

DEMOS

Mozilla | 

THE OPEN DIVIDEND

BUILDING AN AI
OPENNESS STRATEGY
TO UNLOCK THE UK'S
AI POTENTIAL

ELIZABETH SEGER
JAMIE HANCOCK

JUNE 2025

Open Access. Some rights reserved.

Open Access. Some rights reserved. As the publisher of this work, Demos wants to encourage the circulation of our work as widely as possible while retaining the copyright. We therefore have an open access policy which enables anyone to access our content online without charge. Anyone can download, save, perform or distribute this work in any format, including translation, without written permission. This is subject to the terms of the Creative Commons By Share Alike licence. The main conditions are:

- Demos and the author(s) are credited including our web address **www.demos.co.uk**
- If you use our work, you share the results under a similar licence

A full copy of the licence can be found at **<https://creativecommons.org/licenses/by-sa/3.0/legalcode>**

You are welcome to ask for permission to use this work for purposes other than those covered by the licence. Demos gratefully acknowledges the work of Creative Commons in inspiring our approach to copyright. To find out more go to **www.creativecommons.org**



CONTENTS

ACKNOWLEDGEMENTS	PAGE 4
EXECUTIVE SUMMARY	PAGE 6
INTRODUCTION	PAGE 10
SECTION 1: WHAT DO WE MEAN BY A NATIONAL 'AI OPENNESS' STRATEGY?	PAGE 11
SECTION 2: WHY LEAN INTO AI OPENNESS?	PAGE 13
SECTION 3: WHY NOW?	PAGE 19
SECTION 4: SQUARING AI OPENNESS WITH THE UK'S AI SAFETY COMMITMENTS	PAGE 26
SECTION 5: A PURPOSE BUILT AI OPENNESS STRATEGY FOR THE UK	PAGE 32
SECTION 6: HIGH-LEVEL RECOMMENDATIONS	PAGE 41
CONCLUSION	PAGE 45
APPENDICES	PAGE 46

ACKNOWLEDGEMENTS

This report is produced in partnership with Mozilla, who we would like to thank for their support of the work and for lending their expertise to its development.

The report includes important insights from a workshop held on May 12, 2025 which convened experts on AI openness and open-source AI from the UK and Europe to discuss how an open-source AI strategy for the UK might take shape. We would like to thank all participants for generously sharing their time and input to this project.

We would like to extend thanks the following people for offering their input, insight, and feedback at various stages throughout the drafting of this report:

Markus Anderljung, Dave Buckley, Camilla de Coverly Veale, Mark Dansie, Jennifer Ding, Max Gahntz, Naman Goel, Linda Griffin, Lucie-Aimée Kaffee, Yann Lechelle, Nik Marda, Cailean Osborne and Martin Tisne.

Finally we would like to thank colleagues at Demos, specifically Polly Curtis, Sofia Lyall, Sumaya Akthar, and Chloe Burke.

Elizabeth Seger

Jamie Hancock

June 2025

ABOUT THIS REPORT

This paper is part of Demos' strategic focus area on '*Trustworthy Technology*'. With emerging technologies transforming our world at an ever-faster pace, we work to build bridges between politicians, technical experts, and citizens to explore solutions, improve trust, and create policy to ensure our technologies benefit society. Our aim in this report is to help elucidate how open and open-source AI can help the UK achieve its AI ambitions. These include building foundations for a thriving domestic AI ecosystem, reducing dependence on proprietary foreign tech, and driving safe AI adoption for public benefit.

As part of Demos's ongoing efforts to facilitate greater diversity, inclusion, equity and justice in all areas of our work, we assess and publish our approach to meeting our goals in each of our publications. In this report, the subject is fundamentally about securing a more inclusive and just future growth strategy for all, and ensuring the benefits of AI are realised across society.

EXECUTIVE SUMMARY

The UK government has made economic growth its top priority, and the government's success, along with national well-being, depends on the policies it adopts. AI has been made central to this growth agenda, and the UK stands at a pivotal juncture in shaping its AI future.

With global momentum around artificial intelligence accelerating, the country faces both a profound opportunity and a pressing challenge in determining how to secure the long-term economic and public value of AI. While the UK boasts world-leading academic institutions, a growing ecosystem of AI startups, and globally recognised expertise in AI safety and governance, it cannot compete directly with the scale and self-sufficiency of AI superpowers like the US and China.

This report makes the case for a national commitment to AI openness as a strategic move towards achieving the UK's AI ambitions. We set out how this could lead to an 'open dividend' for our growth and innovation ambitions.

