available at: https://chinaaibulletin.substack.com/p/china-ai-bulletin-9).

52 Ryan Fedasiuk and Satvik Pendyala, *Voltcraft: Industrial Competition in the Age of AI*, Working Paper (American Enterprise Institute, 2026), available at: https://www.aei.org/research-products/working-paper/voltcraft-industrial-competition-in-the-age-of-ai-aei/.

53 European Parliament and the Council of the European Union, 'Regulation (EU) 2023/1781 of the European Parliament and of the Council of 13 September 2023 Establishing a Framework of Measures for Strengthening Europe's Semiconductor Ecosystem and Amending Regulation (EU) 2021/694 (Chips Act)', September 2023, available at: https://eur-lex.europa.eu/eli/reg/2023/1781/oj

54 European Commission, 'Proposal for the Chips Act 2.0', June 2026, available at: https://digital-strategy.ec.europa.eu/en/library/proposal-chips-act-20.

55 European Parliament, *Chips Act 2.0*, Briefing (European Parliamentary Research Service, 2026), available at: https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2026)785742.

56 European Commission, 'EU Launches AI Gigafactories Call to Boost Europe's Computing Capacity and Unlock More than €30 Billion in Investment', 30 July 2026, available at: https://ec.europa.eu/commission/presscorner/detail/en/ip_26_1708.

57 Corinna Visser, *AI Gigafactories: Wie sieben Fabriken helfen sollen, Europas KI-Rückstand aufzuholen* (Centrum für Europäische Politik, 2026).

58 Joana Soares, 'The EU's AI Boom Could Undermine Its Own Chip Strategy', 2026, available at: *https://www.techpolicy.press/the-eus-ai-boom-could-undermine-its-own-chip-strategy/*.

59 Arthur Douillard et al., 'DiLoCo: Distributed Low-Communication Training of Language Models', version 3, preprint, arXiv, 2023, available at: *https://doi.org/10.48550/ARXIV.2311.08105*; Sami Jaghouar et al., 'OpenDiLoCo: An Open-Source Framework for Globally Distributed Low-Communication Training', version 1, preprint, arXiv, 2024, available at: *https://doi.org/10.48550/ARXIV.2407.07852*; Arthur Douillard et al., 'Decoupled DiLoCo for Resilient Distributed Pre-Training', version 1, preprint, arXiv, 2026, available at: *https://doi.org/10.48550/ARXIV.2604.21428*.

60 Sevilla and Troynikov, *Could Decentralized Training Solve AI's Power Problem?* (Epoch AI, 2025).

61 Catherine Tucker, 'Network Effects and Market Power: What Have We Learned in the Last Decade?', Antitrust 32, no. 2 (2018): 72–79; Maria Ioannidou and Bingwan Xiong, *Competition Law and Policy in Digital Markets : A Comparative Analysis of the EU and China* (Oxford University Press, 2025); Philipp Staab, *Markets and power in digital capitalism*, First English-language edition (Manchester University Press, 2024); Anselm Küsters, 'EU Competition Law in the Digital Age: Substantive Challenges and Methodological Innovations Using the Example of the European Commission's Consultation Procedure', in *Doing Justice. Procedural Law and Procedural Law Reform in Comparative Perspective*, ed. Henning Glaser, CPG Series of Comparative Constitutional Law, Politics and Governance 7 (Nomos, 2026), available at: *https://doi.org/10.5771/9783845254098-251*.

62 Synergy Research Group, 'European Cloud Providers' Local Market Share Now Holds Steady at 15%', 2025, available at: *https://www.srgresearch.com/articles/european-cloud-providers-local-market-share-now-holds-steady-at-15*.

63 SemiAnalysis, 'The Wild Wild West of LEGO Datacenters'.

64 Synergy Research Group, 'Cloud Market Share Trends: Big Three Together Hold 63% While Oracle and the Neoclouds Inch Higher', 19 November 2025, available at: *https://www.srgresearch.com/articles/cloud-market-share-trends-big-three-together-hold-63-while-oracle-and-the-neoclouds-inch-higher*.

65 See: *https://commission.europa.eu/news-and-media/news/commission-advances-cloud-sovereignty-through-strategic-procurement-2026-04-17_en*.

66 Synergy Research Group, 'European Cloud Providers' Local Market Share Now Holds Steady at 15%'.

67 Filippo Gualtiero Blancato, 'The Cloud Sovereignty Nexus: How the European Union Seeks to Reverse Strategic Dependencies in Its Digital Ecosystem', *Policy & Internet* 16, no. 1 (2024): 12–32, available at: *https://doi.org/10.1002/poi3.358*.

68 Julia Rone, '"The Sovereign Cloud" in Europe: Diverging Nation State Preferences and Disputed Institutional Competences in the Context of Limited Technological Capabilities', *Journal of European Public Policy* 31, no. 8 (2024): 2343–69, available at: *https://doi.org/10.1080/13501763.2024.2348618*.

69 Vaida Gineikyte-Kanclere et al., *European Software and Cyber Dependencies*, PE 778.576 (Policy Department for Transformation, Innovation and Health, European Parliament, 2025), available at: *https://www.europarl.europa.eu/RegData/etudes/STUD/2025/778576/ECTI_STU(2025)778576_EN.pdf*.

70 SemiAnalysis, 'RL Systems: Mind the Gap — Matching Human and Model Capability', SemiAnalysis, 2025, *https://newsletter.semianalysis.com/p/rl-systems-mind-the-gap-matching*.

71 Epoch AI, 'An FAQ on Reinforcement Learning', Epoch AI, 2026, available at: *https://epochai.substack.com/p/an-faq-on-reinforcement-learning*.

72 SemiAnalysis, 'RL Environments and RL for Science: Data Foundries and Multi-Agent Systems', January 2026, available at: *https://newsletter.semianalysis.com/p/rl-environments-and-rl-for-science*.

73 European Commission, 'Common European Data Spaces', European Commission, 2026, available at: *https://digital-strategy.ec.europa.eu/en/policies/data-spaces*.

74 Data Spaces Support Centre, 'Scaling Trusted Data Sharing across Europe', Data Spaces Support Centre, 2026, available at: *https://dssc.eu*.

