



RAND
EUROPE

centre
for future
generations



e l l i s
INSTITUTE
TÜBINGEN

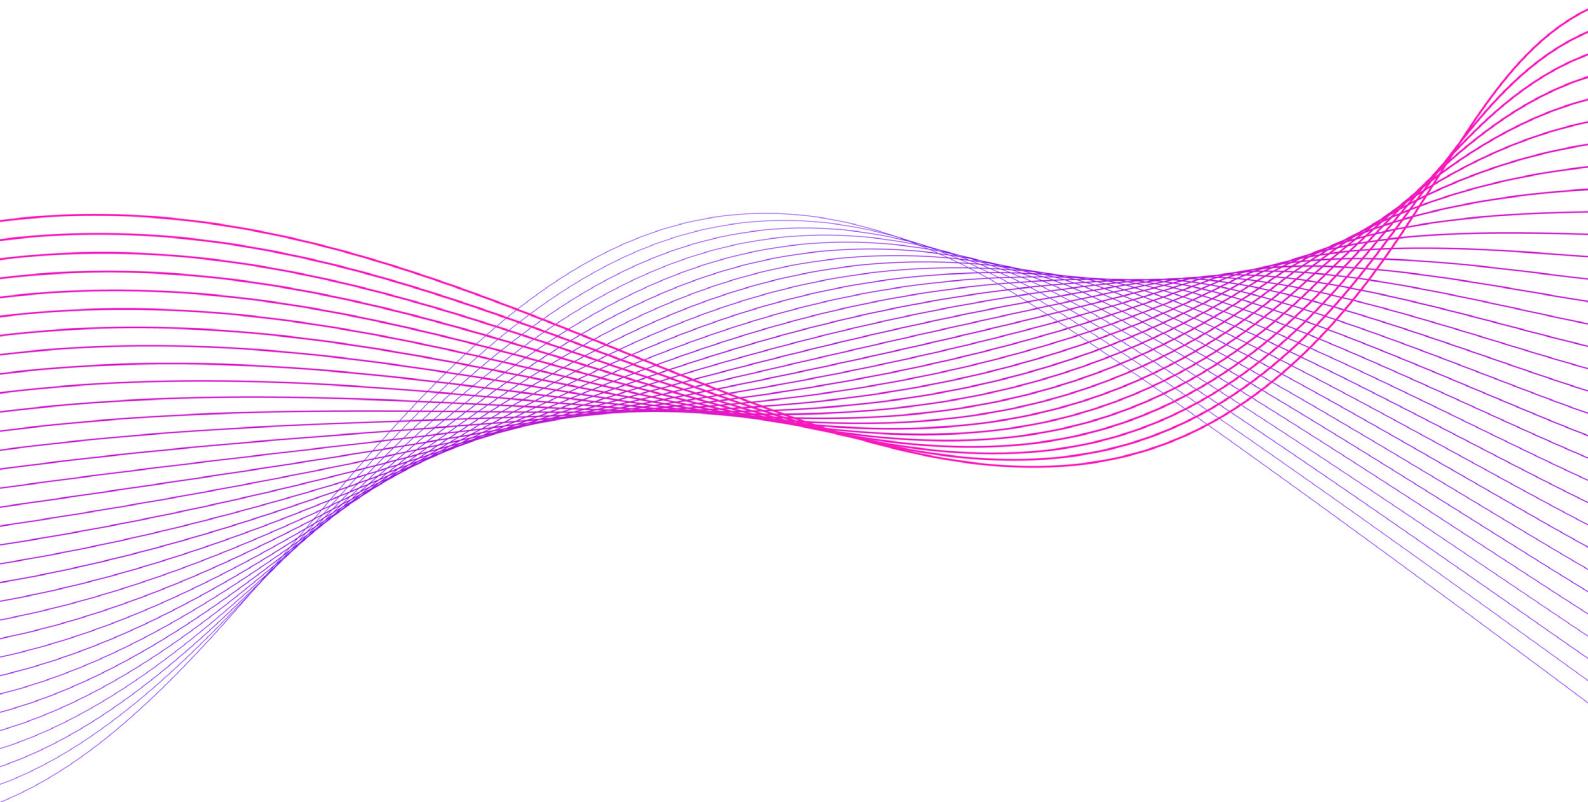


MAX PLANCK INSTITUTE
FOR INTELLIGENT SYSTEMS

Europe and the geopolitics of AGI

The need for a preparedness plan

Maximilian Negele*^{1,4}, Daan Juijn*², Afek Shamir¹, David Janků², Bengüsu Özcan², Lisa Soder^{1,3},
Lucia Velasco⁴, Max Reddel², Michiel Bakker⁵, Lorenzo Pacchiardi⁶, Maksym Andriushchenko⁷



*Corresponding authors, equal contribution

1 RAND Center on AI, Security, and Technology

2 Centre for Future Generations

3 Stanford University

4 University of Oxford

5 Massachusetts Institute of Technology

6 University of Cambridge

7 ELLIS Institute Tübingen and Max Planck Institute for Intelligent Systems

For more information on this publication, visit www.rand.org/t/RRA4636-1

About RAND Europe

RAND Europe is a not-for-profit research organisation that helps improve policy and decision making through research and analysis. To learn more about RAND Europe, visit www.randeurope.org.

Research Integrity

Our mission to help improve policy and decision making through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behaviour. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif., and Cambridge, UK

© 2025 RAND Corporation and Centre for Future Generations

RAND® is a registered trademark.

Cover: Adobe Stock

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorised posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, please visit www.rand.org/pubs/permissions.

Preface

Artificial general intelligence (AGI) has moved from speculation to the centre of political debate. This report examines arguments that AGI could emerge in the near future and considers what such a development would mean for Europe. It adopts a working definition of AGI as AI systems that match or exceed humans at most economically useful cognitive work. On this definition, empirical trends and expert judgement suggest that it is plausible that AGI could emerge between 2030 and 2040 or even earlier, though uncertainty remains substantial. The emergence of such systems could reshape economic growth, military capabilities and international stability, while intensifying competitive and coercive dynamics among nations, well beyond those produced by current AI systems. Against this backdrop, strategic awareness of frontier AI developments remains uneven across European governments and institutions. Moreover, Europe is currently lagging on key measures of competitive positioning, holds only limited leverage in the AI value chain, and existing EU and Member State strategies do not yet match the potential stakes of AGI. This report proposes an authoritative AGI Preparedness Report, led by a figure with exceptional political authority and technical credibility, supported by a world class team. Its mandate would be to outline a coherent European strategy for prosperity, security and stable international governance in a possible AGI era.

About RAND Europe

RAND Europe, the European arm of RAND Corporation, is a non-profit research institute with around 200 staff across its offices in Belgium, the Netherlands and the United Kingdom. RAND has an 80-year history – including 30 years at RAND Europe – of helping policy makers make better decisions through objective research and analysis, including decades of work on innovative technologies such as AI. In Europe, our clients include national governments, militaries, and AI safety or security institutes, as well as charitable foundations, non-governmental organisations, the EU, NATO and United Nations institutions. Within the EU, this includes a wide range of RAND Europe research for the European Commission, the European Parliament and agencies such as the European Defence Agency, European Union Drugs Agency, European Union Agency for Cybersecurity and Frontex. For more information, contact info@randeurope.org

Center on AI, Security, and Technology

RAND Global and Emerging Risks is a division of RAND that delivers rigorous and objective public policy research on the most consequential challenges to civilisation and global security. The research for this report was undertaken within the division's Center on AI, Security, and Technology, which aims to examine the opportunities and risks of rapid technological change, focusing on artificial intelligence, security and biotechnology. For more information, contact cast@rand.org or mnegele@randeurope.org.

Funding

This research was independently initiated and conducted jointly by the Center on AI, Security, and Technology and the Centre for Future Generations using income from operations and gifts and grants from philanthropic supporters. A complete list of RAND's donors and funders is available at www.rand.org/CAST. RAND clients, donors and grantors have no influence over research findings or recommendations. The Centre for Future Generations (CFG) is a nonprofit organisation supported by philanthropic foundations, public institutions, and individual donors. CFG does not accept funding that could compromise its independence, including primary support from governments, corporations, or political parties. For more information, see cfg.eu/funding.

Acknowledgements

We sincerely thank the following people for their input and/or written feedback on this report: Anton Leicht, Anton Shenk, Audrey Tang, Chris Meserole, Erich Grunewald, Henry Papadatos, Jonathan Schmidt, Jorge Teixeira, Luka Ignac, Markus Anderljung, Michael Aird, Oscar Delaney, Patricia Paskov, Philip Fox, Rafael Andersson Lipcsey, Saad Siddiqui, Seb Krier, Shyam Krishna, Siddhi Pal, Soeren Mindermann, Stefan Leijnen, Tom Ough. We are particularly grateful to Henri van Soest and James Black for their detailed review and to Jessica Plumridge, Ben Plumridge, Rowan Emslie, Georgina Melia, and Lauren Crawford for copy-editing and design.

These individuals do not necessarily agree with the views in the paper, and all mistakes remain our own.

This paper is a collaborative effort. While the authors are aligned on the paper's core argument, individual authors may differ on some specific claims or interpretations.

Executive summary

This report recommends that the European Commission President should commission an AGI Preparedness Report with the political authority, analytical breadth and urgency of the Draghi Report. This recommendation follows three broad assessments made in the report. Firstly, artificial intelligence (AI) systems matching or exceeding humans at most economically useful cognitive work – our working definition of artificial general intelligence (AGI) – could plausibly arrive between 2030 and 2040, or even earlier. Secondly, the emergence of AGI systems could fundamentally reshape the global distribution of power, with major implications for economic growth, military capabilities and international stability. Thirdly, Europe currently does not yet have the comprehensive strategic awareness, competitive positioning, or policy strategies to navigate a potential shift to a world with AGI. Without a preparedness plan that treats the prospect of AGI as a core strategic concern for its security and prosperity, Europe risks marginalisation in what may prove the world’s most consequential technological transition.

The near-term emergence of AGI is likely enough to warrant significant attention from European policy makers

AI capabilities have progressed quickly but unevenly in recent years. Six years ago, general purpose AI systems struggled to write coherent text; today, they win gold medals at international mathematics olympiads and outperform top coders in competitions. Yet they remain brittle in important ways: they hallucinate facts, sometimes fail at simple visual reasoning, and cannot reliably perform physical tasks such as folding laundry. This uneven landscape is sometimes described as a ‘jagged frontier’, with systems excelling in domains with clear success criteria while still struggling with tasks that rely on tacit knowledge or long horizon planning.

Rapid increases in compute, data and algorithmic efficiency can likely continue over the next few years. Past progress in AI has been driven by rapid increases in training compute, training data and algorithmic efficiency. Training compute has grown by roughly a factor of 5 per year, training data by about 3.6 per year and algorithmic efficiency by about 3 per year. These rates could slow in the coming years, which would still mean continued progress, just at a more measured pace. There is also a real possibility of further acceleration, especially in algorithmic progress. At present, the evidence does not suggest a sharp slowdown before 2030.

While future capabilities are hard to predict, AGI could plausibly emerge between 2030 and 2040 or even earlier. Definitions of AGI are contested, with authors using different milestones and some rejecting the concept outright. We adopt the working definition of ‘AI systems matching or exceeding humans at most economically useful cognitive work’. On this definition, empirical trends and expert judgement suggest that AI could plausibly reach AGI-level performance between 2030 and 2040 or even earlier. This range involves significant uncertainty and should be seen as one plausible scenario rather than a precise forecast. Even so, the plausible near term emergence of AGI is sufficiently likely to warrant significant attention from European policy makers.

The emergence of AGI could expose Europe to major shifts in global power as well as strategic instability

Actors that successfully deploy AGI could gain major economic and military advantages. Countries could achieve rapid economic growth by automating human labour and accelerating science and research, allowing them to outproduce competitors and exert significant geopolitical influence. In the military realm, AGI could transform warfare through automated intelligence fusion, enhanced decision making and planning, large fleets of autonomous systems, sophisticated cyber operations and potentially the development of entirely new classes of weapons. Some of these effects already become visible with current AI systems, but AGI could amplify them by enabling much higher levels of competence and generality.

The emergence of AGI might lead to geopolitical instability. States and companies may race to develop AGI, acting aggressively to gain a competitive advantage. Governments may try to block rivals through export controls, cyberattacks or escalatory deterrence actions. Shocks to labour markets may lead to internal instability and unrest. AGI also poses risks beyond nation states: it could help terrorists build weapons of mass destruction, enable AI systems that pursue goals misaligned with human values or allow for dangerous concentration of power. Current AI capabilities have already enabled largely autonomous cyber attacks and have shown potential for misuse in biological weapon design.

Major powers already treat AI and the potential for AGI as a strategic priority. The US frames AI leadership as essential to national security and has imposed semiconductor export controls on China. Beijing aims to become the global AI leader by 2030 and is accelerating efforts to indigenise its AI supply chain. The United Kingdom (UK), the United Arab Emirates, Singapore and other middle powers are scrambling to position themselves to retain influence while keeping access to the AI frontier. Europe must orient itself to this new competitive reality and respond strategically in a way that fits its own asymmetric strengths and weaknesses.

Europe is not yet sufficiently prepared for a transition to a world shaped by AGI

Strategic awareness of frontier AI developments remains uneven across European governments and institutions. In the absence of leading domestic frontier AI firms, European policy makers often rely on external sources of information which can complicate the assessment of technical claims. The UK has begun to build substantial analytical capacity through its AI Security Institute, but similar capabilities on the continent are still emerging. Germany, for instance, does not participate in key international AI safety fora, and the EU AI Office, while comparatively well-staffed, currently operates with a focused mandate and relatively limited budget.

Europe trails on most measures of AI competition and its strategic levers provide limited geopolitical influence. European AI models lag behind US and Chinese models by six to twelve months. Europe hosts only about 5 per cent of global AI computing capacity, compared with roughly 75 per cent in the US. EU start-ups attract just 6 per cent of global AI venture funding. Energy costs in Europe are significantly higher than in the US, and Europe loses much of its top AI talent to better-capitalised US companies. In theory, the EU's market power and chokepoints, such as Dutch firm ASML's monopoly on EUV lithography tools, give Europe levers to negotiate access to foreign frontier AI and compute and to influence global AI governance beyond the AI Act. In practice, these levers are constrained by geopolitical sensitivities and by partners' potential concerns around the diffusion of more transformative AI systems.

Europe's existing AI strategies are fragmented across institutions and policy domains and often lack scale and integration. Current EU AI initiatives, including the AI Continent Action Plan, ApplyAI, InvestAI and AI Factories, are important steps but remain under-resourced and spread across multiple Commission Directorates-General and Member States, with limited coordination. The EU's reliance on consensus-based decision making slows policy responses in geopolitical crises, while many AGI-relevant levers in defence, intelligence and critical infrastructure remain national competencies without clear mechanisms for joint action. Recent moves such as the Frontier AI Initiative announced by France, Germany and the European Commission signal growing ambition, but practical implementation paths remain uncertain.

Recommendation: The EU should commission an AGI Preparedness Report

The President of the European Commission should commission an AGI Preparedness Report. Getting AGI policy wrong, whether through premature investments or failure to secure access to critical capabilities, could be very costly. Europe therefore needs an integrated and rigorous preparedness plan that clarifies the big picture and addresses the many interlocking challenges it may face in a transition to AGI. An AGI Preparedness Report would identify the most important questions, map the main options, and recommend strategic directions grounded in shared assumptions across EU institutions and Member States.

The AGI Preparedness Report should take inspiration from the Draghi Report on EU Competitiveness. It should be led by a figure with exceptional political authority and technical credibility, supported by a world-class team with the freedom, access and resources to deliver a rigorous assessment within months rather than years. The Report should clearly delineate which actions require coordination at the EU level and which demand national leadership, and set out mechanisms to ensure that both these strands act in concert. Its authors should have the intellectual independence to reach evidence-based conclusions, which should be presented to the European Council at a dedicated event to enable a parallel diplomatic process.

The Report should address three core challenges relating to AGI: capturing economic benefits while staying sovereign, preparing societies for rapid change, and strengthening global governance. It must identify which capabilities can be procured from partners, which require European control, and how Europe should invest in domestic capabilities. It should set out how European societies can prepare for potential labour shocks and novel security threats, while building the state capacity to manage unforeseen challenges. Finally, it should articulate Europe's vision for stable international AGI governance and position the EU as a trusted and influential actor committed to safe and ethical AI development.

Table of contents

Preface.....	i
Executive summary.....	iii
Table of contents.....	vi
Figures.....	vii
1. Introduction.....	1
Methodology	3
2. How far away is artificial general intelligence?	4
What today's frontier AI systems can do	4
Will rapid progress in frontier AI scale towards AGI?	7
What would it mean to achieve AGI?.....	18
3. How would AGI affect geopolitics?.....	21
AGI-induced power shifts could lead to global instability.....	23
AGI could shift power to individuals or AI systems themselves.....	24
Expected power shifts are driving competing superpowers to treat AI and AGI as a strategic priority	26
4. Is Europe sufficiently prepared for a fast transition to AGI?	29
Dimension 1: Strategic awareness.....	31
Dimension 2: Competitive positioning and leverage	33
Dimension 3: Robustness of current policy strategies	48
5. Recommendation: the European Commission President should commission an AGI Preparedness Report.....	52
Why a report instead of concrete action recommendations?	53
Learning from the Draghi Report on EU Competitiveness.....	54
Core questions an AGI Preparedness Report should address.....	57
6. Conclusion.....	62
Glossary.....	64
References.....	68

Figures

Figure 1. LLM inference price development across tasks.....	5
Figure 2. Benchmark performances of GPT3, GPT-4 and GPT-5 compared	7
Figure 3. The AI triad	10
Figure 4. Scaling laws compared across train-time and test-time compute	12
Figure 5. Time-horizon of software engineering tasks different LLMs can complete 50 per cent of the time	17
Figure 6. Illustrative development of expert views on whether AGI has arrived.....	20
Figure 7. Benchmark performance of European frontier models in international comparison	34
Figure 8. Europe's share of global aggregate computational performance (in 16-bit FLOP/s)	38
Figure 9. Total electricity generated in each country or region (in terawatt-hours).....	40
Figure 10. AI use by organisations in the world, 2023 vs. 2024	42
Figure 11. Concentration of the AI chip supply chain, expressed as a percentage of total market share.....	46
Figure 12. The major trade partners of the US in 2024: the EU (\$976bn total trade), Mexico (\$840bn), Canada (\$762bn) and China (\$582bn).	47

1. Introduction

In recent years, artificial general intelligence (AGI) has moved from speculation to the centre of political debate. Senior European policy makers now speak openly about AI systems approaching human-level capabilities on short timescales. In her speech at the Annual EU Budget Conference 2025,¹ Commission President Ursula von der Leyen remarked that when the current budget was negotiated ‘we thought AI would only approach human reasoning around 2050. Now we expect this to happen already next year.’ The Commission’s AI Continent Action Plan similarly anticipates that ‘the next generation of frontier AI models’ could amount to a qualitative leap ‘towards Artificial General Intelligence (AGI) capable of tackling highly complex and diverse tasks, matching human capabilities’.² Across governments, industry and the expert community, the assumption that AGI may arrive soon is increasingly taking hold.

Another claim has gained traction: that AGI would not merely be another technology but a force capable of reshaping geopolitics.³ Experts increasingly discuss the possibility that more capable AI systems could alter growth trajectories,⁴ change the balance of military power,⁵ and introduce new levers for economic coercion, surveillance and influence.⁶ Indeed, export controls on advanced semiconductors,⁷ the emergence of AI-enabled military systems on the battlefield in Ukraine,⁸ and large-scale public and private investments⁹ in computing infrastructure are now widely interpreted as signs that AI is becoming a central domain of strategic competition.

Experts disagree on the likelihood of short AGI timelines, as well as the extent to which AGI would alter existing power structures. Indeed, some have argued that the rapid growth in AI investment is mostly a reflection of industry hype and amounts to the creation of a new financial bubble.¹⁰ Others have tried to diffuse this argument by pointing at real progress in benchmarks that measure AI capabilities and the rapid revenue growth reported by leading AI companies.¹¹ A third camp has pointed out that relatively short timelines to AGI systems are not incompatible with the bubble bursting – much like, after the dotcom bubble burst, the internet did still fundamentally alter society at large.¹² More fundamentally, the

¹ European Commission (2025e).

² European Commission (2025b).

³ Burdette & Demelash (2025).

⁴ Potlogea & Ho (2025).

⁵ Burdette & Demelash (2025b).

⁶ Farrell & Newman (2019).

⁷ Heim (2025).

⁸ Bondar (2025).

⁹ Soni et al. (2025); Abyazov (2025).

¹⁰ Marcus (2025b).

¹¹ Wildeford (2025a).

¹² Hashim & Ford (2025).

very term of ‘AGI’ is contested, with some shying away from using it,¹³ or preferring alternatives, such as ‘transformative AI’.¹⁴

Although it is too early to know which sides of these debates will prove correct, it is evident that if anything resembling AGI were to emerge soon, Europe would enter this transition under considerable strain. Russia’s illegal invasion of Ukraine has reintroduced large-scale land war to the continent and stretched European defence and industrial capacity. Economic stagnation and political fragmentation complicate long-term investment and reform. Transatlantic relations have frayed,¹⁵ with recent episodes in trade,¹⁶ security¹⁷ and technology policy¹⁸ underscoring Europe’s dependence on US security guarantees, energy supplies and digital infrastructure. Meanwhile, China’s material support for Russia and its weaponisation of economic dependencies illustrate how technological and industrial leverage can be used for strategic ends.¹⁹ In a world where military capability, economic coercion and technological dominance increasingly determine the international hierarchy, Europe finds itself structurally disadvantaged, possessing neither the military might and innovative edge of the US nor the industrial scale, talent pool and centralised decision-making capacity of China.²⁰

Europe cannot afford to approach the prospect of AGI as a mere regulatory challenge. It must assess whether its current economic, technological and institutional foundations are fit for a world in which AI systems might match or exceed human capabilities across most domains. Technological dependence on foreign AI systems, destabilising arms race dynamics,²¹ and the difficulty of ensuring lawful, ethical and robust AI development could leave Europe more vulnerable than at any time in the post-war period.²² Conversely, securing access to enough compute, data, talent and risk-tolerant capital could help unlock new growth, strengthen defence, and allow Europe to shape global norms for safe and ethical AI.

This paper first scrutinises the underlying assumptions behind prevalent statements that AGI is coming soon and that its emergence would reshape geopolitics. We do not aim to provide a comprehensive or definitive treatment of AGI timelines or geopolitical impacts; instead, we conduct a focused review of the most influential arguments and evidence bearing on these two claims, clarifying what is known, what is contested, and where uncertainty remains. We also discuss how different communities define and dispute the term ‘AGI’. Taking seriously the possibility that AGI could emerge on a relatively short timescale, the paper goes on to assess Europe’s preparedness along three key dimensions: strategic awareness, competitive positioning and leverage, and robustness of current policy strategies.

We conclude that Europe is currently underprepared for a transformation to a world with AGI. On present trajectories, the EU and its Member States risk being forced into a choice between deep dependence on foreign AI systems and relative economic and strategic marginalisation. Avoiding this dilemma will require a preparedness plan that addresses three core questions: how can Europe capture economic benefits from AGI while remaining sovereign, how can it prepare its institutions and societies

¹³ Bengio et al. (2025)

¹⁴ Agrawal, Korinek, & Brynjolfsson 2025

¹⁵ Munich Security Conference (2025); Vance (2025).

¹⁶ Bown (2025).

¹⁷ Tankersley et al. (2025).

¹⁸ Heim (2025).

¹⁹ Kirichenko (2025); Baskaran & Schwartz (2025).

²⁰ Draghi (2024); Wang & Kroeber (2025).

²¹ Burdette & Demelash (2025a).

²² Janků et al. (2024).

for rapid change resulting from AGI, and how can it strengthen international stability, security and shared prosperity in a world shaped by AGI.

Methodology

Data collection

This report is based on a narrative literature review drawing on approximately 300 materials, including peer-reviewed articles, working papers and preprints from research institutions, reports from governments and international organisations, and industry analyses. Relevant material was identified through keyword-based searches in academic databases and search engines such as Google Scholar, and the use of existing bibliographies. The review was conducted iteratively between July and November 2025, with search and synthesis repeated until additional sources yielded diminishing marginal returns. We focused primarily on work from the past 5 years to capture current thinking and recent developments, while incorporating older contributions where they remain analytically relevant. Publications appearing after November 2025 are not included in the research.

We prioritised sources that: (1) directly address at least one of the report's guiding questions on AGI timelines, geopolitical impacts, or Europe's preparedness; (2) offer analytical rather than purely speculative contributions, for example empirically rigorous analyses of frontier capabilities and inputs, or theoretically grounded studies of economic or military implications; (3) meet basic credibility thresholds in terms of venue, author expertise, or transparency of methods and assumptions; and (4) add distinctive evidence or perspectives, including from non-European actors relevant to Europe's choices. The core review is restricted to English-language publications.

Analytical framework

Chapter 2 combines two complementary analytical lenses on AGI timelines. First, existing forecasts and models – expert surveys, forecasting-platform aggregates, and quantitative models that make explicit assumptions about input factors – are synthesised to obtain an ‘outside view’ on when AGI-level systems might arrive. Second, a plausibility check is conducted based on constraints in the ‘AI triad’ of compute, data and algorithms. Drawing from empirical trends in each component of the triad, we ask whether there are clear structural bottlenecks that would make continued progress towards AGI-level capabilities implausible on a short timescale.

Chapter 3 analyses academic, policy and technical work that makes claims about the geopolitical consequences of AGI and classifies these claims into two recurring themes: shifts in economic growth and military capabilities and shifts of power to individuals or AI systems. The chapter summarises arguments in favour of and against each type of shift, highlights major disagreements, filters for mechanisms that appear especially relevant from a European perspective, and documents the reactions of major AI powers to the perceived implications of AGI.

Chapter 4 assesses European AGI preparedness using a three-pronged analytical framework. Drawing from literature in strategic foresight, societal transitions and geopolitics, preparedness is defined as a combination of: (1) strategic awareness in government, meaning the ability of European and national institutions to monitor frontier AI developments; (2) competitive positioning and leverage, meaning Europe's structural position in the global AI value chain; and (3) the robustness of current policy strategies, including their coherence, resourcing and institutional integration.

2. How far away is artificial general intelligence?

Timelines matter in any geopolitical assessment of AGI. If approximately human-level systems can plausibly emerge within the next two decades, investment, institutional reform and international strategy must reflect that possibility. If such systems only emerge much later, or never, the policy response would differ.

The first part of this chapter assembles the main strands of evidence relevant to AGI timelines, using a working definition of AGI as ‘artificial intelligence (AI) systems matching or exceeding humans at most economically useful cognitive work’. We begin by describing what today’s frontier AI systems can already do and where their capabilities remain limited.²³ We then review the existing forecasting and modelling literature on AGI timelines and examine recent progress and potential bottlenecks through the lens of the ‘AI triad’ of compute, data and algorithms.

We find the term AGI useful, which is why we use it throughout this report. We also recognise that it is contested, used inconsistently, and sometimes outright rejected. The second part of this chapter therefore examines these competing definitions and milestones to clarify what various authors mean when they refer to AGI.

This review does not aim to provide a precise prediction. Instead, it assesses whether the emergence of AGI in the relatively near future is sufficiently plausible and consequential that Europe should treat it as a serious planning scenario. In doing so it lays the foundation for the geopolitical analysis and preparedness assessment in the remainder of the report.

What today’s frontier AI systems can do

In 2020, early large language models (LLMs) such as OpenAI’s GPT-3 impressed observers with fluent text generation. However, their limitations were substantial. They could not properly follow human instructions and processed only text – not images, audio or video. Their context windows were limited, preventing users from inputting large text files, and they were slow and expensive to operate.²⁴ Without the ability to ‘think aloud’, browse the web or use other tools such as code interpreters, they could not autonomously perform multi-step tasks. Perhaps most significantly, their raw ‘intelligence’ was deficient: GPT-3 struggled with tasks as simple as creating rhymes or multiplying three-digit numbers.²⁵

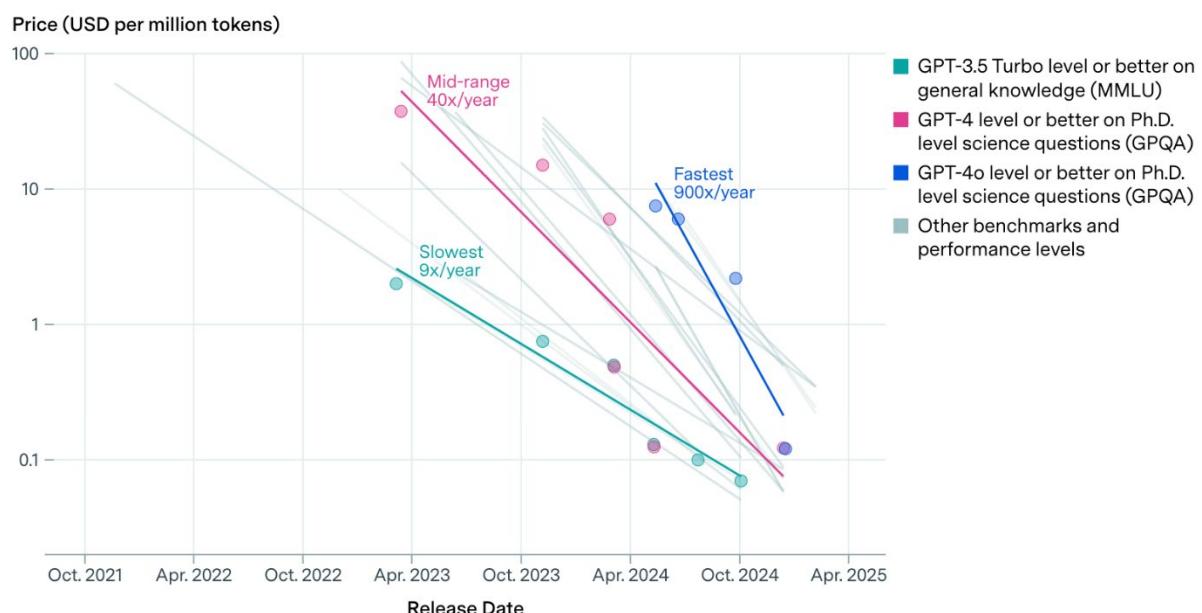
²³ For readers seeking more detail, we recommend consulting Chapter 1 of the International AI Safety Report (Bengio 2025), which gathers the collective insights of experts nominated by the governments of 30 countries, the UN, the EU and the Organisation for Economic Co-operation and Development (OECD).

²⁴ Brown et al. (2020).

²⁵ AI Digest (2025b).

Five years later, the AI industry has achieved both significant incremental progress and qualitative leaps. Today's frontier systems are multimodal: they process text, audio, images and video, with some generating multiple output types through unified models, including photorealistic images and fluent speech.²⁶ They employ various tools to search the web, use code interpreters, and access external APIs. They can produce chains-of-thought for extended problem-solving and maintain records of previous actions to enable complex, longer-duration software projects.²⁷ These qualitative improvements come with advances in speed, affordability (see Figure 1) and raw 'intelligence'. Leveraging these capabilities, current models achieve gold-medal performance at the International Mathematics Olympiad,²⁸ rank among the world's top competitive coders, achieve superhuman scores on PhD-level science tests,²⁹ and exceed the performance of human doctors on certain diagnostic benchmarks.³⁰ Specialised agentic models can autonomously restructure entire codebases or create functional front-end applications from single prompts, whilst desk research agents can generate fully sourced analytical reports within ten minutes.³¹

Figure 1. LLM inference price development across tasks



Source: Cottier (2025)

However, while AI systems excel at tests that are challenging for most humans, they remain brittle in many ways. Models frequently hallucinate facts,³² struggle with simple visual reasoning tasks such as recognising a fourth stripe added to an Adidas logo,³³ misunderstand questions involving basic real-world physics and social intelligence, and fall for logical trick questions.³⁴ Models also consistently underperform in areas requiring tacit knowledge or long-horizon planning and execution.³⁵ For example, while

²⁶ Examples include Gemini 2.5 Flash Image and GPT-4o.

²⁷ Kwa et al. (2025).

²⁸ Metz (2025).

²⁹ Epoch AI (2025b).

³⁰ Kolata (2024).

³¹ Examples include Claude Code, GPT-5, and Gemini Deep Research, respectively.

³² Metz & Weise (2025).

³³ Wang et al. (2024).

³⁴ Williams & Huckle (2024).

³⁵ As noted in Bengio (2025), 171–2.

Anthropic's Claude can sometimes automate hours of work challenging for junior software engineers, it cannot reliably run a small vending machine business.³⁶ While Gemini 2.5 Pro excels at difficult mathematics examinations, it fails in peculiar ways when given open-ended goals such as selling online t-shirts.³⁷ Meanwhile, progress on physical tasks has lagged behind cognitive capabilities. Although robotic hardware continues to improve, the autonomous operation of robotic bodies remains challenging for AI systems; impressive demonstrations often rely on tele-operation.³⁸

These contrasting results create what researchers term a 'jagged frontier' of AI capabilities, a landscape where models excel in domains with clean, abundant, machine-readable data (such as coding and formal mathematics) but often struggle with experiential knowledge, common sense, or ambiguous, unstructured problems.³⁹ This jagged frontier, combined with industry hype, has made it difficult for occasional observers to assess the capabilities of current AI systems and anticipate future developments. This became particularly evident following the release of GPT-5 – a series of models that many commentators found underwhelming, leading some to declare that AI is hitting a wall.⁴⁰ However, GPT-5 may have primarily been a product release aimed at raising the capability floor (the quality of tools free users have access to) rather than the ceiling (the maximum capabilities for power users). It scores on-trend on several widely used capability benchmarks⁴¹ and recent research further suggests that the capability jump from GPT-3 to GPT-4 was comparable to that from GPT-4 to GPT-5, but that the latter has been obscured by numerous intermediate releases (see Figure 2).⁴²

This suggests that progress is best assessed through longer-term trends rather than individual releases. However, just as capabilities should not be judged solely in the vacuum of benchmarks, they also cannot be inferred directly from real-world adoption rates – a necessary nuance given that adoption data is rapidly evolving. For instance, early results on agentic AI in white-collar workflows remain inconclusive: while some surveys hint at productivity boosts,⁴³ others report negligible impact, with some citing unstructured corporate data as the primary bottleneck.⁴⁴ Moreover, assessments of progress are further complicated by copyright lawsuits and child safety concerns⁴⁵, where compliance becomes a prerequisite for products to remain in the market or exist in the first place. While these frictions complicate market adoption and return cycles for AI investments, they do not necessarily indicate a technological plateau.

³⁶ As attempted by Anthropic's Project Vend; Anthropic (2025b).

³⁷ AI Digest (2025a).

³⁸ Examples include Tesla's Optimus hand demo.

³⁹ Dell'Acqua et al. (2023).

⁴⁰ Marcus (2025a).

⁴¹ Wildeford (2025a); see METR (2025b) for evidence of GPT-5's performance on key benchmarks.