Here 'AI openness' is understood as the broad public availability and ease of access to key artefacts and documentation from AI across the AI stack including AI models (weights), code, datasets, documentation, safety tooling and compute resources (Section 1).

NATIONAL AI OPENNESS BENEFITS

For a country looking to drive the domestic AI industry growth and reap the public benefits of widespread AI adoption, the advantages of supporting AI openness are numerous (Section 2). We identify them as:

- Driving Innovation
- Supporting AI Industry Growth
- Enabling Flexible AI Adoption
- Acting as an Economic Multiplier
- Delivering Public Benefit with Public AI
- Strengthening Tech Sovereignty

WHY NOW?

We make a three-pronged argument for why the UK should make its commitment to AI Openness now (Section 3):

- 1. The AI Game is Changing:** The capabilities of open models are rapidly catching up to frontier proprietary models, while the market is shifting toward smaller, specialised, and efficient models.
- 2. The Geopolitical Landscape is Shifting:** Global tech rivalries and export controls have exposed vulnerabilities in AI supply chains. As major powers consolidate AI infrastructure, the UK must protect its interests.
- 3. The UK Needs a Practical Strategy:** With limited resources to compete in frontier AI through brute-force investment, the UK must play to its strengths in research, safety, public sector innovation, and open knowledge. Embracing openness aligns with these strengths and provides a credible path to digital sovereignty.

SQUARING AI OPENNESS WITH THE UK'S AI SAFETY COMMITMENTS

The UK has positioned itself as a global leader in AI safety. However there is some concern that AI openness could be in tension with the UK's safety commitments if sharing models and systems more broadly may increase the risk of misuse or enable the spread of harmful capabilities. In this section, we aim to address this tension (Section 4).

Key Points

- Openness enables broad community oversight, helping identify and fix vulnerabilities faster than closed development.
- For most models, the benefit of openness outweighs the risks, but the balance is less certain for future frontier models, where risks and mitigation are harder to assess.
- Model sharing restrictions are a fallible risk mitigation strategy. We should therefore look to supporting a robust net of safety intervention throughout the AI life cycle.
- Restricting model sharing without a clear understanding of the threat model being guarded against may have negative consequences, limiting the safety and positive innovation benefits of openness.

EMBEDDING AN AI OPENNESS STRATEGY THROUGH THE AI OPPORTUNITIES ACTION PLAN

This section works to articulate what a bespoke AI Openness strategy for the UK could look like utilising the AI Opportunities Action Plan as the framework (Section 5). Our reason for utilising the Action Plan are threefold:

- The Action Plan articulates three high-level goals that are well suited to enhancement by AI openness.
- The Action Plan already alludes to openness but there is opportunity to more strongly embed it.
- The government has already bought into the Action Plan, accepting all proposals.

We map an openness strategy for the UK directly to the Plan's proposals, identifying opportunities to expand ambitions or embed openness more explicitly. A summary according to each of the high-level goals is as follows:

1. Laying the foundations for AI

The UK should treat AI openness - across compute, data, talent, and regulation - as a core component of national AI infrastructure. Openness can amplify each of these AI foundations.

Compute: The UK should actively pursue international computer collaboration, such as through the EuroHPC or EuroStack initiatives, prioritise compute access for public-interest open-source developers, and invest in open-source hardware.

Data: The National Data Library should adopt an “open by default” model, making high-quality datasets broadly available under open-access licences. Where openness isn’t appropriate, intermediate models such as localised access, synthetic data, and structured transparency tools can lower barriers to innovation while preserving privacy and legal protections.

Talent: Excessive restrictions on open-source development could deter top researchers, even if openness isn’t the main draw. To stay competitive, the UK must support international collaboration and commercial use of open-source tools, fostering a culture of collaborative, values-driven innovation.

Regulation: Openness should guide the UK AI Bill and the AI Security Institute’s (AISI’s) role, with regulatory exemptions for open models and strong transparency standards for proprietary systems. AISI can boost global influence by open-sourcing safety tools. Meanwhile, embedding participatory processes like citizens’ assemblies in AI governance will help align decisions with public values and build trust.

2. Driving cross-economy AI adoption (in the public interest)

The UK’s strategy for AI adoption across public and private sectors must shift from a “solutions first” to a “needs first” approach, embedding AI within broader public service reform rather than treating it as a standalone fix. While AI can significantly enhance service quality, its full potential depends on systemic integration.