75 European Commission, 'Common European Data Spaces'.

76 See: *https://manufacturingdata-space-csa.eu*.

77 See: *https://epoch.ai/data/ai-models-documentation/records*.

78 Current AI, 'Open-Source AI Map', 2026, available at: *https://github.com/currentai-org/os-ai-map*.

79 Using manual process, supplemented where possible by matches to the European Tech Map company registry.

80 Expert commentary on leading systems, including GPT, Gemini, Llama and Mistral, stresses that "open" and "closed" are not binary labels. Access to weights, training data, code, evaluations and commercial use rights varies, and these differences determine who can realistically build on a model. Nathan Lambert, 'Reading Today's Open and Closed Performance', Interconnects, 2026, available at: *https://www.interconnects.ai/p/reading-todays-open-closed-performance*.

81 James Evans et al., 'Agentic AI and the next Intelligence Explosion', version 1, preprint, arXiv, 2026, *https://doi.org/10.48550/ARXIV.2603.20639*.

82 See: *https://www.cep.eu/de/eu-themen/details/openai-modell-greift-hugging-face-an-ein-reward-hacking-vorfall-der-alignment-theorie-zur-praxis-macht.html*.

83 Market-style allocation mechanisms, for example, in which agents bid for tasks based on their own assessment of their capabilities, can outperform hierarchical orchestration by frontier models, particularly for brittle reasoning problems that benefit from multiple independent attempts and diverse strategies. Rohit Krishnan, 'Why Coase Needs Hayek: Sometimes Smart Planners Lose to Simple Markets', Strange Loop Canon, 2026, available at: *https://www.strangeloopcanon.com/p/why-smart-planners-lose-to-simple*.

84 Rohit Krishnan and Andrey Fradkin, 'Agent, Know Thyself! (And Bid Accordingly)', Strange Loop Canon, 2026, available at: *https://www.strangeloopcanon.com/p/agent-know-thyself-and-bid-accordingly*.

85 For this proposal, see: *https://ai-frontiers.org/articles/we-need-better-infrastructure-to-govern-ai-agents*.

86 Again, all listed organisations were assigned to countries, supplemented by matching against European Tech Map records.

87 Emerging US proposals would require major developers of frontier AI technologies to retain licensed independent verification organisations. Gabriel Weil, 'Don't Let AI Developers Hire Their Own Referees', AI Frontiers, 29 July 2026, *https://ai-frontiers.org/articles/dont-let-ai-developers-hire-their-own-referees*. Policy debates, including in Europe, are also considering liability and insurance-based approaches, which would make evaluation and containment commercially valuable.

88 Anu Bradford, *The Brussels Effect: How the European Union Rules the World* (Oxford University Press, 2020).

89 For the AI Act and its relationship to the Brussels Effect, see: Marco Almada and Anca Radu, 'The Brussels Side-Effect: How the AI Act Can Reduce the Global Reach of EU Policy', *German*

Law Journal 25, no. 4 (2024): 646–63, available at: https://doi.org/10.1017/glj.2023.108; Ben Crum, 'Brussels Effect or Experimentalism? The EU AI Act and Global Standard-Setting', *Internet Policy Review* 14, no. 3 (2025), available at: https://doi.org/10.14763/2025.3.2032.

90 See: https://digitalpolicyalert.org/activity-tracker.

91 The DPA's jurisdiction field records the jurisdictions to which a measure applies, rather than the body that enacted it. Of the 1,946 policy changes targeting ML and AI development, the tracker records 9,625 jurisdiction entries, approximately 4.95 per policy change, and the field extends to dependent territories mapped to their metropole. European instruments are consequently attributed to all 27 member states, creating a common baseline for each of them. This study esti-mates this minimum value across the EU-27 to be 182 measures (Malta and Romania), with a median of 191, and twenty of the twenty-seven member states lying within ten per cent of the minimum. National increments are calculated as the raw count minus this floor for member states and as the raw count for jurisdictions outside the Union. This estimate is corroborated by Switzerland, which is covered by the tracker for the G20, the EU-27 and Switzerland, but is not covered by the acquis. Switzerland records 37 measures. Note, the procedure assumes that the floor is entirely acquis-derived and that non-EU counts are purely national.

92 The confidence intervals are narrow and the 2018 and 2021 intervals do not overlap.

93 For this, see: Küsters and Waber, 'How SMEs Can Prepare for the EU's AI Regulations'.

94 Emerging Technology Observatory, 'Agora State AI Governance Tracker', Center for Security and Emerging Technology, 2026, available at: https://eto.tech/blog/agora-state-ai-governance-tracker-1/; Gary Marcus, 'US AI Policy Is a Clumsy Mess. Here's What to Do about It', Gary Marcus Substack, 2026, available at: https://garymarcus.substack.com/p/us-ai-policy-is-a-clumsy-mess-heres; Ashley Taylor et al., 'State AGs Fill the AI Regulatory Void', Reuters, 19 May 2025, available at: https://www.reuters.com/legal/legalindustry/state-ags-fill-ai-regulatory-void-2025-05-19/.

95 See, e.g.: Rogier Creemers et al., 'Translation: Measures for the Management of Generative Artificial Intelligence Services (Draft for Comment) – April 2023', *DigiChina*, 11 April 2023, available at: https://digichina.stanford.edu/work/translation-interim-measures-for-the-management-of-generative-artificial-intelligence-services/.

96 Wang, *Breakneck*.

97 Angela Huyue Zhang, *Chinese Antitrust Exceptionalism: How the Rise of China Challenges Global Regulation* (Oxford University Press, 2021).

98 Anselm Küsters, 'Ordoliberalism Goes China? A Comparison of Recent Developments in EU and Chinese Competition Law Considering the Digital Economy', *Constitutional Political Economy* 36, no. 4 (2025): 444–68, available at: https://doi.org/10.1007/s10602-023-09407-y; Ioannidou and Xiong, Competition Law and Policy in Digital Markets.

99 Annalisa Ferrando et al., *Adoption and Investment in AI Across the Euro Area: Insights from Harmonised Firm-Level Data*, Occasional Paper Series no. 395 (European Central Bank, 2026), available at: https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op395.en.pdf.