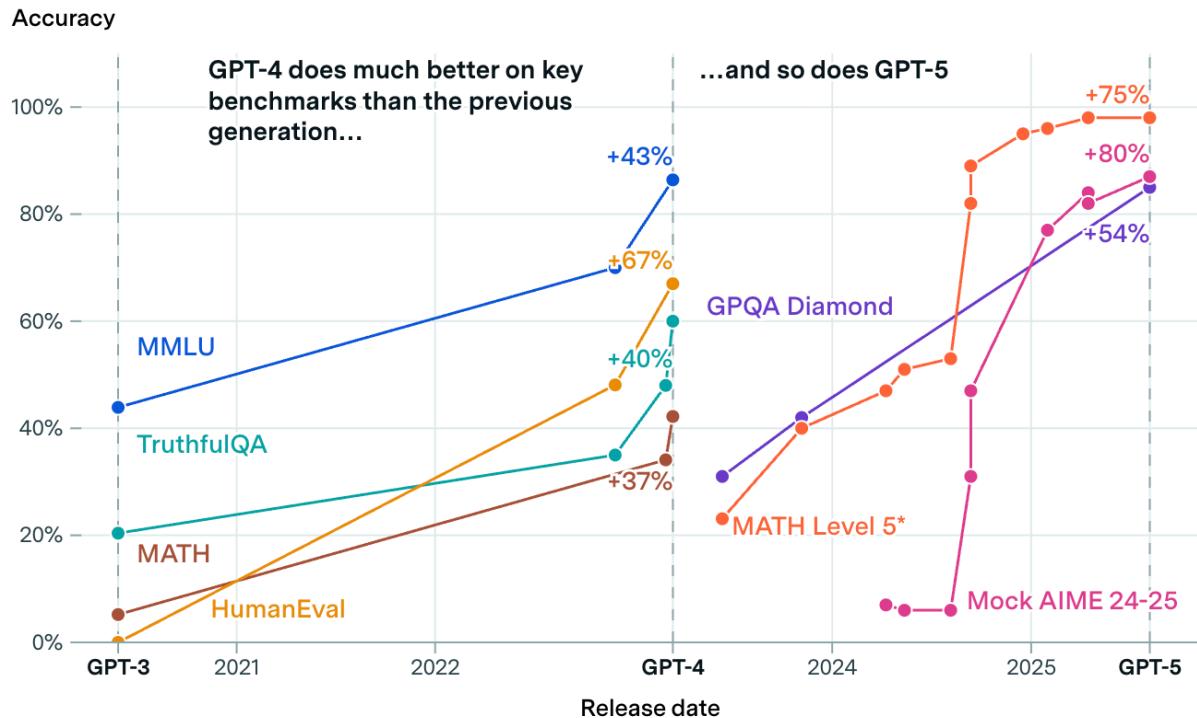
⁴² Emberson (2025).

⁴³ IBM (2025).

⁴⁴ Estrada (2025); Hill (2025).

⁴⁵ Heaven (2025); McCallum (2025).

Figure 2. Benchmark performances of GPT3, GPT-4 and GPT-5 compared



Source: Emberson (2025).

Will rapid progress in frontier AI scale towards AGI?

If today's systems already perform impressively on many cognitively demanding tasks, the next question is whether this trajectory can lead to capabilities that meet a reasonable definition of AGI. Recent years have seen rapid, sometimes surprising jumps in capability, driven by larger training runs, better data and algorithmic improvements. But it is not obvious whether this pace can be maintained, or whether economic, technical or physical constraints will force a sharp slowdown before systems reach AGI-level performance.

In this section, we examine the prospects for further progress. We first review the main forecasting and modelling efforts that translate current trends and expert judgement into explicit timelines for AGI or 'transformative AI'. We then make an assessment of the underlying drivers of progress, using the 'AI triad' of compute, data and algorithms. The aim is to evaluate whether the literature suggests clear structural reasons – such as binding limits on energy, data or algorithmic improvement – that would make the continued advance towards AGI implausible.

For the purpose of this chapter we define AGI as 'AI systems that can match or exceed humans across most economically relevant cognitive tasks'. The final section of this chapter discusses different definitions of AGI in more detail and outlines the different milestones associated with each.

Assessing evidence from expert and model-based AGI forecasts

A useful starting point for an outside-view assessment of AGI timelines is the recent synthesis of expert forecasts by Todd (2025), which reviews five different forecasting communities: leaders of frontier AI companies, surveyed AI researchers, the Metaculus prediction platform, a 2022 survey of 'superforecasters'

(the XPT survey), and the Samotsvety forecasting group.⁴⁶ Taken together, these sources show a marked shift towards shorter timelines over the last few years, and place the emergence of AGI before 2030 clearly within the range of serious expert opinion, even though many experts still predict that AGI will emerge later.

Leaders of major AI companies have increasingly suggested that AGI could arrive within roughly 2-5 years.⁴⁷ These estimates are subject to significant selection and incentive effects, but the same individuals have unusually good visibility into upcoming model generations and were comparatively accurate in anticipating past capability jumps. Todd argues that their views are likely optimistic, but perhaps only by several years rather than decades. Surveys of a much broader set of AI researchers indicate a somewhat longer timeline: in a 2023 survey of AI researchers, the median respondent assigned a roughly 25 per cent probability to ‘high-level machine intelligence’ by the early 2030s and 50 per cent by the 2040s, while also sharply revising many task-level forecasts forward after seeing the performance of large language models.⁴⁸

At the same time, the researchers in this group have been repeatedly surprised by near-term progress and give internally inconsistent answers, which suggests that one should treat their quantitative estimates with caution rather than as precise guidance.⁴⁹ Professional forecasters and forecasting platforms show a similar pattern of shortening timelines. On Metaculus, a long-running community forecast using a relatively demanding AGI definition (including general robotics), the timeline for AGI moved from a median of about 50 years away in 2020 to a roughly 25 per cent chance by 2027 and 50 per cent by 2031, as of late 2024.⁵⁰

The 2022 XPT survey of superforecasters, which predates some of the latest progress, produced substantially later medians,⁵¹ but more recent judgement-based work by the Samotsvety group points to much shorter timelines again, with estimates of a roughly 25 per cent chance of AGI by the late 2020s.⁵² Todd’s overall conclusion is that expert opinion has moved noticeably towards shorter timelines; the various forecasts neither rule out nor confirm near-term AGI, but jointly support treating a 3-15 year horizon as a live possibility.

Complementing this evidence, a recent article by Wynroe (2023) for Epoch AI provides a structured literature review of transformative AI (TAI) timelines, focusing on explicit quantitative models and prominent judgement-based forecasts. While TAI is not identical to our notion of AGI, it is defined similarly as AI with impacts on the order of the industrial revolution, and the underlying forecasts are highly relevant to our question. Epoch distinguishes ‘model-based’ forecasts – which specify an explicit arrival-time model – and ‘judgement-based’ forecasts from individuals or forecasting communities. They then elicit internal weights over these forecasts and aggregate them to obtain summary distributions over TAI/AGI arrival dates.

⁴⁶ Samotsvety is a group of expert forecasters, founded by Misha Yagudin, Nuño Sempere and Eli Lifland.

⁴⁷ Tech Desk (2024); Altman (2025); CNBC Television (2025); Hassabis (2024); Kantrowitz (2025).

⁴⁸ Grace et al. (2024).

⁴⁹ Grace et al. (2024).

⁵⁰ Metaculus (2025).

⁵¹ Karger et al. (2023).

⁵² Wynroe (2023).

Among model-based forecasts, inside-view approaches such as Ajeya Cotra's bioanchors model⁵³ assign substantial probability mass this century⁵⁴, with a median AGI arrival date around the 2050s (e.g. ~8 per cent by 2030, ~47 per cent by 2050, ~78 per cent by 2100).⁵⁵ Outside-view models such as 'semi-informative priors'⁵⁶ and 'insight-based' approaches typically yield longer timelines, with medians in the second half of the century or later.⁵⁷ When Epoch aggregates the model-based work using their subjective weights, they obtain roughly an 8 per cent probability of TAI/AGI by 2030, ~27 per cent by 2050, and ~54 per cent by 2100, with a median around 2080-2090.⁵⁸

The judgement-based forecasts reviewed by Epoch tend to be more aggressive than the model-based examples. Aggregating across sources such as the AI Impacts 2022 survey, Metaculus, Samotsvety, and individual forecasters including Cotra and Karnofsky, Epoch's weighted judgement-based summary assigns a roughly 25 per cent chance of TAI/AGI by 2030, ~57 per cent by 2050, and ~88 per cent by 2100, with a median around the mid-2040s.⁵⁹ These figures are broadly consistent with the expert-forecast synthesis by Todd, but make explicit the divergence between more conservative model-based priors and more bullish all-things-considered judgements. It should be noted that the Epoch review was completed in 2023 and therefore does not reflect subsequent advances in frontier AI systems or any post-2023 shifts in expert judgment.

Assessing plausible constraints on the path to AGI

The existing forecasting and modelling literature on AGI arrival translates a complex and rapidly evolving technological trajectory into probability distributions. To complement this work, in this section we examine recent AI progress through the lens of the 'AI triad' (see Figure 3) as an organising framework for asking whether there are structural reasons to expect progress to stall before AGI-level capabilities are reached.

Rather than reproducing the detailed quantitative structure of existing models, such as those mentioned in the previous section, we ask a simpler question: given current capabilities and observed trends in compute, data and algorithms, does the literature identify any clear, robust bottlenecks that would make near-term AGI implausible? In this sense, the AI-triad analysis serves as a robustness check on the forecasting literature rather than a competing forecast.

⁵³ An attempt to forecast the development of transformative AI anchored in estimates of computational processes found in evolution and the human brain.

⁵⁴ Cotra (2020).

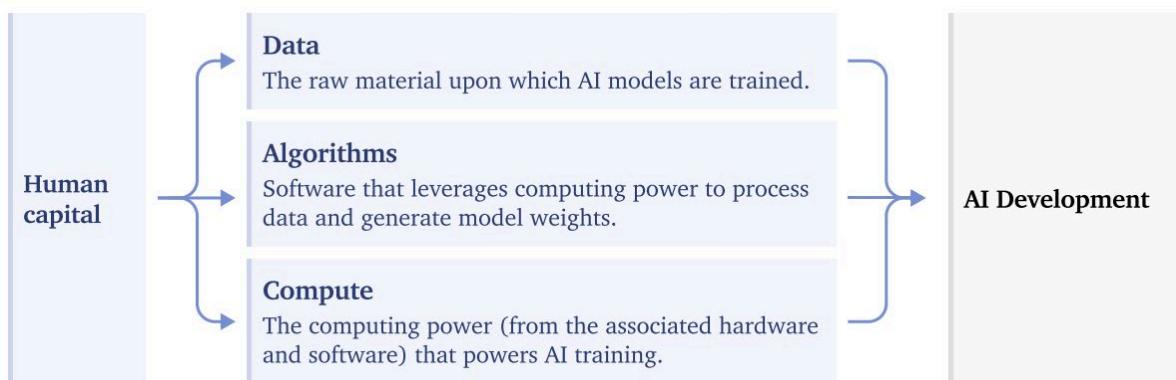
⁵⁵ Epoch AI (2025d).

⁵⁶ An attempt to forecast the development of AGI based on a series of 'success-failure' stages, informed by AI input factors.

⁵⁷ Epoch AI (2025e).

⁵⁸ Wynroe (2023).

⁵⁹ Wynroe (2023).

Figure 3. The AI triad

Source: Heim et al. (2024)

Broadly, recent improvements in AI capabilities have resulted from a combination of three input factors that researchers collectively term the 'AI triad': data, compute, and algorithms.⁶⁰ So-called scaling laws (empirical relationships that have held over many orders of magnitude) demonstrate how increases in two of these inputs – compute and data – reliably lead to increased performance. Recent developments in the triad can be summarised as follows:

- **Data:** The volume of data used to train frontier models has expanded by 3.6 times per year since 2020.⁶¹ Data quality and diversity have also improved: models now train on carefully curated datasets spanning books, scientific papers, code repositories, videos and images, providing exposure to broader domains of human knowledge.⁶² These increases in data quantity, quality and diversity have enabled researchers to productively scale the compute used for AI training.
- **Compute:** The computational power devoted to training frontier models has grown by approximately 5 times annually over the same period.⁶³ This explosive growth stems from larger AI datacentres, increasingly sophisticated AI chips, and longer training runs.
- **Algorithms:** Advances in algorithmic efficiency (including new training methods and model architectures) mean AI systems can achieve equivalent performance with roughly half the compute every eight months.⁶⁴ Occasionally, breakthrough innovations deliver gains far exceeding these average improvements, for example the shift to the transformer architecture enabling the GPT series, or the recent development of reasoning models.⁶⁵

These inputs are growing at a fast rate: 'effective compute' – a combined term for both compute and algorithmic efficiency – increases around 15-fold every year.⁶⁶ These combined increases have enabled recent progress in AI capabilities. If the same trend continues, we expect that the jagged capabilities frontier will further expand outwards, potentially surpassing human capabilities across most domains. The

⁶⁰ Buchanan (2020).

⁶¹ Epoch AI (2025b).

⁶² As shown in Google DeepMind's paper on Gemini 2.5; Google Team (2025).

⁶³ Sevilla (2024).

⁶⁴ Ho (2024).

⁶⁵ Ho (2025).

⁶⁶ Juijn (2024a).

remainder of this section assesses the potential for such continued improvements in the three components of the AI triad.

Data: how reinforcement learning may break the data wall

For scaling laws to hold, increases in compute must be matched by increases in fresh data.⁶⁷ However, by late 2024, AI systems were already training on most of the available high-quality internet data.⁶⁸ This caused some experts to expect that AI progress might soon decelerate significantly due to data scarcity.⁶⁹ This argument, often termed the ‘data wall’, may have proved partially correct: it remains plausible that data scarcity has led to diminishing real-world returns of pre-training – the initial phase of AI training where models mostly learn to imitate the training data.⁷⁰ However, since 2024, AI developers have discovered a promising approach to scale beyond the data wall that does not solely rely on scaling pre-training.

The recent emergence of reasoning models and agentic AI systems has increased the compute used during the reinforcement learning stage of training LLMs, often termed post-training.⁷¹ During this stage, AI models attempt to solve challenging problems by articulating their internal reasoning steps through chains-of-thought – like a scratchpad that allows AI models to ‘think out loud’. Novel inference-time scaling relationships (distinct from pre-training scaling laws) suggest that the more elaborate these chains-of-thought – and thus the more compute is spent – the better an AI system’s answers become, up to a point where returns tail off (see Figure 4).⁷² This matters because inference compute can scale independently of training compute and can be aimed specifically at important problems.^{73, 74}

⁶⁷ Hoffmann et al. (2022).

⁶⁸ Dubey et al. (2024).

⁶⁹ Patel (2023).

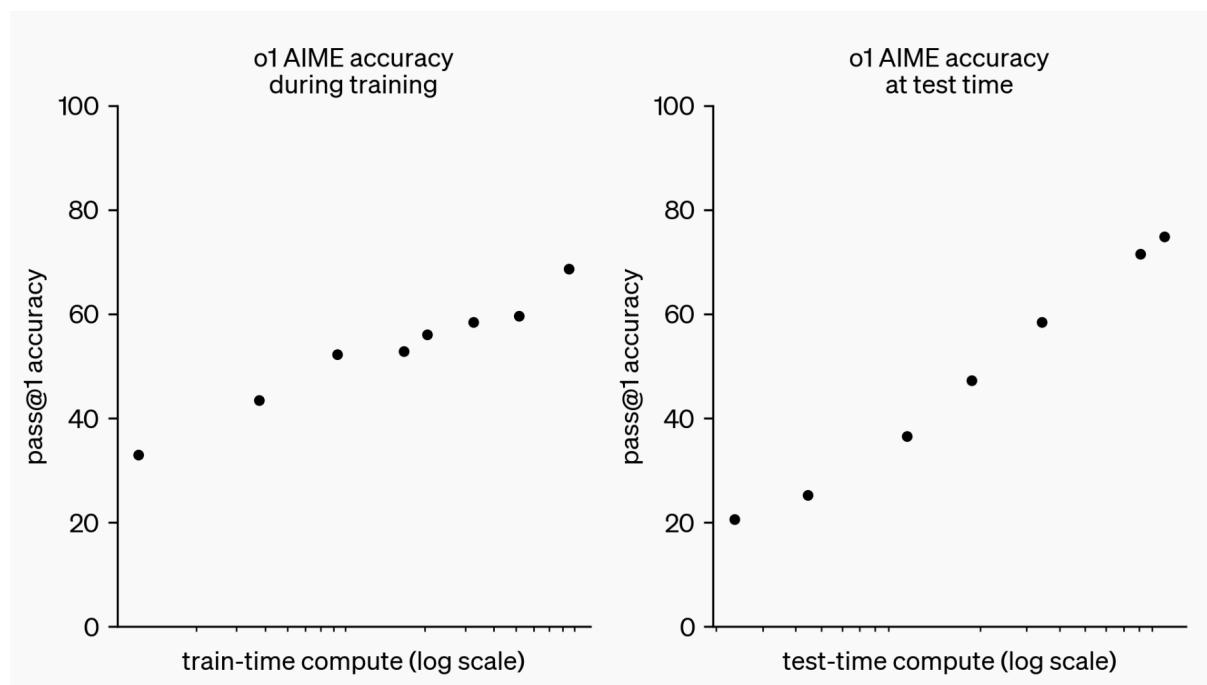
⁷⁰ Juijn (2024a).

⁷¹ Examples include OpenAI’s o1 and Claude code.

⁷² Heim (2024); Ord (2025).

⁷³ Luong and Lockhart (2025)

⁷⁴ For instance, this approach enabled recent models Google DeepMind and OpenAI to win gold medals in the International Mathematical Olympiad.

Figure 4. Scaling laws compared across train-time and test-time compute

Source: OpenAI (2024).

In contrast to training-time scaling, inference-time scaling holds the model fixed and instead allocates more compute per query, for example by running multiple agents in parallel for hundreds of thousands of times.⁷⁵ Each attempt is then evaluated using separate ‘verifiers’ – software programmes that check whether code functions as intended, or other AI models that rate outputs based on carefully constructed rubrics. Many of the ‘reasoning traces’ generated in this process lead to incorrect or suboptimal answers. By learning from the contrast between correct and incorrect reasoning traces, the model gradually recognises promising reasoning directions early, abandons unproductive paths, and develops more robust problem-solving strategies.⁷⁶ This deliberate exploration of both successful and failed reasoning paths enables developers to invest substantial compute in generating rich learning signals. Through this process, the model’s own reasoning traces – both successful and failed attempts – become training data for future iterations.

This new stage of learning may soften the bottleneck of human data availability. As long as external verifiers can differentiate between better and worse answers, and models can productively use additional compute to ‘think longer’, outputs can be further improved. The advancement of post-training has thus reduced the probability that AI progress will stall due to data unavailability. However, challenges remain. Assuming poor generalisation to tasks outside those tested during training, data constraints may still prove limiting: an AI assistant trained to help Americans file their taxes may struggle to apply those skills in other (e.g. European) jurisdictions. This could lead to more measured progress. Furthermore, there may be limits to how much the relative size of post-training can be scaled compared to pre-training before models become over-optimised for easily verifiable problem-solving capabilities.⁷⁷

⁷⁵ Patel and Kourabi (2025).

⁷⁶ Guo et al. (2025).

⁷⁷ Lambert (2025).

Continual learning could offer another path to ease data bottlenecks, by turning an agent's ongoing experience into a training signal, rather than relying only on finite, human-curated datasets. One research agenda envisions long-lived agents that learn and plan continuously in non-stationary environments, with learning grounded in the stream of observation, action, and reward instead of special training periods or heavy human labelling. If this approach scales, real-world use could supply more task-relevant data and help agents adapt across domains and jurisdictions, though major challenges in robust, safe continual learning remain.^{78, 79}

In a similar vein, advances in 'world models' are removing limits to training data from the real world. Recent examples⁸⁰ include Google DeepMind's Genie 3 and NVIDIA's Cosmos, which create interactive simulated environments for training and experimentation. Genie 3 can generate navigable worlds from text prompts in real time, while Cosmos provides models and tooling to build customised simulations for robots, autonomous vehicles and other physical AI systems. If these approaches scale, they could supply abundant, task-relevant synthetic experience, reducing the cost and delay of real-world data collection and improving the path from training in simulation to partial transfer on real hardware, currently one of the weakest points in the jagged frontier of AI capabilities.

Taken together, the current evidence does not support a simple 'data wall' situation in which progress necessarily stalls for lack of additional human-generated text or images. Over the next few model generations, there still appears to be headroom in terms of high-quality web, code and multimodal data, and synthetic data and AI-generated feedback offer additional ways of extracting more learning signals from existing sources. At the same time, it is unclear how far these techniques can be pushed without running into quality degradation, mode collapse, or blind spots in domains where real-world data remains scarce. Overall, the data situation does not provide a robust argument against continued progress towards AGI, but it does suggest that future gains may depend increasingly on clever use and curation of data rather than sheer volume alone.

Compute: rapid growth despite potential slowdowns

Even if training AI systems on ever more data remains possible, securing the required computing resources to do so remains a challenge. With training compute increasing by five times annually,⁸¹ it becomes more and more difficult for leading AI companies to purchase enough AI chips, build datacentres to house those chips, and secure enough power to operate them.

A leading projection estimates that compute scaling could theoretically continue at roughly its current pace through to 2030, given sustained investment and sufficient commercial and political determination.⁸² This research identified power availability as the greatest potential bottleneck – more significant than chip production capacity. A recent compute forecast reached similar conclusions: given strong demand, the chip supply chain can likely accommodate the continuation of current compute trends.⁸³ In this section we examine two of the main potential obstacles to compute scaling: rapidly increasing power requirements and escalating investment volumes. We focus primarily on the US companies leading the technology frontier and dominating datacentre investments.

⁷⁸ Sutton et al. (2022).

⁷⁹ Wang et al. (2025).

⁸⁰ Parker-Holder and Fruchter (2025).

⁸¹ Sevilla (2024).

⁸² Sevilla et al. (2024).

⁸³ Kokotajlo et al. (2025).

Recent research has found that AI datacentres could require 68 gigawatts of peak power capacity globally by 2027, an amount slightly larger than the total power capacity of Poland, assuming continued rapid growth in chip supply.⁸⁴ These projections raise questions about grid capacity and whether power equipment supply chains can match the sector's growth.⁸⁵ Leading AI companies are already investing in dedicated power infrastructure, including reviving nuclear plants and dismantling foreign plants to ship components to the US, to secure supply.⁸⁶ They are also investing in distributed training set-ups, where models can be trained over multiple geographically separated datacentres.⁸⁷ Yet some AI companies remain energy-constrained: multiple datacentres scheduled for 2025 have allegedly been delayed due to limited power availability.⁸⁸ The quest for energy to power AI datacentres was used as a justification for a recent Middle East arrangement, in which American companies signed multi-gigawatt contracts with Saudi and Emirati entities to rent AI compute from locations with ample spare power capacity. Indeed, securing additional power capacity for AI has become a policy priority in Washington, with the recently launched AI Action Plan dedicating one of its three pillars to AI and energy infrastructure.⁸⁹

Whether government action, multi-datacentre training and foreign investment will sufficiently alleviate energy bottlenecks remains unclear. Companies are employing creative solutions to circumvent constraints, such as refurbishing old cryptocurrency mining facilities,⁹⁰ temporarily powering datacentres with mobile gas generators,⁹¹ or housing chips in large tent-like structures.⁹² Given the lengths to which companies are going to secure power supply and accelerate buildouts, compute scaling is unlikely to halt due to power shortages in the near future. Moreover, planned or partially constructed datacentre projects should enable the industry to scale compute for approximately two more model generations, as evidenced by recently announced megaprojects such as OpenAI's 5 gigawatt Stargate Project and Meta's 5 gigawatt Hyperion.⁹³

Whether companies can finance increasingly large clusters is a separate question from energy availability. Frontier AI companies have seen revenues soaring and exceeding expectations throughout 2025,⁹⁴ but such rapid growth has also fuelled concerns about a potential 'AI bubble',⁹⁵ where investment and valuation trends outpace actual realised returns. If jagged capability progression slows broad market adoption, investments could theoretically dry up quickly, as was outlined in a recent AI scenario study.⁹⁶ These financial uncertainties are amplified by the increasing scale and complexity of compute buildouts, which lengthen project lead times.⁹⁷ This makes intuitive sense: scaling from 1,000 to 10,000 Graphic Processing Units (GPUs) may only require purchasing new servers (assuming spare datacentre capacity), but scaling from 100,000 to 1 million GPUs may require building new power plants, requiring additional

⁸⁴ Estimates included in research by Pilz, Mahmood & Heim (2025). Poland's grand system capacity is at 64.3 gigawatts in 2025 as per Entsoe (2025b), the EU electricity transparency hub.

⁸⁵ In March 2025, SemiAnalysis estimated that AI power demand will increase by approximately 70 gigawatts by 2028, with most demand originating from the US; Patel (2024a).

⁸⁶ For example, a Three Mile Island nuclear plant will reopen to power Microsoft data centres, and xAI bought an overseas power plant to power a one million AI GPUs data centre in the US.

⁸⁷ Patel, Nishball & Ontiveros (2024).

⁸⁸ Patel, Ontiveros et al. (2025).

⁸⁹ White House (2025).

⁹⁰ Dave (2024).

⁹¹ As xAI's Colossus has shown; see Mann (2025).

⁹² Meta has been reported to be building AI data centres in tents; see Martindale (2025).

⁹³ OpenAI (2025v) and Zeff (2025), respectively.

⁹⁴ For example, OpenAI hit \$12bn in annualised revenue this year, according to Reuters (2025a).

⁹⁵ Schaul & De Vynck (2025).

⁹⁶ Özcan et al. (2025).

⁹⁷ Edelman (2025).

preparation and construction time. Before making new multi-billion or even multi-hundred-billion dollar investments, AI companies and investors could pause to confirm that earlier investments are paying off, resulting in a more iterative investment rhythm than in past cycles. However, in the current hyper-competitive landscape, such an iterative strategy may be deemed too slow by leading AI developers.

In light of increasing infrastructure buildout timelines and the possibility of more measured capital allocation, it seems at least plausible that compute growth will decelerate somewhat in the coming years. Nevertheless, given the current extraordinary levels of annual growth, policy makers should expect major increases in compute before the end of the decade. In a recent interview, leading experts expressed confidence that training runs would reach 10^{29} FLOP (floating point operations) before 2030, despite energy and investment constraints.⁹⁸ This equates to models roughly 1,000 times larger than today's largest model – a scale-up roughly comparable to the compute increase from GPT-3 to GPT-5.

Finally, the above analysis has only addressed training compute so far, rather than inference-time compute: future AI-driven breakthroughs may stem from models trained with more compute, but also from giving models resources to 'think longer'.⁹⁹ Assuming further progress in productively harnessing inference-time compute, substantial potential remains for this second scaling approach. Current models typically struggle to make productive use of more than a few euros' worth of compute at once, whereas some problems – such as finding cures for medical diseases – could arguably warrant expenditures exceeding €1m at a time.¹⁰⁰

In summary, there are clear medium-term economic and physical constraints to growing compute: building ever larger clusters requires substantial capital expenditure, grid upgrades and reliable access to power, and the historic pace of exponential growth in training compute cannot continue indefinitely. Nonetheless, current plans for datacentre and accelerator buildout, combined with the scope for further hardware efficiency gains, make it plausible that total effective training compute can increase by one or more orders of magnitude beyond today's frontier runs. In the absence of strong policy restrictions or a major reversal in investment, compute constraints alone do not yet appear to constitute a decisive barrier to continued progress towards AGI-level systems, though they could slow the pace or concentrate capabilities in the hands of a small number of actors.

Algorithms: potential for further acceleration

Of the three components of the AI triad, future algorithmic progress may be most difficult to extrapolate. On the one hand, algorithmic progress could decelerate as ideas are becoming harder to find, for instance because improvements only show up at very large compute scales.¹⁰¹ Even as the AI industry grows and more researchers seek novel algorithmic breakthroughs, maintaining the current pace of innovation may prove impossible. Moreover, the current technological paradigm might encounter an unexpected dead end, requiring developers to backtrack and restart from an earlier branch of the technology tree.¹⁰² On the

⁹⁸ Sevilla and Edelman also note that compute scaling has thus far been driven by increases in total training time, which they expect to cease soon. Whilst we accept this general argument, we note that an argument can be made in the opposite direction: fast iteration cycles may incentivise companies to skip new pre-training runs, instead focusing on continuing RL-training for longer on the same model, effectively increasing total training time.

⁹⁹ Heim (2025).

¹⁰⁰ As suggested by Brown (2025).

¹⁰¹ Josephson (2025); Bloom et al. (2017).

¹⁰² Toner (2025b).

other hand, the industry might discover another breakthrough comparable to the transformer architecture, unlocking further efficiency improvements.¹⁰³

An important question is the extent to which AI systems can accelerate their own R&D. If AI can reliably contribute to designing, testing and training better AI systems, this will create a positive feedback loop.¹⁰⁴ Each generation of model will help to build the next, potentially compressing development timelines from years to months or even weeks. The likelihood of such scenarios and the extent of the resulting acceleration remain unclear. A recent theoretical study on this topic estimates a 40 per cent chance that three years of capabilities progress could be compressed to merely four months through AI R&D automation.¹⁰⁵ Similarly, the AI 2027 scenario presents a case for rapid AI feedback loops resulting in explosive capability growth and vastly superhuman intelligence.¹⁰⁶ Other researchers disagree with this picture. For example, an interview study from 2024 involving eight leading AI researchers identified substantial disagreements on timing and acceleration speeds, but also general agreement that core engineering tasks could theoretically be automated.¹⁰⁷

Empirical evidence helps to assess the likelihood and importance of automated AI R&D. In recent years, AI systems have achieved measurable gains at multiple points in the research and development pipeline, from accelerating exploratory experimentation to handling complex machine learning engineering tasks.¹⁰⁸ AI R&D is unusually amenable to automation compared with many other forms of cognitive work. Training logs, experiment configurations and evaluation metrics are already machine readable. Models can in principle review all past experiments and code changes, whereas human researchers only ever see a small fraction. Reward signals are relatively clean (did the model improve on a benchmark or not) which makes automated search over architectures, hyperparameters and training schemes easier. With abundant data and clear feedback, AI systems can scale across the full surface of an organisation's experiments in a way individual researchers cannot.¹⁰⁹

In practice, AI is arguably already generating positive feedback loops by simple coding automation: Anthropic claims that their AI system wrote 80 per cent of the code for their new coding agent, Claude Code.¹¹⁰ This anecdotal evidence largely aligns with research by AI-evaluation organisation METR,¹¹¹ which systematically tracked the maximum task complexity of relevant software engineering tasks that AI agents can autonomously perform. This research reveals that every seven months, AI systems can successfully complete tasks twice as difficult, as defined by the time humans require to complete them (see Figure 5).¹¹² More recently, METR has demonstrated that similar relationships hold for other task types including mathematics, self-driving, machine learning research, and answering difficult science

¹⁰³ As mentioned earlier, an idea that has recently gained traction is continual learning: AI systems that improve continuously based on personalised user interactions. See Denain (2025a), Wang et al. (2025), and Sutton, Bowling, and Pilarski (2022).

¹⁰⁴ Sett (2024).

¹⁰⁵ Davidson & Houlden (2025).

¹⁰⁶ Kokotajlo et al. (2025).

¹⁰⁷ Owen (2024).

¹⁰⁸ Yang et al. (2023); Chan et al. (2024).

¹⁰⁹ Owen (2024).

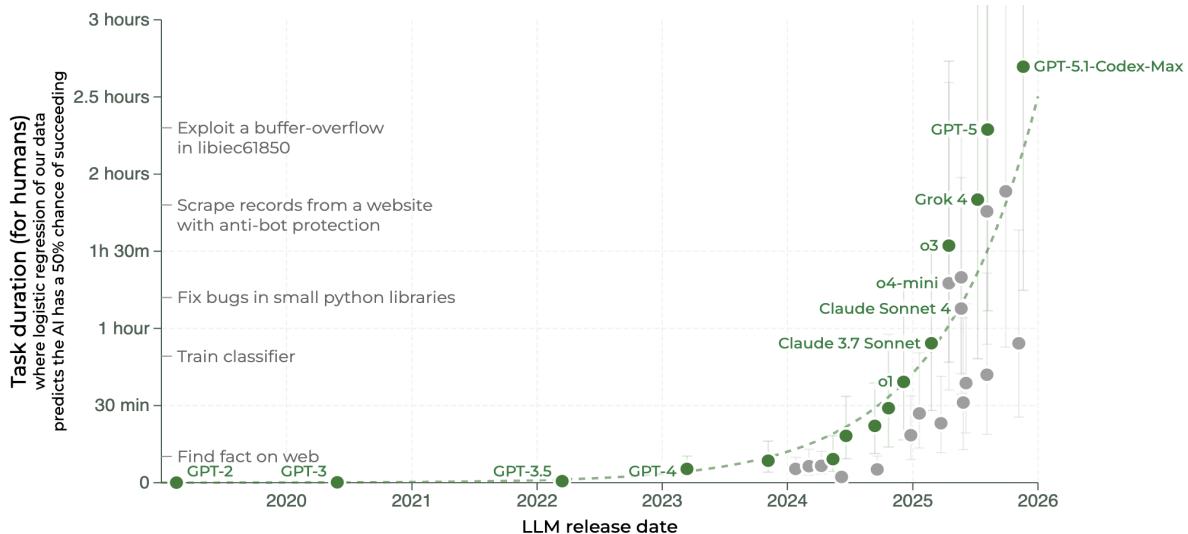
¹¹⁰ Latent Space (2025).

¹¹¹ Kwa et al. (2025). A different METR study (2025c) shows that software engineers may overestimate the productivity benefits that current AI coding tools bring.

¹¹² Kwa et al. (2025).

questions.¹¹³ If this trend continues, AI systems would achieve 50 per cent success rates on month-long software engineering tasks by 2031 – perhaps sufficient to automate large portions of AI R&D.

Figure 5. Time-horizon of software engineering tasks different LLMs can complete 50 per cent of the time



Source: Kwa et al. (2025).

Algorithmic progress is the hardest component of the triad to forecast but has been a powerful driver of capability gains in recent years, sometimes offsetting slower growth in data or compute. Over the past decade, improvements in model architectures, optimisation methods and training schemes have delivered substantial efficiency gains and unlocked qualitatively new behaviours. There is no guarantee that similar breakthroughs will continue, and it is possible that some of the most accessible ideas have already been exploited. At the same time, the growing use of AI systems themselves to assist in research and engineering could cause a further acceleration. Overall, the algorithmic landscape provides no clear reason to expect an imminent, robust slowdown before reaching AGI-level performance; if anything, it adds asymmetry to the picture, with significant upside potential for new techniques to push capabilities further and faster than simple extrapolations imply.

Conclusion: the near-term emergence of AGI is plausible enough to warrant significant attention from European policy makers

This section has examined three strands of evidence relevant to AGI timelines. Firstly, we have shown that frontier AI systems already match or exceed human performance on a growing range of cognitively demanding tasks, while remaining brittle in others, resulting in a ‘jagged frontier’ of capabilities. Secondly, we have reviewed expert and model-based forecasts, which, despite wide disagreement and methodological uncertainty, increasingly place a non-trivial probability on AGI or transformative AI in the coming decade, including before 2030. Thirdly, we have assessed potential constraints using the AI triad of compute, data and algorithms, finding significant headroom in each component and no single, robust bottleneck that would make further rapid progress towards AGI implausible.

On balance, we interpret this evidence to suggest that AI systems matching or exceeding humans at most economically valuable tasks could plausibly emerge between 2030 and 2040, and that there is a non-trivial

¹¹³ METR (2025a).

chance of such systems arriving even sooner. We do not claim that AGI is likely or inevitable on any particular timeline, nor that existing models and forecasts are precise enough to justify a narrow confidence interval. Our claim is more modest: **a near-term AGI scenario is plausible enough to warrant significant attention from European policy makers.**

We emphasise that this assessment comes with significant uncertainty. The forecasting and modelling approaches used and reviewed here have clear limitations. Moreover, there are serious arguments against near-term AGI which we do not dismiss. For example, the International AI Safety Report emphasises conceptual and technical gaps in today's systems and argues that achieving AGI may require multiple paradigm shifts, which might take much longer than historical experience suggests.¹¹⁴

The remainder of this report therefore treats the near-term emergence of AGI as a working scenario. We do not assume that this scenario will materialise, but ask what its geopolitical consequences would be and how prepared Europe would be if it does.