Open-source AI procurement is key, offering cost efficiency, vendor independence, and better system compatibility. To support this, the UK should foster data-rich, open experimentation environments, develop an AI Knowledge Hub for shared practices, mandate open-source procurement when feasible, ensure interoperability, implement transparent evaluation with public benchmarks, and engage citizens in prioritising AI use and identifying unacceptable applications.

3. Securing a future for homegrown AI (AI Sovereignty)

The UK’s third high-level Action Plan goal aims to underpin UK AI sovereignty. In the context of the AI Opportunities Action Plan, AI sovereignty could be conceptualised as aiming to ensure that the UK has reliable access to AI capabilities, that the value of AI-led economic transformation is captured in the UK, and that sources of influence over the global development and deployment of AI technologies.

AI openness will aid each of the three goals in turn:

Access: Through open-source initiatives, the UK can contribute to and benefit from shared AI resources that remain permanently accessible—unlike proprietary models that can be restricted, taxed, or withdrawn by their owners. This collaborative approach pools international expertise and resources to create AI capabilities that no single nation could achieve alone.

Value: Ultimately value will accrue to AI industries that succeed at putting AI tools in consumer hands. Laying foundations for a thriving AI openness ecosystem in the UK will help the UK capture value by lowering barriers to entry for new business and by acting as an economic multiplier across industries.

Influence: The UK can influence the future direction of AI development by sharing this expertise widely and remaining a powerful contributor to the global AI research and safety environments.

RECOMMENDATIONS

This report makes five high-level recommendations to Government (Section 6):

- 1. Commit to an AI openness strategy for the UK and look to deploy it through integrations with the existing AI Opportunities Action Plan agenda.** The UK is primed to extract maximum value from the open ecosystems that it helps build and maintain, capitalising on our deep bench of scientific and AI expertise.
- 2. Use a commitment to AI openness to demonstrate dedication to building and deploying AI in the public interest.** This would help build trust and underpin realised public benefit from AI.
- 3. Pursue AI Sovereignty through outward collaboration and resource sharing, and promoting open development guidelines.** Outward collaboration is also about contributing to a thriving global open-source counterpoint to proprietary big tech. The larger collective action, the larger and more viable the counterpoint.
- 4. Influence a positive future for AI globally by openly sharing AISI's AI safety research insights and tools.** Through AISI, the UK government has the opportunity to influence the direction of development for frontier AI globally. This effect will be more profound the more widely AISI shares its research insights.
- 5. Use the forthcoming UK AI Bill as an opportunity to promote greater transparency and openness in AI development across the board.** The UK's forthcoming AI Bill offers a prime opportunity to enshrine the UK's commitment to AI safety while fostering a more transparent, open, and innovation-friendly development environment. Creating and implementing regulation for AI along the spectrum from fully-open to proprietary models is difficult. Therefore, it must be an objective of regulation not to accidentally disincentivise open development.

INTRODUCTION

The UK government has made economic growth its number one mission. Its success as a government, and our collective well-being as a nation, will depend on the policies and actions it takes in this area. AI is central to this growth agenda and the UK now stands at a pivotal juncture in shaping its AI future. With global momentum building around artificial intelligence, the promise of AI as a transformative economic and societal force is undeniable. The UK faces both a challenge and an opportunity in defining a distinctive strategy that will secure long-term economic benefits of AI and deliver public value.

Home to world-leading universities, a growing base of AI startups, and strong AI safety and governance expertise, the UK is a leading AI player with strong foundations for growth. However, the US and China currently dominate the AI landscape with unmatched scale, industrial ecosystems, data resources, and compute infrastructure that enable a high degree of AI self-sufficiency. The UK cannot compete on identical terms, so how will the UK capture a meaningful share of the emerging AI opportunity?

The question is not just about staying relevant at the forefront of AI development, but also about ensuring that its citizens and economy meaningfully benefit from the forthcoming AI transition through safe and effective AI applications, a competitive domestic AI sector, and the ability to shape the development and deployment of the technology in line with national priorities.

The *AI Opportunities Action Plan*, released in January 2025, sets a strong direction for the UK's AI ambitions. It highlights the need for innovation, strategic investment, talent development and, most notably, digital sovereignty with the establishment of a Sovereign AI Unit. Yet while the Plan articulates important goals, it lacks clarity on the 'how'. How will the UK support the rollout of the Action Plan to achieve its stated goals, particularly in a global market dominated by foreign private platforms and proprietary systems?