100 While 47% of German firms reported moderate or significant AI use (equal to Austria and above the euro-area average of 38%), the Netherlands and Belgium lead on significant use alone, at 23% and 18% respectively.

101 The Monopolkommission reaches a closely aligned diagnosis, which is used throughout this section: Monopolkommis-sion, *Wettbewerb 2026: Für eine wettbewerbsorientierte Wirtschaftspolitik* (Monopolkommission, 2026), available at: https://www.monopolkommission.de/publikationen/hauptgutachten/wettbewerb-2026-fuer-eine-wettbewerbsorientierte-wirtschaftspolitik.

102 Greta L. Polites and Elena Karahanna, 'Shackled to the Status Quo: The Inhibiting Effects of Incumbent System Habit, Switching Costs, and Inertia on New System Acceptance1', *MIS Quarterly* 36, no. 1 (2012): 21–42, available at: https://doi.org/10.2307/41410404.

103 Jung-Kuei Hsieh et al., 'Post-Adoption Switching Behavior for Online Service Substitutes: A Perspective of the Push-Pull-Mooring Framework', *Computers in Human Behavior* 28, no. 5 (2012): 1912–20, available at: https://doi.org/10.1016/j.chb.2012.05.010.

104 For the different components of switching costs, see: Soumya Ray et al., 'Online Users' Switching Costs: Their Nature and Formation', *Information Systems Research* 23, no. 1 (2012): 197–213, available at: https://doi.org/10.1287/isre.1100.0340.

105 For the links between organisational components, procedures, and embedded knowledge in general, see: Rebecca M. Henderson and Kim B. Clark, 'Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms', *Administrative Science Quarterly* 35, no. 1 (1990): 9, available at: https://doi.org/10.2307/2393549.

106 Monopolkommission, *Wettbewerb 2026: Für eine wettbewerbsorientierte Wirtschaftspolitik*, 380.

107 Monopolkommission, *Wettbewerb 2026: Für eine wettbewerbsorientierte Wirtschaftspolitik*, 381–82.

108 Heidrun C. Hoppe, 'The Timing of New Technology Adoption: Theoretical Models and Empirical Evidence', *The Manchester School* 70, no. 1 (2002): 56–76, available at: *https://doi.org/10.1111/1467-9957.00283*.

109 Monopolkommission, *Wettbewerb 2026: Für eine wettbewerbsorientierte Wirtschaftspolitik*, 387.

110 See the remarks on ordoliberalism in Section 2.1 above.

111 Monopolkommission, *Wettbewerb 2026: Für eine wettbewerbsorientierte Wirtschaftspolitik*, 391.

112 Monopolkommission, *Wettbewerb 2026: Für Eine Wettbewerbsorientierte Wirtschaftspolitik*, 391.

113 Deirdre Ahern, 'Operationalising AI Regulatory Sandboxes under the EU AI Act: The Triple Challenge of Capacity, Coordination and Attractiveness to Providers', *Cambridge Forum on AI: Law and Governance* 1 (2025): e35, available at: *https://doi.org/10.1017/cfl.2025.10023*; Thomas Buocz et al., 'Regulatory Sandboxes in the AI Act: Reconciling Innovation and Safety?', *Law, Innovation and Technology* 15, no. 2 (2023): 357–89, available at: *https://doi.org/10.1080/17579961.2023.2245678*.

114 See: *https://www.cep.eu/eu-topics/details/european-ai-sovereignty-instead-of-chinese-models.html*.

115 Monopolkommission, *Wettbewerb 2026: Für eine wettbewerbsorientierte Wirtschaftspolitik*, 368–69.

116 In line with the literature on switching costs: Polites and Karahanna, 'Shackled to the Status Quo'.

117 SemiAnalysis, 'Can AMD Break the CUDA Moat? AMD Advancing AI 2026', SemiAnalysis, 25 July 2026, available at: *https://newsletter.semianalysis.com/p/can-amd-break-the-cuda-moat-amd-advancing*.

118 Dana Petcu, 'Portability and Interoperability between Clouds: Challenges and Case Study', *Proceedings of the 4th European Conference on Towards a Service-Based Internet* (Berlin, Heidelberg), ServiceWave'11, 2011, 62–74; Louise Amoore, *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others* (Duke University Press, 2020).

119 Tomaso Duso and Martin Peitz, 'Wettbewerbspolitik und Industriepolitik unter einem Hut', *Perspektiven der Wirtschaftspolitik* 26, no. 3 (2025): 323–43, available at: *https://doi.org/10.1515/pwp-2025-0016*.

120 The term "acqui-hire" refers to an acquisition where the buyer's main objective is to hire the target company's team rather than to obtain its products or existing customers.

121 Monopolkommission, *Wettbewerb 2026: Für eine Wettbewerbsorientierte Wirtschaftspolitik,* 397–401.

122 See: *https://newsletter.semianalysis.com/p/agentx-inferencexv3-does-cuda-moat*.

123 Notably vLLM, SGLang, KV-cache and routing layers.

124 Anselm Küsters et al., 'Towards a Rule-Bound EU Transformation Policy: Can Ordnungspolitik and Industrial Strategy Be Reconciled?', *ORDO. Jahrbuch für die Ordnung von Wirtschaft und Gesellschaft* 76, no. 1 (2026): 1–30, available at: *https://doi.org/10.1515/ordo-2026-2012*.

125 This paragraph draws on the translated sources described in Emmie Hine's China Bulletin, issue 9, available at: *https://chinaaibulletin.substack.com/p/china-ai-bulletin-9*.

126 Francisco Javier Varela Sandoval et al., *How Middle Powers Can Weather US and Chinese AI Dominance : The Case for 'Sovereign AI' Strategies* (Royal Institute of International Affairs, 2026), available at: *https://doi.org/10.55317/9781784136710*.

127 Here, I follow a recent expert report on frontier AI that distinguishes between "placing bets within the current paradigm" and backing alternative approaches. See: European Commission, *Key Insights from the First Meeting of the European Expert Forum on Frontier AI* (European Commission, 2026), 18–21, available at: *https://digital-strategy.ec.europa.eu/en/library/ai-office-publishes-frontier-ai-expert-findings-eu-competitiveness-sovereignty-and-security*.