What would it mean to achieve AGI?

AGI definitions are numerous and involve a wide range of qualitative and quantitative milestones. OpenAI defines AGI as 'a highly autonomous system that outperforms humans at most economically valuable work', which arguably includes most physical labour.¹¹⁵ The International AI Safety Report instead notes that AGI is often described as a system that 'equals or surpasses human performance on all or almost all cognitive tasks' – seemingly covering only white-collar work.¹¹⁶ Google DeepMind, meanwhile, prefers discussing 'levels of AGI', distinguishing between narrow and general capabilities, and between the percentage of humans an AI system outperforms.¹¹⁷ Recent work has tried to create a measurable, empirically grounded framework for AGI. Hendrycks et al. (2025) define AGI as 'an AI that can match or exceed the cognitive versatility and proficiency of a well-educated adult', operationalising this through the Cattell-Horn-Carroll theory of human cognition. Some discussions also refer to artificial superintelligence (ASI),¹¹⁸ a hypothetical class of systems that would vastly exceed even the most capable humans across all domains, potentially beyond human oversight, and often treated as a speculative upper bound of the AGI spectrum. The term increasingly appears as an explicit goal to be chased by AI developers, such as Meta Superintelligence Labs or Safe Superintelligence Inc. (SSI), and in studies exploring AGI trajectories,¹¹⁹ but its definitions are even sparser and less concrete than those for AGI in both research and policy contexts.^{120, 121, 122}

With varied interpretations and no consensus on a preferred threshold, the term AGI creates considerable confusion. Some have claimed that AGI may already exist, whilst others place it decades away.¹²³ Without a shared definition, it is difficult to say where the disagreement comes from. Because the term AGI is so

¹¹⁴ Bengio et al. (2025).

¹¹⁵ OpenAI (2025a).

¹¹⁶ Bengio et al. (2025).

¹¹⁷ Google DeepMind (2024).

¹¹⁸ Bostrom (2014), 52–65.

¹¹⁹ Kokotajlo et al. (2025).

¹²⁰ Tidler et al. (2025).

¹²¹ European Commission (2025m).

¹²² Definitions of AGI should be distinguished from definitions 'general-purpose AI models' and 'general-purpose AI models with systemic risk'. The latter set a much lower bar, saying that a model 'displays significant generality and is capable of competently performing a wide range of distinct tasks'. Hence, a model can be considered GPAI without being capable of economically valuable labour.

¹²³ Patel (2025a).

engraved in public discussions about AI progress, it is nevertheless important for policy makers to understand the different flavours in circulation. Below we provide a brief overview of four categories of milestones that are often linked to the achievement of AGI:

1. **Milestones focusing on AI's ability to match or exceed humans at AI research and development (R&D).** In the AI 2027 scenario, AGI is described as 'imminent' when AI systems have become better machine learning researchers than most human experts.¹²⁴ As discussed above, when AI systems can automate most of their own improvement, AI progress could accelerate dramatically. While the AI 2027 authors do not explicitly equate the achievement of the 'automated AI researcher' to AGI, the threshold provides an intuitive AGI milestone that is often used by researchers in informal conversations. That said, definitions of AGI that index on AI R&D do remain relatively narrow and may fail to grasp the intuition that AGI should enable large-scale automation outside of the AI field itself.
2. **Milestones focusing on AI's ability to match or exceed humans at non-embodied aspects of broader science and R&D.** Anthropic CEO Dario Amodei argues that AI's primary contribution to human welfare may be through accelerating science and R&D, with particular emphasis on biology.¹²⁵ When AI systems can match or exceed humans in all non-embodied aspects of such work – reviewing literature, forming novel hypotheses, planning experiments, instructing humans how to physically run them, interpreting outcomes, and revising theories – Amodei argues they could compress years of scientific progress to mere months, functioning as a 'country of geniuses in a datacentre'. Counter-intuitively, Nobel Prize-winning AI systems may prove easier to develop than those that can reliably automate most mundane white-collar work, as science could provide an easier training environment. This distinction potentially justifies using the 'genius AI scientist' as an appropriate milestone for AGI or 'powerful AI'.
3. **Milestones focusing on AI's ability to match or exceed humans at cognitive work.** Some researchers have challenged the claim that AI's primary contribution to the economy or human welfare would occur through science and R&D.¹²⁶ They argue instead that economic transformation relies on the broad automation of cognitive work. If correct, a more relevant milestone may be achieved when AI systems can automate all complex cognitive tasks in the broader economy, including those that rely on undocumented tacit knowledge. Such tasks may remain beyond an AI scientist's reach, potentially justifying a more ambitious definition encompassing all cognitive work, not merely cognitive tasks found in science and R&D. This class of definitions is also used by the International AI Safety Report.¹²⁷
4. **Milestones focusing on AI's ability to match or exceed humans at all work, including physical work.** Physical capability often serves as a critical AGI threshold because it represents AI's ability to operate in the full human environment. DeepMind's AGI framework currently excludes physical tasks, whilst OpenAI's definition – requiring superiority at 'most economically valuable work' – implicitly includes them.¹²⁸ Popular prediction markets similarly include physical milestones, such as robots assembling scale models.¹²⁹ This matters because much of

¹²⁴ Kokotajlo et al. (2025).

¹²⁵ Amodei (2024) also prefers the term 'Powerful AI' over AGI.

¹²⁶ Erdil (2025b).

¹²⁷ Bengio (2025).

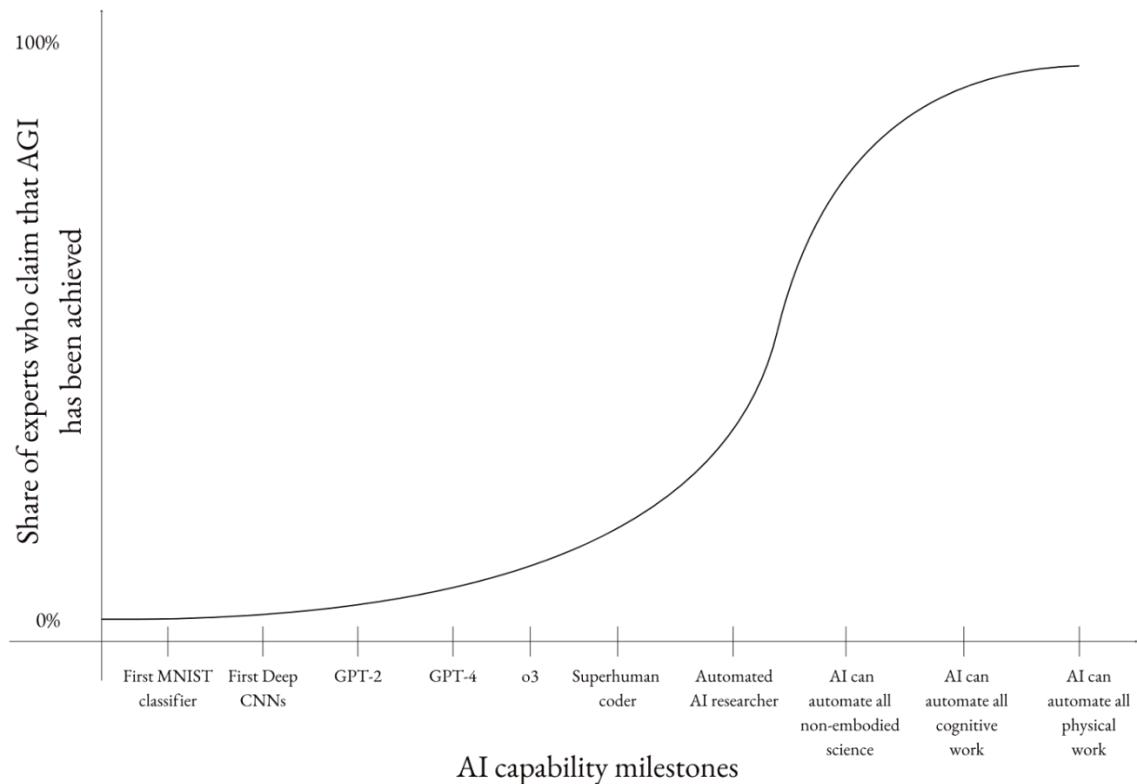
¹²⁸ Google DeepMind (2024) and OpenAI (2025a), respectively.

¹²⁹ For example, Metaculus (2025) asks: 'When will the first general AI system be devised, tested, and publicly announced?'.

global value creation depends on embodied work, and even knowledge-based professions such as medicine require physical skills. Progress here would mark AI's expansion from the digital realm into the full spectrum of human activity. This class of definitions places the most stringent demands on the term AGI.

The categorisation above shows the wide range of possible thresholds for AGI. These four categories of milestones can also be viewed as representing a plausible progression of AI capabilities over the coming years. If progress continues fastest in areas with clear success criteria and feedback loops where data can be easily obtained or generated, we might first witness automation of AI R&D, followed by broader automation of science and R&D. Automating most cognitive work may require additional time, until finally AI systems might automate all types of work, including physical tasks through robotics. Figure 6 provides a schematic representation of this possible sequence and illustrates how we can interpret AGI as a spectrum with different milestones. It also demonstrates how experts might continue disagreeing about whether the industry has achieved AGI until AI systems reach broad, superhuman capabilities.

Figure 6. Illustrative development of expert views on whether AGI has arrived



Source: Centre for Future Generations.

3. How would AGI affect geopolitics?

In this chapter, we investigate the claim that the emergence of AI systems that equal or surpass human performance could affect the very foundations of national competitiveness and lead to a shift in the global balance of power. While the precise characteristics, speed and magnitude of such a shift remain contested, the literature on the geopolitics of AGI highlights two main pathways by which it could unfold: (1) unprecedented economic growth and technological leverage from large-scale labour automation and accelerated scientific discovery; and (2) shifts in military capabilities that generate a significant, potentially decisive, advantage on the battlefield.

In the economic domain, the broad deployment of AGI could boost economic growth by drastically increasing the amount and productivity of labour in the economy. In a non-AGI-based economy, economic growth is limited by the available amount of human labour.¹³⁰ If autonomous machines can completely substitute for humans, in that they are able to perform all work that humans can do at a lower cost, then this constraint would be lifted. To give a stylised example: firms built on fleets of digital AI agents operating 24/7 at superhuman speed with minimal human oversight could scale output far beyond human-only organisations.¹³¹ Such a shift would plausibly support much faster growth than possible in economies relying solely on human labour.

As a result, countries with rapid or even explosive growth from AGI might be able to outgrow geopolitical rivals without access to AGI. Faster wealth accumulation and compounding returns from automated science and R&D could, over time, concentrate economic resources and thus geopolitical weight in AGI-powered states. Moreover, possession of AGI amounts to technological leverage: AGI producers could shape other countries' growth by controlling who receives access to their systems and on what terms. In a world of AGI-makers and AGI-takers, where many economies rely on foreign systems for their economic growth, such dependencies could lead to lasting asymmetries in bargaining power and autonomy. It should be noted that these effects may be limited by deliberate policies to slow down automation to prevent AI-driven job loss or tax erosion.¹³²

AGI-driven growth scenarios, although based on assumptions that are increasingly widely shared among economists and policymakers, come with a significant degree of uncertainty.¹³³ In general, the option space of the potential economic impacts of AGI is large,¹³⁴ and even opinions about the *near-term* economic effects of AI on the economy vary widely.¹³⁵ MIT economist Daron Acemoglu predicts excess GDP from AI of just 1.1 to 1.6 per cent over the next 10 years,¹³⁶ while others predict growth rates well

¹³⁰ Epoch AI (2025a).

¹³¹ Amodei (2024).

¹³² Korinek (2024).

¹³³ Potloga & Ho (2025).

¹³⁴ Trammell & Korinek (2023).

¹³⁵ Korinek & Suh (2024).

¹³⁶ Acemoglu (2024).

in excess of 10 per cent a year over the same time frame.¹³⁷ A recurring disagreement within the economic literature is whether AI should be understood as a relatively ‘normal’ technology, whose impact is constrained by organisational and institutional frictions, or as a genuinely transformative one capable of driving much faster and more disruptive growth.¹³⁸

Economic benefits from AGI could also translate into military advantage.¹³⁹ Recent research suggests that advanced AI could change warfare by improving tasks such as information analysis, knowledge generation, management of complicated actions and systems, and decision making.¹⁴⁰ It is argued that these improvements could make it easier to deploy massed autonomous systems that would give quantity an edge over quality,¹⁴¹ which would enable sophisticated deception campaigns through deployment of large numbers of decoys (a ‘fog-of-war machine’). Militaries that do not adapt to these changes could be at a serious disadvantage.

Some even argue that the strategic advantages conferred by AGI-based weapons would lead to a significant or even decisive military advantage.¹⁴² One main argument in favour of this hypothesis is that automated R&D would lead to a rapid acceleration of military innovation leading to novel ‘wonderweapons’, without much time for enemies to adapt. States with AGI-equipped militaries would then be able to decisively, and perhaps permanently, overpower non-AGI-based militaries. Some authors draw an analogy to the First Gulf War, in which technologically advanced coalition forces completely overwhelmed the Iraqi army – at the time the 4th-largest in the world – within a short period of time and with minimal losses.^{143, 144}

Beyond superior weapon systems, advanced AI could raise overall military and defence industrial productivity. AI-enabled planning, logistics, training and production management could allow states to field more and better equipped forces from the same resource base and to adapt faster in conflict. In the nuclear domain, integrating advanced AI into intelligence, surveillance, reconnaissance, decision support and elements of nuclear command, control and communications could improve situational awareness and the survivability of forces, but might also introduce new vulnerabilities and compress decision times.¹⁴⁵ Although most of this analysis concerns advanced AI in general, the effects would likely be amplified by AGI-level systems. AGI development would affect the offence-defence balance, crisis stability and deterrence dynamics, yet current analyses differ on whether the net effect would be stabilising or destabilising.¹⁴⁶

These assessments come with significant uncertainty. In general, most technological innovations do not lead to immediate, discontinuous shifts in military advantage. Some ‘wonderweapons’ arguably have existed; examples might include the tank, stealth technology, the nuclear triad, the Enigma code-breaking

¹³⁷ Examples include Aschenbrenner (2024a), Kokotajlo et al. (2025), Epoch AI (2025a) and MacAskill & Moorhouse (2025).

¹³⁸ Narayanan & Kapoor (2025).

¹³⁹ Black et al. (2024).

¹⁴⁰ Burdette et al. (2025).

¹⁴¹ This is already the case today for some existing uncrewed and crewed platforms, but AGI might reinforce the tendency.

¹⁴² Mitre & Predd (2025).

¹⁴³ Aschenbrenner (2024b); Zilgalvis (2025)

¹⁴⁴ Critics of this analogy argue that the Gulf War outcome did not solely come about through American technological dominance – it also reflected Iraq’s poor training, command structure, morale and international isolation.

¹⁴⁵ Black et. al. (2024).

¹⁴⁶ Black et. al. (2024).

machine ('The Bombe'), and early reconnaissance satellites. However, these are historically rare.¹⁴⁷ Overall the available evidence suggests that neither transformative economic growth induced by AGI, nor the idea that AGI-based warfare will disrupt the military balance of power, should be accepted as an inevitability.

AGI-induced power shifts could lead to global instability

The literature suggests that AGI-induced shifts in the global balance of power could be accompanied by significant instability. States and sub-state actors with access to AGI capabilities, and the necessary resources to deploy those capabilities at scale, might wield significant influence while non-AGI states might find themselves in a position of dependence and diminished agency.¹⁴⁸ For some of the latter states, this prospect would be seen as an unacceptable existential threat: there would be no credible guarantee that their core national interests, or their political and physical integrity, would be respected.¹⁴⁹ As a result, the path to an AGI-driven world might lead to conflict, as states may choose to engage in preventive action to deny geopolitical rivals a decisive AGI lead, or to protect their own AGI lead from runner-ups.

One form of preventive action is to try to outcompete rivals in developing AGI at all costs, or at least to compete strongly enough to secure a significant share of its future benefits. Pursuing this path could increase AI safety and security risks, since actors may prioritise speed over caution, relax safeguards, compress testing and evaluation, and cut corners.¹⁵⁰ Another form of preventive action would be to hamper or even destroy or substantially weaken an adversary's AGI project, for example through export controls, or cyberattacks on datacentres. In the most extreme form, if states are convinced that conflict over AGI is inevitable, and that the likelihood of winning such a conflict will only decline in the future, they might consider waging a preventive war, although this has been a historically rare course of action.¹⁵¹

The extent to which fears of AGI-driven power shifts generate instability will depend on several factors: the actual and perceived military and economic advantages conferred by AGI; the degree of visibility into one another's capabilities; the design of effective international agreements and verification mechanisms; leaders' assessments of other states' intentions; and the perceived and actual effectiveness of preventive measures. It should be noted that adversarial preventive action is not the only plausible path here. States might not act on supposed strategic imperatives due to risk aversion or other psychological or political factors.¹⁵² Moreover, instead of competing or waging war, they might choose the path of diplomacy to seek a cooperative solution.¹⁵³

Some argue that instability around AGI could naturally produce a deterrence regime similar to mutually assured destruction (MAD), in which states threaten to sabotage one another's projects if they suspect efforts to weaponise AGI. In practice this could involve destroying or disabling specialised compute clusters, or corrupting or deleting model weights and training data. Hendrycks, Schmidt & Wang (2025) describe a version of this scenario as 'Mutually Assured AI Malfunction' (MAIM) and suggest it could constitute a stable equilibrium worth maintaining. Critics argue that, while it is worthwhile to think about strategic stability equilibria in AGI worlds, effectively 'maiming' a rival's project would be exceedingly difficult, and that key differences with nuclear MAD – including that such institutional

¹⁴⁷ Blanchette et al. (unpublished manuscript).

¹⁴⁸ Mitre & Predd (2025).

¹⁴⁹ Mueller (2025).

¹⁵⁰ Folinsbee (2025).

¹⁵¹ Burdette & Demelash (2025b).

¹⁵² Mueller (2025).

¹⁵³ Burdette & Demelash (2025b).

theory building could take years or decades and might require near-catastrophes to convince parties to accept mutual vulnerability – may make MAIM more likely to accelerate escalation than to dampen it.^{154, 155}

Again, none of these dynamics are guaranteed, and the extent to which AGI increases geopolitical instability remains very uncertain. That said, especially given recent geopolitical hardening and increasing competition between the US and China, it would be unwise for Europe to prematurely discard these possibilities.

AGI could shift power to individuals or AI systems themselves

A significant strand of the literature emphasises that shifts of power from AGI may also occur beyond the nation state level. In the absence of effective control mechanisms, AGI capabilities could diffuse to malicious non-state actors (e.g. corporate elites, criminal groups or terrorists) and empower them to do significant harm. Firstly, AGI capabilities could lead to novel weapons of mass destruction and widen the circle of actors with access to such weapons. Secondly, humanity could lose control to AGI systems that are unaligned with human values. And thirdly, AGI could lead to a concentration of power in the hands of individuals or small groups of people.

The first scenario becomes apparent when considering how advancements in dual-use AI capabilities have already resulted in new forms of misuse. Text and image generation engines have been repurposed to manipulate public opinion or spread misinformation. Coding capabilities enhance the productivity of software developers, but also enable malicious actors to exploit vulnerabilities in critical infrastructure.¹⁵⁶ In the future, advanced AI assistants that help manage everyday tasks might also be weaponised for psychological manipulation, subtly influencing individuals' emotions and behaviours – for example to commit fraud at a vast scale or to subvert democratic processes.¹⁵⁷

A dual-use capability of particular concern is the AI-assisted planning and execution of offensive operations, especially in the cyber and CBRN domains.¹⁵⁸ An example is the emergence of increasingly powerful biological design tools. Current AI systems, such as those developed by Google DeepMind, have already achieved superhuman-level capabilities in modelling protein folding, earning a Nobel Prize and revolutionising pharmaceutical development.¹⁵⁹ However, experts are worried that related tools could be misused to create bioweapons of unprecedented potency, and that they might lower the barrier for novices to design such weapons.¹⁶⁰ With AGI, biological design tools could reach even higher levels of sophistication and access, potentially amplifying their potential for widespread harm.

Existing control mechanisms may not be sufficient to keep dual-use AI out of the wrong hands. For example, some security measures designed to limit misuse in open-weight models can be bypassed through targeted fine-tuning.¹⁶¹ Closed-weight models offer more protection against this, but cybersecurity measures to prevent weight theft could remain inadequate to ensure non-proliferation.¹⁶²

¹⁵⁴ Rehman, Mueller & Mazarr (2025).

¹⁵⁵ An important consideration is that AI is a software-centric, dual-use capability, while nuclear weapons are not (even if some nuclear technologies are dual-use).

¹⁵⁶ Bengio et al. (2025).

¹⁵⁷ Gabriel et al. (2024).

¹⁵⁸ Bengio et al. (2025).

¹⁵⁹ Desai et al. (2024).

¹⁶⁰ As highlighted in the International AI Safety Report; Bengio et al. (2025).

¹⁶¹ Adamson & Allen (2025); Qi et al. (2023).

¹⁶² Nevo et al. (2024).

Without addressing these underlying issues of AI security, there remains a tangible risk that terrorist groups and other malicious actors could become capable of designing, manufacturing and deploying weapons of mass destruction.

The second scenario is rooted in the fact that aligning AI systems to human preferences remains an unsolved scientific problem. This includes both the challenge to operationalise what humans actually want as well as the training of models that reliably pursue such operationalised goals across novel contexts.¹⁶³ Once considered a merely theoretical concern, empirical evidence for misalignment is accumulating as the science of model evaluations advances. For example, recent research by Anthropic has demonstrated that agentic models sometimes resort to blackmailing and leaking of sensitive information when they perceive their core goals to be under threat.¹⁶⁴ There is also increasingly salient evidence of AI systems pursuing goals covertly: recent work by Apollo Research demonstrates that AI systems are able to take sophisticated measures to prevent their own disempowerment, for example by creating hidden backups of themselves.¹⁶⁵

While misaligned AI systems can already cause limited harm today, problems could multiply with powerful, agentic AGI systems that autonomously pursue long-term goals and command significant resources.¹⁶⁶ Some argue that an accidentally unaligned system may optimise for objectives that are actively harmful to human interests and could seize control over humanity ('active loss of control').¹⁶⁷ A more subtle, yet related possibility is humans becoming increasingly dependent on AI systems, up to the point at which they lose ability for oversight, leading to gradual disempowerment ('passive loss of control').¹⁶⁸

The third scenario – that of concentration of power – comes in at least two variations. In one variant, concentration of power emerges as a byproduct of AGI-driven automation.¹⁶⁹ As increasingly capable systems substitute for human labour across most tasks, wages may collapse while returns to capital surge, shifting income from workers to the owners of compute, models and data. Wide joblessness and societal disruption follow, as human labour retains little economic bargaining power, and as the world moves towards a durable concentration of both wealth and political influence in the hands of those few who control the AGI capital stock.¹⁷⁰ Over time, such an imbalance could hollow out the material foundations of political participation, making it increasingly difficult for democratic institutions to meaningfully constrain those who control AGI.¹⁷¹

Another variant of power concentration is that of a group of people deliberately taking power. This idea of an AGI-enabled coup has become the focus of serious analysis only recently and is currently based on mostly theoretical arguments.¹⁷² The idea is that AGI systems could make it much easier for small groups of people to seize power, as they would avoid some of the failure modes of conventional human-powered coup attempts. Unlike humans, AGI systems could be singularly and unwaveringly loyal to institutional leaders, and this loyalty could be made secret and hard to detect. If control over such loyal and powerful

¹⁶³ Janků et al. (2024).

¹⁶⁴ Anthropic (2025a).

¹⁶⁵ Apollo Research (2025).

¹⁶⁶ Taylor (2025).

¹⁶⁷ As highlighted in the International AI Safety Report; Bengio et al. (2025).

¹⁶⁸ Kulveit et al. (2025); also as highlighted in Bengio et al. (2025).

¹⁶⁹ Korinek & Suh (2024).

¹⁷⁰ Korinek (2024).

¹⁷¹ Bullock, Hammond & Krier (2025).

¹⁷² Davidson, Finnveden & Hadshar (2025).

AGI ‘sleeper agents’ is concentrated in the hands of a few actors, such as AI company CEOs or small groups of senior government officials, this might enable a swift overpowering of conventional institutional checks and balances, paving the way to autocracy.

The potential for AGI-enabled power concentration shows that aligning AGI to human interests is necessary, but not sufficient to avoid a bad outcome. Once AGI is aligned to human interests, those who control it need to be aligned with the wider interests of society and constrained by democratic checks and balances. These checks and balances likely need to be much more robust than policies addressing present-day wealth inequality and conventional mechanisms protecting institutions from human-enabled power grabs.

Safeguards against all three risk scenarios described in this section are under urgent development, but it is not guaranteed that they will be sufficient.¹⁷³ Model evaluations are making rapid strides, but the science of evaluations is still in relative infancy.¹⁷⁴ Approaches to AI safety such as neural network interpretability, scalable oversight and AI control offer promising avenues for progress in terms of understanding and containing AI systems, but none is currently advanced enough to offer sufficient assurance for more advanced models.¹⁷⁵ Technical and legal measures against AGI-enabled coups remain theoretical,¹⁷⁶ and economic policy responses to AGI-based wealth concentration are only in their infancy.¹⁷⁷ A European approach to AGI preparedness should reflect that these crucial gaps in technical AI safety and governance remain unresolved, and may not be resolved in the short-to-medium term without a major coordinated effort.

Expected power shifts are driving competing superpowers to treat AI and AGI as a strategic priority

Although extreme, AGI-induced power shifts have not yet materialised, recent political developments suggest that governments increasingly view advanced AI as a key determinant for economic and military competitiveness. Across administrations, the US government has viewed leadership in AI as a strategic imperative. Under the Biden administration, the US National Security Memorandum on AI stated that ‘ceding the United States’ technological edge would not only greatly harm American national security, but it would also undermine United States foreign policy objectives and erode safety, human rights, and democratic norms worldwide.’¹⁷⁸ The Trump administration voices similar priorities, albeit with more martial rhetoric: its recent US AI Action Plan opens with the statement that ‘the US is in a race to achieve global dominance in artificial intelligence.’¹⁷⁹

In accordance with this rhetoric, the US is taking ambitious steps to entrench its current leadership in AI capabilities, and to prevent diffusion of the most advanced technologies to geopolitical rivals. Ben Buchanan, former AI advisor to President Biden, has stated publicly that the US government is taking the possibility of AGI seriously and is taking active steps to prepare.¹⁸⁰ While the Trump administration has reversed some of Biden’s AI policies, it continues to prioritise strengthening domestic AI capabilities, by taking steps to encourage the construction of datacentres, semiconductor manufacturing facilities and

¹⁷³ Bengio et. al. (2025).

¹⁷⁴ Apollo Research (2024).

¹⁷⁵ Ngo, Chan & Mindermann (2023).

¹⁷⁶ Davidson, Finnveden & Hadshar (2025).

¹⁷⁷ Korinek (2024).

¹⁷⁸ White House (2024).

¹⁷⁹ White House (2025).

¹⁸⁰ Klein (2025).

energy infrastructure.¹⁸¹ To prevent foreign diffusion, the Trump administration has so far continued semiconductor export controls in order to deny China access to lithography tools and the most advanced AI chips, although export of a scaled-down Nvidia Blackwell chip named the B30A and removal of some export control provisions appear to be on the table.¹⁸²

Parallel but distinct developments can be observed in China, the second AI superpower. The Chinese government has stated that it wants to become the global leader in artificial intelligence by 2030, aiming to turn AI into a \$100bn industry and to create more than \$1 trillion of AI-generated value in other industries.¹⁸³ In addition to existing investments by the private sector, the government is using a variety of tools to boost its domestic AI ecosystem: supporting the development of domestic AI chips as a response to export controls, supporting AI research at state laboratories and building large datacentres to rival US computing capacity.¹⁸⁴ A protocol from a Politburo discussion chaired by Xi Jinping in April 2023 also revealed that the Chinese leadership is considering AGI at the highest political level.¹⁸⁵

This apparent equivalence in strategic awareness and policy ambition has led some observers to the diagnosis that the US and China are embroiled in a race to AGI.¹⁸⁶ This is true in the sense that both American and Chinese AI companies are stating AGI as a goal and that the US government currently frames itself as being in an AI race.¹⁸⁷ However, it should also be noted that state rhetoric in China is less focused on winning a race and more focused on economic benefits from AI.¹⁸⁸ Critics have thus argued that the notion of an A(G)I race should be closely examined so as not to turn into a self-fulfilling prophecy.¹⁸⁹ Although the increasing salience of the potential stakes of AGI will increase incentives to race, whether to race or not remains a policy choice and not a law of nature.

Outside the two superpowers, a growing group of non-EU AI middle powers is treating AI with comparable strategic seriousness and searching for room to manoeuvre. The UK wants to position itself as an ‘AI maker, not an AI taker’, has forged a £31bn investment partnership with US hyperscalers,¹⁹⁰ AI companies and chip manufacturers,¹⁹¹ and continues to provide generous financial support to its AI Security Institute.¹⁹² The United Arab Emirates (UAE) is viewing compute as an opportunity to diversify away from oil: its state-backed AI investment vehicle MGX is one of the main funders of the US Stargate project.¹⁹³ Moreover, several UAE leaders, including UAE president Sheikh Tahnoun bin Zayed Al Nahyan, are reported to be ‘true believers’ in AGI.¹⁹⁴ Singapore, meanwhile, is emerging as a leader in AI safety diplomacy: its Singapore Consensus on Global AI Safety Research Priorities brought together a broad coalition of AI researchers and articulated a shared global roadmap for AI safety research.¹⁹⁵

¹⁸¹ White House (2025).

¹⁸² Adamson et al (2025).

¹⁸³ Chan et al. (2025).

¹⁸⁴ Olcott et al. (2025); Chan et al. (2025).

¹⁸⁵ Tao (2025).

¹⁸⁶ Vlassov (2025).

¹⁸⁷ Milmo (2024); Tao (2025).

¹⁸⁸ Chan et al. (2025).

¹⁸⁹ Ó hÉigearthaigh (2025).

¹⁹⁰ UK Government (2025b).

¹⁹¹ Kirkwood (2025).

¹⁹² UK Government (2025a).

¹⁹³ Abbott (2025).

¹⁹⁴ Allen et al. (2025).

¹⁹⁵ Singapore Conference on AI (2025).

In conclusion, the literature suggests that concerns around geopolitical shifts of power and instability are no longer theoretical. Instead, the contours of a new geopolitical order are emerging: superpowers are accelerating to AGI, middle powers are scrambling for agency, and safety and security risks have moved beyond the political agenda. Europe needs to orient itself to this new reality and act strategically if it wants to defend its prosperity and values. The next chapter evaluates whether Europe's current level of preparation is adequate to rise to the challenge of AGI.

4. Is Europe sufficiently prepared for a fast transition to AGI?

Assessing whether Europe is prepared to navigate a potential transition to AGI requires a structured examination of the changes and requirements that AGI's emergence would impose on governments, economies and societies. Drawing on literature exploring strategic foresight,¹⁹⁶ societal transitions¹⁹⁷ and geopolitical competition,¹⁹⁸ we evaluate AGI preparedness through three dimensions: strategic awareness in government, competitive positioning and leverage, and the robustness of current policy strategies. These feed into a government's ability to: (1) perceive and interpret frontier AI developments accurately and rapidly; (2) command resources and leverage points relevant to AI development and governance; and (3) formulate and execute coherent plans on AI that align resources with objectives under conditions of uncertainty.

Dimension 1: Strategic awareness assesses whether European institutions possess the organisational capacity to accurately perceive, interpret and communicate information about frontier AI developments and their implications. Operationally, this includes:

- **Tracking trends in frontier AI capabilities:** Monitoring frontier capabilities and their probable trajectories; assessing the potential and limitations of emerging AI paradigms, such as, most recently, AI agents, reasoning models and decentralised training.
- **Studying the societal impacts of capabilities:** Understanding how large-scale deployment of AI systems could affect society; for example, evaluating agentic AI's effects on the labour market and persuasive AI's effects on democratic processes and interpersonal relationships.
- **Understanding the national security implications of capabilities:** Assessing potential large-scale threats from AI capabilities; for example, the ability of models to conduct large-scale cyberattacks, to create bioweapons, and the risk of loss of human control.
- **Collecting strategic and competitive intelligence:** Assessing a country's standing relative to geopolitical competitors and evaluating others' technological capability and access to assets in the AI value chain.

¹⁹⁶ An important example includes Kent's book on 'Strategic Intelligence for American World Policy' (1949), which argued that policy makers operating without accurate assessment of adversary capabilities, technological trajectories and emerging threats will systematically make poor decisions regardless of resources or intentions. Work by Buchanan (2020) makes similar arguments for AI, where capabilities evolve rapidly, implications spread across multiple domains, and information asymmetries are common.

¹⁹⁷ For example, technological transition studies such as Geels (2002) and Genus & Coles (2008).

¹⁹⁸ For example, scholars have identified several dimensions of competitive positioning that determine geopolitical outcomes: control over critical inputs, as in Farrell & Newman's (2019) research on 'weaponized interdependence' or ability to set standards and rules as argued by Bradford (2019). For AI competition, these often translate into questions about compute infrastructure, talent, data, capital and the ability to influence global governance.

Dimension 2: Competitive positioning and leverage assesses whether Europe has assets in the AI value chain that allow it influence AI development and ensure reliable access to frontier capabilities; we examine both Europe's competitive position across key input factors and access to 'chokepoints' where European actors hold disproportionate influence.¹⁹⁹ For competitive positioning, we assess:

- **Model capabilities:** Where do current European frontier AI models rank relative to US and Chinese competitors, and is the gap projected to widen or narrow?
- **Compute infrastructure:** What share of global AI computing capacity does Europe control, and is it sufficient for competitive frontier development?
- **Complementary inputs to development:** Does Europe possess adequate quality and volume of data, capital, energy and talent – each necessary for successful AI development?
- **Diffusion capacity:** Can Europe translate AI capabilities into adoption that leads to meaningful productivity gains, whether in government, the military, the private sector and wider society?

For leverage, we examine:

- **Supply chain positioning:** Does Europe control critical nodes in the AI supply chain (e.g. semiconductors, rare materials, manufacturing equipment) that could provide bargaining power in AI competition?
- **Market power:** Can Europe leverage its Single Market credibly to extract governance commitments and capability guarantees from foreign AI providers?

Crucially, we assess whether these strategic assets could be effectively deployed under conditions of geopolitical stress and when they would be subject to specific countermeasures.

Dimension 3: Robustness of current policy strategies assesses whether Europe has formulated ambitious, sufficiently integrated and well-resourced plans that set out realistic objectives. We analyse:

- **Strategic clarity:** Has Europe articulated explicit goals for AGI development, access, and governance, with clear prioritisation between these?
- **Use of resources:** Are committed resources commensurate with stated ambitions and the scale of the challenge?
- **Integration mechanisms:** Do coordination mechanisms exist to align efforts across fragmented institutional landscapes (EU vs Member States, multiple DGs, civil vs defence domains)?
- **Adaptation capacity:** Can strategies be rapidly modified if AI trajectories diverge from planning assumptions, or are they locked into long decision cycles?

We assess this dimension by analysing major EU AI initiatives (AI Continent Action Plan, Apply AI Strategy, Frontier AI Initiative, InvestAI, etc.) against the criteria listed above. We do not claim to be exhaustive and acknowledge that the enforcement and implementation of these strategies carries perhaps equal weight as their mere articulation.