This paper argues that the UK should clarify and more fully commit to an "AI openness" strategy as a defining feature of its national approach. From models to data to compute infrastructure, open and open-source AI offers a practical and strategic pathway to delivering on the Action Plan's goals. It offers a route to AI sovereignty through cooperation, and enables the UK industry to build and deploy AI systems that are transparent, agile, and aligned with public interest. It also reflects the changing dynamics of the AI economy, where smaller, domain-specific models and the open ecosystems that support them are becoming more competitive and attracting investment. Together with plans for courting private partnership, leaning into an AI openness strategy will unlock the UK's opportunity to become, as articulated in the Action Plan, an "AI maker, not an AI taker".

In September 2024, we explored the technical possibilities for open sourcing AI in our paper, *Open Horizons*.¹ This present report, *The Open Dividend*, builds on the first to explore how the UK can use open and open-source AI to maximise the impact of the AI Opportunities Action Plan. The paper draws on international examples, evolving AI trends, and the UK's own strengths to make the case for a strategic commitment to AI openness. We make the case that there is an open dividend to be achieved by embedding AI openness in our growth agenda.

1 Seger & O'Dell (2024). Open Horizons: Exploring Nuanced Technical and Policy Approaches to Openness in AI. https://demos.co.uk/wp-content/uploads/2024/08/Mozilla-Report_2024.pdf

SECTION 1

WHAT DO WE MEAN BY A NATIONAL 'AI OPENNESS' STRATEGY?

A '**national AI openness strategy**' for the UK would be a coherent agenda that seeks to support and facilitate greater openness in AI development as a strategic mechanism toward realising the UK's broader AI goals. In this paper, we take the high level goals of bolstering AI innovation, driving AI adoption (for public benefit), and striving for greater AI sovereignty as a guiding framework for the UK from the *AI Opportunities Action Plan* (hereafter referred to as the *Action Plan*).

A national AI openness strategy would work by employing measures facilitating greater openness in AI development and deployment in the UK. As presented at the Columbia/Mozilla Convening on AI openness, **AI openness** can be generally understood as the broad public availability and ease of access² to key artefacts and documentation from AI across the AI tech stack including AI models (weights and code), datasets, documentation, safety tooling, and compute resources

(See Figure 1).³ We propose that the concept of AI openness is expanded further to incorporate accessibility of compute resources and support for participating in development communities and collaborative initiatives.

'AI openness' is distinct from '**open-source AI**'.⁴ A prominent definition of open-source AI developed in a co-design process coordinated by the Open Source Initiative (OSI) refers to the availability of a model for public download under open-source licence such that anyone is able to freely use, study, modify, and share the model.⁵ Open-source model distribution and licensing is only one component of AI openness.