128 OECD, *OECD Artificial Intelligence Review of Germany* (OECD Publishing, 2024), available at: *https://doi.org/10.1787/609808d6-en*.

129 *The Design and Implementation of Mission-Oriented Innovation Policies: A New Systemic Policy Approach to Address Societal Challenges*, OECD Science, Technology and Industry Policy Papers no. 100, vol. 100, OECD Science, Technology and Industry Policy Papers (2021), available at: *https://doi.org/10.1787/3f6c76a4-en*.

130 Elizabeth Gibney, '"Another DeepSeek Moment": Chinese AI Model Kimi K2 Stirs Excitement', *Nature*, ahead of print, 16 July 2025, available at: *https://doi.org/10.1038/d41586-025-02275-6*. See also: *The Design and Implementation of Mission-Oriented Innovation Policies*, vol. 100.

131 Wang, *Breakneck*.

132 See: *https://www.interconnects.ai/p/kimi-k3-the-open-weights-escalation*.

133 Yann LeCun, 'A Path Towards Autonomous Machine Intelligence', *OpenReview*, 27 June 2022, available at: *https://openreview.net/forum?id=BZ5a1r-kVsf*.

134 Anthropic, 'Project Fetch: Phase Two', 18 June 2026, available at: https://www.anthropic.com/research/project-fetch-phase-two.

135 Richard Sutton, 'The Bitter Lesson', 13 March 2019, available at: http://www.incompleteideas.net/IncIdeas/BitterLesson.html.

136 Cristina Caffarra, 'The Sovereignty Narratives That Don't Help Europe', 2026, available at: https://www.techpolicy.press/the-sovereignty-narratives-that-dont-help-europe/.

137 Tamara Carleton and William R. Cockayne, 'Inspired by DARPA: A Global Comparison of Radical Innovation Government Agencies', in *Inventing the Almost Impossible*, ed. Tamara Carleton et al., Future of Business and Finance (Springer Nature Switzerland, 2023), available at: https://doi.org/10.1007/978-3-031-36224-8_3.

138 See the "Moonshot in Artificial Intelligence" proposal, available at: https://cairne.eu/wp-content/uploads/2023/11/Moonshot-proposal.pdf.

139 SemiAnalysis, Korea's Trillion-Dollar Sovereign AI Investment, 1.9.2026, available at: https://newsletter.semianalysis.com/p/koreas-trillion-dollar-sovereign.

140 Research by Interface suggests that Germany's AI talent broadly matches the competitive niches identified in the study, particularly industrial AI: https://www.interface-eu.org/publications/germany-case-study.

141 See: Tagesspiegel, Open Source strategisch für den Staat gestalten, 2.9.2026, available at: https://background.tagesspiegel.de/digitalisierung-und-ki/briefing/open-source-strategisch-fuer-den-staat-gestalten.

142 See: https://www.forbes.com/sites/josipamajic/2026/06/30/world-model-startups-raise-3-billion-vcs-bet-beyond-llms/.

143 See: European Chips Act (2022), Staff Working document, available at: https://digital-strategy.ec.europa.eu/en/library/european-chips-act-staff-working-document.

144 Note that the €100 billion figure was presented as public funding alone, without any assumption of leverage. Here, however, it is treated it as a total investment benchmark, using the Chips Act's 1:1 ratio to calculate the implied public share.

145 Monopolkommission, *Wettbewerb 2026: Für eine Wettbewerbsorientierte Wirtschaftspolitik*.

146 See: https://digital-markets-act.ec.europa.eu/commission-fines-google-eur890-million-breaches-digital-markets-act-2026-07-23_en.

147 See: https://ec.europa.eu/commission/presscorner/detail/en/ip_26_1634.

148 Fali Wang et al., 'A Comprehensive Survey of Small Language Models in the Era of Large Language Models: Techniques, Enhancements, Applications, Collaboration with LLMs, and Trustworthiness', version 1, preprint, arXiv, 2024, available at: *https://doi.org/10.48550/ARXIV.2411.03350*.

149 Jon Saad-Falcon et al., 'Intelligence per Watt: Measuring Intelligence Efficiency of Local AI', version 2, preprint, arXiv, 2025, available at: *https://doi.org/10.48550/ARXIV.2511.07885*.

6. Literature list

Ahern, Deirdre. 'Operationalising AI Regulatory Sandboxes under the EU AI Act: The Triple Challenge of Capacity, Coordination and Attractiveness to Providers'. Cambridge Forum on AI: Law and Governance 1 (2025): e35. Available at: *https://doi.org/10.1017/cfl.2025.10023*.

Almada, Marco, and Anca Radu. 'The Brussels Side-Effect: How the AI Act Can Reduce the Global Reach of EU Policy'. German Law Journal 25, no. 4 (2024): 646–63. Available at: *https://doi.org/10.1017/glj.2023.108*.

Amoore, Louise. Cloud Ethics: Algorithms and the Attributes of Ourselves and Others. Duke Univer-sity Press, 2020.

Anthropic. 'Project Fetch: Phase Two'. 18 June 2026. Available at: *https://www.anthropic.com/research/project-fetch-phase-two*.

ASCII UK. 'Grid Interconnection Delays Bottleneck AI Data Center Buildout'. July 2026. Available at: *https://ascii.co.uk/news/article/news-20260710-25ce5e50/grid-interconnection-delays-bottleneck-ai-data-center-buildo*.

Biebricher, Thomas, Werner Bonefeld, and Peter Nedergaard, eds. The Oxford Handbook of Ordolib-eralism. Oxford University Press, 2022.

Blancato, Filippo Gualtiero. 'The Cloud Sovereignty Nexus: How the European Union Seeks to Reverse Strategic Dependencies in Its Digital Ecosystem'. Policy & Internet 16, no. 1 (2024): 12–32. Available at: *https://doi.org/10.1002/poi3.358*.

Bradford, Anu. Digital Empires: The Global Battle to Regulate Technology. Oxford University Press, 2023.

Bradford, Anu. The Brussels Effect: How the European Union Rules the World. Oxford University Press, 2020.