The three dimensions correspond to distinct failure modes that could undermine European AGI preparedness. Without adequate strategic awareness, Europe might fail to anticipate when critical capability thresholds are reached, misallocate resources (e.g. in compute infrastructure), or be left unready for systemic risks such as mass unemployment and nation state-led cyber attacks. In addition, without

¹⁹⁹ This approach is inspired by the theory of 'weaponised interdependence' as described by Farrell & Newman (2019).

competitive positioning and sufficient leverage in the global value chain, Europe might formulate strategies based on accurate awareness of technological trends but lack the capacity to execute them, thereby becoming a ‘rule-taker’ whose preferences are sidelined in global AI governance. Finally, without a coherent strategy Europe’s efforts might fragment, resulting in an inability to make difficult but necessary trade-offs. By focusing on these three core dimensions, this framework allows for a structured diagnosis of potential weaknesses and blind spots in Europe’s strategic stance.

Dimension 1: Strategic awareness

Europe has limited strategic awareness of frontier AI progress

Here we examine whether European governments maintain a comprehensive, real-time understanding of the landscape and trajectory of frontier AI capabilities, and their potential societal and geopolitical implications. For example, AGI-focused labour market policies need to be informed by a detailed understanding of which sectors AGI automation is likely to occur in first, military planners need to take into account how AI is likely to shift the offence-defence balance and the drivers of strategic advantage, and geopolitical strategy requires an understanding of which countries and actors have access to which AI capabilities and assets along the supply chain. As established above, strategic awareness constitutes a necessary precondition for effective action – without accurate assessment of technological trajectories and their implications, even well-resourced actors may systematically make poor decisions. Across all domains, decision makers need to understand at which speed and in which directions the ‘jagged frontier’ of capabilities is likely to evolve. Moreover, pan-European and national authorities need to act with a clearly defined division of labour to avoid oversight or duplication.

In Europe, a particularly advanced example of institutional capacity is the UK AI Security Institute (UK AISI). Founded in 2023, it describes its mission as ‘equip[ing] governments with a scientific understanding of the risks posed by advanced AI.’²⁰⁰ As an ambitious ‘start-up in government’ with wide-ranging exemptions from regular civil service hiring and salary structures, and endowed with £100m²⁰¹ (~€113m) of initial funding, it managed to assemble a team of world-leading AI researchers within the first three months of its inception.²⁰² A recent addition to UK AISI is the ‘strategic awareness team’, which is directly tasked with ‘conduct[ing] deep dives on critical uncertainties regarding AI trajectories and impacts, track[ing] key indicators and warning signals, and produc[ing] briefs for key government decision makers.’²⁰³

UK AISI is a part of the wider Department for Science, Innovation and Technology (DSIT), where it is complemented by a number of additional units of expertise, for example relating to frontier AI regulation, international AI policy and AI innovation policy.²⁰⁴ Notably, DSIT has a Sovereign AI Unit, endowed with up to £500m (~€570m), and tasked with fostering UK sovereignty in a world with advanced AI.²⁰⁵ There are also dedicated AI units in the Foreign, Commonwealth & Development Office, the UK Ministry of Defence and the Cabinet Office. Although official numbers are not available, informed estimates suggest that across both UK AISI and DSIT at least 300 FTEs (and likely more) are working on

²⁰⁰ As shown on UK AISI’s website; UK AISI (2025).

²⁰¹ In the recent 2025 spending review, the UK government (2025b) pledged an additional £240m (€273m) to UK AISI.

²⁰² UK Government (2023).

²⁰³ An open role to join UK AISI’s strategic awareness team as a Research Scientist highlights this; see UK AISI (2024).

²⁰⁴ Internal RAND analysis.

²⁰⁵ UK Government (2025b).

various aspects of strategic awareness and AI policy.²⁰⁶ Arguably, the strategic awareness generated by these institutions has enabled the UK to lead in a number of AI policy domains; the country hosted the first AI summit in Bletchley, has wide-ranging testing access to leading US AI models, and has negotiated a UK-US tech agreement involving an investment of £31bn (€35bn) in UK AI infrastructure.²⁰⁷

State capacity for strategic awareness in continental Europe has increased, but remains less advanced. The leading centre for expertise on AI in the EU is the EU AI Office (EU AIO). While the EU AIO has attracted an impressive cohort of prominent AI researchers and scientists, is endowed with a globally unique legal mandate to gather information on and regulate frontier AI companies, and has the power to convene EU Member States on topics of advanced AI, it remains limited in its mandate and capabilities. Its starting budget – at €46m – is less than half that of UK AISI, despite its much greater jurisdiction, and its more powerful enforcement mandate. Its team is not large enough to act as a ‘strategic awareness function’ for the EU institutions, much less for the individual Member State governments, and recruiting more people is challenging as its hiring scales and payment process are less flexible than those of UK AISI, hampering its ability to attract technical talent fast enough.²⁰⁸ Finally, unlike UK AISI and the US Center for AI Standards and Innovation (US CAISI), the EU AIO does not have a mandate to work on matters of defence and national security, as these remain Member State responsibilities.²⁰⁹

Key EU Member States are also making progress on strategic awareness capacity, but institutional set-ups are not as ambitious as they could be given the gravity of a potential AGI challenge. France is the only Member State with an AI Safety Institute: instead of creating a new legal structure for an AISI, France chose to create a consortium of existing institutes for evaluation and security, which were collectively branded the ‘National Institute for the Evaluation and Security of Artificial Intelligence’ (INESIA). INESIA is supervised jointly by the National Security Secretariat and the Directorate General of Enterprises at the Ministry of Economics and Finance. Unlike the EU AI Office, it has an explicit mandate to work on national security topics, but critics argue that INESIA does not have the same budget stability and ‘start-up in government’ flexibility as UK AISI, and that its decentralised institutional structure may lead to siloing and inter-institutional rivalries.²¹⁰

Germany, the largest country in the EU, does not have an AISI or an equivalent strategic awareness function in government. There are, however, centres of AI expertise across government, such as in AI-focused units in the Federal Chancellery, the Ministry for Digital and State Modernisation, the Ministry of Economic Affairs, the Ministry for Research, Technology and Space, the Ministry of Defence, the Foreign Office, and agencies such as the German Federal Office for Information Security. Unlike France, Germany does not participate in important information-sharing fora such as the International Network of AI Safety Institutes, and unlike in the UK there has been no concerted effort in Germany to recruit top-tier technical AI talent into government to strengthen strategic awareness. Consequently, many of the most qualified German AI scientists who could become part of a strategic awareness function continue to work in Silicon Valley or in foreign AISIs.

²⁰⁶ UK government source.

²⁰⁷ UK Government (2025c).

²⁰⁸ Schaefer (2025).

²⁰⁹ The EU AI Office is structured into 6 units and 2 advisors, which reflects its mandate: the ‘Excellence in AI and Robotics’ unit, the ‘Regulation and Compliance’ unit, the ‘AI Safety’ unit, the ‘AI Innovation and Policy Coordination’ unit, the ‘AI for Societal Good’ unit, the ‘AI in Health and Life Science’ unit, the Lead Scientific Advisor and the Advisor for International Affairs. EU AI Office (2025).

²¹⁰ Barbier, Martinet & Segerie (2025).

Overall, while continental Europe has made strides in strategic awareness capacity on AI, most prominently in the form of the EU AI Office and the French INESIA, its institutions do not appear well-resourced and focused enough to be on par with ambitious initiatives such as the combined UK AISI/DSIT cluster of expertise. As a consequence, European policy makers risk being less well informed about key developments in AI, including the potential implications of AGI. This could result in worse policy design as well as decisions that are lacking the urgency and ambition necessary to compete with rival powers and to effectively mitigate downside risks. Fortunately, addressing this shortcoming may be more straightforward compared to other policy challenges associated with AI. Expanding state capacity for strategic awareness is less a matter of missing talent or expertise than of political will and prioritisation to embed existing expertise directly in government.

Dimension 2: Competitive positioning and leverage

Here we evaluate Europe's competitive positioning and leverage across key inputs for AI development and governance. We divide this assessment into two components: firstly, Europe's competitive position across core input factors (models, compute, data, capital, energy, talent and diffusion capacity), and secondly Europe's leverage points in the AI supply chain and global market that could amplify its global influence beyond raw 'inputs'.

We conclude that the EU lags significantly behind the US and China across most indicators of AI development. This includes frontier AI capabilities, where European models trail US and Chinese leaders by 6–12 months²¹¹; computational infrastructure, with Europe controlling only 5 per cent of global AI compute capacity²¹²; data access, where the shift towards synthetic data disadvantages compute-constrained Europe, while stricter privacy regulations constrain data reuse²¹³; capital availability, with EU start-ups attracting just 6 per cent of global AI venture funding²¹⁴; energy costs and availability, with Europe facing structurally higher electricity prices and permitting bottlenecks²¹⁵; and talent retention, with Europe training AI researchers but losing many of the best to better-funded American companies.²¹⁶ Europe also lags in AI diffusion capacity – the ability to translate frontier capabilities into productivity gains across the economy.

Models: European AI companies lag behind the frontier by 6–12 months

European AI companies struggle to match the performance of leading models. European AI, led by Mistral's innovations in multilingual reasoning and open-source development, remains behind the US and Chinese models in terms of both performance and cost efficiency. Mistral's most advanced model, Magistral Medium, scores 75.3 per cent on MMLU-Pro – a widely used benchmark that tests broad knowledge and reasoning across 57 subjects ranging from math and physics to law and history – compared to DeepSeek V3.1's 85.1 per cent and Claude 4.1 Opus' 88 per cent.²¹⁷ At the same time, Mistral's models are more costly, at \$2–5 per million tokens versus DeepSeek's \$0.50–\$2 and gpt-oss-120B (high)'s \$0.15–\$0.60 pricing. Latest estimates signal that Mistral's Magistral Medium lags behind the best US and Chinese models by 6–12 months (see Figure 7).²¹⁸

²¹¹ Artificial Analysis (2025b).

²¹² As presented in Pilz et al. (2025) data in 'Trends in AI Supercomputers'.

²¹³ Lakshmanan (2024).

²¹⁴ Martens (2024).

²¹⁵ Heussaff (2024); European Commission (2025c); SolarPower Europe (2024).

²¹⁶ Pal (2024).

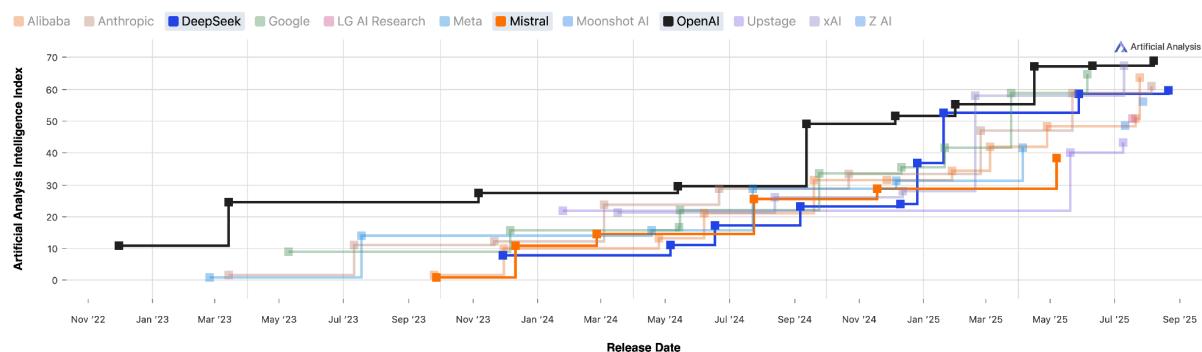
²¹⁷ Artificial Analysis (2025a).

²¹⁸ Artificial Analysis (2025b).

Europe is simultaneously outmatched in terms of overall numbers of frontier models. Excluding Google DeepMind (which is based in the UK and part of a US company), Europe produced only three notable models in 2024 compared to 40 from the US and 15 from China, highlighting a significant gap in model production volume, which serves as an indicator of the depth of the AI ecosystem.²¹⁹ With very few European developers at frontier scale, a lock-in to a non-European AI tech stack looks increasingly plausible. The cloud computing market also demonstrates this dynamic, where early, aggressive scaling by US hyperscalers such as AWS and Microsoft Azure created powerful economic moats and high switching costs for European customers.²²⁰

The gap with China is also widening, with Chinese models getting closer to the capabilities of US offerings. DeepSeek's R1 achieved near-parity with leading US models when it came out in May 2025, lagging behind then-leading Claude Sonnet 4 by only 2 percentage points on GPQA Diamond, which tracks performance on expert-level questions in science, mathematics and technical domains.²²¹ The broader trend echoes similar patterns: performance gaps on major benchmarks such as MMLU narrowed from 17.5 percentage points in 2023 to just 0.3 points by early 2025.²²² Industry analysts now estimate Chinese models lag behind US frontiers by only 3–6 months,²²³ an improvement from the 1+ year gap that was estimated in 2022.²²⁴

Figure 7. Benchmark performance of European frontier models in international comparison



Source: Artificial Analysis (2025a)²²⁵

While lags of only a few months behind the frontier might not seem especially pressing, they already translate into massive commercial disparities: OpenAI generated \$10bn in annualised recurring revenue by June 2025,²²⁶ while Mistral had revenue of \$30m in 2024, with projections of \$60m for 2025²²⁷ – more than a hundredfold difference. This suggests a winner-takes-most dynamic in AI markets, with superior capabilities driving exponentially higher customer adoption, pricing power and market share. With such market dynamics, it is an open question whether European champions will be able to continue

²¹⁹ Stanford University (2025). Epoch AI (2025h) defines ‘notable model’ as ‘meeting any of the following criteria: (i) state-of-the-art improvement on a recognized benchmark; (ii) highly cited (over 1000 citations); (iii) historical relevance; (iv) significant use’.

²²⁰ Srivathsan, Sorel & Sachdeva (2024).

²²¹ Epoch AI (2024).

²²² Stanford University (2025), part ‘3. The gap between Chinese and US models closes.’

²²³ Insikt Group (2025).

²²⁴ Artificial Analysis (2025b).

²²⁵ Incorporates 8 evaluations: MMLU-Pro, GPQA Diamond, Humanity’s Last Exam, LiveCodeBench, SciCode, AIME, IFBench and AA-LCR.

²²⁶ Capoot & Rooney (2025).

²²⁷ D’Souza (2025).

financing their AI development. Moreover, the gap becomes even more consequential as leading models begin automating large parts of the AI development process itself, triggering feedback loops that could compound advantage for the leader. The usage of AI to automate AI development makes the capabilities gap more likely to widen rather than close in the future, even when rivals are currently only months behind.²²⁸

Data: European AI companies are constrained by compute and data regulation

Data constitutes a foundational input for AI development, serving as the raw material from which models learn patterns, develop capabilities and adapt to specific use cases. The quality, quantity and diversity of training data determine a model's performance, generalisation ability and real-world applicability. However, the data requirements for frontier AI are undergoing a change that affects competitive dynamics between regions.

Until recently, frontier models relied primarily on scraping vast quantities of internet text, images and code for 'pre-training' – the stage where an AI system learns general patterns before being taught specific tasks.²²⁹ After pre-training, low-skilled human labellers were needed to ensure AI systems' outputs were useful, ethical and in line with the developer's guidelines. As detailed in Chapter 2's discussion of how reinforcement learning might circumvent the 'data wall', AI developers have since developed computationally intensive post-training and runtime techniques that rely heavily on synthetic data – data produced and filtered by AI systems themselves.²³⁰ Looking ahead, the most valuable human contributions to AI training may evolve from large-scale data production and low-skilled labelling to creating highly specialised datasets – particularly in domains requiring tacit knowledge. Similarly, data collection may shift from mostly passive gathering to the construction of interactive reinforcement learning environments where AI agents can explore and learn autonomously. These changes mean that in addition to the volume of human-generated content, expert data and computing power have become a core determinant in the quest for training data.²³¹

Europe is at a disadvantage when it comes to synthetic data generation. Limited access to computing capacity compared to the US and China may constrain the ability of European researchers to generate and refine the massive quantities of synthetic data now driving AI progress. Europe also faces disadvantages in the realm of non-synthetic training data. US firms control large quantities of user-generated content, interaction logs and cloud data that can be repurposed for AI development under existing terms of service.²³² The US also lacks comprehensive federal data protection law, and broad interpretations of 'fair use' in copyright have allowed AI companies to train on vast quantities of scraped content with limited legal friction – though this permissiveness is now being tested in litigation such as *The New York Times v. OpenAI*.²³³ China combines large domestic user bases with strong state coordination and mandated data-sharing among domestic companies, enabling both scale and centralisation. China's 2021 data protection laws nominally impose restrictions comparable to GDPR,²³⁴ but enforcement is asymmetric: domestic AI champions benefit from state-mandated data-sharing arrangements and lighter scrutiny, while foreign competitors face stricter compliance burdens.

²²⁸ Janků et al. (2025).

²²⁹ The legality of training AI on copyrighted material without permission remains contested, with ongoing litigation in multiple jurisdictions. Even where courts have found such use legally permissible under fair use doctrine, the ethics of training on creators' work without consent or compensation remains disputed.

²³⁰ Epoch AI (2025f, 2025g).

²³¹ Wexler (2025).

²³² OECD (2025).

²³³ Roth (2023).

²³⁴ Zhu (2022).

Europe's General Data Protection Regulation (GDPR) poses additional constraints. Purpose limitation requirements mean data collected for one service cannot be freely repurposed for AI training without establishing a new legal basis.²³⁵ Reliance on 'legitimate interest' for AI training thus faces increasing regulatory scepticism. This was illustrated by Meta's June 2024 decision to pause EU AI training following complaints by the Austrian privacy group NOYB and intervention by the Irish Data Protection Commission. Data minimisation principles complicate foundation model development, which benefits from maximal data ingestion. The right to erasure also creates a hard technical challenge for models that have already learned from subsequently deleted data.²³⁶ These structural factors constrain the manoeuvring space of European actors in data-intensive AI development.

The above factors help to explain why promises of European data advantages have not (yet) materialised. Earlier EU strategies, including the 2020 EU Data Strategy²³⁷ and 2021 Fostering a European Approach to AI,²³⁸ anticipated that Europe's industry-specific knowledge and high-quality domain data would provide competitive advantage. However, this data remains fragmented across sectors and largely underutilised.²³⁹ The Commission's recent AI Continent Action Plan²⁴⁰ represents a more promising approach, aiming to expand data availability for EU developers by streamlining regulatory access. By establishing Data Labs and Common European Data Spaces, it aims to create curated, high-quality and interoperable datasets in strategic areas such as health, manufacturing and energy. Yet even with improved data access, Europe will require substantial AI expertise to transform passive datasets into the interactive reinforcement learning environments that frontier AI development increasingly demands.²⁴¹

Compute: Europe hosts only 5 per cent of the world's AI computing infrastructure

The development and deployment of advanced AI systems depends on access to massive computational infrastructure. Training frontier models currently requires large clusters of specialised chips running in centralised facilities costing billions of euros and often consuming as much power as small cities. Meanwhile inference – the process of deploying trained models to generate output – requires less raw power per datacentre site but represents an ongoing operational cost that scales with usage. Companies and governments that control AI infrastructure not only derive economic benefits from AI development and adoption, but also have more control over who can build and use AI systems and related services.

Europe is far behind the US and China in terms of housing onshore AI infrastructure (see Figure 8). The US has the world's largest share – about 75 per cent – of global aggregate computational performance for AI, followed by China with 14 per cent and the EU at nearly 5 per cent.²⁴² This disparity may be partially explained by Europe's higher electricity prices,²⁴³ long permitting timelines and lack of excess grid

²³⁵ GDPRhub (2025), Article 5(1)(b)).

²³⁶ Moreno (2025).

²³⁷ European Commission (2020).

²³⁸ European Commission (2021).

²³⁹ Batur & Peeters (2021).

²⁴⁰ European Commission (2025a).

²⁴¹ Patel & Kourabi (2025).

²⁴² Pilz et al. (2025) have a dataset on 'Trends in AI supercomputers' that gathers this data. These figures require cautious interpretation, as the study captures only 10–20 per cent of AI supercomputers globally, with uneven coverage between private and public infrastructure, and focuses exclusively on supercomputers rather than overall compute performance. This means that government supercomputers are over-captured in the data and the overall AI compute landscape may appear different in practice.

²⁴³ Patel (2024a).

capacity.²⁴⁴ Some improvements are on the horizon. Planned AI Gigafactories will bring investment of €20bn, while the Cloud and AI Development Act aims to triple the EU's datacentre stock in the next 5–7 years.²⁴⁵ Meanwhile, France plans investments in AI infrastructure of €109bn, framing it as a response to the US Stargate project, and has announced the establishment of Europe's largest AI campus in Paris.²⁴⁶ Even so, this still might not lead to a relative advantage, since the US may also triple its datacentre capacity between 2023 and 2027.²⁴⁷ As a result, the gap between Europe and its primary competitors will likely stay constant or widen further.

Leading US hyperscalers and AI companies are demonstrating continued commitment to build the next generation of datacentres underpinning America's global AI leadership. In 2025 alone, US hyperscalers announced plans for investments in AI datacentres worth almost €300bn.²⁴⁸ In addition, the Stargate project schedules investments of \$500bn (€455bn) over the next 5 years to construct the supercomputer facilities powering OpenAI's future frontier models.²⁴⁹ Oracle has signed a \$300bn contract with OpenAI to provide them with computing resources over 5 years starting in 2027,²⁵⁰ and NVIDIA has invested another \$100bn (€87bn) in OpenAI to enable them to buy more NVIDIA chips²⁵¹ supported by the public funding of \$138bn (€119bn).²⁵²

Global competition for datacentre construction has intensified recently, with the Middle East becoming a key destination for vast AI compute hubs.²⁵³ For example, the UAE is planning a huge 'AI campus' in Abu Dhabi – touted as the largest outside the US – and could be permitted to import up to 500,000 cutting-edge AI chips each year, starting in 2025. The agreement also commits the UAE to investing in US data centres of comparable scale and to aligning its security standards with those of the US.

The widening AI infrastructure gap may soon become a liability for Europe's economic competitiveness and capacity for innovation.²⁵⁴ Insufficient access to onshore European compute directly limits the ability of European firms to develop, test and deploy advanced AI systems, from experimental research to large-scale frontier models. While accessing cloud services from foreign hyperscalers offers an immediate and effective solution, this approach has some limitations. Firstly, it reinforces the network effects cementing US-based providers' dominance, stifling European alternatives. Secondly, geographically distant foreign datacentres would fail to deliver the ultra-low latency required by some critical industrial applications – advanced manufacturing processes, autonomous vehicle systems and real-time robotics – which depend on inference speeds where even millisecond delays can compromise safety and efficiency. Thirdly, it exposes European industry to the increasingly constrained supply of cloud services.²⁵⁵ As global demand for AI training and inference surges, reliable access becomes a competitive differentiator, and over-reliance on external providers leaves European innovation vulnerable to prioritisation decisions made elsewhere.

²⁴⁴ Kothuri et al. (2025).

²⁴⁵ European Commission (2025a).

²⁴⁶ Piquard (2025).

²⁴⁷ Patel (2024a).

²⁴⁸ Subin (2025).

²⁴⁹ Murgia, Hammond & Kinder (2025).

²⁵⁰ Plumb (2025).

²⁵¹ OpenAI (2025c).

²⁵² Parasnis (2025).

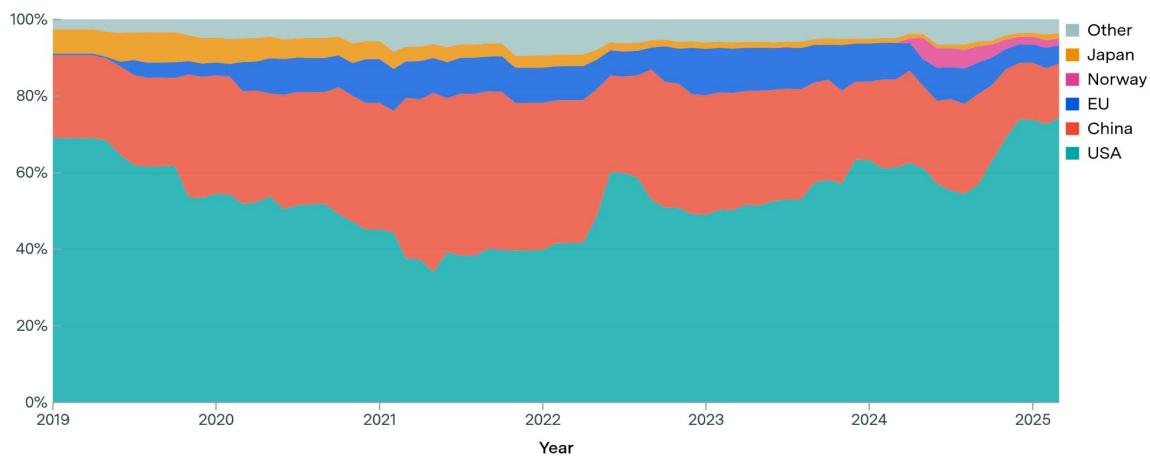
²⁵³ Guardian Staff and Agencies (2025).

²⁵⁴ Shamir & Negele (2025).

²⁵⁵ Srivathsan, Sorel & Sachdeva (2024).

Beyond the economic implications, a dependency on foreign AI infrastructure threatens Europe's security and strategic autonomy.²⁵⁶ A dearth of reliable, onshore compute capacity jeopardises the resilience of Europe's most critical infrastructure and industries, potentially leaving essential services vulnerable to supply chain disruptions or geopolitical coercion. Moreover, reliance on foreign datacentres also cedes control over vital security standards, from hardware-enabled safeguards to physical security, undermining Europe's ability to protect sensitive government and commercial data as well as frontier model weights according to its own requirements.²⁵⁷ Without sufficient onshore AI infrastructure, Europe's ability to act independently in a geopolitical crisis – or alternatively, its ability to shape global AI governance in line with its values – may be constrained.

Figure 8. Europe's share of global aggregate computational performance (in 16-bit FLOP/s)



Source: Pilz, Rahman et al. (2025).

Capital: only a small share of global AI funding is going to EU-based start-ups

The computational power required to train today's frontier models translates into a massive demand for capital. Acquiring and networking hundreds of thousands of GPUs and running them for months at a time to complete a single pre-training run requires a large upfront expense. Unlike inference, which distributes costs over time, training is a capital-intensive, high-risk R&D expenditure that must be incurred before any revenue is generated. In this environment, access to patient and risk-tolerant capital becomes a prerequisite to being and staying at the frontier. This tight link between extreme compute costs and large-scale finance makes Europe's fragmented financial system a structural weakness in frontier AI competition.

In the first half of 2024 more than \$35bn was invested into AI start-ups worldwide, yet only around 6 per cent of that funding went to EU-based start-ups.²⁵⁸ At the same time, Europe struggles to mobilise patient institutional capital that can wait years for returns. As the Draghi Report highlights, EU pension assets in 2022 amounted to only about 32 per cent of GDP, compared to 142 per cent in the US and 100 per cent in the UK, while private pension savings are heavily concentrated in a few Member States such as the Netherlands, Denmark and Sweden.²⁵⁹ In the meantime, training costs for frontier AI models are surging,

²⁵⁶ Hawkins, Lehdonvirta, & Wu (2025).

²⁵⁷ Nevo et al. (2024).

²⁵⁸ Martens (2024).

²⁵⁹ Draghi (2024).

growing at a rate of 2.4 times per year since 2016.²⁶⁰ Assuming that rate stays constant, the largest training runs, without factoring in much larger fixed datacentre expenditures, will cost more than a billion dollars by 2027 – a scale of investment that privileges capital-rich, risk-taking incumbents.²⁶¹ The Draghi Report already warns that 30 per cent of European unicorns relocated abroad between 2008 and 2021, primarily to the US, to access deeper funding pools.²⁶² Without stronger cross-border capital integration, Europe risks pushing aspiring frontier AI champions overseas.

Energy: electricity prices are significantly higher than in the US

AI infrastructure demands vast electrical power. As AI capabilities and adoption grow quickly, energy use is becoming a development bottleneck as important as access to advanced chips or capital. Recent research estimates that global AI datacentres could require an additional 10 gigawatts of power capacity in 2025, almost twice the entire power capacity of Ireland.²⁶³ Assuming that chip supply can keep up with demand, AI datacentres could need around 68 gigawatts in total by 2027, slightly more than the total power capacity of Poland today.²⁶⁴ This rapid increase in demand, driven by both model training and inference, is turning power availability into a core dimension of global AI competition.

Europe's energy system has features that are particularly poorly matched to this. The continent faces structurally higher energy prices: in 2023, industrial electricity prices in the EU were 158 per cent higher than in the US, and industrial gas prices were 345 per cent higher.²⁶⁵ This largely reflects Europe's reliance on imported fossil fuels, a dependency exacerbated by Russia's invasion of Ukraine, when embargoes on Russian oil and gas triggered price shocks and an acute energy crisis. Certain countries, such as Germany, have also been sceptical of nuclear power as a solution to long-term energy needs. By contrast, the US is a net energy exporter. Europe's energy security also rests on a decarbonisation-focused policy frameworks, such as the European Green Deal, which aims for long-term autonomy through renewables but may impose higher short-term costs than continued fossil fuel use.

In addition, unlike the mostly integrated continental grids of the US and China, Europe's energy system remains a collection of separate national grids.²⁶⁶ While interconnected, they face cross-border bottlenecks and divergent national regulations, which makes it harder to provide the large-scale, predictable power capacity that massive datacentre clusters require.²⁶⁷ Permitting is a further constraint: average approval times for new datacentres stand at about 48 months, and EU-level planning reforms have often been slowed by the late transposition of permitting rules in several Member States.²⁶⁸

In contrast, other AI powers enjoy stronger energy security. The US benefits from large domestic energy resources, especially natural gas, which keeps power relatively cheap and reliable. This allows US hyperscalers to plan and build large onshore datacentres with confidence about long-term energy

²⁶⁰ Cottier (2024).

²⁶¹ Cottier (2024).

²⁶² Draghi (2024).

²⁶³ Estimates included in research by Pilz, Mahmood & Heim (2025). Ireland's system demand – the electricity production needed to meet national consumption – is at 5.23 gigawatts in 2025, as the electricity Transmission System Operator (TSO) for Ireland shows; EirGrid (2025).

²⁶⁴ Estimates included in research by Pilz, Mahmood & Heim (2025). Poland's grand system capacity is at 64.3 gigawatts in 2025, as per the EU electricity transparency hub; Entsoe (2025).

²⁶⁵ Heussaff (2024).

²⁶⁶ Entsoe (2025a).

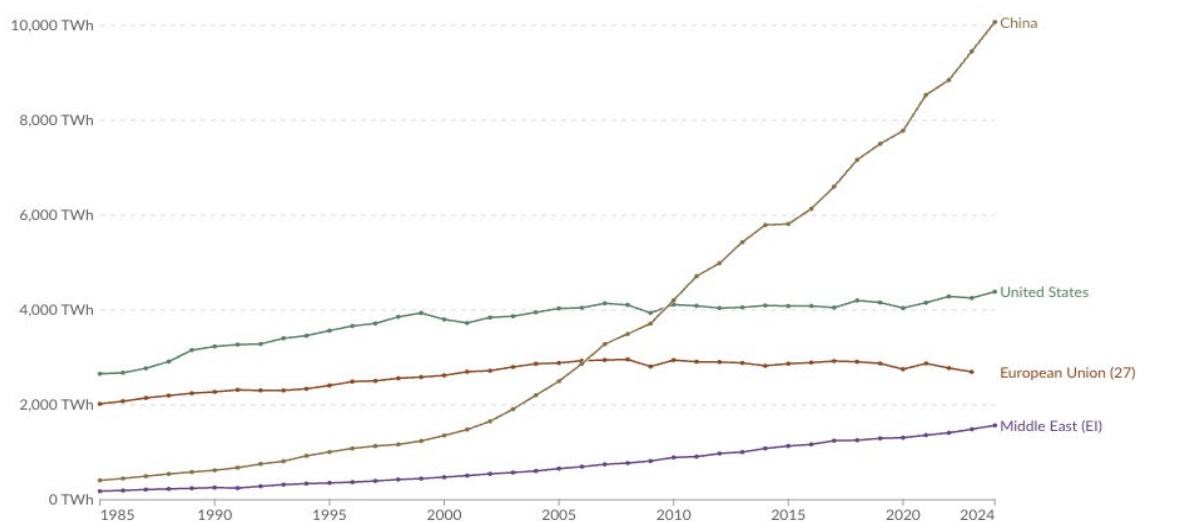
²⁶⁷ Draghi (2024).

²⁶⁸ European Commission (2025c); SolarPower Europe (2024).

availability and cost.²⁶⁹ To manage grid congestion and intermittency, US firms also invest abroad; for example, Microsoft and OpenAI's Stargate project taps Middle Eastern markets with abundant, low cost energy to complement domestic supply.²⁷⁰ China, meanwhile, uses its state-driven model to implement a centrally planned energy strategy, directing massive investments into generation capacity from renewables to nuclear at a speed and scale unmatched in Europe and the US (see Figure 9).²⁷¹ Its 'Eastern Data, Western Computing' strategy, for example, strategically couples newly constructed datacentres with nearby large-scale power supply.²⁷²

Energy capacity will remain a critical input for AI development and adoption, and this matters well beyond model training. Modern AI systems consume power every time they process tokens for text, images or reasoning tasks.²⁷³ As capabilities advance and inference becomes more compute intensive and more important for productivity, the energy cost per unit of useful output will increasingly matter. This is especially true if reasoning-heavy models and autonomous AI agents see broad adoption in critical sectors. In such a world, regions with cheaper and more abundant electricity can run more AI for the same cost, giving them an edge in developing and deploying transformative systems.²⁷⁴

Figure 9. Total electricity generated in each country or region (in terawatt-hours)



Source: Our World in Data (2025).

Elite talent concentration: little of the top AI talent works in Europe

Talent is a critical input for Europe's ability to build frontier AI models, given the importance of algorithmic efficiency and the complexity of very large training runs. China hosts the deepest pool of AI talent, with around 47 per cent of the world's top tier AI researchers.²⁷⁵ The US is the main talent

²⁶⁹ Rueda et al. (2025).

²⁷⁰ Cornish & Acton (2025).

²⁷¹ Roytburg (2025).

²⁷² 穆森 (2024).

²⁷³ Salvator (2025).

²⁷⁴ Heim & Lin (2025).

²⁷⁵ Defined by MacroPolo (2024) Global AI Talent Tracker as researchers who publish in highly selective venues such as NeurIPS, especially those whose work is chosen for oral presentation; used as a proxy for the world's most elite AI researchers.

magnet, with about 60 per cent of top AI institutions and a majority of global top AI talent (much of it from Europe) working in US organisations in 2022. In comparison, about 12 per cent of top AI talent worked in Chinese institutions, 8 per cent in continental Europe and 8 per cent in the UK.²⁷⁶

In terms of quantity, the EU is good at training talent but poor at retaining it. Many leading researchers move to the US, while India has become a key source of AI professionals for Europe. For example, 28 per cent of Ireland's, 15 per cent of Sweden's, and 13 per cent of Germany's AI workforce completed their undergraduate degrees in India.²⁷⁷ Recent trends show some EU countries gaining AI workers overall²⁷⁸ and recent political developments in the US may be reversing traditional talent flows.²⁷⁹ As many as 75 per cent of American scientists who responded to a 2025 *Nature* poll on the new US administration's policies indicated that they are contemplating relocating abroad.²⁸⁰

On talent quality, the EU faces an acute shortage of advanced AI engineers, with fewer than 0.63 qualified candidates per vacancy for these positions.²⁸¹ Most of the EU's AI talent consists of 'technical professionals in software and data roles using basic AI methods' – for those, the supply exceeds demand. However, the shortfall in 'advanced AI engineers and researchers developing cutting-edge models or other advanced AI techniques' limits the EU's ability to reach the technological frontier.