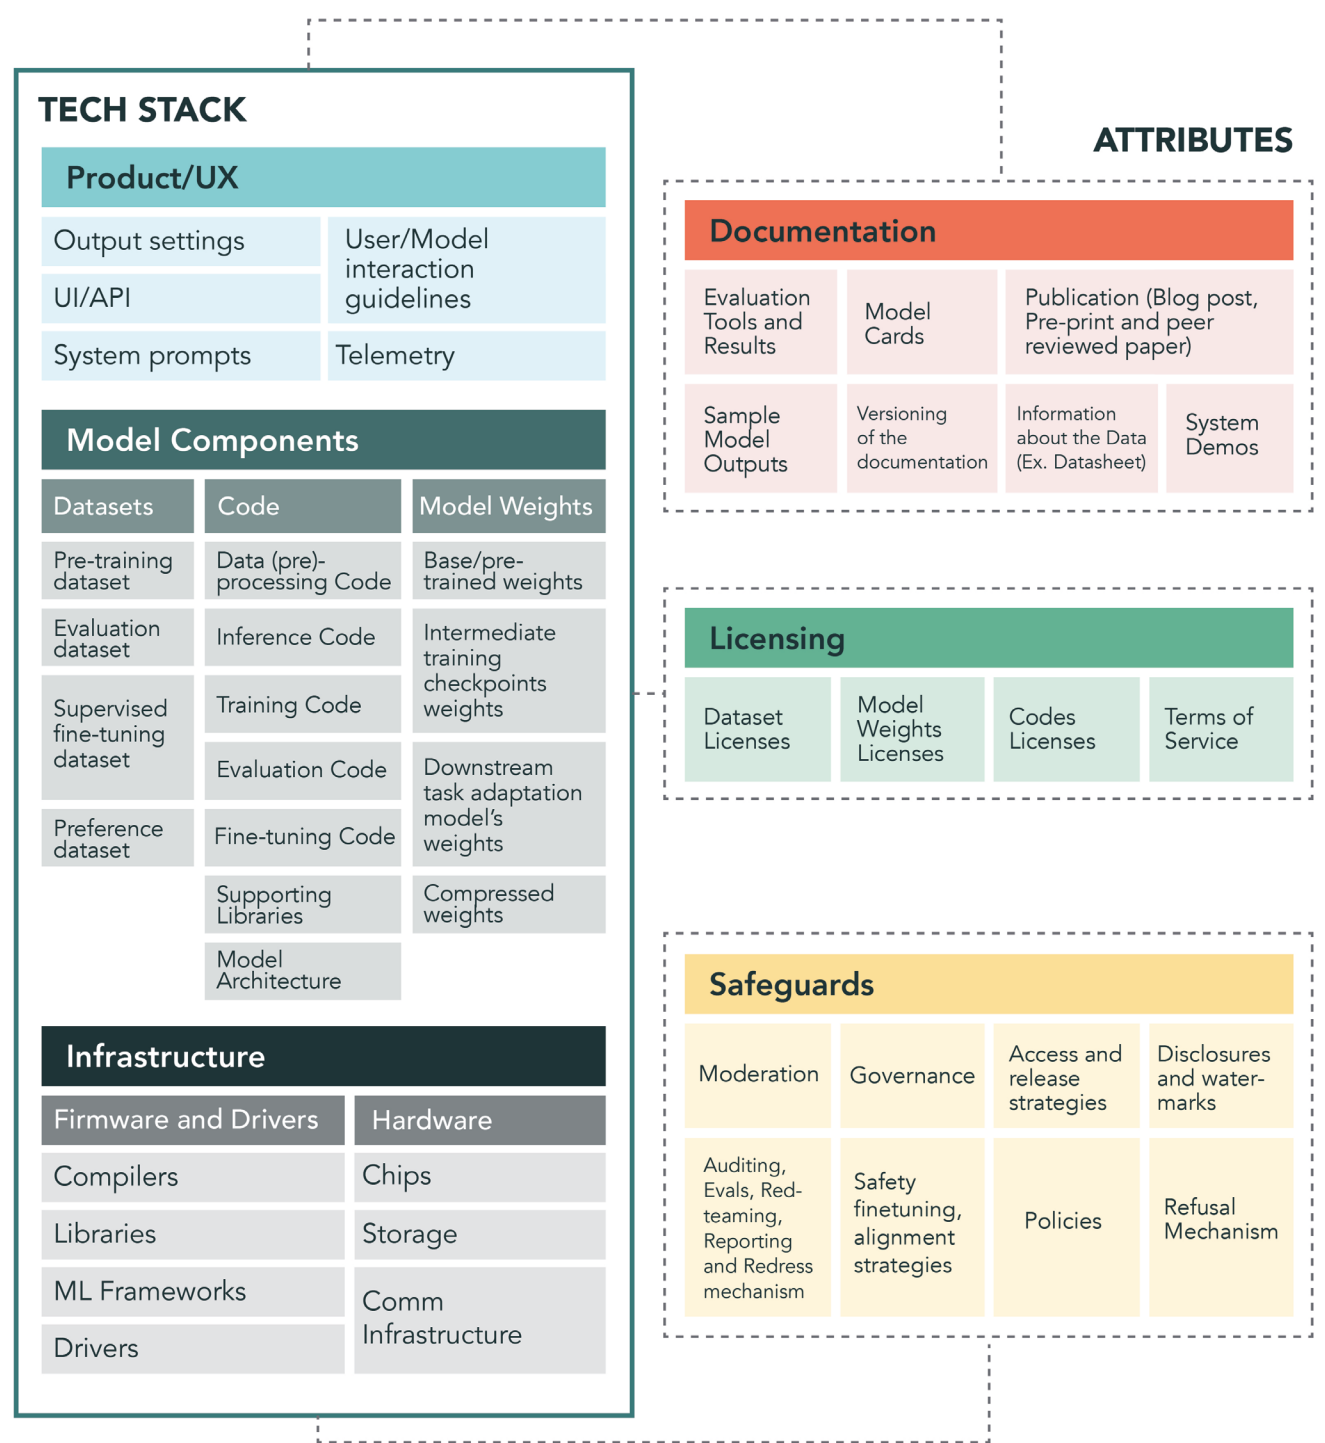
2 We add "ease of access" to the original Columbia Convening definition of AI openness building on Solaiman et al.'s (2025) recent paper, *Beyond Release: Access Considerations for Generative AI Systems*. <https://arxiv.org/abs/2502.16701>