Bundesministerium für Digitales und Staatsmodernisierung (BMDS). Nationale Rechenzentrumsstrategie. Bundeskabinett, 2026. Available at: *https://bmds.bund.de/fileadmin/BMDS/Dokumente/Publikationen/260312_BMDS_Template-RZ-Strategie_V2_barrierefrei_final.pdf*.

Bundesnetzagentur. 'Bundesnetzagentur veröffentlicht Daten zum Strommarkt: negative Preise und Preisspitzen 2025'. 2026. Available at: *https://www.bundesnetzagentur.de/1087156*.

Bundesnetzagentur. 'Growth in Renewable Energy in 2024'. January 2025. Available at: *https://www.bundesnetzagentur.de/SharedDocs/Pressemitteilungen/EN/2025/20250108_EE.html*.

Buocz, Thomas, Sebastian Pfotenhauer, and Iris Eisenberger. 'Regulatory Sandboxes in the AI Act: Reconciling Innovation and Safety?' Law, Innovation and Technology 15, no. 2 (2023): 357–89. Available at: *https://doi.org/10.1080/17579961.2023.2245678*.

Caffarra, Cristina. 'The Sovereignty Narratives That Don't Help Europe'. 2026. Available at: *https://www.techpolicy.press/the-sovereignty-narratives-that-dont-help-europe/*.

Carleton, Tamara, and William R. Cockayne. 'Inspired by DARPA: A Global Comparison of Radical Innovation Government Agencies'. In Inventing the Almost Impossible, edited by Tamara Carleton, Shaun West, and William R. Cockayne. Future of Business and Finance. Springer Nature Switzerland, 2023. Available at: *https://doi.org/10.1007/978-3-031-36224-8_3*.

Chen, Zhengyu, Siqi Wang, Teng Xiao, et al. 'Revisiting Scaling Laws for Language Models: The Role of Data Quality and Training Strategies'. In Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), edited by Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar. Association for Computational Linguistics, 2025. Available at: *https://doi.org/10.18653/v1/2025.acl-long.1163*.

Creemers, Rogier, Graham Webster, Zac Haluza, Helen Toner, and Seaton Huang. 'Translation: Measures for the Management of Generative Artificial Intelligence Services (Draft for Comment) – April 2023'. DigiChina, 11 April 2023. Available at: *https://digichina.stanford.edu/work/translation-interim-measures-for-the-management-of-generative-artificial-intelligence-services/*.

Crum, Ben. 'Brussels Effect or Experimentalism? The EU AI Act and Global Standard-Setting'. Internet Policy Review 14, no. 3 (2025). Available at: *https://doi.org/10.14763/2025.3.2032*.

Current AI. 'Open-Source AI Map'. 2026. Available at: *https://github.com/current-tai-org/os-ai-map*.

Data Spaces Support Centre. 'Scaling Trusted Data Sharing across Europe'. Data Spaces Support Centre, 2026. Available at: *https://dssc.eu*.

Douillard, Arthur, Qixuan Feng, Andrei A. Rusu, et al. 'DiLoCo: Distributed Low-Communication Training of Language Models'. Version 3. Preprint, arXiv, 2023. Available at: *https://doi.org/10.48550/ARXIV.2311.08105*.

Douillard, Arthur, Keith Rush, Yani Donchev, et al. 'Decoupled DiLoCo for Resilient Distributed Pre-Training'. Version 1. Preprint, arXiv, 2026. Available at: *https://doi.org/10.48550/ARXIV.2604.21428*.

Duso, Tomaso, and Martin Peitz. 'Wettbewerbspolitik und Industriepolitik unter einem Hut'. Perspektiven der Wirtschaftspolitik 26, no. 3 (2025): 323–43. Available at: *https://doi.org/10.1515/pwp-2025-001*6.

Electric Perspectives / Edison Electric Institute. 'EEI Report: Annual Electricity Generation Jumps 3 Percent'. September 2025. Available at: *https://www.electricperspectives.com/eei-statistical-yearbook-electricity-generation/*.

Emerging Technology Observatory. 'Agora State AI Governance Tracker'. Center for Security and Emerging Technology, 2026. Available at: *https://eto.tech/blog/agora-state-ai-governance-tracker-1/*.

Enerdata. 'China Installs Record Capacity for Solar (+45%) and Wind (+18%) in 2024'. February 2025. Available at: *https://www.enerdata.net/publications/daily-energy-news/china-installs-record-capacity-solar-45-and-wind-18-2024.html*.

Epoch AI. 'An FAQ on Reinforcement Learning'. Epoch AI, 2026. Available at: *https://epochai.substack.com/p/an-faq-on-reinforcement-learning*.

European Commission. 'Common European Data Spaces'. European Commission, 2026. Available at: *https://digital-strategy.ec.europa.eu/en/policies/data-spaces*.

European Commission. 'EU Launches AI Gigafactories Call to Boost Europe's Computing Capacity and Unlock More than €30 Billion in Investment'. 30 July 2026. Available at: *https://ec.europa.eu/commission/presscorner/detail/en/ip_26_1708*.

European Commission. Key Insights from the First Meeting of the European Expert Forum on Frontier AI. European Commission, 2026. Available at: *https://digital-strategy.ec.europa.eu/en/library/ai-office-publishes-frontier-ai-expert-findings-eu-competitiveness-sovereignty-and-security*.

European Commission. 'Proposal for the Chips Act 2.0'. June 2026. Available at: *https://digital-strategy.ec.europa.eu/en/library/proposal-chips-act-20*.

European Commission. Directorate-General for Communications Networks, Content and Technology, ICF, CARSA, VTT, and Timelex. Study on the next Data Frontier. Publications Office of the European Union, 2025. Available at: *https://doi.org/10.2759/3847027*.

European Data Protection Board. 'Guidelines 2/2026 on Anonymisation of Personal Data'. July 2026. Available at: *https://www.edpb.europa.eu/system/files/2026-07/edpb_guidelines_202602_anonymisation_v1_en_0.pdf*.

European Data Protection Board. 'Guidelines 3/2026 on Web Scraping in the Context of Generative AI'. July 2026. Available at: *https://www.edpb.europa.eu/system/files/2026-07/edpb_guidelines_2020603_webscraping_v1_en_0.pdf*.