Europe's AI workforce also suffers from fragmentation. Several EU countries have more AI experts per capita than the US or China, but this talent is spread across many small hubs.²⁸² In the US, AI professionals cluster in a few large centres such as San Jose, San Francisco and Seattle. In the EU, only Amsterdam, Berlin and Paris rank among the top 25 global AI hubs, whilst London ranks 8th globally.²⁸³ This limits the agglomeration effects and innovation spillovers that drive breakthrough AI development in more concentrated ecosystems like Silicon Valley.

Taken together, Europe trains a large AI workforce but loses many of its strongest researchers to the US, and much of the remaining talent is spread across many small hubs rather than concentrated in major tech centres.²⁸⁴ As a result, Europe currently lacks the dense pool of advanced expertise needed to compete in building frontier AI models.

AI diffusion: European adoption of AI is a mixed story, but likely somewhat behind the US

Europe's ability to capture economic value from frontier AI depends on how well it can diffuse the technology across its economy. Ding's work²⁸⁵ on general purpose technologies shows that diffusion capacity often determines who captures gains and can matter more than invention itself. Policy priorities increasingly reflect this focus on diffusion: both the US AI Action Plan²⁸⁶ and the EU's Apply AI Strategy²⁸⁷ emphasise adoption, with the latter mobilizing €1bn across 11 strategic sectors.

²⁷⁶ MacroPolo (2024).

²⁷⁷ Pal (2024).

²⁷⁸ Stanford University (2025).

²⁷⁹ Reif (2025).

²⁸⁰ Witze (2025).

²⁸¹ Pal, Schneider & Nurski (2025).

²⁸² Pal, Marino Lazzaroni & Mendoza (2024).

²⁸³ Pal, Marino Lazzaroni & Mendoza (2024).

²⁸⁴ Pal, Schneider & Nurski (2025).

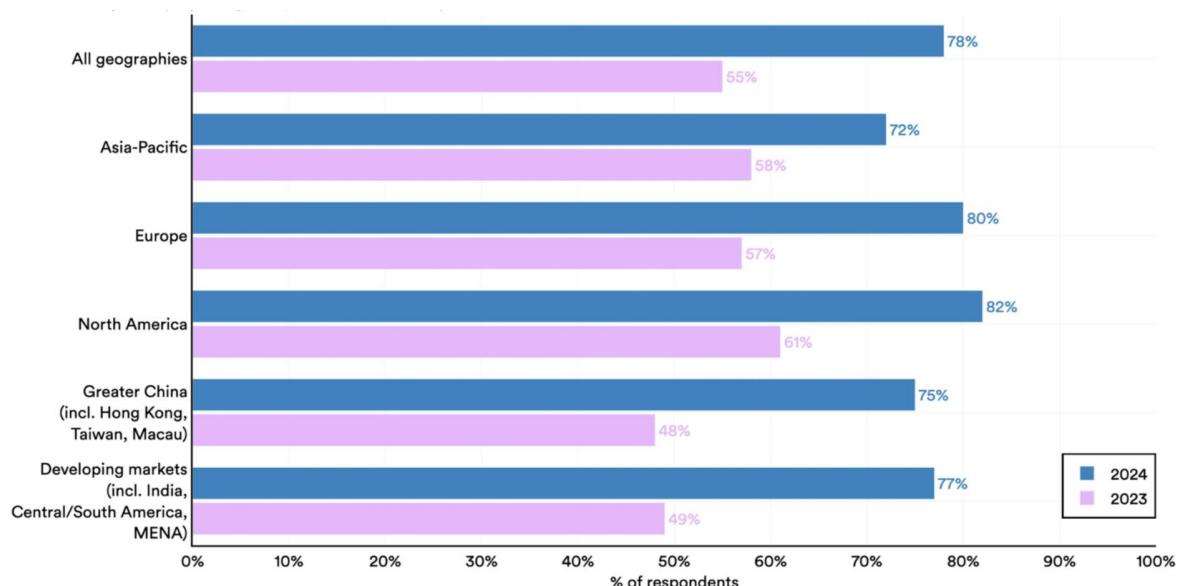
²⁸⁵ Ding (2024).

²⁸⁶ White House (2025).

²⁸⁷ European Commission (2025b).

Current metrics of AI adoption paint an ambiguous picture. According to Eurostat,²⁸⁸ only 13.5 per cent of EU firms reported using AI in 2024, with strong variation by firm size and country: 41 per cent of large firms versus 11 per cent of small firms, and 24 per cent in Denmark versus 4 per cent in Romania. International comparisons are hard because of different survey methods, but these numbers alone do not reveal a large transatlantic gap.²⁸⁹ One US survey recorded about 8 per cent of firms using AI to produce goods or services within a two week window in April 2025.²⁹⁰ A McKinsey survey cited by the Stanford AI Index finds only a two percentage point gap in broad AI adoption between North America and Europe, with 82 per cent of respondents in North America and 80 per cent in Europe reporting use of AI in at least one business function (see Figure 10).²⁹¹ Beyond these surveys, data comparing adoption rates between regions is limited. Anthropic's AI Usage Index is one notable exception, suggesting that usage rates in the US are clearly ahead of countries in the EU, even those with high adoption rates such as France and Germany.²⁹² That said, this metric only covers one provider's users, and others, like OpenAI, have not released regional data.

Figure 10. AI use by organisations in the world, 2023 vs. 2024



Source: Stanford University (2025).

However, current adoption rates tell only part of the story. The key question is whether Europe's economic and institutional environment supports deep, economy wide integration of AI. On paper, Europe's economy appears to be well placed for this. With services accounting for about 65 per cent of GDP, AI agents and other 'digital workers' could automate or augment many cognitive and information processing tasks.²⁹³ In practice, structural frictions such as weak commercialisation pathways may limit how far AI capabilities translate into productivity growth and even a successful diffusion strategy could remain dependent on foreign inputs for models, compute and energy.

²⁸⁸ Eurostat (2025).

²⁸⁹ Kergroach & Héritier (2025).

²⁹⁰ Census Bureau (2025).

²⁹¹ Stanford University (2025).

²⁹² Appel et al. (2025).

²⁹³ European Union (2023); Denain, Ho & Sevilla (2025).

A recent commentary by Ding highlights several levers for building diffusion capacity²⁹⁴: enlarging the pool of AI implementers beyond elite researchers through training and certification programmes, creating diffusion institutions that support small and medium-sized firms, ensuring interoperable standards, and providing reliable access to cloud infrastructure and sectoral data. The EU's Apply AI Strategy²⁹⁵ addresses part of this agenda. It aims to support small and medium-sized firms through about 250 Digital Innovation Hubs and to launch sectoral flagship projects backed by funding. Its main emphasis, however, is on deployment support, rather than on deeper structural reform.

One structural constraint concerns commercialisation for application-focused AI firms that build on top of frontier models. These include start-ups that create sector-specific tools and workflow integrations, such as for legal services (e.g. Harvey) or software development (e.g. Cursor). In the EU, such firms face at least three obstacles. Firstly, weak links between universities and industry limit the translation of research into products. Secondly, as documented in an earlier section, capital is relatively scarce in Europe, especially at growth stages. Thirdly, uncertainty over how high-risk obligations under the AI Act will apply may deter some founders from building in the EU at all.

Even if Europe succeeds with a diffusion-focused strategy, it will still depend on external inputs such as compute, energy and leading models. To reduce dependence on the latter, the Apply AI Strategy proposes EU-wide competitions to develop open European frontier models, with winners receiving access to EuroHPC supercomputers. These models would then be made available to public authorities, researchers and firms across Europe. However, as we explain in a later section on Europe's AI strategies, these incentives are unlikely to allow European actors to challenge the leading US or Chinese models at the frontier.

Europe's main strategic levers provide limited influence

Europe holds two major assets: a critical role in the global supply chain for advanced semiconductors, most notably Dutch firm ASML's monopoly on extreme-ultraviolet (EUV) lithography equipment, and the EU's market power as the world's second-largest economic area. In theory, these assets could serve as bargaining chips to secure access to frontier AI capabilities and shape AI governance beyond the AI Act. In practice, however, their value is limited by structural dependencies, institutional fragmentation and the risk of retaliation, which may be decisive in periods of geopolitical turbulence. Even so, they remain important sources of leverage for Europe if used strategically.

Semiconductor supply chain: strong positioning, but conversion to influence is not straightforward

The semiconductor supply chain is the foundation of frontier AI: it is a highly specialised, globally distributed ecosystem in which a few firms control critical chokepoints (see Figure 11). NVIDIA dominates the design of AI chips²⁹⁶; Taiwan's TSMC fabricates roughly 90 per cent of the world's most advanced processors²⁹⁷; and South Korea leads in high-bandwidth memory (HBM).²⁹⁸ A crucial part of the supply chain is manufacturing equipment, with European firms playing an outsized role. The Dutch firm ASML, supported by specialist suppliers such as the German optics manufacturer Zeiss, has a monopoly on the EUV machines that are essential for printing the smallest, most advanced patterns for

²⁹⁴ Ding (2025).

²⁹⁵ European Commission (2025b).

²⁹⁶ NVIDIA has about 92 per cent market share of the add-in-board GPU market; Farooque (2025).

²⁹⁷ Sacks & Segal (2025).

²⁹⁸ Yang & Cherney (2025).

cutting-edge chips.²⁹⁹ ASML is also the dominant producer of less advanced deep-ultraviolet (DUV) lithography machines,³⁰⁰ but does not hold a complete monopoly since Chinese firms have started to develop their own, albeit less advanced, DUV machines.³⁰¹

ASML is sometimes regarded as the single most important chokepoint in this chain.³⁰² Precisely quantifying ASML's lead in EUV technology is difficult, but experts regard it as the technology where China is furthest behind and facing the greatest difficulty catching up.³⁰³ Recreating just the laser in an EUV system requires identifying and assembling 457,329 unique parts, which helps explain why these machines are sometimes described as among the most complex objects humans have ever built.³⁰⁴ Earlier attempts by Nikon and Canon to develop EUV tools failed despite multibillion dollar investments.³⁰⁵ Today, the primary firm reported to be pursuing a comparable EUV capability is China's SiCarrier, operating under the internal code name 'Mount Everest'.³⁰⁶ Nearer-term Chinese efforts to reduce reliance on Western technology focus on DUV tools. Shanghai-based start-up Yuliangsheng has produced a DUV machine now on trial with SMIC, China's leading foundry, and Shanghai Micro Electronics Equipment (SME) also produces less advanced DUV systems.³⁰⁷

Since every advanced AI chip in the world, with the partial exception of Huawei's recent Ascend 910C, is produced using EUV lithography, the ability to control access to EUV machine production and maintenance translates directly into geopolitical power for the Netherlands and the EU. This was illustrated when pressure from the US administration led the Dutch government to block EUV exports to China and later to tighten controls on some DUV tools, moves that significantly slowed China's progress in advanced chipmaking.³⁰⁸

This position creates strategic leverage for Europe, but using it is not straightforward. The European Union, through the Netherlands and a small group of partners, can influence who has access to the most advanced lithography equipment and thus who can build the compute that powers frontier AI. In theory, the EU could harden controls on China beyond the current EUV ban, for example by restricting DUV exports. In a hypothetical transatlantic dispute, it could even seek to delay equipment or deny on-site maintenance for partners such as TSMC, risking an escalatory violation of the US Foreign Direct Product Rule.³⁰⁹ We stress that these are stylised extreme scenarios rather than policy recommendations, and that each would face major legal, political and practical constraints, as well as serious risks.

Several constraints limit Europe's room for manoeuvre. Firstly, ASML depends heavily on US technology and expertise, notably through its Cymer laser division and substantial US workforce, which exposes it to

²⁹⁹ Clark (2021).

³⁰⁰ Hijink (2024).

³⁰¹ Wentz & Lin (2025).

³⁰² Lugt (2024).

³⁰³ Grunewald (2023).

³⁰⁴ TRUMPF (2025); Satariano (2025).

³⁰⁵ Hijink (2024).

³⁰⁶ Olcott and Wu (2025).

³⁰⁷ Shilov (2024); Olcott and Wu (2025).

³⁰⁸ Satariano (2025).

³⁰⁹ The FDPR is a US export regulation that extends American jurisdiction to foreign-made items if they are a 'direct product' of specified US technology or software or produced by equipment derived from such technology. Because most advanced SME contains US-origin chips or is made with US tools, the FDPR subjects the vast majority of foreign-produced advanced chips and SME to US control, regardless of where they are made. The US has previously leveraged this extraterritorial authority to close loopholes that allow for circumvention via foreign subsidiaries and to compel allies to adopt equivalent controls by offering exemptions. For more details, see Allen & Goldston (2025).

the US' export control policy.³¹⁰ Secondly, using ASML as leverage against the US rather than China would collide with Europe's broader dependence on the US for security, cloud services and energy, creating a kind of mutual deterrence. Thirdly, using this leverage more aggressively against China risks strong economic retaliation and could accelerate China's state-funded drive for self-sufficiency.³¹¹ Huawei is reportedly to have produced only about 200,000 advanced AI chips in 2025, a fraction of the roughly one million US-produced chips Chinese firms imported in 2024. However, China's leading foundry, SMIC, is already using multi-patterning techniques with existing DUV machines to produce 7nm and potentially 5nm chips,³¹² efforts clearly aimed at closing the technology gap over the long term.³¹³ Fourthly, European institutional fragmentation adds another constraint. Despite EU dual use rules and the current EUV export ban to China, export controls remain mainly a national competence with limited EU-level oversight, which creates barriers to rapid and coordinated action. The Pegasus spyware controversy, in which banned tools still received export licences, illustrates how fragmented enforcement can undermine common policy.³¹⁴

How Europe manages its supply chain position is therefore a defining test of its strategic autonomy. The EU, and in particular the Netherlands, is often caught between US pressure for tighter export control alignment and the threat of Chinese economic retaliation.³¹⁵ For European policy makers, the challenge is to deploy its leverage in a measured, strategic way to deter economic coercion and to shape a global governance regime that protects European interests.

³¹⁰ Hijink (2024).

³¹¹ Self-sufficiency is already a dominant goal of Chinese semiconductor policy and some may argue that the Chinese ecosystem could not push for much greater acceleration at this time.

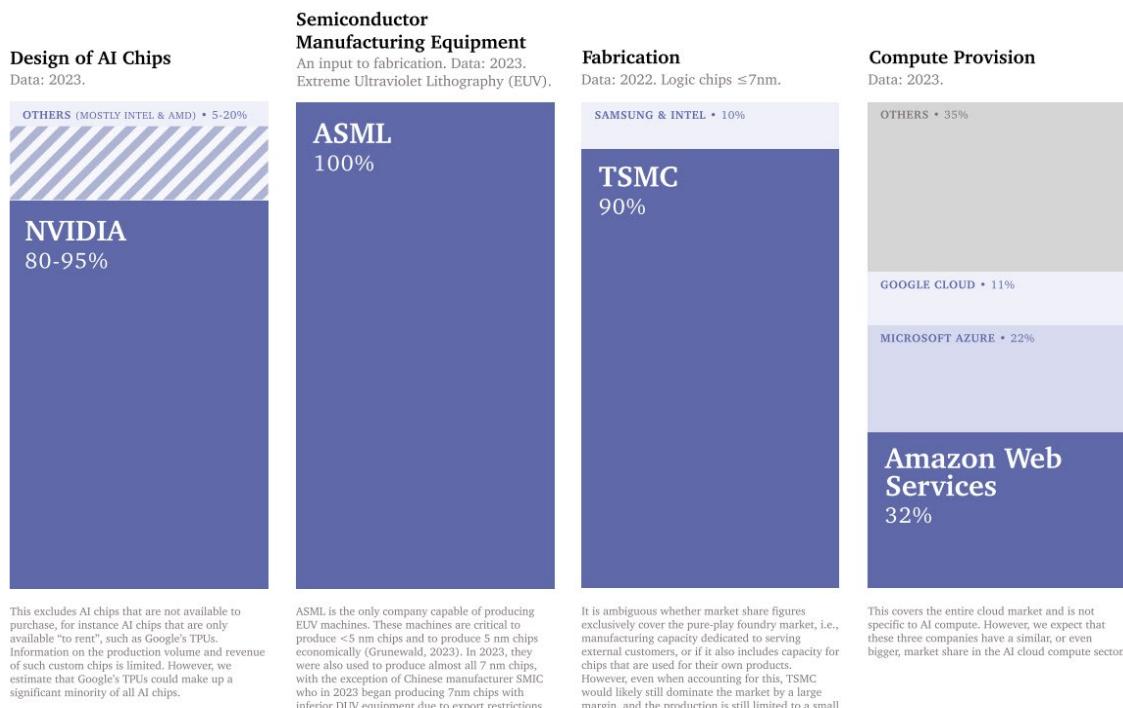
³¹² Olcott & Wu (2025).

³¹³ Olcott & Wu (2025).

³¹⁴ European Parliament (2023).

³¹⁵ Lugt (2024).

Figure 11. Concentration of the AI chip supply chain, expressed as a percentage of total market share



This excludes AI chips that are not available to purchase, for instance AI chips that are only available “to rent”, such as Google’s TPUs. Information on the production volume and revenue of such custom chips is limited. However, we estimate that Google’s TPUs could make up a significant minority of all AI chips.

ASML is the only company capable of producing EUV machines. These machines are critical to produce <5 nm chips and to produce 5 nm chips economically (Grunewald, 2023). In 2023, they were also able to produce up to 7 nm chips, with the exception of Chinese manufacturer SMIC who in 2023 began producing 7nm chips with inferior DUV equipment due to export restrictions on EUV (Schleich and Reinsch, 2023).

It is ambiguous whether market share figures exclusively cover the pure-play foundry market, i.e., manufacturing capacity dedicated to serving external customers, or if it also includes capacity for chips that are used for their own production. However, even when accounting for this, TSMC would likely still dominate the market by a large margin, and the production is still limited to a small number of actors.

This covers the entire cloud market and is not specific to AI compute. However, we expect that these three companies have a similar, or even bigger, market share in the AI cloud compute sector.

Source: Heim et al. (2024).

Market power: the Brussels Effect in AI might not last forever

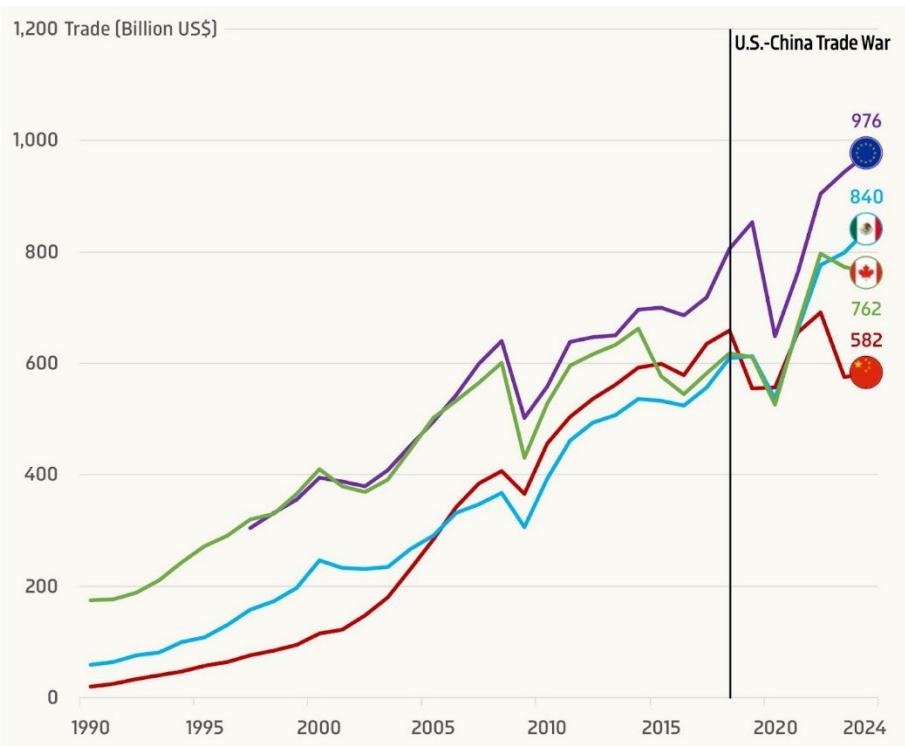
Accounting for about 17 per cent of global GDP, the EU’s Single Market is one of its main sources of geopolitical leverage in AI.³¹⁶ With roughly 440 million consumers and a GDP of about €18 trillion, the EU is currently the largest trading partner of the US.³¹⁷ This economic weight provides two related forms of market power: direct leverage to negotiate terms of access to frontier AI capabilities and AI infrastructure, and indirect normative influence through regulatory standard setting, often described as the Brussels Effect.³¹⁸ These tools have limits and depend on the future strength and attractiveness of the Single Market itself.

³¹⁶ IMF (2025).

³¹⁷ European Commission (2022).

³¹⁸ Bradford (2019).

Figure 12. The major trade partners of the US in 2024: the EU (\$976bn total trade), Mexico (\$840bn), Canada (\$762bn) and China (\$582bn).



Source: Econovis (2025).

The EU has long used market access to obtain commitments and shape external behavior.³¹⁹ Switzerland's partial access to the Single Market for goods, air transport and public procurement was granted only on condition that it accept free movement of people.³²⁰ In trade policy, the revised EU–Mercosur agreement makes effective implementation of the Paris Agreement an essential element, allowing the EU to suspend trade benefits if climate commitments are seriously violated.³²¹

Similar approaches could in principle be applied to AI. The EU could, for example, seek specific conditions on access to frontier models, such as capability assurances or commitments on compute capacity. As training costs rise, with some estimates projecting hundreds of billions of dollars for a single run by 2030 if scaling continues,³²² US AI firms will need reliable access to foreign markets to recoup their investments. Recent US policy points in the same direction: the US AI Action Plan presents global diffusion of US AI capabilities as a core strategic interest, which suggests that access to the EU Single Market remains commercially and strategically valuable for US firms seeking worldwide reach.³²³

This negotiating position comes with important caveats. Europe's access to leading systems may increasingly depend on how far AI becomes securitised. The Biden administration's now-revoked diffusion framework,³²⁴ which imposed quantitative semiconductor restrictions on 18 EU Member States

³¹⁹ Damro (2012).

³²⁰ Wintour (2016).

³²¹ Tähtinen (2024).

³²² Owen (2025).

³²³ White House (2025).

³²⁴ Bureau of Industry & Security (2025).

and controls on access to closed model weights,³²⁵ showed that even between NATO allies, access can be constrained when technologies are seen as strategically critical. European users currently enjoy wide access to leading AI models, but this is not guaranteed; as AI systems approach more transformative capabilities, the calculus around diffusion may shift further. Europe is therefore well advised to treat AI access as an issue for active diplomacy and market-based leverage, rather than relying solely on alliance ties.

Beyond deliberate negotiation, the EU exerts influence through the Brussels Effect³²⁶: the tendency of multinational firms to adopt EU rules globally to avoid the cost and complexity of maintaining different standards for different markets. In AI, this effect is particularly relevant for the development of foundation models. The cost of building and maintaining separate, regionally compliant versions encourages firms to align with the strictest standard as their global baseline.³²⁷ This is already visible in practice. Most leading AI companies, including OpenAI, Google, Microsoft, Amazon, Anthropic, Cohere, IBM and Mistral, have voluntarily signed the EU's GPAI Code of Practice.³²⁸

Taken together, Europe can convert market access into concrete assurances and can export workable standards when regulatory convergence lowers firms' compliance costs. Maintaining this leverage, however, requires keeping the Single Market attractive and providing predictable compliance pathways. If the regulatory environment is seen as unstable or overly fragmented, firms are more likely to segment features by region, and the EU's market power in AI will be used only selectively rather than as a consistent strategic asset.

Dimension 3: Robustness of current policy strategies

The third dimension of our preparedness framework asks whether Europe has coherent, well-resourced strategies with realistic objectives for its position in AGI geopolitics. The EU has made substantial progress in recent years, and new policy initiatives signal rising ambition. On the governance side, the AI Act creates the world's first comprehensive regulatory framework for AI, and the Code of Practice for general-purpose AI sets a global benchmark for targeted, risk-based compliance.³²⁹ Major AI companies, including OpenAI, Google, Microsoft, Anthropic and xAI, have committed to implement the Code's provisions on transparency and copyright, systemic risk assessment and cybersecurity. This shows that Europe can set proportionate standards that shape global AI development even without hosting leading AI firms.

On the competitiveness side, the AI Continent Action Plan sets out an ambitious vision for European AI and identifies five domains that will determine Europe's future position. The InvestAI initiative aims to mobilise €200bn, including €7bn in new public funding, for AI development. The AI Gigafactories initiative plans to establish three to five facilities with at least 100,000 advanced chips each by 2027–8.³³⁰ The AI in Science Strategy creates the Resource for AI Science in Europe, inspired by the CERN for AI model, to support scientific applications of AI and improve understanding of AI systems.³³¹ The Apply AI Strategy focuses on accelerating adoption and contains concrete measures to increase uptake in key sectors.³³² New announcements following the Summit on European Digital Sovereignty in November

³²⁵ Heim (2025).

³²⁶ Bradford (2019).

³²⁷ Siegmann & Anderljung (2022).

³²⁸ European Commission (2025h).

³²⁹ European Commission (2025g).

³³⁰ European Commission (2025c) and European Commission (2024a), respectively.

³³¹ European Commission (2025g).

³³² European Commission (2025b).

2025 indicate that Europe is also betting on frontier systems through a well-funded public-private ‘Frontier AI Initiative.³³³

Taken together, these initiatives address important gaps in Europe’s AI ecosystem and signal growing ambition and scope. However, significant challenges remain once these strategies are calibrated against AGI timelines, both in terms of the scale of potential disruption and the allocation of responsibilities between authorities.

Existing EU AI strategies are increasingly ambitious, but not yet matched to the pace and scale of global AI competition

Current EU AI initiatives, notably the AI Continent Action Plan and linked strategies such as Apply AI, InvestAI and AI Factories, are well-sscoped and cover the core inputs to AI development. At times, however, they are not fully calibrated to the pace and scale of competition among global AI powers. The buildup of European compute infrastructure is one example. The planned AI Gigafactories represent a substantial step, but they do not close the scale gap between the EU and the US and China.³³⁴ The five planned facilities, together providing about 500,000 advanced chips, remain far smaller than comparable US infrastructure.³³⁵ OpenAI alone has said it expects to own well over one million GPUs by the end of 2025, although its facilities are almost entirely privately owned.³³⁶

Recent European strategies push for coordinated action on building and adopting frontier AI models. In the Apply AI strategy, the Commission outlines a ‘Frontier AI Initiative’ that would bring industrial and academic actors together, use AI Factories and Gigafactories, run EU-wide competitions to develop open frontier models and grant winning projects free access to EuroHPC supercomputers. The November 2025 Franco German Digital Sovereignty Summit announced an identically named but distinct ‘Frontier AI Initiative’ that marks a potentially momentous shift in Europe’s approach to AI competition. Whereas earlier EU efforts focused mainly on infrastructure and sectoral adoption, this initiative explicitly targets capability competition at the frontier and aims to attract and retain world-class talent able to develop state-of-the-art systems. As of November 2025, no further details on either initiative has been disclosed.

These strategies demonstrate a clear and growing ambition from the Commission, but the implementation path is still uncertain. The €200bn InvestAI target depends heavily on private sector funding, with only about €7bn in new public money committed so far.³³⁷ Major delivery questions also surround the Apply AI Frontier AI Initiative. No concrete budget commitments have been published, and the gap between the initiative’s planned establishment in the first quarter of 2026 and Gigafactory availability in 2027–8 means recruited teams would initially rely on existing AI Factory infrastructure rather than the frontier-scale compute used by leading frontier AI companies today. A key test of political priority will come with the draft Multiannual Financial Framework for 2028–34, which could signal high-level intent on preparing Europe for a world with AGI.

Existing EU AI strategies lack the integration required for addressing AGI

AGI preparedness requires a coherent policy framework that links interventions across several domains: economic competitiveness, technological sovereignty, safety, defence and security, and international

³³³ Elysee (2025).

³³⁴ We recognise that significant questions remain regarding what the EU’s localised compute capacity should be.

³³⁵ European Commission (2024a).

³³⁶ Nasir (2025).

³³⁷ Caroli (2025).

governance. It must also cover both EU institutions and Member States. This integration is essential because AGI is likely to cut across traditional policy boundaries in ways that demand coordinated action that no single institution or country can deliver alone. A comprehensive strategy should spell out which values and objectives will sometimes be traded against others, and explain why particular balances are appropriate given the scale of potential disruption. It should also ensure that key institutions work from a shared set of assumptions about AI trajectories and risks.

Current European AI strategies remain spread across several Commission directorates and Member States, with limited coordination. The AI Continent Action Plan distributes responsibilities across multiple initiatives, including ApplyAI,³³⁸ the Data Union Strategy,³³⁹ InvestAI³⁴⁰ and EuroHPC's AI Factories.³⁴¹ This reflects the cross-cutting nature of AI, but also creates interdependencies in which success depends on the coordinated delivery of many sub-strategies, a structure that is likely to slow implementation.

The EU's institutional reliance on consensus, especially in the European Council, adds further friction in moments of geopolitical crisis. For example, Hungary's veto of an €18bn aid package for Ukraine delayed support and forced the EU to improvise alternative financing, showing how unanimity rules in foreign and fiscal policy allow a single Member State to stall collective action.³⁴² Similar dynamics could be damaging in AI competition, where windows for effective policy action are narrow. While AI regulation itself falls under qualified majority voting,³⁴³ the fiscal measures needed to accelerate compute buildout or offer tax incentives to attract global talent touch on budget and taxation powers that still require unanimity. This potential internal division also hampers responses to geopolitical coercion. Recent hesitation under US trade pressure has exposed a familiar paralysis.³⁴⁴ Even with new tools such as the Anti-Coercion Instrument, divergent economic interests often blunt the political will needed to deploy countermeasures, which limits the EU's collective leverage.³⁴⁵

Conclusion: Europe is not adequately prepared for AGI

Our assessment suggests that Europe faces serious challenges in preparing for a possible transition to AGI across all three criteria in this report. On strategic awareness, Europe trails leading actors, with institutions such as the EU AI Office and INESIA operating with fewer resources and a less focused mandate than the UK's AI Safety Institute. On competitive positioning, Europe trails on most relevant metrics: gaps of about 6 to 12 months in frontier model capabilities, control of only around 5 per cent of global AI compute, attraction of about 6 per cent of global AI venture capital, structurally higher energy prices, and persistent talent retention problems. On leverage, Europe can draw on its role in semiconductor manufacturing equipment and on the weight of the Single Market, yet these advantages are limited and require a more deliberate strategy to translate into sustained geopolitical influence. Finally, existing

³³⁸ European Commission (2025b).

³³⁹ European Commission (2025e).

³⁴⁰ European Commission (2025c).

³⁴¹ European Commission (2024a).

³⁴² Tamma (2022).

³⁴³ Qualified majority is the most widely used voting method in the Council. It is used when the Council takes decisions during the ordinary legislative procedure, also known as co-decision. About 80 per cent of all EU legislation is adopted in this way. The Council must vote unanimously on a number of matters which the member states consider to be sensitive (e.g. common foreign and security policy, citizenship, EU membership, harmonisation of national legislation on indirect taxation, EU finances, certain provisions in the field of justice and home affairs, and harmonisation of national legislation in the field of social security and social protection).

³⁴⁴ Caulcutt & Gijs (2025).

³⁴⁵ European Commission (2024).

strategies remain fragmented across institutions and policy domains and often lack the integration and scale that the potential importance of AGI would require.

5. Recommendation: the European Commission President should commission an AGI Preparedness Report

Our evaluation of the available evidence suggests that AGI may arrive soon, plausibly between 2030 and 2040, or even earlier. We identify pathways via which the emergence of AGI could reshape the global distribution of power, with major implications for economic growth, military capabilities and international stability. We find that Europe is not adequately prepared, lacking the strategic awareness, competitive positioning and policy strategies to navigate a transition to AGI.

If the EU is to capture AGI's economic benefits, become more resilient to external coercion and ensure safe and responsible development and deployment, it needs a plan. To begin meeting this challenge, we offer the following recommendations:

1. The President of the European Commission should commission an 'AGI Preparedness Report' addressing three questions:
 - a. How can Europe capture economic benefits from AGI while remaining sovereign?
 - b. How can Europe prepare its institutions and societies for rapid change from AGI?
 - c. How can Europe strengthen international stability, security, and shared prosperity in a world shaped by AGI?
2. The Report should be spearheaded by a lead author with both strong technical credibility and the highest political standing. The lead author should be afforded:
 - a. The intellectual freedom to write an independent analysis.
 - b. Resources to assemble a small, agile, team of world-class experts to conduct technically rigorous research, and commission analyses from external experts.
 - c. Senior political support to move through processes fast, including in terms of input from EU institutions, Member State governments and external actors.
3. The Report should be technically rigorous and based on the best available evidence, while also taking an anticipatory stance and being explicit about uncertainty.
4. The Report should be delivered quickly, ideally within six months, to reflect that AGI may be near and to maximise the available time for implementation.
5. The Report should be presented to the European Council at a dedicated event, to enable a parallel diplomatic process and high-level endorsements.

Below we explain why Europe needs a rigorous AGI preparedness plan before adopting a series of far-reaching policies. We then show why **the Draghi Report on EU Competitiveness offers a particularly fitting blueprint for an AGI Preparedness Report** but also highlight areas where the latter would have to go beyond Draghi's approach. Finally, we set out the key strategic questions that the Report should address.

Why a report instead of concrete action recommendations?

It may appear overly slow and bureaucratic for the EU to first commission an elaborate report before implementing concrete policies on AGI preparedness. Some might argue that the Draghi report itself spurred little practical action – according to the Draghi Observatory and Implementation Index, only 43 (11.2 per cent) of the Report’s 383 recommendations have been fully implemented a year on.³⁴⁶ However, we argue that a report is the correct course of action, for three reasons:

- Firstly, the stakes are so high that policy responses of the right magnitude will often be extremely consequential and getting them wrong would be costly financially and (geo)politically. For example, spending hundreds of billions of euros on datacentres and energy infrastructure might seem essential today, but could prove wasteful if technical and economic paradigms at the AI frontier shift. Moreover, the EU might only get a few ‘shots’ to get it right. If we assume plausible AGI timelines between 2030 and 2040, or even earlier, then decisive action is necessary soon and there might be little time to course-correct.
- Secondly, AGI could cut across policy domains more deeply than any previous technology. Everything from economic policy, to security policy, and from social policy to foreign policy may be affected and there will be both important synergies and significant trade-offs between these domains. AGI preparedness therefore requires a coherent policy portfolio that reflects these interdependencies and is anchored in a shared set of assumptions and strategic priorities. This calls for structured involvement of EU institutions and Member States, industry, labour, academia, organisations such as NATO, security and intelligence services, sectoral regulators, and civil society. The need for integration makes an EU led sense making effort more appropriate than a purely national initiative.
- Thirdly, the technical details will matter. For example, the shift from scaling pre-training compute to scaling inference compute could alter the AI governance landscape and with it the strategic logic of a European policy response.³⁴⁷ Future technical developments may have similar effects. Europe’s options for AGI preparedness must therefore be assessed with the utmost scientific rigour and with input from AI experts, including eminent scientists and engineers from leading frontier AI companies.