3 Basdevant, A. et al. (2024). *Towards a Framework for Openness in Foundation Models: Proceedings from the Columbia Convening on Openness in Artificial Intelligence*. <https://arxiv.org/abs/2405.15802>

4 For a thorough discussion of the difference between often confused terms such as open access, open-source, open science, open licence, open knowledge, and open collaboration in relation to AI see White et al. (2024). *The Model Openness Framework: Promoting Completeness and Openness for Reproducibility, Transparency, and Usability in Artificial Intelligence*. Retrieved from <https://arxiv.org/abs/2403.13784>

5 For specific conditions and model component access requirements see: OSI (2024). *The Open Source AI Definition - 1.0*. Retrieved 23 April, 2025. <https://opensource.org/ai/open-source-ai-definition> ; for a discussion of other definitions, see OpenUK (2024). *State of Open: The UK in 2024 Phase Four*. <https://openuk.uk/wp-content/uploads/2024/12/State-of-Open-The-UK-in-2024-Phase-Four.pdf>

FIGURE 1
GENERAL-PURPOSE AI SYSTEM STACK & DIMENSIONS OF OPENNESS⁶



‘Public AI’ is another related term that will feature in this report. Public AI refers to AI models and supporting infrastructure including datasets, compute, and safety tooling that are developed and maintained as a public good.⁷ They are accessible to the public (and therefore usually open or open-source) and accountable to the public for their function and impact.

⁶ Figure reproduced from Basdevant, A. et al. (2024). Towards a Framework for Openness in Foundation Models: Proceedings from the Columbia Convening on Openness in Artificial Intelligence. <https://arxiv.org/abs/2405.15802>
⁷ Surman, M., Marda, N. and Sun, J. (September 30, 2024). Public AI. Mozilla. <https://www.mozillafoundation.org/en/research/library/public-ai/>

SECTION 2

WHY LEAN INTO AI OPENNESS?

For a country looking to drive domestic AI industry growth and reap the public benefits of widespread AI adoption, the advantages of supporting AI openness are numerous. Those we explore here include:

- Driving AI innovation
- Supporting AI industry
- Supporting flexible AI adoption
- Acting as an economic multiplier
- Driving public benefit through public AI
- Facilitating greater tech sovereignty and influence over AI futures

2.1 DRIVING AI INNOVATION

Modern AI is the product of decades of open research, public collaboration, and community-driven experimentation. Foundational breakthroughs - such as the transformer architecture that underpins today's large language models and training methodologies that are now standard practice in model development - emerged from openly shared academic research and open-source prototypes. The rapid pace of AI progress has been fuelled by this culture of openness, where researchers and developers build on each other's work, refining methodologies and iterating faster.

Efficiency gains are a particularly important outcome of this process. Faced with limited access to compute due to high cost and geopolitical barriers (e.g. chip export controls on China), open developers have had strong incentive to create smarter, leaner approaches. Breakthroughs in fine-tuning - specifically Low Rank Adaptation (LoRA)⁸ - were driven by open-source communities out of necessity. It is a process by which the performance of smaller models can be significantly improved by optimising model weights using the outputs of more high-capable

8 T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer (2023). QLoRA: Efficient Finetuning of Quantized LLMs. DOI: 10.48550/arXiv.2305.14314. arXiv: 2305.14314 [cs]

models as training data. More recent advances include Ai2's (Allen AI's) OLMo 2,⁹ a relatively small model which outperforms OpenAI's GPT 3.5 and GPT4o mini, and OlympicCoder¹⁰ which outperforms Anthropic's much larger Claude Sonnet 3.7 model at complex coding tasks. Overall, tasks which once required models with over 100 billion parameters just two years ago can now be accomplished with models under 2 billion parameters.¹¹ The 2025 AI Index report similarly notes that the performance gap between open and closed weight models has decreased from 8% to 1.7% in just the past year.¹²

Given the escalating costs of compute and energy, this improved efficiency brought by open innovation is not just technically useful, it is economically essential. For AI developers outside the US and China and not partnered with a major compute provider (e.g. AWS, Microsoft Azure, or Google Cloud), model efficiency is the key to competitive market participation. Reducing compute and energy usage will also be key to minimising environmental impacts of AI as development and use surges. A recent report from the International Energy Agency (IEA) predicts that energy usage by AI-optimised data centres will quadruple between now and 2030, more than doubling global data centre energy consumption, and accounting for more than 20% of total growth in global energy demand.¹³