European Parliament. Chips Act 2.0. Briefing. European Parliamentary Research Service, 2026. Available at: *https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2026)785742*.

European Parliament and the Council of the European Union. 'Regulation (EU) 2023/1781 of the European Parliament and of the Council of 13 September 2023 Establishing a Framework of Measures for Strengthening Europe's Semiconductor Ecosystem and Amending Regulation (EU) 2021/694 (Chips Act)'. September 2023. Available at: *https://eur-lex.europa.eu/eli/reg/2023/1781/oj*.

Evans, James, Benjamin Bratton, and Blaise Agüera y Arcas. 'Agentic AI and the next Intelligence Explosion'. Version 1. Preprint, arXiv, 2026. Available at: *https://doi.org/10.48550/ARXIV.2603.20639*.

Fedasiuk, Ryan, and Satvik Pendyala. Voltcraft: Industrial Competition in the Age of AI. Working Paper. American Enterprise Institute, 2026. Available at: *https://www.aei.org/research-products/working-paper/voltcraft-industrial-competition-in-the-age-of-ai-aei/*.

Ferrando, Annalisa, Sara Lamboglia, Judit Rariga, and Maurice Schmidt. Adoption and Investment in AI Across the Euro Area: Insights from Harmonised Firm-Level Data. Occasional Paper Series No. 395. European Central Bank, 2026. Available at: *https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op395.en.pdf*.

Germany Insolvencies. 'Germany's Industrial Electricity Prices: Double the US'. May 2026. Available at: *https://germanyinsolvencies.com/blog/germany-industrial-electricity-prices/*.

Gibney, Elizabeth. '"Another DeepSeek Moment": Chinese AI Model Kimi K2 Stirs Excitement'. Nature, ahead of print, 16 July 2025. Available at: *https://doi.org/10.1038/d41586-025-02275-6*.

Gineikyte-Kanclere, Vaida, Militsa Eggert, and Goda Skiotyte. European Software and Cyber Dependencies. PE 778.576. Policy Department for Transformation, Innovation and Health, European Parliament, 2025. Available at: *https://www.europarl.europa.eu/RegData/etudes/STUD/2025/778576/ECTI_STU(2025)778576_EN.pdf*.

Goldschmidt, Nils, and Michael Wohlgemuth, eds. Grundtexte zur Freiburger Tradition der Ordnungsökonomik. Untersuchungen zur Ordnungstheorie und Ordnungspolitik 50. Mohr Siebeck, 2008.

Henderson, Rebecca M., and Kim B. Clark. 'Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms'. Administrative Science Quarterly 35, no. 1 (1990): 9. Available at: *https://doi.org/10.2307/2393549*.

Hooker, Sara. 'On the Slow Death of Scaling'. Preprint, SSRN, 2025. Available at: *https://doi.org/10.2139/ssrn.5877662*.

Hoppe, Heidrun C. 'The Timing of New Technology Adoption: Theoretical Models and Empirical Evidence'. The Manchester School 70, no. 1 (2002): 56–76. Available at: *https://doi.org/10.1111/1467-9957.00283*.

Hsieh, Jung-Kuei, Yi-Ching Hsieh, Hung-Chang Chiu, and Yi-Chu Feng. 'Post-Adoption Switching Behavior for Online Service Substitutes: A Perspective of the Push–Pull–Mooring Framework'. Computers in Human Behavior 28, no. 5 (2012): 1912–20. Available at: *https://doi.org/10.1016/j.chb.2012.05.010*.

IDC and Lisbon Council and CARSA. D2.1 First Report on Facts and Figures: European Data Market Study 2024-2026. European Commission, Directorate-General for Communications Net-works, Content and Technology, 2025. Available at: *https://ec.europa.eu/newsroom/repository/document/2025-13/EDM_2024__2026__First_Report_on_Facts_and_Figures_9cU0iIjcuSZhjmwI9TchwqenidQ_114043.pdf*.

Ioannidou, Maria, and Bingwan Xiong. Competition Law and Policy in Digital Markets : A Comparative Analysis of the EU and China. Oxford University Press, 2025.

Jaghouar, Sami, Jack Min Ong, and Johannes Hagemann. 'OpenDiLoCo: An Open-Source Framework for Globally Distributed Low-Communication Training'. Version 1. Preprint, arXiv, 2024. Available at: *https://doi.org/10.48550/ARXIV.2407.07852*.

Kaplan, Jared, Sam McCandlish, Tom Henighan, et al. 'Scaling Laws for Neural Language Models'. Version 1. Preprint, arXiv, 2020. Available at: *https://doi.org/10.48550/ARXIV.2001.08361*.

Krishnan, Rohit. 'Why Coase Needs Hayek: Sometimes Smart Planners Lose to Simple Markets'. Strange Loop Canon, 2026. Available at: *https://www.strangeloopcanon.com/p/why-smart-planners-lose-to-simple*.

Krishnan, Rohit, and Andrey Fradkin. 'Agent, Know Thyself! (And Bid Accordingly)'. Strange Loop Canon, 2026. Available at: *https://www.strangeloopcanon.com/p/agent-know-thyself-and-bid-accordingly*.

Küsters, Anselm. 'EU Competition Law in the Digital Age: Substantive Challenges and Methodological Innovations Using the Example of the European Commission's Consultation Procedure'. In Doing Justice. Procedural Law and Procedural Law Reform in Comparative Perspective, edited by Henning Glaser. CPG Series of Comparative Constitutional Law, Politics and Governance 7. Nomos, 2026. Available at: *https://doi.org/10.5771/9783845254098-251*.

Küsters, Anselm. 'Ordoliberalism Goes China? A Comparison of Recent Developments in EU and Chinese Competition Law Considering the Digital Economy'. Constitutional Political Economy 36, no. 4 (2025): 444–68. Available at: *https://doi.org/10.1007/s10602-023-09407-y*.

Küsters, Anselm. Small is beautiful 2.0:

Mit digitaler Dezentralisierung zu einer menschlicheren Wirtschaft. 1. Auflage. Verlag Herder, 2026.

Küsters, Anselm, and Céline Nauer. Strengthening the EU's Position in the AI Industry: The Importance of Software and Research. Analysis. Friedrich Naumann Foundation for Freedom, Global Innovation Hub, 2024. Available at: *https://shop.freiheit.org/\#!/Publikation/1940*.