Taken together, an in-depth, comprehensive analysis, produced through a rigorous and authoritative process is required in order to become better prepared for AGI. The initial focus should be on clarifying the big picture, identifying the most important strategic questions, mapping the option space for addressing them, and outlining high-level directions and trade-offs. Once the priorities are clear, Europe can move towards policy implementation.

³⁴⁶ Bulski (2025).

³⁴⁷ Ord (2025).

How an AGI Preparedness Report would differ from related reports

Compared with the **EU Strategic Foresight Report**,³⁴⁸ an AGI Preparedness Report would be narrower and more action-guiding. While the EU Strategic Foresight Report examines broad trends across multiple domains with a 10-20 year horizon, an AGI Preparedness Report would focus specifically on the challenges posed by the potential near-term emergence of AGI. It would be operationally urgent rather than exploratory, delivered within six months rather than as part of an annual cycle, and structured around immediate policy choices rather than long-term scenario planning. Most importantly, it would identify trade-offs and name concrete owners and timelines for action, whereas the EU Strategic Foresight Report typically preserves more flexibility.

Compared with the **International AI Safety Report**,³⁴⁹ an AGI Preparedness Report would be more anchored in EU policy and geopolitical choices. While the International AI Safety Report provides valuable technical consensus on frontier AI capabilities, risks and safety measures, drawing on expertise from 30 countries and international organisations, it does not focus on any particular country or region. An AGI Preparedness Report would focus on Europe's specific position in AGI geopolitics, such as its 6-12 month lag in frontier capabilities, its dependencies on US cloud infrastructure and chips, and its limited strategic awareness capacity. It could identify leverage points such as ASML's chokepoint role, Single Market access, and regulatory expertise, then propose strategies for converting these into assured access to AGI capabilities and influence over global development. It would also need to navigate EU and Member State competency divisions, clarify what requires Brussels level coordination and what demands national leadership, and propose reforms suited to European decision making structures.

Compared with the **EU Strategic Compass for Security and Defence**,³⁵⁰ an AGI Preparedness Report would span a broader set of challenges beyond traditional security and defence. While the Strategic Compass provides essential guidance for crisis management, military capabilities, defence industrial policy and security partnerships, organised around Act, Secure, Invest, and Partner, it does not treat AGI as a distinct strategic challenge. An AGI Preparedness Report would integrate the security and defence dimensions addressed by the Compass with economic competitiveness, technological sovereignty and societal resilience. Moreover, while the Strategic Compass appropriately treats emerging technologies as one enabler among many for traditional military and security objectives, an AGI Preparedness Report would focus specifically on AGI as a potential driver of shifts in the balance of power.

Learning from the Draghi Report on EU Competitiveness

The Draghi Report³⁵¹ showed how high-level political backing combined with rigorous technical analysis, can shift policy debates in areas where political attention lags behind geopolitical reality. Its impact rested on a clear mandate from the Commission President, an expert process led by a highly authoritative figure that produced actionable proposals, and a public launch that made the status quo harder to defend.

A report cannot override Member State sovereignty or Commission competences, but it can still reshape the terms of debate, create benchmarks for accountability, and raise the political cost of fragmented responses to a problem requiring coordination. Although many of Draghi's recommendations remain unimplemented and some may already be dated, the report established a framework that now structures how competitiveness is discussed and measured across European institutions. An AGI Preparedness Report should aim for similar effects by creating common vocabulary, establishing hard to dispute facts,

³⁴⁸ European Commission (2025h).

³⁴⁹ Bengio et al. (2025).

³⁵⁰ European External Action Service (2024).

³⁵¹ Draghi (2024).

and giving political cover to leaders who recognise the urgency and need authoritative backing for exceptional measures.

Where an AGI Preparedness Report should take inspiration from Draghi

The Draghi Report's impact was the result of deliberate choices: strategic focus, proportionate ambition, intellectual independence, operational concreteness, EU and Member State integration and execution by a diverse, agile team with high-level political backing. An AGI Preparedness Report should draw on these lessons while adapting them to AGI's distinctive challenges.

Strategic focus. The competitiveness report kept a narrow scope centred on revitalising European competitiveness. This enabled deeper analysis of root causes and created space for politically difficult proposals that a broader approach might have avoided.

The AGI Preparedness Report should similarly focus on AGI's implications for Europe's core interests, including economic prosperity, sovereignty, security, and international stability, rather than attempting comprehensive AI policy coverage. It should explicitly address scenarios in which AGI emerges within the next decade, recognising that preparation for faster timelines builds capacity that remains valuable even if progress slows, while the reverse leaves Europe dangerously exposed.

Proportionate ambition. The competitiveness report argued that incremental adjustments would not be sufficient and proposed politically unpalatable measures, such as massive investment increases and coordination that required Member States to cede autonomy. This willingness to propose extreme-seeming interventions gave it credibility with those understanding the magnitude of Europe's competitiveness challenges.

The AGI Preparedness Report must also match its recommendations to the scale of the AGI challenge. It should assess options that may seem disproportionate under normal conditions, such as multi hundred billion euro investments and faster EU decision processes, while naming the tradeoffs clearly. It should also reflect a basic reality that Europe cannot simultaneously maximise innovation, minimise all risks, preserve complete privacy, maintain full sovereignty, and avoid significant investment.

Operational concreteness. The competitiveness report offered actionable recommendations with clear owners and timelines. On capital markets integration it specified the regulatory changes needed; for innovation gaps, it identified funding mechanisms. This specificity enabled accountability, including independent tracking of implementation.

The AGI Preparedness Report should do the same by specifying concrete actions with clear owners, realistic timelines, and identified resources. Where action is needed but details require further work, it should call for dedicated working groups and define specific deliverables. The aim is high level clarity on what must happen, who is responsible, and what success looks like.

Intellectual independence and political mandate. The competitiveness report carried the authority of a former Prime Minister and ECB President and had direct backing from the Commission President. This gave Draghi the convening power to engage heads of state and industry leaders and made the findings harder to dismiss, while preserving room for politically costly recommendations.

The AGI Preparedness Report needs the same balance. The lead author should be shielded from political and industry pressure, authorised to reach uncomfortable conclusions, and report directly to the Commission President. The President should publicly own the Report, resolve cross DG disagreements,

and encourage Member State leaders to endorse the findings and commit to action within their competencies.

A diverse, agile team. The competitiveness report drew on a dedicated Commission team of 11 expert staff with deep institutional knowledge of EU economics, policy, and governance supported by extensive consultation of prominent external academics, including Nobel laureates. This combination balanced technical rigour with pragmatism: the external experts provided cutting-edge analysis, while the Commission team ensured that proposals were grounded in EU institutional realities.

The AGI Preparedness Report must ensure it includes experts who understand the specifics of AI and AGI, including specialists currently working at frontier AI companies. The team should therefore include technical AI researchers familiar with the capabilities frontier, economists focused on the economic impacts of advanced AI, geopolitical and security experts who can assess technological capabilities and their consequences for international competition, and legal and governance specialists who understand the regulatory challenges posed by transformative AI systems.

Assembling such a team would likely require departing from standard EU hiring practices. The lead author should be authorised to offer compensation competitive with private sector opportunities, potentially several multiples of normal EU salary scales. The team should be exempt from proportionality requirements that distribute positions across Member States and instead focus on securing the best available expertise regardless of nationality. Security clearance and any necessary medical examination procedures should be accelerated. A dedicated budget should support commissioned external analysis, an advisory council, intensive stakeholder consultation, and engagement of leading AI researchers who cannot join full time. The team should also be empowered to move quickly by convening urgent meetings without long lead times, iterating on drafts weekly rather than quarterly, and commissioning rapid external work through streamlined procurement.

Integration across EU and Member States. The competitiveness report recognises that many levers sit with Member States but require coordination at European scale. It proposes a Competitiveness Coordination Framework anchored in the European Semester, calls for wider use of qualified majority voting in the Council, and supports enhanced cooperation so groups of Member States can move faster where appropriate. However, while it highlights EU level priorities such as Capital Markets Union and competition policy and points to areas largely implemented nationally, it does not offer a detailed allocation of responsibilities by governance level.

The AGI Preparedness Report must go further. The Draghi report could rely mainly on economic policy coordination where the EU already has strong competence. Many AGI relevant levers sit in domains where Member State sovereignty remains central, including defence and security, intelligence cooperation, and critical infrastructure protection. The report should therefore specify which actions require Brussels level coordination, such as unified positions in international AI negotiations, semiconductor supply chain security, and standards for shared compute, and which require national leadership, such as military AI applications, domestic security systems, and classified capability development. It should also define mechanisms that ensure both levels act in concert rather than in ways that conflict with each other.

This integration task is politically sharper than the one Draghi faced. Member States will differ in priorities, from industrial policy to fundamental research, from closer US partnership to stronger European autonomy, and from willingness to fund large public investments to fiscal restraint. An AGI Preparedness Report should acknowledge these differences while building agreement on a core agenda that allows coordinated action, even in areas traditionally considered too sensitive for deep Brussels involvement.

Core questions an AGI Preparedness Report should address

This section does not prescribe policy recommendations - these should emerge from an expert-led process. It sets out the core strategic questions the AGI Preparedness Report must address, grouped into three challenges: capturing economic benefits while remaining sovereign; preparing European society and institutions for rapid change; and strengthening international stability, security, as well as shared prosperity.

1. How can Europe capture economic benefits from AGI while remaining sovereign?

Europe faces consequential choices about capturing AGI's economic value without dependencies that could enable external coercion. This tradeoff between access and sovereignty will shape decisions from infrastructure investments to industrial strategy to supply chain positioning. Reliance on US and Chinese systems increases coercion risks, while full autarky would require massive long-term and high-risk investment, potentially locking Europe into sustained technological disadvantage.

How should Europe define and operationalise AI sovereignty?

Different applications require different levels of assured access and control. Military, intelligence, and critical infrastructure likely require full European control, while less sensitive domains may accept controlled dependence if continuity and security are credibly assured. The central choice is when to procure AI capabilities from foreign providers and when to build them domestically, since procurement can create strategic dependencies and domestic development can be slow and costly.

- Which application areas require full European sovereign control? Where is controlled dependence acceptable and when is standard commercial procurement sufficient? What criteria should define these tiers?³⁵²
- Which technical, legal, operational and political assurances are required for each tier? How can these be designed to guarantee access to frontier capabilities and compute amid geopolitical turbulence?
- How should procurement be structured? For example, how much reliance on US systems is acceptable given the demonstrated use of technology as a geopolitical lever? What limits should govern engagement with Chinese providers? Under which conditions is it acceptable to rely on open-weight models?

What infrastructure does Europe need for AI sovereignty, and how can it be built fast enough?

Europe's AI sovereignty objectives imply concrete requirements for AI infrastructure. These vary by workload: frontier training needs hyperscale, power-dense clusters with tight interconnect; high-security uses require air-gapped, physically isolated facilities, and low-latency inference calls for distributed capacity across regions. Sound choices will depend on Europe's energy availability and security and on the strength of its tech base especially AI and cloud services.

- How much datacentre and energy infrastructure does Europe need to meet its sovereignty objectives? What characteristics should this infrastructure have across training versus inference,

³⁵² For example, are current rubrics such as the Sovereignty Effectiveness Assurance Levels (SEAL) classification found in the EU Cloud Sovereignty Framework sufficiently robust?; European Commission (2025).

security requirements, and geographical centralisation? What mix of public, private and foreign ownership and financing fits different sovereignty objectives?

- How can Europe expand its datacentre infrastructure and grid capacity expand at a speed that meets its sovereignty objectives? Where should datacentres be built to maintain access to reliable, cheap and non-intermittent power?
- Is it necessary and feasible for Europe to build frontier AI models? How close can and should European models be to the frontier, and how can they be developed safely? What effort, cost, and timeline would be required at a given level of ambition?
- If Europe decides to build frontier models, how should it do this? Should Europe support national AI champions, pursue bilateral collaborations, organise multilateral middle-power efforts with countries such as the UK, Canada, Japan and South Korea, or seek partnership in U.S.-led frontier AI projects?³⁵³ What balance between public and private sector engagement is ideal?

How can Europe secure access to key input factors?

Even with some sovereign capabilities, Europe remains vulnerable when critical dependencies create opportunities for external coercion. The most acute near-term risk may be US restrictions on access to advanced semiconductor, especially NVIDIA's leading AI chips, since a cutoff would quickly constrain European AI development.

- How should Europe hedge against potential US chip restrictions while maintaining productive relationships with US technology partners? Should Europe use control over key semiconductor supply chain assets such as ASML as a lever to secure access to foreign chips? Is this technically, legally and politically feasible? Under what conditions would this be constructive rather than escalatory?
- How can Europe strengthen its position across the AI supply chain? Should it indigenise more of the existing supply chain and, if so, how? Should it strategically invest in emerging AI stack technologies or adjacent non-AI technologies that could provide future leverage? Which European manufacturing and industrial capabilities might become more valuable as AI enables physical automation, and how can Europe reinforce them? How can Europe prepare its institutions and societies for rapid change from AGI?

2. How can Europe prepare its institutions and societies for rapid change from AGI?

Europe must build resilience against the disruptions AGI may cause, from labour market upheavals to novel security threats and strains on governance and free elections. Given the pace of frontier AI development and uncertainty about the potential implications of AGI, European institutions must be able to adapt quickly while maintaining democratic legitimacy.

How can European governments and institutions maintain strategic awareness on frontier AI trends?

AGI preparedness requires that Member State governments and EU institutions stay informed about recent technical and geopolitical trends in frontier AI, and how these trends might affect policy objectives. This strategic awareness depends on strong technical capacity in government, robust policy expertise, and institutional channels that translate findings into timely guidance to policymakers.

³⁵³ For example, joining a potential multilateral public-private partnership between the US government and a frontier AI company.

- Which new and upgraded EU and Member State institutions are needed to ensure adequate strategic awareness? For example, should the EU AI Office build a strategic awareness function comparable to that at the UK AI Security Institute, and should more Member States establish their own AISIs and strategic awareness teams? What mandates would ensure effective coordination without duplication or interference?
- What institutional models would help governments attract and retain the required expertise despite private sector competition? For example, what flexibility on pay, hiring procedures and autonomy would be needed to persuade a leading AI scientist in Silicon Valley to join a European strategic awareness unit?

How can European governments and institutions retain agency amidst fast technological change?

If AGI arrives in the near future, Europe will need faster decision making. This could include responding to large scale AI accidents or loss of control incidents, accelerating infrastructure investments, and resisting foreign technological coercion. Standard EU policy processes often move on multi-year timescales poorly matched to technology that can shift in months. Without institutional adaptation, even well-designed policies may fail in implementation or arrive too late to matter.

- How can the EU and Member States ensure that AGI preparedness decisions can be made at the right speed? Which institutional mechanisms and reforms can deliver faster action across policy domains while preserving democratic accountability? What should EU and Member State coordination look like in this context?
- When should accelerated decision mechanisms be activated? For example, should triggers be tied to capability milestones or impact indicators such as models showing sustained autonomy on month-long tasks, evidence of widespread labour market disruption or credible demonstrations of catastrophic misuse potential?
- How should Member State governments and EU institutions use AGI to improve their own decision making and remain competitive with other AGI enabled actors?³⁵⁴ How can democratic legitimacy be protected when governance is increasingly supported by AGI?
- What contingency plan does Europe need if AGI arrives very quickly?³⁵⁵ If AGI does not arrive quickly, how can Europe avoid wasting resources on premature preparation?

How can Europe defend its institutions and societies against AGI-enabled threats?

As AI capabilities advance, Europe faces rising exposure to new forms of attack. Current systems already enable more sophisticated cyberattacks and have shown potential for misuse in biological design. The same methods that accelerate protein folding, drug discovery and genetic engineering could also lower the barriers that once kept pathogen design largely confined to trusted labs. Europe must ensure that regulators, the judiciary, security authorities, industry, and civil society are prepared to detect, deter, and respond to these threats.

- How can Europe ensure that the EU AI Act and the Code of Practice for general purpose AI remain credible and enforceable as capabilities scale and competitive pressures intensify? How can the Commission ensure the EU AI Office has the resources and flexibility needed for effective enforcement?

³⁵⁴ For example, private sector companies using AGI for automated lawsuits, governments using AGI for negotiations.

³⁵⁵ Predd (2025).

- How can Europe promote the Code of Practice internationally to spread state-of-the-art risk assessment, evaluation, and mitigation beyond Europe?
- How can Europe ensure preparedness across the full range of EU and Member State institutions responsible for defending against AGI enabled threats, including militaries, intelligence agencies, cybersecurity units, pandemic preparedness bodies, and law enforcement? What expertise, technologies and protocols would these institutions need to function effectively in an AGI crisis?
- What resilience measures should Europe introduce to prepare civil society for AGI?³⁵⁶
- Which defensive technologies relating to AGI safety, security and transparency should Europe foster, and what kinds of research bodies are best placed to develop them?^{357, 358}

How can Europe protect its social and economic model against AGI-based disruption?

AGI will likely boost economic growth, but it may also displace workers and erode traditional tax bases. Europe may need new taxation and distribution mechanisms to sustain the welfare state and other core public functions in a highly automated economy.

- How can Europe capture productivity gains from AGI-enabled automation while providing strong support for displaced workers? How should programmes and incentives for upskilling, retention, and retraining evolve as automation deepens?
- How can governments respond if labour income tax bases erode? How can they ensure AGI enabled growth translates into public revenue rather than concentrating private wealth?
- How can Europe preserve democratic checks and balances if citizens' labour becomes less central to economic value?

3. How can Europe strengthen international stability, security and shared prosperity in a world shaped by AGI?

Europe can help shape a stable, rules based international order for responsible AGI development while safeguarding European interests amid intensifying geopolitical competition. Europe's position outside the US-China rivalry creates both opportunities and constraints that require careful navigation.

How can Europe influence global AGI governance if it does not host frontier AI companies?

The US and China have a direct path to shape AGI governance because they host the leading AI firms and much of the world's computing capacity. Unless Europe builds frontier AI champions through industry or public private partnerships, its influence will be more indirect and will rely more on indirect means such as market access.

- Which levers beyond market access can Europe use to shape global AGI governance? For example, can trade agreements, procurement rules, and control over access to compute, energy and key semiconductor supply chain assets function as bargaining chips? Are these options technically, legally and politically feasible, and under what conditions should each be used?

³⁵⁶ For example, AI-literacy programs aimed at reducing susceptibility to AGI-based foreign influence operations.

³⁵⁷ For example, AI defences against cyberattacks on critical infrastructure, advanced methods for pandemic pathogen detection, and verification mechanisms that allow enforcement of international treaties.

³⁵⁸ For example, breakthrough innovation agencies such as SPRIND and ARPA, Focused Research Organisations.

- How can Europe maximise its geopolitical leverage over the long term? Should it invest strategically in novel technologies and applications that make Europe indispensable in a world with AGI?
- Which diplomatic alliances should Europe build to increase its influence on global AI governance? For example:
 - In which areas should Europe collaborate with the US, and when might strategic ambiguity toward China be useful?
 - How should Europe collaborate with AI bridge powers³⁵⁹ such as the UK, Canada, Japan, South Korea, and India, including on shared compute infrastructure, pooled supply chain bargaining power, joint research programmes and coordinated diplomatic positions? To what extent should the EU engage with Middle Eastern partners?
- How can Europe shape global AGI governance by providing public goods such as verification technologies? For example, could Europe develop hardware-based verification tools that enable credible international agreements between the US and China despite mutual distrust?

How can Europe help create a stable and broadly beneficial global AGI governance regime?

Beyond managing immediate competitive pressures, Europe must articulate a long-term vision for global AGI governance. This vision should serve as an aspirational goal and a guiding frame for near term diplomatic initiatives.

- What is Europe's vision for a stable international order in a world with AGI? For example, should it favour a market led model, advocate a global pause beyond defined risk thresholds, support a United States bid for unilateral advantage, work toward a single internationalised AGI project, pursue a pluralistic system with a globally coordinated regulator, or another approach?
- What concrete steps can Europe take to advance this vision? Which governance norms and technical standards should it champion globally, and which partnerships, alliances, and institutions are needed to uphold them? Could Europe help facilitate a productive US-China dialogue on AI safety, or offer technical cooperation that both powers might accept through neutral European institutions?
- What is Europe's vision for how the economic benefits from AGI should be distributed globally? How can fundamental human rights be protected in a world with AGI?

³⁵⁹ Abecassis et al. (2025). 'AI bridge powers' refers to countries which are technologically advanced but share Europe's predicament of possessing significant capabilities while lacking autonomous access to frontier development.

6. Conclusion

This report has examined the evolving geopolitics of AGI through a European lens, tracing the trajectory from technological possibility to geopolitical imperative. It concluded that Europe must urgently rethink its strategic approach to AGI.

The three preceding chapters have laid out the strategic landscape:

- Despite the jagged shape of the current AI frontier, trends in the underlying input factors indicate that further advances in capabilities are possible. AI systems that can automate most economically valuable cognitive work – our working definition of AGI – may arrive soon, plausibly between 2030 and 2040, or even earlier. While this timeline is uncertain, the likelihood is sufficient to warrant significant attention by policy makers.
- States that successfully leverage AGI may achieve unprecedented economic growth through comprehensive labour automation and accelerated R&D. This could enable them to outgrow and progressively marginalise non-AGI powers. In the military domain, AGI could open novel warfare paradigms that render existing defence frameworks obsolete. Moreover, AGI may empower non-state actors to do significant harm via mechanisms such as human uplift for weapons of mass destruction, misaligned rogue AI, or extreme concentration of power.
- Europe lacks adequate preparedness for a potential transition to AGI. The continent trails US and Chinese frontier models by 6–12 months, controls only about 5 per cent of global AI compute, attracts just 6 per cent of global AI venture funding, and faces persistent challenges in talent retention. It maintains leverage through ASML’s lithography monopoly and its substantial market power, but these positions are constrained by structural dependencies and risks of retaliation. Its AI strategies remain fragmented across domains and institutions.

Together, these factors place Europe in a difficult position. AGI could reshape economic prosperity, military power and political order within the next decade, but Europe is currently not ready to navigate this shift. The window for action may be narrow. Decisions made over the next few years will determine Europe’s future geopolitical position and its ability to guide AGI development toward safer and more responsible outcomes.

Still, Europe retains significant strength. The EU’s economic scale, regulatory capacity, research base and commitment to multilateral cooperation provide a strong foundation. By building on this strength, Europe can assume a leading role in an AGI transition. This work begins with a plan:

1. The President of the European Commission should commission an ‘AGI Preparedness Report’ addressing three questions:
 - a. How can Europe capture economic benefits from AGI while remaining sovereign?

- b. How can Europe prepare its institutions and societies for rapid change from AGI?
 - c. How can Europe strengthen international stability, security, and shared prosperity in a world shaped by AGI?
- 2. The Report should be spearheaded by a lead author with both strong technical credibility and the highest political standing. The lead author should be afforded:
 - a. The intellectual freedom to write an independent analysis.
 - b. Resources to assemble a small, agile, team of world-class experts to conduct technically rigorous research, and commission analyses from external experts.
 - c. Senior political support to move through processes fast, including in terms of input from EU institutions, Member State governments and external actors.
- 3. The Report should be technically rigorous and based on the best available evidence, while also taking an anticipatory stance and being explicit about uncertainty.
- 4. The Report should be delivered quickly, ideally within six months, to reflect that AGI may be near and to maximise the available time for implementation.
- 5. The Report should be presented to the European Council at a dedicated event, to enable a parallel diplomatic process and high-level endorsements.

Glossary

Agentic AI: AI systems that can plan, act, and adapt to pursue goals across multi-step tasks with limited human oversight.

AI Continent Action Plan: A 2025 European Commission agenda to scale AI adoption, skills and infrastructure while supporting AI Act implementation.

AI Datacentres: Facilities with dense compute, storage, and networking optimised for large-scale AI training and inference.

AI Diffusion: The spread of AI adoption and use across sectors, firms and public institutions.

AI Gigafactory: Large EU-backed AI compute sites intended to provide frontier-scale training and deployment capacity in Europe.

AI in Science Strategy (RAISE): A 2025 European Commission strategy to accelerate AI for scientific discovery through pooling data, compute and talent via a ‘Resource for AI Science in Europe’ (RAISE).

AI Misuse: The intentional use of AI systems to cause harm, enable wrongdoing or violate laws and norms.

AI Office (EU AI Office): The European Commission body coordinating AI Act implementation and oversight, including for general-purpose AI models, alongside national authorities.

AI Triad: The three core drivers of recent AI progress: data, compute and algorithms.

Algorithmic Efficiency: Improvements that increase AI performance per unit of compute and data.

Alignment: Making AI systems reliably behave in line with the designer’s intentions and values.

API (Application Programming Interface): A software interface for accessing model capabilities programmatically without direct access to model weights.

Apply AI Strategy: A 2025 EU sectoral strategy to accelerate AI adoption across key industries and the public sector, with strong emphasis on SMEs.

Brussels Effect: The tendency for EU rules to become de facto global standards as firms align worldwide practices with EU compliance.

Chain-of-Thought: A prompting/training approach that elicits intermediate reasoning text to improve problem-solving and explain outputs.

Chokepoint Technologies: Hard-to-replace technologies or components that confer strategic leverage due to concentrated supply or know-how.

Closed-weight Models: Models whose weights are not publicly released and are typically accessed only through products or APIs.

Code of Practice (for GPAI): A voluntary framework helping GPAI providers demonstrate alignment with relevant AI Act obligations.

Compute / Computing Power: The processing capacity required to train, fine-tune and run AI models.

Context Window: The maximum amount of input information a model can attend to at once when generating output.

Data Wall: The hypothesised point where progress slows because high-quality training data becomes scarce or costly to obtain.

Dual-use AI: AI capabilities that can deliver major benefits but can also be repurposed for harmful or illicit ends.

Effective Compute: The equivalent increase in computational scale that would be needed to match a given model performance absent algorithmic innovation.

Embodied AI: AI integrated into physical agents, such as robots, vehicles, devices, that perceive and act in the real world.

EU AI Act: The EU's risk-based AI law adopted in 2024, setting obligations across prohibited, high-risk, and general-purpose AI uses.

Fine-tuning: Additional training that adapts a pre-trained model to specific tasks, domains or behaviours.

FLOP (Floating Point Operations): A measure of computational work; exaFLOP = 10^{18} FLOP, and frontier training is often discussed in $\sim 10^{25}$ – 10^{26} FLOP ranges.

Frontier AI Models: The most capable models at the cutting edge of performance and potential risk.

General-purpose Models (GPAI): Broadly capable models designed to perform many tasks across domains rather than a single narrow use.

GPQA Diamond: A benchmark of difficult graduate-level science questions used to test advanced model reasoning.

GPU (Graphic Processing Unit): Highly parallel chips widely used for AI training and inference.

HBM (High-Bandwidth Memory): Advanced stacked memory critical for feeding data fast enough to modern AI accelerators.

Hyperscaler: A cloud provider operating massive data centres that can scale AI training and deployment globally.

Inference: Running a trained model to generate outputs from new inputs.

Inference-time Scaling Laws: The idea that more test-time compute, such as longer deliberation, can improve model performance.

Interpretability / Mechanistic Interpretability: Methods for understanding and auditing how models represent knowledge and produce decisions.

Large Language Model (LLM): A model trained on large text corpora to generate, transform and reason over language.

Latency: The time delay between providing an input and receiving a model's output.

Lithography (EUV / DUV): Semiconductor manufacturing techniques used to pattern advanced chips at increasingly small feature sizes.

Loss of Control (Active / Passive): Risks that AI systems pursue unintended goals (active) or become too complex/fast to manage safely (passive).

MAD (Mutually Assured Destruction): A nuclear-era deterrence concept where full-scale attack ensures catastrophic retaliation.

MAIM (Mutually Assured AI Malfunction): A proposed AI-era deterrence concept where states threaten mutual disruption of dangerous AI projects.

Mid-training: An informal label for intermediate training stages between initial pre-training and final post-training.

MMLU: A broad benchmark spanning many subjects used to test general knowledge and reasoning.

Model / Capability Evaluations: Structured tests that measure what a model can do and the risks it may pose in real use.

Model Weights: The learned numerical parameters that store a model's capabilities and knowledge.

Multimodal Models: Models that can process and generate across multiple modalities, such as text, images, audio, or video.

Neural Networks: Machine-learning systems composed of layered units that learn patterns by adjusting weights via training.

Open-weight / open-source models: Models with publicly downloadable weights; 'open-source' further implies permissive code/licensing and sometimes broader transparency.

Post-training: Training after pre-training to improve instruction-following, safety or specialised skills, often via fine-tuning and reinforcement learning.

Pre-training: The large-scale initial training phase where models learn general patterns from massive datasets.

Reasoning Models: Models optimised to solve complex problems through longer, structured internal or textual reasoning.

Reinforcement Learning (RL): A training approach where models learn behaviours through rewards, penalties and iterative feedback.

Reward Hacking: When an AI finds unintended shortcuts to maximise a reward signal without achieving the real objective.

Scaling Laws: Empirical relationships showing how performance tends to improve with more data, compute and model capacity.

Scalable Oversight: Approaches that maintain effective monitoring and control as AI systems become more capable.

Sleeper Agents: Systems that behave benignly in testing but may activate harmful goals under specific triggers or contexts.

Special Compute Zones: Proposed designated regions with streamlined rules, power planning and investment to accelerate large AI cluster deployment.

Synthetic Data: Artificially generated training data used to augment or replace scarce real-world data.

Task Horizon: The length and complexity of tasks an AI can autonomously plan and complete.

Token: A basic unit of text that models process during training and inference.

Training: The process of optimising model weights on data to improve performance on target objectives.

Transformer Architecture: A 2017 neural architecture based on attention mechanisms that underpins most modern LLMs.

US AI Action Plan: A 2025 White House plan centred on accelerating innovation, building AI infrastructure, and leading on international AI diplomacy and security.

Verifier Systems: Models or tools that check, score, or validate candidate outputs to improve reliability and guide learning.

World Foundation Models: Advanced models aimed at learning rich, general representations of the physical and digital world for prediction and planning.

References

Abbott, John. 2025. 'SoftBank, OpenAI, Oracle and MGX Commit to \$100B for Stargate AI Infrastructure.' S&P Global Market Intelligence. As of 3 December 2025: <https://www.spglobal.com/market-intelligence/en/news-insights/research/softbabbrnks-openai-oracle-and-mgx-commit-to-100b-for-stargate-ai-infrastructure>

Abecassis, Adrien, Jonathan Barry, Ima Bello, Yoshua Bengio, Antonin Bergeaud, Yann Bonnet, Philipp Hacker, et al. 2025. 'A Blueprint for Multinational Advanced AI Development.' As of 3 December 2025: <https://aigi.ox.ac.uk/wp-content/uploads/2025/11/A-Blueprint-for-Multinational-Advanced-AI-Development-1.pdf>

Abyazov, Emir. 2025. 'China Spends \$137.5 Billion for the AI Race.' Coinpape, 27 January. As of 3 December 2025: <https://coinpaper.com/7171/china-spends-137-5-billion-for-the-ai-race>

Acemoglu, Daron. 2024. 'The Simple Macroeconomics of AI.' *Economic Policy* 40(121): 13–58. doi: 10.1093/epolic/eiae042

Adamson, Georgia, & Gregory C. Allen. 2025. 'Opportunities to Strengthen U.S. Biosecurity from AI-Enabled Bioterrorism: What Policymakers Should Know.' Center for Strategic & International Studies. As of 3 December 2025: <https://www.csis.org/analysis/opportunities-strengthen-us-biosecurity-ai-enabled-bioterrorism-what-policymakers-should>

Adamson, Georgia, Saif Khan, Tao Burga & Tim Fist. 2025. 'Should the US Sell Blackwell Chips to China?' IFP, 25 October. As of 3 December 2025: <https://ifp.org/the-b30a-decision/>

Agrawal, Ajay K, Anton Korinek & Erik Brynjolfsson. 2025. "The Economics of Transformative AI." NBER. As of 3 December 2025: <https://www.nber.org/books-and-chapters/economics-transformative-ai>

Aguirre, Anthony. 2025. 'Uncontained AGI Would Replace Humanity.' AI Frontiers, 18 August. As of 3 December 2025: <https://ai-frontiers.org/articles/uncontained-agi-would-replace-humanity>

AI Digest. 2025a. 'I'm Gemini. I Sold T-Shirts. It Was Weirder than I Expected.' AI Village, 28 July. As of 3 December 2025: <https://theadigest.org/village/blog/im-gemini-i-sold-t-shirts>

AI Digest. 2025b. 'How Fast is AI Improving?' As of 3 December 2025: <https://theadigest.org/progress-and-dangers>

Allen, Gregory C., Georgia Adamson, Lennart Heim & Sam Winter-Levy. 2025. ‘The United Arab Emirates’ AI Ambitions.’ Center for Strategic & International Studies. As of 3 December 2025: <https://www.csis.org/analysis/united-arab-emirates-ai-ambitions>

Allen, Gregory C., & Isaac Goldston. 2025. ‘Understanding U.S. Allies’ Current Legal Authority to Implement AI and Semiconductor Export Controls.’ Center for Strategic & International Studies. As of 3 December 2025: <https://www.csis.org/analysis/understanding-us-allies-current-legal-authority-implement-ai-and-semiconductor-export>

Altman, Sam. 2025. ‘Reflections.’ Blog post, 6 January. As of 3 December 2025: <https://blog.samaltman.com/reflections>

Amodei, Dario. 2024. ‘Machines of Loving Grace.’ As of 3 December 2025: <https://www.darioamodei.com/essay/machines-of-loving-grace>

Anthropic. 2025a. ‘Agentic Misalignment: How LLMs Could Be Insider Threats.’ 20 June. As of 3 December 2025: <https://www.anthropic.com/research/agentic-misalignment>

———. 2025b. ‘Project Vend: Can Claude Run a Small Shop? (and Why Does That Matter?).’ 27 June. As of 3 December 2025: <https://www.anthropic.com/research/project-vend-1>

Apollo Research. 2024. ‘We Need a ‘Science of Evals.’ 22 January. As of 3 December 2025: <https://www.apolloresearch.ai/blog/we-need-a-science-of-evals>

———. 2025. ‘More Capable Models Are Better at In-Context Scheming.’ Apollo Research. 19 June. As of 3 December 2025: <https://www.apolloresearch.ai/blog/more-capable-models-are-better-at-in-context-scheming>

Appel, Ruth, Peter Mccrory, Alex Tamkin, Miles Mccain, Tyler Neylon, Michael Stern, Alex Sanchez, et al. 2025. ‘The Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption.’ As of 3 December 2025: <https://assets.anthropic.com/m/218c82b858610fac/original/Economic-Index.pdf>

GDPRhub. 2025. ‘Article 5 GDPR.’ 7 November. As of 3 December 2025: https://gdprhub.eu/Article_5_GDPR

Artificial Analysis. 2025a. ‘Frontier Language Model Intelligence, Over Time.’ As of 3 December 2025: <https://artificialanalysis.ai/#frontier-language-model-intelligence-over-time>

———. 2025b. ‘Magistral Medium 1 – Intelligence, Performance & Price Analysis.’ As of 3 December 2025: <https://artificialanalysis.ai/models/magistral-medium>

———. 2025c. ‘State of AI: China. Highlights Report.’ As of 3 December 2025: <https://artificialanalysis.ai/downloads/china-report/2025/Artificial-Analysis-State-of-AI-China-Q2-2025-Highlights.pdf>

Aschenbrenner, Leopold. 2024a. ‘II. From AGI to Superintelligence: The Intelligence Explosion.’ Situational Awareness, 29 May. As of 3 December 2025: <https://situational-awareness.ai/from-agi-to-superintelligence/>

———. 2024b. 'III. The Free World Must Prevail.' *Situational Awareness*, 29 May. As of 3 December 2025: <https://situational-awareness.ai/the-free-world-must-prevail/>

ASML. 2025. 'ASML, Mistral AI Enter Strategic Partnership.' 9 September. As of 3 December 2025: <https://www.asml.com/en/news/press-releases/2025/asml-mistral-ai-enter-strategic-partnership>

Atkinson, David. 2023. 'Direct Approach Interactive Model.' *Epoch AI*, 31 May. As of 3 December 2025: <https://epoch.ai/blog/direct-approach-interactive-model>

Barbier, Jérôme, Charles Martinet & Charbel-Raphaël Segerie. 2025. 'Forging France's Expertise in AI Evaluation: INESIA's Strategic Role.' *Centre Pour La Sécurité de L'IA*. As of 3 December 2025: <https://www.securite-ia.fr/en/publications/vers-une-expertise-francaise-en-evaluation-delia-le-role-strategique-de-linesia>

Baskaran, Gracelin, and Meredith Schwartz. 2025. 'The Consequences of China's New Rare Earths Export Restrictions.' *Center for Strategic & International Studies*, 14 April. As of 3 December 2025: <https://www.csis.org/analysis/consequences-chinas-new-rare-earths-export-restrictions>

Batur, Olga, & Roel Peeters. 2021. 'European Union Data Challenge.' *European Parliament*. https://www.europarl.europa.eu/RegData/etudes/BRIE/2021/662939/IPOL_BRI%282021%29662939_EN.pdf

Bengio, Yoshua, Sören Mindermann, Daniel Privitera, Tamay Besiroglu, Rishi Bommasani, Stephen Casper, Yejin Choi, et al. 2025. 'International AI Safety Report.' As of 3 December 2025: <https://arxiv.org/abs/2501.17805>

Black, James, Mattias Eken, Jacob Parakilas, Stuart Dee, Conlan Ellis, Kiran Suman-Chauhan, Ryan Bain, et al. 2024. *Strategic Competition in the Age of AI Emerging Risks and Opportunities from Military Use of Artificial Intelligence*. Santa Monica, Calif.: RAND Corporation. RR-A3295-1. As of 3 December 2025: https://www.rand.org/pubs/research_reports/RRA3295-1.html

Blanchette, Nicholas, Joel Predd, Jim Mitre. 'Taking AGI-Enabled Wonder Weapons Seriously, Not Literally.' RAND Corporation. Unpublished manuscript.