Open-source ecosystems have also been driving innovation in distributed compute for decentralised model training and use. Decentralised model training works by distributing the computational workload across a network of globally dispersed, often heterogeneous, computing resources including consumer-grade GPUs. For example, Prime Intellect has recently completed a training run for INTELLECT-2, a 32 billion parameter open-source reasoning model with contributions for 20 independent computer providers spanning 3 continents.¹⁴ In a similar vein, researchers out of Abu Dhabi and Chengdu have released Prima.cpp, an open-source software that helps run and use large-scale models like DeepSeekR1 or Llama-3-70b by linking up a collection of home compute clusters (i.e. laptops, desktops, and phones).¹⁵

These recent advances in decentralised compute are lowering the barriers to entry for large-scale AI development, enabling researchers and smaller organisations to collaboratively train and use powerful models without needing massive centralised infrastructure.

2.2 SUPPORTING AI INDUSTRY

Today's AI industry is heavily reliant on open-source foundations. Widely used machine learning and deep learning frameworks such as TensorFlow¹⁶ and PyTorch¹⁷ are maintained as open-source projects and supported by contributions from a mix of independent developers and large corporations. These frameworks power countless commercial applications and academic projects, offering a common, accessible infrastructure for AI development.

Open repositories of datasets and pretrained models such as those found on GitHub or Hugging Face also allow developers to build products and services without needing to train models from scratch. This lowers the entry barrier, especially for smaller companies, startups, and researchers with limited resources.

9 <https://allenai.org/blog/olmo2-32B>

10 <https://huggingface.co/blog/open-r1/update-3>

11 <https://huggingface.co/blog/evijit/smollm-deepseek-bias-eval>

12 Maslej, N. et al. (April 2025). The AI Index 2025 Annual Report. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University. <https://hai.stanford.edu/ai-index/2025-ai-index-report>

13 IEA (April 2025). AI is set to drive surging electricity demand from data centres while offering the potential to transform how the energy sector works. <https://www.iea.org/news/ai-is-set-to-drive-surging-electricity-demand-from-data-centres-while-offering-the-potential-to-transform-how-the-energy-sector-works>

14 (April 2025). INTELLECT-2: Launching the First Globally Distributed Reinforcement Learning Training of a 32B Parameter Model. Accessed 5 May, 2025. <https://www.primeintellect.ai/blog/intellect-2>

15 Li et al. (2025). PRIMA.CPP: Speeding Up 70B-Scale LLM Inference on Low-Resource Everyday Home Clusters. <https://arxiv.org/abs/2504.08791>

16 Why TensorFlow. Accessed May 27, 2025. <https://www.tensorflow.org/about>

17 Zemlin, J. (2022). Welcoming PyTorch to the Linux Foundation. The Linux Foundation. Retrieved August 16, 2024, from <https://www.linuxfoundation.org/blog/blog/welcoming-pytorch-to-the-linux-foundation>

Accordingly, open-source has become a well-established and financially viable category of venture capital, with valuations of companies that develop open-source software sometimes reaching tens of billions of dollars.¹⁸ In AI, this trend can be seen with Mistral, a French AI ‘unicorn’ which open-sources many of its models was valued at \$6 billion in 2024.¹⁹ Meanwhile Meta continues to position itself as all-in on open-source,²⁰ announcing a variety of details at its recent ‘LLamaCon’ event about forthcoming open-source model releases, open APIs, security products for open-source developers, and more.²¹

At first glance, open-source can seem odd as a business plan – why invest in developing a model just to release it for free? But the benefits of open-sourcing to AI companies are numerous. Because the models are freely available, more people are encouraged to adopt them and interact with the company’s wider product ecosystem. Moreover, the companies can benefit from the (often free) labour and expertise of people in wider open-source communities who may add features and suggest improvements. There is a clear business strategy in providing one software component for free and charging for complementary software or services called ‘commoditising the complement’²²: the free software expands the market of potential users by lowering barriers to access and adoption, while the paid-for services make up for the cost.

Aside from open-source software and AI models, open-source hardware has also been essential to the growth of heavy data-processing industries including AI.