Küsters, Anselm, Patrick Stockebrandt, and André Wolf. 'Towards a Rule-Bound EU Transformation Policy: Can Ordnungspolitik and Industrial Strategy Be Reconciled?' ORDO. Jahrbuch Für Die Ordnung von Wirtschaft Und Gesellschaft 76, no. 1 (2026): 1–30. Available at: *https://doi.org/10.1515/ordo-2026-2012*.

Küsters, Anselm, and Ben Waber. 'How SMEs Can Prepare for the EU's AI Regulations'. Harvard Busi-ness Review, 24 September 2025. Available at: *https://hbr.org/2025/09/how-smes-can-prepare-for-the-eus-ai-regulations*.

Lambert, Nathan. 'Reading Today's Open and Closed Performance'. Interconnects, 2026. Available at: *https://www.interconnects.ai/p/reading-todays-open-closed-performance*.

LeCun, Yann. 'A Path Towards Autonomous Machine Intelligence'. OpenReview, 27 June 2022. Available at: *https://openreview.net/forum?id=BZ5a1r-kVsf*.

Lowdown. 'xAI Wins 41 Gas Turbines for Colossus'. April 2026. Available at: *https://lowdown.today/t/data-centres/1/xai-wins-41-gas-turbines-for-colossus/*.

Marcus, Gary. 'US AI Policy Is a Clumsy Mess. Here's What to Do about It'. Gary Marcus Substack, 2026. Available at: *https://garymarcus.substack.com/p/us-ai-policy-is-a-clumsy-mess-heres*.

Miller, Chris. Chip War: The Fight for the World's Most Critical Technology. Simon & Schuster Ltd, 2023.

Monopolkommission. Wettbewerb 2026: Für eine Wettbewerbsorientierte Wirtschaftspolitik. Monopolkommission, 2026. Available at: *https://www.monopolkommission.de/publikationen/hauptgutachten/wettbewerb-2026-fuer-eine-wettbewerbsorientierte-wirtschaftspolitik*.

NVIDIA Corporation. 'NVIDIA H100 Tensor Core GPU'. 2026. Available at: *https://www.nvidia.com/en-us/data-center/h100/*.

OECD. OECD Artificial Intelligence Review of Germany. OECD Publishing, 2024. Available at: *https://doi.org/10.1787/609808d6-en*.

Pandey, Ajey, Jeremie Eliahou Ontiveros, and Dylan Patel. How AI Labs Are Solving the Power Crisis: The Onsite Gas Deep Dive. SemiAnalysis, 2025. Available at: *https://newsletter.semianalysis.com/p/how-ai-labs-are-solving-the-power*.

Pearce, Tim, Tabish Rashid, Dave Bignell,

Raluca Georgescu, Sam Devlin, and Katja Hofmann. 'Scaling Laws for Pre-Training Agents and World Models'. Version 1. Preprint, arXiv, 2024. Available at: *https://doi.org/10.48550/ARXIV.2411.04434*.

Petcu, Dana. 'Portability and Interoperability between Clouds: Challenges and Case Study'. Proceedings of the 4th European Conference on Towards a Service-Based Internet (Berlin, Heidelberg), ServiceWave'11, 2011, 62–74.

Polites, Greta L., and Elena Karahanna. 'Shackled to the Status Quo: The Inhibiting Effects of Incumbent System Habit, Switching Costs, and Inertia on New System Acceptance1'. MIS Quarterly 36, no. 1 (2012): 21–42. Available at: *https://doi.org/10.2307/41410404*.

Rasmussen, E. G. and others. Report from the Extreme Ultraviolet (EUV) Lithography Workshop. NIST SP 1500-208. National Institute of Standards and Technology (NIST), 2023. Available at: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-208.pdf.

Rauh, Christian. 'Validating a Sentiment Dictionary for German Political Language—a Workbench Note'. Journal of Information Technology & Politics 15, no. 4 (2018): 319–43. Available at: *https://doi.org/10.1080/19331681.2018.1485608*.

Ray, Soumya, Sung S. Kim, and James G. Morris. 'Online Users' Switching Costs: Their Nature and Formation'. Information Systems Research 23, no. 1 (2012): 197–213. Available at: *https://doi.org/10.1287/isre.1100.0340*.

Remus, Robert, Uwe Quasthoff, and Gerhard Heyer. 'SentiWS - A Publicly Available German-Language Resource for Sentiment Analysis'. In Proceedings of the Seventh International Conference on Language Resources and Evaluation, edited by Nicoletta Calzolari. European Language Resources Association, 2010.

Rone, Julia. '"The Sovereign Cloud" in Europe: Diverging Nation State Preferences and Disputed Institutional Competences in the Context of Limited Technological Capabilities'. Journal of Euro-pean Public Policy 31, no. 8 (2024): 2343–69. Available at: *https://doi.org/10.1080/13501763.2024.2348618*.

Saad-Falcon, Jon, Avanika Narayan, Hakki Orhun Akengin, et al. 'Intelligence per Watt: Measuring Intelligence Efficiency of Local AI'. Version 2. Preprint, arXiv, 2025. Available at: *https://doi.org/10.48550/ARXIV.2511.07885*.

Sandoval, Francisco Javier Varela, Isabella Wilkinson, Alex Krasodomski, and Rowan Wilkinson. How Middle Powers Can Weather US and Chinese AI Dominance : The Case for 'Sovereign AI' Strategies. Royal Institute of International Affairs, 2026. Available at: *https://doi.org/10.55317/9781784136710*.

SemiAnalysis. 'Can AMD Break the CUDA Moat? AMD Advancing AI 2026'. SemiAnalysis, 25 July 2026. Available at: *https://newsletter.semianalysis.com/p/can-amd-break-the-cuda-moat-amd-advancing*.

SemiAnalysis. 'RL Environments and RL for Science: Data Foundries and Multi-Agent Systems'. January 2026. Available at: *https://newsletter.semianalysis.com/p/rl-environments-and-rl-for-science*.