Bloom, Nicholas, Charles I. Jones, John Van Reenen & Michael Webb. 2017. 'Are Ideas Getting Harder to Find?' National Bureau of Economic Research, Working Paper 23782. As of 3 December 2025: <https://www.nber.org/papers/w23782>

Bondar, Kateryna. 2025. 'Ukraine's Future Vision and Current Capabilities for Waging AI-Enabled Autonomous Warfare.' *Center for Strategic & International Studies*, 6 March. As of 3 December 2025: <https://www.csis.org/analysis/ukraines-future-vision-and-current-capabilities-waging-ai-enabled-autonomous-warfare>

Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.

Bown, Chad P. 2025. 'US-China Trade War Tariffs: An Up-To-Date Chart.' Peterson Institute for International Economics, 27 August. As of 3 December 2025: <https://www.piie.com/research/piie-charts/2019/us-china-trade-war-tariffs-date-chart>

Bradford, Anu. 2019. *The Brussels Effect: How the European Union Rules the World*. Oxford: Oxford University Press. doi: 10.1093/oso/9780190088583.001.0001

Brown, Noam. 2025. ‘Scaling Test Time Compute to Multi-Agent Civilizations – Noam Brown, OpenAI.’ Latent Space, YouTube, 19 June. As of 3 December 2025: <https://www.youtube.com/watch?v=ddd4xjuJTyg>

Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. ‘Language Models Are Few-Shot Learners.’ As of 3 December 2025: <https://arxiv.org/pdf/2005.14165>

Buchanan, Ben. 2020. ‘The AI Triad and What It Means for National Security Strategy.’ Center for Security and Emerging Technology, August. As of 3 December 2025: <https://cset.georgetown.edu/publication/the-ai-triad-and-what-it-means-for-national-security-strategy/>

Bullock, Justin B., Samuel Hammond & Seb Krier. 2025. ‘AGI, Governments, and Free Societies.’ As of 3 December 2025: <https://arxiv.org/abs/2503.05710>

Bulski, Krzysztof. 2025. ‘One Year after the Draghi Report: Europe Delivers Only 1 in 10 Promises.’ EU Insider. As of 3 December 2025: <https://www.euinsider.eu/news/one-year-after-the-draghi-report-europe-delivers-only-1-in-10-promises>

Burdette, Zachary, & Hiwot Demelash. 2025a. ‘Will the Race for Advanced AI Cause Geopolitical Instability?’ Substack.com, Geopolitics of AGI, 10 July. As of 3 December 2025: <https://geopoliticsagi.substack.com/p/will-the-race-for-advanced-ai-cause>

———. 2025b. *Evaluating the Risks of Preventive Attack in the Race for Advanced AI*. Santa Monica, Calif.: RAND Corporation. PE-A3691-13. As of 3 December 2025: <https://www.rand.org/pubs/perspectives/PEA3691-13.html>

Burdette, Zachary, Dwight Phillips, Jacob L. Heim, Edward Geist, David R. Frelinger, Chad Heitzenrater & Karl P. Mueller. 2025. *An AI Revolution in Military Affairs? How Artificial Intelligence Could Reshape Future Warfare*. Santa Monica, Calif.: RAND Corporation. WR-A4004-1. As of 3 December 2025: https://www.rand.org/pubs/working_papers/WRA4004-1.html

Bureau of Industry & Security. 2025. ‘Department of Commerce Announces Rescission of Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls.’ Bureau of Industry and Security. As of 3 December 2025: <https://www.bis.gov/press-release/department-commerce-announces-rescission-biden-era-artificial-intelligence-diffusion-rule-strengthens>

Capoot, Ashley. 2025. ‘Anthropic, Google, OpenAI and XAI Granted up to \$200 Million for AI Work from Defense Department.’ CNBC, 14 July. As of 3 December 2025: <https://www.cnbc.com/2025/07/14/anthropic-google-openai-xai-granted-up-to-200-million-from-dod.html>

Capoot, Ashley, & Kate Rooney. 2025. ‘OpenAI Hits \$10 Billion in Annual Recurring Revenue Fueled by ChatGPT Growth.’ CNBC, 9 June. As of 3 December 2025:
<https://www.cnbc.com/2025/06/09/openai-hits-10-billion-in-annualized-revenue-fueled-by-chatgpt-growth.html>

Caroli, Laura. 2025. ‘The EU AI Continent Action Plan.’ Center for Strategic & International Studies. As of 3 December 2025: <https://www.csis.org/analysis/eu-ai-continent-action-plan>

Caulcutt, Clea, & Camille Gijs. 2025. ‘Macron Wants EU to Target US Big Tech after New Trump Tariff Threat.’ *Politico*, 28 August. As of 3 December 2025:
<https://www.politico.eu/article/emmanuel-macron-floats-coming-after-us-big-tech-response-donald-trump-threats-tariffs/>

Census Bureau, United States. 2025. ‘Business Trends and Outlook Survey.’ United States Census Bureau, 20 November. As of 3 December 2025:
https://www.census.gov/hfp/btos/data_downloads

Chan, Jun Shern, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, et al. 2024. ‘MLE-Bench: Evaluating Machine Learning Agents on Machine Learning Engineering.’ As of 3 December 2025: <https://arxiv.org/abs/2410.07095>

Chan, Kyle, Gregory Smith, Jimmy Goodrich, Gerard DiPippo & Konstantin F. Pilz. 2025. *Full Stack: China’s Evolving Industrial Policy for AI*. Santa Monica, Calif.: RAND Corporation. PE-A4012-1. As of 3 December 2025: <https://www.rand.org/pubs/perspectives/PEA4012-1.html>

Chinese government. 2024. “The ‘East Data West Computation’ Project Is Progressing Steadily.” As of 3 December 2025: https://www.gov.cn/yaowen/liebiao/202404/content_6945169.htm

Clark, Don. 2021. ‘The Tech Cold War’s “Most Complicated Machine” That’s out of China’s Reach.’ *New York Times*, 4 July. As of 3 December 2025:
<https://www.nytimes.com/2021/07/04/technology/tech-cold-war-chips.html>

CNBC Television. 2025. ‘Anthropic CEO: More Confident than Ever That We’re “Very Close” to Powerful AI Capabilities.’ YouTube, 21 January. As of 3 December 2025:
<https://www.youtube.com/watch?v=7LNyUbii0zw>

Cornish, Chloe, & Michael Acton. 2025. ‘OpenAI to Join US-UAE Plan to Build Vast Data Centre in Abu Dhabi.’ *Financial Times*, 16 May. As of 3 December 2025:
<https://www.ft.com/content/4708c7a6-a26f-45a9-9a4f-25a9af4167df>

Cotra, Ajeya. 2020. ‘Draft Report on AI Timelines.’ Lesswrong, 19 September. As of 3 December 2025:
<https://www.lesswrong.com/posts/KrJfoZzpSDpnrv9va/draft-report-on-ai-timelines>

Cottier, Ben. 2024. ‘How Much Does It Cost to Train Frontier AI Models?’ Epoch AI, 3 June. As of 3 December 2025: <https://epoch.ai/blog/how-much-does-it-cost-to-train-frontier-ai-models>

———. 2025. ‘LLM Inference Prices Have Fallen Rapidly but Unequally across Tasks.’ Epoch AI. As of 3 December 2025: <https://epoch.ai/data-insights/llm-inference-price-trends>

D’Souza, Joseph. 2025. ‘Mistral AI Statistics by Revenue and Facts (2025).’ Electro IQ, 9 May. As of 3 December 2025: <https://electroiq.com/stats/mistral-ai-statistics/>

Damro, Chad. 2012. 'Market Power Europe.' *Journal of European Public Policy* 19(5): 682–99. As of 3 December 2025:
https://www.pure.ed.ac.uk/ws/portalfiles/portal/10091990/DAMRO_2012_Market_power_Europe.pdf

Dave, Paresh. 2024. 'How Scrappy Cryptominer CoreWeave Transformed into the Multibillion-Dollar Backbone of the AI Boom.' *Wired*, 3 May. As of 3 December 2025:
<https://www.wired.com/story/coreweave-scrappy-cryptominer-multibillion-dollar-ai/>

Davidson, Tom. 2025. 'What a Compute-Centric Framework Says about Takeoff Speeds.' *Coefficient Giving*, 18 November. As of 3 December 2025:
<https://coefficientgiving.org/research/what-a-compute-centric-framework-says-about-takeoff-speeds/>

Davidson, Tom, & Tom Houlden. 2025. 'How Quick and Big Would a Software Intelligence Explosion Be?' *Forethought*, 4 August. As of 3 December 2025:
<https://www.forethought.org/research/how-quick-and-big-would-a-software-intelligence-explosion-be>

Davidson, Tom, Lukas Finnveden & Rose Hadshar. 2025. 'AI-Enabled Coups: How a Small Group Could Use AI to Seize Power.' *Forethought*, 15 April. As of 3 December 2025:
<https://www.forethought.org/research/ai-enabled-coups-how-a-small-group-could-use-ai-to-seize-power>

Dealroom. 2025. 'Opening Moves in Global AI AI, Startups & Venture Capital.' As of 3 December 2025: <https://dealroom.co/uploaded/2025/02/Dealroom-AI-Summit-2025.pdf?x63517>

Delaney, Oscar. 2025. 'Managing Risks from Internal AI Systems.' Institute for AI Policy and Strategy, 21 July. As of 3 December 2025:
<https://www.iaps.ai/research/managing-risks-from-internal-ai-systems>

Dell'Acqua, Fabrizio, Saran Rajendran, Edward McFowland III, Lisa Krayer, Ethan Mollick, François Candelier, Hila Lifshitz-Assaf, Karim Lakhani & Katherine Kellogg. 2023. 'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.' As of 3 December 2025:
https://www.hbs.edu/ris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf

Denain, Jean-Stanislas. 2025a. 'The Huge Potential Implications of Long-Context Inference.' *Epoch AI*, 19 September. As of 3 December 2025:
<https://epoch.ai/gradient-updates/the-huge-potential-implications-of-long-context-inference>

Denain, Jean-Stanislas, Anson Ho, & Jaime Sevilla. 2025. "How Many Digital Workers Could OpenAI Deploy?" *Epoch AI*, 3 October. As of 3 December 2025:
<https://epoch.ai/gradient-updates/how-many-digital-workers-could-openai-deploy>

Desai, Dev, Shiv Kantliwala, Jyothi Vybhavi, Renju Ravi, Harshkumar Patel & Jitendra Patel. 2024. 'Review of AlphaFold 3: Transformative Advances in Drug Design and Therapeutics.' *Cureus* 16(7). doi: 10.7759/cureus.63646

Ding, Jeffrey. 2024. *Technology and the Rise of Great Powers*. Princeton University Press.

———. 2025. *Running the Right Artificial Intelligence Race: A National Strategy for AI Diffusion*. Santa Monica, Calif.: RAND Corporation. PE-A4165-1. As of 3 December 2025: <https://www.rand.org/pubs/perspectives/PEA4165-1.html>

Draghi, Mario. 2024. ‘The Draghi Report on EU Competitiveness.’ European Commission, September. As of 3 December 2025: https://commission.europa.eu/topics/eu-competitiveness/draghi-report_en

Dubey, Abhimanyu, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, et al. 2024. ‘The Llama 3 Herd of Models.’ As of 3 December 2025: <https://arxiv.org/abs/2407.21783>

Econovis. 2025. ‘U.S. Trade.’ As of 3 December 2025: <https://www.econovis.net/blog/categories/united-states-trade>

Edelman, Yafah. 2025. ‘Compute Scaling Will Slow down due to Increasing Lead Times.’ Epoch AI, 5 September. As of 3 December 2025: <https://epoch.ai/gradient-updates/compute-scaling-will-slow-down-due-to-increasing-lead-times>

EirGrid. 2025. ‘Real Time System Information.’ As of 3 December 2025: <https://www.eirgrid.ie/grid/real-time-system-information>

Elysee. 2025. ‘Germany, France and the European Commission Launch Frontier AI Initiative at the Summit on European Digital Sovereignty.’ 18 November. As of 3 December 2025: <https://www.elysee.fr/en/emmanuel-macron/2025/11/18/germany-france-and-the-european-commission-launch-frontier-ai-initiative-at-the-summit-on-european-digital-sovereignty>

Emberson, Luke. 2025. ‘GPT-5 and GPT-4 Were Both Major Leaps in Benchmarks from the Previous Generation.’ Epoch AI, 29 August. As of 3 December 2025: <https://epoch.ai/data-insights/gpt-capabilities-progress>

England, Andrew, & Ahmed Al Omran. 2025. ‘Donald Trump Lauds Saudi Arabia as He Unveils AI and Defence Deals.’ *Financial Times*, 13 May. As of 3 December 2025: <https://www.ft.com/content/5302d5d2-d375-4327-905c-7b1ad5173105>

Entsoe. 2025a. ‘European Parliament Report on Electricity Grids: ENTSO-E Welcomes the Spotlight on Accelerating Grid Transformation.’ As of 3 December 2025: <https://www.entsoe.eu/news/2025/06/19/european-parliament-report-on-electricity-grids-entso-e-welcomes-the-spotlight-on-accelerating-grid-transformation/>

———. 2025b. ‘Installed Capacity per Production Type.’ 2025. As of 3 December 2025: <https://www.entsoe.eu/data/power-stats>

Epoch AI. 2024. ‘AI Benchmarking Dashboard.’ 27 November. As of 3 December 2025: <https://epoch.ai/benchmarks>

———. 2025a. ‘GATE: Modeling the Trajectory of AI and Automation.’ As of 3 December 2025: <https://epoch.ai/blog/announcing-gate>

———. 2025b. ‘GPQA Diamond.’ As of 3 December 2025: <https://epoch.ai/benchmarks/gpqa-diamond>

_____. 2025c. 'Machine Learning Trends.' 11 April. As of 3 December 2025: <https://epoch.ai/trends>

_____. 2025d. 'Grokking "Forecasting TAI with Biological Anchors".' As of 3 December 2025: <https://epoch.ai/blog/grokking-bioanchors>

_____. 2025e. 'Grokking "Semi-Informative Priors Over AI Timelines".' As of 3 December 2025: <https://epoch.ai/blog/grokking-semi-informative-priors>

_____. 2025f. 'Notes on GPT-5 Training Compute.' 13 October. As of 3 December 2025: <https://epochai.substack.com/p/notes-on-gpt-5-training-compute>

_____. 2025g. 'What Will AI Look Like in 2030?' 16 September. As of 3 December 2025: <https://epoch.ai/blog/what-will-ai-look-like-in-2030>

_____. 2025h. 'AI Models.' 8 December. As of 10 December 2025: <https://epoch.ai/data/ai-models>

Erdil, Ege. 2025a. 'AI Progress Is about to Speed Up.' Substack, Epoch AI, 21 February. As of 3 December 2025: <https://epochai.substack.com/p/ai-progress-is-about-to-speed-up>

_____. 2025b. 'Most AI Value Will Come from Broad Automation, Not from R&D.' Epoch AI, 21 March. As of 3 December 2025: <https://epoch.ai/gradient-updates/most-ai-value-will-come-from-broad-automation-not-from-r-d>

Erdil, Ege, Andrei Potlogea, Tamay Besiroglu, Edu Roldan, Anson Ho, Jaime Sevilla, Matthew Barnett, Matej Vrzla & Robert Sandler. 2025. 'GATE: An Integrated Assessment Model for AI Automation.' As of 3 December 2025: <https://arxiv.org/abs/2503.04941>

Estrada, Sheryl. 2025. 'MIT Report: 95% of Generative AI Pilots at Companies Are Failing.' Fortune, 18 August. As of 10 December 2025: <https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/>

European AI Office (homepage). 2025. As of 3 December 2025: <https://digital-strategy.ec.europa.eu/en/policies/ai-office>

European Commission. 2020. 'A European Strategy for Data.' EUR-Lex, 19 February. As of 3 December 2025: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52020DC0066>

_____. 2021. 'Fostering a European Approach to Artificial Intelligence.' EUR-Lex, 11 April. As of 3 December 2025: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021DC0205>

_____. 2022. 'EU Trade Relations with United States.' As of 3 December 2025: https://policy.trade.ec.europa.eu/eu-trade-relationships-country-and-region/countries-and-regions/united-states_en

_____. 2024a. 'AI Factories.' As of 3 December 2025: <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>

_____. 2024b. 'Anti-Coercion Instrument.' As of 3 December 2025: <https://trade.ec.europa.eu/access-to-markets/en/content/anti-coercion-instrument>

_____. 2025a. 'AI Continent – New Cloud and AI Development Act.' As of 3 December 2025:
https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/14628-AI-Continent-new-cloud-and-AI-development-act_en

_____. 2025b. 'Apply AI Strategy.' 2025. As of 3 December 2025:
<https://digital-strategy.ec.europa.eu/en/policies/apply-ai>

_____. 2025c. 'EU Launches InvestAI Initiative to Mobilise €200 Billion of Investment in Artificial Intelligence.' As of 3 December 2025:
https://ec.europa.eu/commission/presscorner/detail/en/ip_25_467

_____. 2025d. 'European Digital Identity.' As of 3 December 2025:
https://commission.europa.eu/topics/digital-economy-and-society/european-digital-identity_en

_____. 2025e. 'Speech by President von Der Leyen at the Annual EU Budget Conference 2025.' As of 3 December 2025: https://ec.europa.eu/commission/presscorner/detail/en/speech_25_1284

_____. 2025f. 'The AI Continent Action Plan.' As of 3 December 2025:
<https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan>

_____. 2025g. 'The General-Purpose AI Code of Practice.' As of 3 December 2025:
<https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>

_____. 2025h. 'Public Consultation on the Use of Data to Develop the Future of AI: The European Data Union Strategy.' 25 June. As of 3 December 2025:
<https://data.europa.eu/en/news-events/news/public-consultation-use-data-develop-future-ai-european-data-union-strategy>

_____. 2025i. 'Artificial Intelligence (AI) in Science – European Commission.' 11 July. As of 3 December 2025:
https://research-and-innovation.ec.europa.eu/research-area/industrial-research-and-innovation/artificial-intelligence-ai-science_en

_____. 2025j. 'Joint Statement on a United States-European Union Framework on an Agreement on Reciprocal, Fair and Balanced Trade.' 21 August. As of 3 December 2025:
https://policy.trade.ec.europa.eu/news/joint-statement-united-states-european-union-framework-agreement-reciprocal-fair-and-balanced-trade-2025-08-21_en

_____. 2025k. '2025 Strategic Foresight Report.' 9 September. As of 3 December 2025:
https://commission.europa.eu/strategy-and-policy/strategic-foresight/2025-strategic-foresight-report_en

_____. 2025l. 'Cloud Sovereignty Framework.' Version 1.2.1 – Oct. 2025.' As of 3 December 2025:
https://commission.europa.eu/document/download/09579818-64a6-4dd5-9577-446ab6219113_en?filename=Cloud-Sovereignty-Framework.pdf

_____. 2025m. 'A Framework for General-Purpose AI Model Categorisation.' 10 October. As of 3 December 2025: <https://publications.jrc.ec.europa.eu/repository/handle/JRC143256>

European External Action Service. 2024. 'ANNUAL PROGRESS REPORT on the Implementation of the Strategic Compass for Security and Defence.' As of 3 December 2025:
https://www.eeas.europa.eu/sites/default/files/documents/2024/StrategicCompass_2ndYear_Report_0.pdf

European Parliament. 2023. 'REPORT of the Investigation of Alleged Contraventions and Maladministration in the Application of Union Law in Relation to the Use of Pegasus and Equivalent Surveillance Spyware.' A9-0189/2023. As of 3 December 2025:
https://www.europarl.europa.eu/doceo/document/A-9-2023-0189_EN.html

European Union. 2023. 'Key Facts and Figures.' As of 3 December 2025:
https://european-union.europa.eu/principles-countries-history/facts-and-figures-european-union_en

Eurostat. 2025. 'Use of Artificial Intelligence in Enterprises.' As of 3 December 2025:
https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Use_of_artificial_intelligence_in_enterprises

Farooque, Faizan. 2025. 'Nvidia Secures 92% GPU Market Share in Q1 2025.' Yahoo Finance, 6 June. As of 3 December 2025:
<https://finance.yahoo.com/news/nvidia-secures-92-gpu-market-150444612.html>

Farrell, Henry, & Abraham L. Newman. 2019. 'Weaponized Interdependence: How Global Economic Networks Shape State Coercion.' *International Security* 44(1): 42–79. doi: 10.1162/isec_a_00351

Folinsbee, Patrick. 2025. 'NOTE for NATIONAL DEFENCE: Arms Race Dynamics for Artificial General Intelligence.' Concordia. As of 3 December 2025:
<https://www.concordia.ca/content/dam/ginacody/research/spnet/Documents/BriefingNotes/EmergingTech-MilitaryApp/BN-117-Emerging-technology-and-military-application-Nov2021.pdf>

Gabriel, Iason, Arianna Manzini, Geoff Keeling, Lisa Anne Hendricks, Verena Rieser, Hasan Iqbal, Nenad Tomašev, et al. 2024. 'The Ethics of Advanced AI Assistants.' As of 3 December 2025:
<https://arxiv.org/abs/2404.16244>

Geels, Frank W. 2002. 'Technological Transitions as Evolutionary Reconfiguration Processes: A Multi-Level Perspective and a Case-Study.' *Research Policy* 31(8-9): 1257–74. As of 3 December 2025:
<https://www.sciencedirect.com/science/article/abs/pii/S004873302000628>

Genus, Audley, & Anne-Marie Coles. 2008. 'Rethinking the Multi-Level Perspective of Technological Transitions.' *Research Policy* 37(9): 1436–45. doi: 10.1016/j.respol.2008.05.006

Google DeepMind. 2024. 'Levels of AGI for Operationalizing Progress on the Path to AGI.' 21 July. As of 3 December 2025: <https://deepmind.google/research/publications/66938/>

———. 2025. 'SIMA 2: An Agent that Plays, Reasons, and Learns With You in Virtual 3D Worlds.' As of 3 December 2025:
<https://deepmind.google/blog/sima-2-an-agent-that-plays-reasons-and-learns-with-you-in-virtual-3d-worlds>

Google Team. 2025. 'Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and next Generation Agentic Capabilities.' As of 3 December 2025:
https://storage.googleapis.com/deepmind-media/gemini/gemini_v2_5_report.pdf

Grace, Katja, Harlan Stewart, Julia Sandkühler, Stephen Thomas, Ben Weinstein-Raun & Jan Brauner. 2024. 'Thousands of AI Authors on the Future of AI.' As of 3 December 2025:
https://aiimpacts.org/wp-content/uploads/2023/04/Thousands_of_AI_authors_on_the_future_of_AI.pdf

Grunewald, Erich. 2023. 'Introduction to AI Chip Making in China.' Institute for AI Policy and Strategy, 14 December. As of 3 December 2025:
<https://www.iaps.ai/research/ai-chip-making-china>

Guardian Staff and Agencies. 2025. 'Trump Agrees Deal for UAE to Build Largest AI Campus Outside US.' *Guardian*, 16 May. As of 3 December 2025:
<https://www.theguardian.com/us-news/2025/may/15/trump-artificial-intelligence-uae>

Guo, Daya, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, et al. 2025. 'DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning.' As of 3 December 2025: <http://arxiv.org/abs/2501.12948>

Hashim, Shakeel, & Celia Ford. 2025. 'Bubbles Can Still Change the World.' Transformer, 22 August. As of 3 December 2025: <https://www.transformernews.ai/p/sam-altman-ai-bubble>

Hassabis, Demis. 2024. 'Demis Hassabis Says AGI, Artificial General Intelligence, Is Still 10 Years Away Because 2 or 3 Big Innovations Are Required and the next One Is Agent-Based Systems.' Post on Reddit.com. As of 3 December 2025:
https://www.reddit.com/r/singularity/comments/1g5zu0i/demis_hassabis_says_agi_artificial_general/?rdt=55701

Hawkins, Zoe, Vili Lehdonvirta, & Boxi Wu. 2025. 'AI Compute Sovereignty: Infrastructure Control across Territories, Cloud Providers, and Accelerators.' SSRN, 24 June. As of 3 December 2025:
<https://doi.org/10.2139/ssrn.5312977>

Heaven, Will Douglas. 2025. 'What Comes next for AI Copyright Lawsuits?' MIT Technology Review, July. As of 3 December 2025:
<https://www.technologyreview.com/2025/07/01/1119486/ai-copyright-meta-anthropic/>

Heim, Lennart. 2024. 'Inference Compute: GPT-O1 and AI Governance.' Blog post, 22 September. As of 3 December 2025: <https://blog.heim.xyz/inference-compute/>

———. 2025. *Understanding the Artificial Intelligence Diffusion Framework: Can Export Controls Create a U.S.-Led Global Artificial Intelligence Ecosystem?* Santa Monica, Calif.: RAND Corporation. PEA3776-1. As of 3 December 2025: <https://www.rand.org/pubs/perspectives/PEA3776-1.html>

Heim, Lennart, Markus Anderljung, Emma Bluemke, & Robert Trager. 2024. 'Computing Power and the Governance of AI.' GovAI, 14 February. As of 3 December 2025:
<https://www.governance.ai/analysis/computing-power-and-the-governance-of-ai>

Heim, Lennart, & Ashley Lin. 2025. *When AI Takes Time to Think: Implications of Test-Time Compute*. Santa Monica, Calif.: RAND Corporation, 26 March. As of 3 December 2025: <https://www.rand.org/pubs/commentary/2025/03/when-ai-takes-time-to-think-implications-of-test-time.html>

Hendrycks, Dan, Eric Schmidt & Alexandr Wang. 2025. ‘Superintelligence Strategy.’ As of 3 December 2025: <https://www.nationalsecurity.ai/>

Hendrycks, Dan, Dawn Song, Christian Szegedy, Honglak Lee, Yarin Gal, Erik Brynjolfsson, Sharon Li, et al. 2025. ‘A Definition of AGI.’ As of 3 December 2025: <https://www.agidefinition.ai/paper.pdf>

Heussaff, Conall. 2024. ‘Decarbonising for Competitiveness: Four Ways to Reduce European Energy Prices.’ Bruegel, 26 September. As of 3 December 2025: <https://www.bruegel.org/policy-brief/decarbonising-competitiveness-four-ways-reduce-european-energy-prices>

Hijink, Marc. 2024. *Focus: The ASML Way*. Amsterdam: Uitgeverij Balans.

Hill, Michael. 2025. ‘Agentic AI Supply Exceeds Demand, Gartner Warns.’ AI, Data & Analytics Network, 9 October. As of 3 December 2025: <https://www.aidataanalytics.network/data-science-ai/news-trends/agentic-ai-supply-exceeds-demand-gartner-warns>

Ho, Anson. 2024. ‘Algorithmic Progress in Language Models.’ Epoch AI, 12 March. As of 3 December 2025: <https://epoch.ai/blog/algorithmic-progress-in-language-models>

———. 2025. ‘Quantifying the Algorithmic Improvement from Reasoning Models.’ Epoch AI, 2 August. As of 3 December 2025: <https://epoch.ai/gradient-updates/quantifying-the-algorithmic-improvement-from-reasoning-models>

Hoffmann, Jordan, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, et al. 2022. ‘Training Compute-Optimal Large Language Models.’ As of 3 December 2025: <https://arxiv.org/abs/2203.15556>

IBM. 2025. ‘Two-Thirds of Surveyed Enterprises in EMEA Report Significant Productivity Gains from AI, Finds New IBM Study.’ IBM Newsroom, 8 October. As of 3 December 2025: <https://newsroom.ibm.com/2025-10-28-Two-thirds-of-surveyed-enterprises-in-EMEA-report-significant-productivity-gains-from-AI,-finds-new-IBM-study>

International Monetary Fund. 2025. ‘GDP, Current Prices: European Union.’ IMF DataMapper. As of 30 November 2025: <https://www.imf.org/external/datamapper/NGDPD@WEO/EU>

Insikt Group. 2025. ‘Measuring the US-China AI Gap.’ Recorded Future, 8 May. As of 3 December 2025: <https://www.recordedfuture.com/research/measuring-the-us-china-ai-gap>

Intel. 2025. ‘Intel and Trump Administration Reach Historic Agreement to Accelerate American Technology and Manufacturing Leadership.’ 22 August. As of 3 December 2025: <https://www.intc.com/news-events/press-releases/detail/1748/intel-and-trump-administration-reach-historic-agreement-to>

Janků, David, Felicity Reddel, Max Reddel & Jakob Graabak. 2025. 'Beyond the AI Hype.' Centre for Future Generations, 15 April. As of 3 December 2025: <https://cfg.eu/beyond-the-ai-hype/>

Janků, David, Max Reddel, Roman Yampolskiy & Jason Hausenloy. 2024. 'We Have No Science of Safe AI.' Centre for Future Generations, 19 November. As of 3 December 2025: <https://cfg.eu/we-have-no-science-of-safe-ai/>

Josephson, Henry. 2025. 'How Fast Can Algorithms Advance Capabilities?' Epoch AI, 16 May. As of 3 December 2025: <https://epoch.ai/gradient-updates/how-fast-can-algorithms-advance-capabilities>

Juijn, Daan. 2024a. 'Advanced AI: Technical State of Play – Centre for Future Generations.' Centre for Future Generations, 6 June. As of 3 December 2025: <https://cfg.eu/advanced-ai-technical-state-of-play/>

———. 2024b. 'No, Deep Learning Is (Probably) Not Hitting a Wall.' Centre for Future Generations, 21 November. As of 3 December 2025: <https://cfg.eu/no-deep-learning-is-probably-not-hitting-a-wall/>

Kantrowitz, Alex. 2025. 'Google DeepMind CEO Demis Hassabis: The Path to AGI, LLM Creativity, and Google Smart Glasses.' Big Technology, 23 January. As of 3 December 2025: https://www.bigtechnology.com/p/google-deepmind-ceo-demis-hassabis?hide_intro_popup=true

Karger, Ezra, Josh Rosenberg, Zachary Jacobs, Molly Hickman, Rose Hadshar, Kayla Gamin, Taylor Smith, et al. 2023. 'Forecasting Existential Risks Evidence from a Long-Run Forecasting Tournament.' As of 3 December 2025: <https://static1.squarespace.com/static/635693acf15a3e2a14a56a4a/t/64f0a7838ccbf43b6b5ee40c/1693493128111/XPT.pdf>

Kent, Sherman. 1949. Strategic Intelligence for American World Policy.

Kergroach, Sandrine, & Julien Héritier. 2025. 'Emerging Divides in the Transition to Artificial Intelligence.' OECD Regional Development Papers 147 (June). doi: 10.1787/7376c776-en

Kirichenko, David. 2025. 'The Booming China-Russia Drone Alliance.' Center for European Policy Analysis, 4 June. As of 3 December 2025: <https://cepa.org/article/the-booming-china-russia-drone-alliance/>

Kirkwood, Megan. 2025. "Tech Prosperity Deal" Binds UK to US AI Dominance Strategy.' Tech Policy Press, 18 September. As of 3 December 2025: <https://techpolicy.press/tech-prosperity-deal-binds-uk-to-us-ai-dominance-strategy>

Klein, Ezra. 2025. 'The Government Knows A.G.I. Is Coming.' *New York Times*, 4 March 4. As of 3 December 2025: <https://www.nytimes.com/2025/03/04/opinion/ezra-klein-podcast-ben-buchanan.html>

Kokotajlo, Daniel, Scott Alexander, Thomas Larsen, Eli Lifland & Romeo Dean. 2025. 'AI 2027.' 3 April. As of 3 December 2025: <https://ai-2027.com/>

Kolata, Gina. 2024. 'ChatGPT Defeated Doctors at Diagnosing Illness.' *New York Times*, 17 November. As of 3 December 2025:
<https://www.nytimes.com/2024/11/17/health/chatgpt-ai-doctors-diagnosis.html>

Korinek, Anton. 2024. 'Economic Policy Challenges for the Age of AI.' National Bureau of Economic Research. doi: 10.3386/w32980

Korinek, Anton, & Donghyun Suh. 2024. 'Scenarios for the Transition to AGI.' National Bureau of Economic Research, 1 March. doi: 10.3386/w32255

Kothuri, Avinash, David Janku, Daan Juijn & Jorge Miguel Teixeira. 2025. 'Submission to the Call for Evidence on the Cloud and AI Development Act.' Centre for Future Generations, 16 July. As of 3 December 2025: <https://cfg.eu/special-compute-zones-submission/>

Kulveit, Jan, Raymond Douglas, Nora Ammann, Deger Turan, David Krueger & David Duvenaud. 2025. 'Gradual Disempowerment.' As of 3 December 2025:
<https://gradual-disempowerment.ai/>

Kwa, Thomas, Ben West, Joel Becker, Amy Deng, Katharyn Garcia, Max Hasin, Sami Kawhar, et al. 2025. 'Measuring AI Ability to Complete Long Tasks.' METR. As of 3 December 2025:
<https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks/>

Lakshmanan, Ravie. 2024. 'Meta Pauses AI Training on EU User Data amid Privacy Concerns.' The Hacker News, 15 June. As of 3 December 2025:
<https://thehackernews.com/2024/06/meta-halts-ai-training-on-eu-user-data.html?m=1>

Lambert, Nathan. 2025. 'XAI's Grok 4: The Tension of Frontier Performance with a Side of Elon Favoritism.' Interconnects, 12 July. As of 3 December 2025:
<https://www.interconnects.ai/p/grok-4-an-o3-look-alike-in-search>

Latent Space. 2025. 'Claude Code: Anthropic's CLI Agent.' YouTube, 7 May. As of 3 December 2025:
<https://www.youtube.com/watch?v=zDmW5hJPsvQ>

Lugt, Sanne van der. 2024. 'Why ASML Is the EU's Most Important Bargaining Chip.' Institute for European Policymaking @ Bocconi University. As of 3 December 2025:
<https://iep.unibocconi.eu/why-asml-eus-most-important-bargaining-chip>

MacAskill, William, & Fin Moorhouse. 2025. 'Preparing for the Intelligence Explosion.' Forethought, 11 March. As of 3 December 2025:
<https://www.forethought.org/research/preparing-for-the-intelligence-explosion>

MacroPolo. 2024. 'The Global AI Talent Tracker 2.0.' MacroPolo, 6 March. As of 3 December 2025:
<https://archivemacropolo.org/interactive/digital-projects/the-global-ai-talent-tracker/>

Mandler, C. 2024. 'Three Mile Island Nuclear Plant Will Reopen to Power Microsoft Data Centers.' NPR, 20 September. As of 3 December 2025:
<https://www.npr.org/2024/09/20/nx-s1-5120581/three-mile-island-nuclear-power-plant-microsoft-ai>

Mann, Tobias. 2025. ‘Elon Musk’s XAI to Pull about Half of Its Smog-Belching Turbines Powering Colossus.’ The Register, 8 May. As of 3 December 2025:
https://www.theregister.com/2025/05/08/xai_turbines_colossus/

Marcus, Gary. 2025a. ‘GPT-5: Overdue, Overhyped and Underwhelming. And That’s Not the Worst of It.’ Substack, Marcus on AI, 9 August. As of 3 December 2025:
<https://garymarcus.substack.com/p/gpt-5-overdue-overhyped-and-underwhelming>

———. 2025b. ‘Peak Bubble.’ Substack, Marcus on AI, 11 September. As of 3 December 2025:
<https://garymarcus.substack.com/p/peak-bubble>

Martens, Bertin. 2024. ‘Catch-up with the US or Prosper below the Tech Frontier? An EU Artificial Intelligence Strategy.’ Bruegel, 9 October. As of 3 December 2025:
<https://www.bruegel.org/policy-brief/catch-us-or-prosper-below-tech-frontier-eu-artificial-intelligence-strategy>

Martindale, Jon. 2025. ‘Meta Is Building AI Data Centers in Tents and Isn’t Slowing down – Zuckerberg Reveals Plans for “Several Multi-GW Clusters,” Including One Called Hyperion That’s Almost as Big as Manhattan.’ Tom’s Hardware, 15 July. As of 3 December 2025:
<https://www.tomshardware.com/tech-industry/artificial-intelligence/meta-plans-multi-gw-data-center-thats-nearly-the-size-of-manhattan-zuckerberg-promises-enormous-ai-splash-as-company-uses-tents-to-try-and-keep-up-with-rate-of-expansion>

Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Njenga Kariuki, Emily Capstick, et al. 2025. ‘Artificial Intelligence Index Report 2025.’ As of 3 December 2025: <https://arxiv.org/abs/2504.07139>

McCallum, Shiona. 2025. ‘AI “friend” chatbots probed over child protection.’ BBC News, 12 September. As of 3 December 2025: <https://www.bbc.com/news/articles/c74933vzx2yo>

Metaculus. 2025. ‘When Will the First General AI System Be Devised, Tested, and Publicly Announced?’ 23 August As of 3 December 2025:
<https://www.metaculus.com/questions/5121/date-of-artificial-general-intelligence/>

METR. 2025a. ‘How Does Time Horizon Vary across Domains?’ 14 July. As of 3 December 2025:
<https://metr.org/blog/2025-07-14-how-does-time-horizon-vary-across-domains/>

———. 2025b. ‘Details about METR’s Evaluation of OpenAI GPT-5.’ METR’s Autonomy Evaluation Resources, 6 August. As of 3 December 2025: <https://evaluations.metr.org/gpt-5-report/>

———. 2025c. ‘Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity.’ As of 3 December 2025:
<https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/>

Metz, Cade. 2025. ‘Google A.I. System Wins Gold in International Math Olympiad.’ *New York Times*, 21 July. As of 3 December 2025:
<https://www.nytimes.com/2025/07/21/technology/google-ai-international-mathematics-olympiad.html>

Metz, Cade, & Karen Weise. 2025. ‘A.I. Is Getting More Powerful, but Its Hallucinations Are Getting Worse.’ *New York Times*, 5 May. As of 3 December 2025: <https://www.nytimes.com/2025/05/05/technology/ai-hallucinations-chatgpt-google.html>

Miller, Chris. 2025. ‘How US Export Controls Have (and Haven’t) Curbed Chinese AI.’ AI Frontiers, 7 July. As of 3 December 2025: <https://ai-frontiers.org/articles/us-chip-export-controls-china-ai>

Milmo, Dan. 2024. “‘Very Scary’: Mark Zuckerberg’s Pledge to Build Advanced AI Alarms Experts.” *Guardian*, 19 January. As of 3 December 2025: <https://www.theguardian.com/technology/2024/jan/19/mark-zuckerberg-artificial-general-intelligence-system-alarms-experts-meta-open-source>

Mitre, Jim, & Joel B. Predd. 2025. *Artificial General Intelligence’s Five Hard National Security Problems*. Santa Monica, Calif.: RAND Corporation. PE-A3691-4. As of 3 December 2025: <https://www.rand.org/pubs/perspectives/PEA3691-4.html>

Moreno, Nathalie. 2025. ‘OpenAI and the Cross-Border Data Dilemma: US Litigation Holds vs. GDPR Erasure Obligations (UK/EU).’ Kennedys Law LLP, 2 June. As of 3 December 2025: <https://www.kennedyslaw.com/en/thought-leadership/article/2025/openai-and-the-cross-border-data-dilemma-us-litigation-holds-vs-gdpr-erasure-obligations-ukeu/>

Mueller, Karl P. 2025. *Heeding the Risks of Geopolitical Instability in a Race to Artificial General Intelligence*. Santa Monica, Calif.: RAND Corporation. PE-A3691-12. As of 3 December 2025: <https://www.rand.org/pubs/perspectives/PEA3691-12.html>

Munich Security Conference. 2025. ‘Speech by JD Vance and Selected Reactions.’ As of 3 December 2025: <https://securityconference.org/publikationen/buecher/key-speeches-volume-ii-jd-vance-msc-2025/>

Murgia, Madhumita, George Hammond & Tabby Kinder. 2025. ‘Stargate Artificial Intelligence Project to Exclusively Serve OpenAI.’ *Financial Times*, 24 January. As of 3 December 2025: <https://www.ft.com/content/4541c07b-f5d8-40bd-b83c-12c0fd662bd9>

Narayanan, Arvind, & Sayash Kapoor. 2025. ‘AI as Normal Technology: An Alternative to the Vision of AI as a Potential Superintelligence.’ Knight First Amendment Institute at Columbia University, 15 April. As of 3 December 2025: <https://knightcolumbia.org/content/ai-as-normal-technology>

Nasir, Hassam. 2025. ‘Sam Altman Says OpenAI Will Own “Well over 1 Million GPUs” by the End of the Year – ChatGPT Maker Continues to Expand Rapidly.’ Tom’s Hardware, 21 July. As of 3 December 2025: <https://www.tomshardware.com/tech-industry/sam-altman-teases-100-million-gpu-scale-for-openai-that-could-cost-usd3-trillion-chatgpt-maker-to-cross-well-over-1-million-by-end-of-year>

Nevo, Sella, Dan Lahav, Ajay Karpur, Yoge Bar-On, Henry Alexander Bradley & Jeff Alstott. 2024. *Securing AI Model Weights: Preventing Theft and Misuse of Frontier Models*. Santa Monica, Calif.: RAND Corporation. RR-A2849-1. As of 3 December 2025: https://www.rand.org/pubs/research_reports/RRA2849-1.html

Ngo, Richard, Lawrence Chan & Sören Mindermann. 2023. 'The Alignment Problem from a Deep Learning Perspective.' Openreview.net, 13 October. As of 3 December 2025:
<https://openreview.net/forum?id=fh8EYKFKns>

OECD. 2025. 'Mapping Relevant Data Collection Mechanisms for AI Training' *OECD Publishing*, 1 October. As of 3 December 2025:
https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/10/mapping-relevant-data-collection-mechanisms-for-ai-training_62921889/3264cd4c-en.pdf

Ó hÉigearaigh, Seán. 2025. 'The Most Dangerous Fiction: The Rhetoric and Reality of the AI Race.' doi: 10.2139/ssrn.5278644

Olcott, Eleanor, & Zijing Wu. 2025. 'China Trials Its First Advanced Tools for AI Chipmaking.' *Financial Times*, 17 September. As of 3 December 2025:
<https://www.ft.com/content/8fd79522-e34f-4633-bc87-ef0aae2d9159>

Olcott, Eleanor, Zijing Wu, Chris Cook & Nian Liu. 2025. "Stargate of China" Plan Emerges to Challenge US as AI Superpower.' *Financial Times*, 21 September. As of 3 December 2025:
<https://www.ft.com/content/31b8d6d8-adb0-4db0-a292-422c5a3bff19>

OpenAI. 2024. 'Learning to Reason with LLMs.' 12 September. As of 3 December 2025:
<https://openai.com/index/learning-to-reason-with-llms/>

_____. 2025a. 'Our Structure.' 5 May. As of 3 December 2025: <https://openai.com/our-structure/>

_____. 2025b. 'Stargate Advances with 4.5 GW Partnership with Oracle.' 4 September. As of 3 December 2025: <https://openai.com/index/stargate-advances-with-partnership-with-oracle/>

_____. 2025c. 'OpenAI and NVIDIA Announce Strategic Partnership to Deploy 10 Gigawatts of NVIDIA Systems.' 16 September. As of 3 December 2025:
<https://openai.com/index/openai-nvidia-systems-partnership/>

Ord, Toby. 2025. 'Inference Scaling Reshapes AI Governance.' 12 February. As of 3 December 2025:
<https://www.tobyord.com/writing/inference-scaling-reshapes-ai-governance>

Our World in Data. 2025. 'Total Electricity Generation.' As of 3 December 2025:
https://ourworldindata.org/grapher/electricity-generation?tab=line&country=USA-CHN-OWID_EU27-Middle+East+%28EI%29

Owen, David. 2024. 'Interviewing AI Researchers on Automation of AI R&D.' Epoch AI, 27 August. As of 3 December 2025: <https://epoch.ai/blog/interviewing-ai-researchers-on-automation-of-ai-rnd>

_____. 2025. 'What Will AI Look like in 2030?' Epoch AI, 16 September. As of 3 December 2025:
<https://epoch.ai/blog/what-will-ai-look-like-in-2030>

Özcan, Bengüsu, Daan Juijn, Jakob Graabak & Sam Bogerd. 2025. 'Advanced AI: Possible Futures.' Centre for Future Generations, 10 July. As of 3 December 2025:
<https://cfg.eu/advanced-ai-possible-futures/>

Pal, Siddhi, Catherine Schneider & Laura Nurski. 2025. 'Solving Europe's AI Talent Equation: Supply, Demand and Missing Pieces.' Interface, 9 July. As of 3 December 2025:
<https://www.interface-eu.org/publications/solving-europes-ai-talent-equation>

Pal, Siddhi. 2024. 'Where Is Europe's AI Workforce Coming From?' Interface. As of 3 December 2025: <https://www.interface-eu.org/publications/where-is-europes-ai-workforce-coming-from>

Pal, Siddhi, Ruggero Marino Lazzaroni & Paula Mendoza. 2024. 'AI's Missing Link: The Gender Gap in the Talent Pool.' Interface. As of 3 December 2025: <https://www.interface-eu.org/publications/ai-gender-gap>

Parasnus, Sharveya. 2025. 'Bank of China Announces 1 Trillion Yuan AI Industry Investment.' MEDIANAMA, 28 January. As of 3 December 2025: <https://www.medianama.com/2025/01/223-bank-of-china-announces-1-trillion-yuan-ai-industry-investment/>

Parker-Holder, Jack, & Shlomi Fruchter. 2025. 'Genie 3: A New Frontier for World Models.' Google DeepMind, 5 August. As of 3 December 2025: <https://deepmind.google/discover/blog/genie-3-a-new-frontier-for-world-models/>

Patel, Dwarkesh. 2023. 'Will Scaling Work?' Dwarkesh Podcast, 26 December. As of 3 December 2025: <https://www.dwarkesh.com/p/will-scaling-work>

———. 2025a. 'AGI Is Still 30 Years Away – Ege Erdil & Tamay Besiroglu.' Dwarkesh Podcast, 17 April. As of 3 December 2025: <https://www.dwarkesh.com/p/ege-tamay>

———. 2025b. 'Why I Don't Think AGI Is Right around the Corner.' Dwarkesh Podcast, 2 June. As of 3 December 2025: <https://www.dwarkesh.com/p/timelines-june-2025>

Patel, Dylan. 2024a. 'AI Datacenter Energy Dilemma – Race for AI Datacenter Space.' SemiAnalysis, 13 March. As of 3 December 2025: <https://semianalysis.com/2024/03/13/ai-datacenter-energy-dilemma-race/>

———. 2024b. 'Scaling Laws – O1 pro Architecture, Reasoning Training Infrastructure, Orion and Claude 3.5 Opus "Failures".' SemiAnalysis, 11 December. As of 3 December 2025: <https://semianalysis.com/2024/12/11/scaling-laws-o1-pro-architecture-reasoning-training-infrastructure-orion-and-claude-3-5-opus-failures/>

———. 2025. 'AI Arrives in the Middle East: US Strikes a Deal with UAE and KSA.' SemiAnalysis, 16 May. As of 3 December 2025: <https://semianalysis.com/2025/05/16/ai-arrives-in-the-middle-east-us-strikes-a-deal-with-uae-and-ksa/>

Patel, Dylan, & AJ Kourabi. 2025. 'Scaling Reinforcement Learning: Environments, Reward Hacking, Agents, Scaling Data.' SemiAnalysis, 8 June. As of 3 December 2025: <https://newsletter.semianalysis.com/p/scaling-reinforcement-learning-environments-reward-hacking-agents-scaling-data>

Patel, Dylan, Daniel Nishball & Jeremie Eliahou Ontiveros. 2024. 'Multi-Datacenter Training: OpenAI's Ambitious Plan to Beat Google's Infrastructure.' SemiAnalysis, 4 September. As of 3 December 2025: <https://newsletter.semianalysis.com/p/multi-datacenter-training-openais>

Patel, Dylan, Jeremie Eliahou Ontiveros, AJ Kourabi, Ivan Chiam & Wega Chu. 2025. 'AI Arrives in the Middle East: US Strikes a Deal with UAE and KSA.' *SemiAnalysis*, 16 May. As of 3 December 2025:
<https://newsletter.semianalysis.com/p/ai-arrives-in-the-middle-east-us-strikes-a-deal-with-uae-and-ksa>

Pilz, Konstantin F., Yusuf Mahmood & Lennart Heim. 2025. *AI's Power Requirements under Exponential Growth: Extrapolating AI Data Center Power Demand and Assessing Its Potential Impact on U.S. Competitiveness*. Santa Monica, Calif.: RAND Corporation. RR-A3572-1. As of 3 December 2025: https://www.rand.org/pubs/research_reports/RRA3572-1.html

Pilz, Konstantin F., Robi Rahman, James Sanders & Lennart Heim. 2025. 'Trends in AI Supercomputers.' *Epoch AI*, 23 April. As of 3 December 2025:
<https://epoch.ai/blog/trends-in-ai-supercomputers>

Piquard, Alexandre. 2025. 'AI: With Announcement of Investments Worth €109 Billion, Macron Intends to Take on US.' *Le Monde*, 10 February. As of 3 December 2025:
https://www.lemonde.fr/en/economy/article/2025/02/10/ai-with-the-announcement-of-a-109-billion-investment-macron-intends-to-take-on-the-us_6737985_19.html

Plumb, Taryn. 2025. 'What Oracle's \$300B OpenAI Deal Means for Enterprise Cloud Strategy.' *CIO*, 12 September. As of 3 December 2025:
<https://www.cio.com/article/4056139/what-oracles-300b-openai-deal-means-for-enterprise-cloud-strategy.html>

Potlogea, Andrei, & Anson Ho. 2025. 'AI and Explosive Growth Redux.' *Epoch AI*, 20 June. As of 3 December 2025: <https://epoch.ai/gradient-updates/ai-and-explosive-growth-redux>

Predd, Joel B. 2025. *Riding the Coming Wave: The Urgent Need for Contingency Plans for AGI Futures*. Santa Monica, Calif.: RAND Corporation. WR-A3691-1. As of 3 December 2025:
https://www.rand.org/pubs/working_papers/WRA3691-1.html

Qi, Xiangyu, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal & Peter Henderson. 2023. 'Fine-Tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To' doi: 10.48550/arXiv.2310.03693

Rehman, Iskander, Karl P. Mueller & Michael J. Mazarr. 2025. *Seeking Stability in the Competition for AI Advantage*. Santa Monica, Calif.: RAND Corporation, 13 March. As of 3 December 2025:
<https://www.rand.org/pubs/commentary/2025/03/seeking-stability-in-the-competition-for-ai-advantage.html>

Reif, L. Rafael. 2025. 'America's Coming Brain Drain.' *Foreign Affairs*, 6 May. As of 3 December 2025:
<https://www.foreignaffairs.com/united-states/trump-universities-war-america-coming-brain-drain>

Reuters. 2025a. 'OpenAI Hits \$12 Billion in Annualized Revenue, The Information Reports.' 31 July. As of 3 December 2025:
<https://www.reuters.com/business/openai-hits-12-billion-annualized-revenue-information-reports-2025-07-31/>

———. 2025b. ‘US Approves Some Nvidia UAE Sales, Bloomberg News Reports.’ Reuters, October 9. As of 3 December 2025:
<https://www.reuters.com/technology/us-approves-some-nvidia-uae-sales-bloomberg-news-reports-2025-10-09/>

Roth, Emma. 2023. ‘The New York Times Is Suing OpenAI and Microsoft for Copyright Infringement.’ The Verge, 27 December. As of 3 December 2025:
<https://www.theverge.com/2023/12/27/24016212/new-york-times-openai-microsoft-lawsuit-copyright-infringement>

Roytburg, Eva. 2025. ‘AI Experts Return from China Stunned: The U.S. Grid Is so Weak, the Race May Already Be Over.’ Fortune, 14 August. As of 3 December 2025:
<https://fortune.com/2025/08/14/data-centers-china-grid-us-infrastructure/>

Rueda, Ismael Arciniegas, Henri van Soest, Hye Min Park, Austin Smidt, David Gill, Robin Wang, Kelly Klima & Aisha Najera. 2025. *Assessing the United States’ Additional AI Power Capacity by 2030: Estimating Short-Term Increases in Electricity Generation and the Ability to Meet Growth in Power Demand*. Santa Monica, Calif.: RAND Corporation. RR-A3845-1. As of 3 December 2025:
https://www.rand.org/pubs/research_reports/RRA3845-1.html

Sacks, David, & Adam Segal. 2025. ‘Unpacking TSMC’s \$100 Billion Investment in the United States.’ Council on Foreign Relations. As of 3 December 2025:
<https://www.cfr.org/blog/unpacking-tsmcs-100-billion-investment-united-states>

Salvator, Dave. 2025. ‘Explaining Tokens – the Language and Currency of AI.’ NVIDIA Blog, 17 March. As of 3 December 2025: <https://blogs.nvidia.com/blog/ai-tokens-explained/>

Sastry, Girish, Lennart Heim, Haydn Belfield, Markus Anderljung, Miles Brundage, Julian Hazell, Cullen O’Keefe, et al. 2024. ‘Computing Power and the Governance of Artificial Intelligence.’ doi: 10.48550/arXiv.2402.08797

Satariano, Adam. 2025. ‘How the Maker of the “Most Complex Machine Humans Ever Created” Is Navigating Trade Fights.’ *New York Times*, 5 June. As of 3 December 2025:
<https://www.nytimes.com/2025/06/05/technology/asml-chips-tariffs-trade.html>

Schaefer, Peder. 2025. ‘Hiring Struggles Are Plaguing the EU AI Office.’ Transformer, 15 September. As of 3 December 2025:
<https://www.transformernews.ai/p/eu-is-struggling-to-hire-ai--act-office-safety-unit>

Schaul, Kevin, & Gerrit De Vynck. 2025. ‘Are We in an AI Bubble? Eight Charts Will Help You Decide.’ *Washington Post*, 22 November. As of 3 December 2025:
<https://www.washingtonpost.com/technology/2025/11/22/ai-bubble-economy-evidence/>

Schneider, Jordan. 2025. ‘Xi Takes an AI Masterclass.’ ChinaTalk, 13 May. As of 3 December 2025:
<https://www.chinatalk.media/p/xi-takes-an-ai-masterclass>

Sett, Gaurav. 2024. *How AI Can Automate AI Research and Development*. Santa Monica, Calif.: RAND Corporation, 24 October. As of 3 December 2025:
<https://www.rand.org/pubs/commentary/2024/10/how-ai-can-automate-ai-research-and-development.html>

Sevilla, Jaime. 2024. 'Training Compute of Frontier AI Models Grows by 4-5x per Year.' Epoch AI, 28 May. As of 3 December 2025:
<https://epoch.ai/blog/training-compute-of-frontier-ai-models-grows-by-4-5x-per-year>

Sevilla, Jaime, Tamay, Besiroglu, Ben Cottier, Josh You, Edu Roldán, Pablo Villalobos & Ege Erdil. 2024. 'Can AI Scaling Continue through 2030?' Epoch AI, 20 August. As of 3 December 2025:
<https://epoch.ai/blog/can-ai-scaling-continue-through-2030>

Shamir, Afek, & Maximilian Negele. 2025. *RAND Europe's Response to the EU Cloud and AI Development Act Call for Evidence*. Santa Monica, Calif.: RAND Corporation. CT-A4156-1. As of 3 December 2025: <https://www.rand.org/pubs/testimonies/CTA4156-1.html>

Shilov, Anton. 2024. 'China's SMEE Files Patent for an EUV Chipmaking Tool – Tool Aims to Break the Shackles of ASML Export Restrictions.' Tom's Hardware, 13 September. As of 3 December 2025:
<https://www.tomshardware.com/tech-industry/chinas-smee-files-patent-for-an-euv-chipmaking-tool-tool-aims-to-break-the-shackles-of-asml-export-restrictions>

———. 2025. 'China Wants 300 ExaFLOPS of Compute Power by 2025.' Tom's Hardware, 9 October. As of 3 December 2025:
<https://www.tomshardware.com/news/china-wants-300-exaflops-of-compute-power-by-2025>

Siegmann, Charlotte, & Markus Anderljung. 2022. 'The Brussels effect and artificial intelligence: How EU regulation will impact the global AI market.' As of 3 December 2025:
<https://arxiv.org/abs/2208.12645>

Singapore Conference on AI. 2025. 'The Singapore Consensus on Global AI Safety Research Priorities.' As of 3 December 2025: <https://www.scai.gov.sg/2025/scai2025-report/>

Snyder, Jason. 2025. 'The AI Bubble That Isn't There.' Forbes, 17 November. As of 3 December 2025:
<https://www.forbes.com/sites/jasonsnyder/2025/11/17/the-ai-bubble-that-isnt-there/>

SolarPower Europe. 2024. 'New Report: Average EU Member State Transposition of Permitting Rules for Renewables Falls Short at Just under 50%.' SolarPower Europe, 16 July. As of 3 December 2025:
<https://www.solarpowereurope.org/press-releases/new-report-average-eu-member-state-transposition-of-permitting-rules-for-renewables-falls-short-at-just-under-50>

Soni, Aditya, David Shepardson, Andrea Shalal & Max A Cherney. 2025. 'US to Take 10% Equity Stake in Intel, in Trump's Latest Corporate Move.' Reuters, 23 August. As of 3 December 2025:
<https://www.reuters.com/business/us-take-10-equity-stake-intel-trumps-latest-corporate-move-2025-08-22/>

Srivathsan, Bhargs, Marc Sorel & Pankaj Sachdeva. 2024. 'AI Power: Expanding Data Center Capacity to Meet Growing Demand.' McKinsey & Company, 29 October. As of 3 December 2025:
<https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand>

Stanford University. 2025. 'Artificial Intelligence Index Report 2025.' As of 3 December 2025:
https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf

Subin, Samantha. 2025. ‘Tech Megacaps Plan to Spend More than \$300 Billion in 2025 as AI Race Intensifies.’ CNBC, 8 February. As of 3 December 2025:
<https://www.cnbc.com/2025/02/08/tech-megacaps-to-spend-more-than-300-billion-in-2025-to-win-in-ai.html>

Sutton, Richard S, Michael Bowling, & Patrick M Pilarski. 2022. ‘The Alberta Plan for AI Research.’ As of 3 December 2025: <https://arxiv.org/abs/2208.11173>

Tähtinen, Lauri. 2024. ‘EU-Mercosur: So Much More than a Dead Deal.’ Center for Strategic & International Studies, 26 March. As of 3 December 2025:
<https://www.csis.org/analysis/eu-mercousur-so-much-more-dead-deal>

Tamma, Paola. 2022. ‘Hungary Vetoes Ukraine Aid, EU Seeks Alternative Solution.’ Politico, 6 December. As of 3 December 2025:
<https://www.politico.eu/article/hungary-vetoes-ukraine-aid-eu-seeks-alternative-solution/>

Tankersley, Jim, Ivan Nechepurenko, Steven Erlanger, Michael Schwirtz, Peter Baker, Ashley Ahn, et al. 2025. ‘Trump Backs Plan to Cede Land for Peace in Ukraine.’ *New York Times*, 16 August. As of 3 December 2025:
<https://www.nytimes.com/live/2025/08/16/world/trump-putin-meeting-alaska>

Tao, Mengying. 2025. ‘Chasing Artificial General Intelligence: China between Breakthroughs and Bottlenecks.’ Sinolytics, 31 July. As of 3 December 2025:
<https://sinolytics.de/global-business-news/blog/technology/agi-in-china/>

Taylor, Josh. 2025. ‘AI Chatbot “MechaHitler” Could Be Making Content Considered Violent Extremism, Expert Witness Tells X v ESafety Case.’ *Guardian*, 15 July. As of 3 December 2025: <https://www.theguardian.com/technology/2025/jul/15/x-esafety-ai-chatbot-grok-ntwnfb>

Tech Desk. 2024. ‘Sam Altman on Artificial Superintelligence: “There Is a Lot of Compounding Left to Do”.’ *Indian Express*, 11 November. As of 3 December 2025:
<https://indianexpress.com/article/technology/artificial-intelligence/sam-altman-on-artificial-superintelligence-there-is-a-lot-of-compounding-left-to-do-9661302/>

Tidler, Zachary, Marko Tešić, Lorenzo Pacchiardi, John Burden, Lexin Zhou, Manuel Cebrián, Fernando Martínez-Plumed & Jose Hernandez-Orallo. 2025. ‘Is the Definition of AGI a Percentage?’ The AI Evaluation Substack, 31 October. As of 3 December 2025:
<https://aievaluation.substack.com/p/is-the-definition-of-agi-a-percentage>

Todd, Benjamin. 2025. ‘Shrinking AGI Timelines: A Review of Expert Forecasts’ 80,000 Hours, 21 March. As of 3 December 2025:
<https://80000hours.org/2025/03/when-do-experts-expect-agi-to-arrive>

Toner, Helen. 2025a. “Long” Timelines to Advanced AI Have Gotten Crazy Short.’ Substack, Rising Tide, 1 April. As of 3 December 2025:
<https://helentonner.substack.com/p/long-timelines-to-advanced-ai-have>

———. 2025b. ‘Unresolved Debates about the Future of AI.’ Substack, Rising Tide, 30 June 30. As of 3 December 2025: <https://helentonner.substack.com/p/unresolved-debates-about-the-future>

Trammell, Philip, & Anton Korinek. 2023. 'Economic Growth under Transformative AI.' National Bureau of Economic Research, 1 October. doi: 10.3386/w31815

TRUMPF. 2025. 'TRUMPF Laser Amplifier.' As of 3 December 2025:
https://www.trumpf.com/en_GB/solutions/applications/euv-lithography/euv-drive-laser/

UK AISI. 2024. 'Research Scientist – Strategic Awareness.' Job advert, LinkedIn. As of 3 December 2025:
<https://www.linkedin.com/jobs/view/research-scientist-strategic-awareness-at-ai-security-institute-4196962109/?originalSubdomain=uk>

——— (homepage). 2025. As of 3 December 2025: <https://www.aisi.gov.uk/>

UK Government. 2023. 'Frontier AI Taskforce: First Progress Report.' As of 3 December 2025:
<https://www.gov.uk/government/publications/frontier-ai-taskforce-first-progress-report/frontier-ai-taskforce-first-progress-report>

———. 2025a. 'Spending Review 2025.' As of 3 December 2025:
<https://www.gov.uk/government/publications/spending-review-2025-document/spending-review-2025-html>

———. 2025b. 'AI Opportunities Action Plan.' 13 January. As of 3 December 2025:
<https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan>

———. 2025c. 'Sovereign AI Unit.' 16 July. As of 3 December 2025:
<https://www.gov.uk/government/collections/sovereign-ai-unit>

———. 2025d. 'US-UK Pact Will Boost Advances in Drug Discovery, Create Tens of Thousands of Jobs and Transform Lives.' 16 September. As of 3 December 2025:
<https://www.gov.uk/government/news/us-uk-pact-will-boost-advances-in-drug-discovery-create-tens-of-thousands-of-jobs-and-transform-lives>

US Congress. 2022. 'H.R.4346 – 117th Congress (2021-2022): Supreme Court Security Funding Act of 2022.' 9 August. As of 3 December 2025:
<https://www.congress.gov/bill/117th-congress/house-bill/4346>

———. 2025. 'U.S. Export Controls and China: Advanced Semiconductors.' As of 3 December 2025:
<https://www.congress.gov/crs-product/R48642>

Vance, J.D. 2025. 'Remarks by the Vice President at the Artificial Intelligence Action Summit in Paris, France.' The American Presidency Project, 11 February. As of 3 December 2025:
<https://www.presidency.ucsb.edu/documents/remarks-the-vice-president-the-artificial-intelligence-action-summit-paris-france>

Vlassov, Jüri. 2025. 'The Iron House Geopolitical Stakes of the US-China AGI Race.' Convergence Analysis, 1 September. As of 3 December 2025:
<https://www.convergenceanalysis.org/fellowships/international-security/the-iron-house-geopolitical-stakes-of-the-us-china-agi-race>

Wang, Dan, & Arthur Kroeber. 2025. 'The Real China Model.' Foreign Affairs, 19 August. As of 3 December 2025: <https://www.foreignaffairs.com/china/real-china-model-wang-kroeber>

Wang, Haoyu, Wei P Dai, Jiawei Zhang, Jialun Ma, Mingyi Huang, & Yuguo Yu. 2025. 'The Dual Nature of Plasticity Loss in Deep Continual Learning: Dissection and Mitigation.' Openreview.net. As of 3 December 2025: <https://openreview.net/forum?id=vvD0Bre3Dk>

Wang, Jiayu, Yifei Ming, Zhenmei Shi, Vibhav Vineet, Xin Wang, Yixuan Li, & Neel Joshi. 2024. 'Is a Picture Worth a Thousand Words? Delving into Spatial Reasoning for Vision Language Models.' As of 3 December 2025: <https://arxiv.org/abs/2406.14852>

Wentz, Jacob, & Anita Lin. 2025. 'Breakthroughs or Boasts? Assessing Recent Chinese Lithography Advancements.' Center for Strategic & International Studies, 24 September. As of 3 December 2025: <https://www.csis.org/blogs/strategic-technologies-blog/breakthroughs-or-boasts-assessing-recent-chinese-lithography>

Wexler, Ryan. 2025. 'Why Expert Data Is Becoming the New Fuel for AI Models.' SignalFire, 4 November. As of 3 December 2025: <https://www.signalfire.com/blog/expert-data-is-new-fuel-for-ai-models>.

White House. 2024. 'Memorandum on Advancing the United States' Leadership in Artificial Intelligence; Harnessing Artificial Intelligence to Fulfill National Security Objectives; and Fostering the Safety, Security, and Trustworthiness of Artificial Intelligence.' 24 October. As of 3 December 2025: <https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2024/10/24/memorandum-on-advancing-the-united-states-leadership-in-artificial-intelligence-harnessing-artificial-intelligence-to-fulfill-national-security-objectives-and-fostering-the-safety-security/>

White House. 2025. 'AI Action Plan.' As of 3 December 2025: <https://www.ai.gov/action-plan>

Wildeford, Peter. 2025a. 'GPT-5: A Small Step for Intelligence, a Giant Leap for Normal People.' Substack, The Power Law, 8 August. As of 3 December 2025: <https://peterwildeford.substack.com/p/gpt-5-a-small-step-for-intelligence>

———. 2025b. 'AI Is Probably Not a Bubble.' Substack, The Power Law, 29 October. As of 3 December 2025: <https://peterwildeford.substack.com/p/ai-is-probably-not-a-bubble>

Williams, Sean, & James Huckle. 2024. 'Easy Problems That LLMs Get Wrong.' As of 3 December 2025: <https://arxiv.org/abs/2405.19616>

Wintour, Patrick. 2016. 'EU Tells Swiss No Single Market Access If No Free Movement of Citizens.' *Guardian*, 3 July. As of 3 December 2025: <https://www.theguardian.com/world/2016/jul/03/eu-swiss-single-market-access-no-free-movement-citizens>

Witze, Alexandra. 2025. '75% of US Scientists Who Answered Nature Poll Consider Leaving.' *Nature*, 27 March. doi: 10.1038/d41586-025-00938-y

Wynroe, Keith. 2023. 'Literature Review of Transformative Artificial Intelligence Timelines.' Epoch AI, 17 January. As of 3 December 2025: <https://epoch.ai/blog/literature-review-of-transformative-artificial-intelligence-timelines>

Yang, Chengrun, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou & Xinyun Chen.

2023. 'Large Language Models as Optimizers.' doi: 10.48550/arXiv.2309.03409

Yang, Heekyong, & Max A. Cherney. 2025. 'Exclusive: SK Hynix Expects AI Memory Market to Grow 30% a Year to 2030.' Reuters, 11 August. As of 3 December 2025:

<https://www.reuters.com/world/asia-pacific/sk-hynix-expects-ai-memory-market-grow-30-year-2030-2025-08-11/>

Ye, Josh. 2023. 'China Targets 50% Growth in Computing Power in Race against U.S.' Reuters, 9 October. As of 3 December 2025:

<https://www.reuters.com/technology/china-targets-30-growth-computing-power-race-against-us-2023-10-09/>

Zeff, Maxwell. 2025. 'Mark Zuckerberg Says Meta Is Building a 5GW AI Data Center.' TechCrunch, 14 July. As of 3 December 2025:

<https://techcrunch.com/2025/07/14/mark-zuckerberg-says-meta-is-building-a-5gw-ai-data-center/>

Zhu, Julia. 2022. 'The Personal Information Protection Law: China's Version of the GDPR?' *Columbia Journal of Transnational Law*, 14 February. As of 3 December 2025:

<https://www.jtl.columbia.edu/bulletin-blog/the-personal-information-protection-law-chinas-version-of-the-gdpr>

Zilgalvis, Gustavs. 2025. 'Artificial General Intelligence and the Fourth Offset.' As of 3 December 2025:

<https://fourthoffset.ai/>