SemiAnalysis. 'RL Systems: Mind the Gap — Matching Human and Model Capability'. SemiAnalysis, 2025. Available at: *https://newsletter.semianalysis.com/p/rl-systems-mind-the-gap-matching*.

SemiAnalysis. 'The Wild Wild West of LEGO Datacenters'. 30 July 2026. Available at: *https://newsletter.semianalysis.com/p/the-wild-wild-west-of-lego-datacenters*.

Sevilla, Jaime, and Anton Troynikov. Could Decentralized Training Solve AI's Power Problem? Epoch AI, 2025. Available at: *https://epoch.ai/publications/could-decentralized-training-solve-ais-power-problem*.

Sevilla, Jaime, and Anton Troynikov. Could Decentralized Training Solve AI's Power Problem? Epoch AI, 2025. Available at: *https://epoch.ai/publications/could-decentralized-training-solve-ais-power-problem*.

Shibata, Tomoatsu, Mohammad N. Nasiri, Jun Suzuki, and Yasunori Baba. 'The Co-Evolution of Exploitation and Exploration toward Industrialising Frontier Technology: Insights from ASML's Pursuit of EUV Lithography'. Technology Analysis & Strategic Management 38, no. 6 (2026): 832–46. Available at: *https://doi.org/10.1080/09537325.2025.2501740*.

Soares, Joana. 'The EU's AI Boom Could Undermine Its Own Chip Strategy'. 2026. Available at: *https://www.techpolicy.press/the-eus-ai-boom-could-undermine-its-own-chip-strategy/*.

Staab, Philipp. Markets and power in digital capitalism. First English-Language edition. Manchester University Press, 2024.

Sutton, Richard. 'The Bitter Lesson'. 13 March 2019. Available at: *http://www.incompleteideas.net/IncIdeas/BitterLesson.html*.

Synergy Research Group. 'European Cloud Providers' Local Market Share Now Holds Steady at 15%'. 2025. Available at: *https://www.srgresearch.com/articles/european-cloud-providers-local-market-share-now-holds-steady-at-15*.

Taylor, Ashley, Clayton Friedman, and Gene Fishel. 'State AGs Fill the AI Regulatory Void'. Reuters, 19 May 2025. Available at: *https://www.reuters.com/legal/legalindustry/state-ags-fill-ai-regulatory-void-2025-05-19/*.

The Design and Implementation of Mission-Oriented Innovation Policies: A New Systemic Policy Approach to Address Societal Challenges. OECD Science, Technology and Industry Policy Papers No. 100. Vol. 100. OECD Science, Technology and Industry Policy Papers. 2021. Available at: *https://doi.org/10.1787/3f6c76a4-en*.

Tucker, Catherine. 'Network Effects and Market Power: What Have We Learned in the Last Decade?' Antitrust 32, no. 2 (2018): 72–79.

Van Der Velden, Susanne, Mohammad Nasir Nasiri, Niels G. Noorderhaven, and Henk Akkermans. 'Incumbent Success in the Era of Ferment: Navigation of Intergenerational Transition of Technology'. Academy of Management Proceedings 2020, no. 1 (2020): 20381. Available at: *https://doi.org/10.5465/AMBPP.2020.20381abstract*.

Visser, Corinna. AI Gigafactories: Wie sieben Fabriken helfen sollen, Europas KI-Rückstand aufzuholen. Centrum für Europäische Politik, 2026.

Wang, Dan. Breakneck: China's Quest to Engineer the Future. First edition. W. W. Norton & Company, Inc., 2025.

Wang, Fali, Zhiwei Zhang, Xianren Zhang, et al. 'A Comprehensive Survey of Small Language Models in the Era of Large Language Models: Techniques, Enhancements, Applications, Collaboration with LLMs, and Trustworthiness'. Version 1. Preprint, arXiv, 2024. Available at: *https://doi.org/10.48550/ARXIV.2411.03350*.

Weil, Gabriel. 'Don't Let AI Developers Hire Their Own Referees'. AI Frontiers, 29 July 2026. Available at: *https://ai-frontiers.org/articles/dont-let-ai-developers-hire-their-own-referees*.

Zhang, Angela Huyue. Chinese Antitrust Exceptionalism: How the Rise of China Challenges Global Regulation. Oxford University Press, 2021.

Zhang, Irene, and Nick Corvino. '"Eastern Data, Western Compute" Is Fake'. ChinaTalk, 10 July 2026. Available at: *https://www.chinatalk.media/p/eastern-data-western-compute-is-fake*.

About the Author



Source: Centre for European Policy

Dr. Anselm Küsters, LL.M.

is Scientific Director for the Nordic Countries and Head of the Department of Digitalisation and AI at the Centre for European Policy (cep). He also works in the field of Digital Humanities as an affiliated researcher at the Max Planck Institute for Legal History and Legal Theory in Frankfurt am Main and as a post-doctoral researcher (Habilitand) at the Humboldt University in Berlin. From April 2025 to March 2026, he held an Interim Professorship for Digital Humanities at the University of Stuttgart. Anselm holds a doctorate from Goethe University Frankfurt am Main (Dr. Phil.) and studied Economic and Social History at the University of Oxford, UK (M.Phil., Distinction), European Law at the University of Würzburg (LL.M.), and History, German Literature, and Economics at the University of Heidelberg (B.A. and B.Sc.).

Centrum für Europäische Politik

FREIBURG
Kaiser-Joseph-Straße 266,
D-79098 Freiburg

BERLIN
Raum 4205
Schiffbauerdamm 40,
D-10117 Berlin
Phone + 49 761 38693-0

Centre de Politique Européenne / PARIS
17, rue Saint Fiacre
F-75002 Paris
Phone +33 1 88 80 72 00

Centro Politiche Europee / ROMA
Via A. Brunetti, 60
I-00186 Roma
Phone +390684388433

The Centrum für Europäische Politik FREIBURG | BERLIN, the Centre de Politique Européenne PARIS, and the Centro Politiche Europee ROMA form the Centres for European Policy Network FREIBURG | BERLIN | PARIS | ROMA.

Free of vested interests and party-politically neutral, the Centres for European Policy Network provides analysis and evaluation of European Union policy, aimed at supporting European integration and upholding the principles of a free-market economic system.



freiheit.org