As demand for computational power has surged, especially in cloud services and machine learning, open hardware standards have played a crucial role in enabling scalable, cost-effective infrastructure. A key driver of this has been the Open Compute Project Foundation (OCP), launched in 2011 after Facebook, facing mounting infrastructure demands, developed new software, servers, racks, power supplies, and data centre designs tailored for energy and operational efficiency and scalability.²³ Facebook open-sourced the designs giving rise to international collaboration among now 400 + companies including Microsoft, Google, Intel, Dell, Cisco, NVIDIA, AMD, IBM, OVH Cloud, Tencent, and ARM to create commodity hardware designs that are more efficient and flexible for scalable computing.²⁴

The realised benefits of these open standards are clear. Microsoft, for example, has reported over 40% cost savings in server deployment using OCP designs,²⁵ Meta reported saving \$1.2 billion,²⁶ and large AI players like Baidu and Alibaba have adopted OCP-inspired architectures to build more flexible, thermally efficient data centres.²⁷ However, the rise of highly specialized AI hardware like Google’s TPUs²⁸ and NVIDIA’s integrated AI platforms²⁹ is starting to challenge the open ethos by encouraging vertically integrated, proprietary stacks. In response, and to help avoid vendor lock-in, OCP is introducing projects such as its Open Systems for AI (OSAI) initiative working to standardise scalable AI infrastructures,³⁰ its Open Chiplet Economy

18 Laverne (2025). ‘The Open Source Payoff’. Serena. <https://blog.serenacapital.com/the-open-source-payoff-5e835c54c0f1>

19 Lunden (2024). ‘Sources: Mistral AI raising at a \$6B valuation, SoftBank ‘not in’ but DST is’. TechCrunch. <https://techcrunch.com/2024/05/09/sources-mistral-ai-raising-at-a-6b-valuation-softbank-not-in-but-dst-is/>

20 Meta (2025). ‘Everything we announced at our first-ever LlamaCon’. <https://ai.meta.com/blog/llamacon-llama-news/>; see also Section 4 of OpenUK (2024). State of Open: The UK in 2024 Phase Four. <https://openuk.uk/wp-content/uploads/2024/12/State-of-Open-The-UK-in-2024-Phase-Four.pdf>

21 Meta (2025). ‘Everything we announced at our first-ever LlamaCon’. <https://ai.meta.com/blog/llamacon-llama-news/>

22 E.g., Angular Ventures Weekly (2024). ‘Commoditize Your Complement: Meta AI Edition’. Medium. <https://medium.com/angularventures/commoditize-your-complement-meta-ai-edition-f81e44498aed>

23 Open Compute Project: About. (Accessed May 21, 2025). <https://www.opencompute.org/about>

24 Open Compute Project: Membership Directory. (accessed May 21, 2025). <https://www.opencompute.org/membership/membership-directory>

25 Microsoft Blog (Jan 27, 2014). Microsoft contributes cloud server designs to the Open Compute Project. (Accessed May 21, 2025). https://blogs.microsoft.com/blog/2014/01/27/microsoft-contributes-cloud-server-designs-to-the-open-compute-project/?utm_source=chatgpt.com

26 Miller, R. (2014) Facebook: Open Compute Has Saved Us \$1.2 Billion. Accessed 27 May, 2025. <https://www.datacenterknowledge.com/data-center-chips/facebook-open-compute-has-saved-us-1-2-billion>

27 Morgan, T. (September 14, 2021). Taking the Long View on Open Computing. The Next Platform. https://www.nextplatform.com/2021/09/14/taking-the-long-view-on-open-computing/?utm_source=chatgpt.com

28 Cloud Tensor Processing Units: Accelerate AI development with Google Cloud TPUs. Accessed 25 May, 2025. <https://cloud.google.com/tpu>

29 Nvidia DGX Platform. Accessed 25 May, 2025. https://www.nvidia.com/en-us/data-center/dgx-platform/?utm_source=chatgpt.com

30 Coyle, R. (January, 2025). Open Systems for AI. <https://www.opencompute.org/projects/open-systems-for-ai>

subproject working to facilitate integration of specialised components from different vendors,³¹ and its OCP Marketplace AI Portal to serve as a hub for designers and builders to share the latest AI infrastructure products.³²

2.3 SUPPORTING FLEXIBLE AI ADOPTION

One of the central goals of the AI Opportunities Action Plan is to push hard on AI adoption across the private and public sectors. It has a specific emphasis on piloting and scaling AI tools for deployment in public services to improve service efficiency and quality and to cut costs. However, public service bodies operate under tight financial constraints, and if AI adoption is to be a route to cost savings the AI solutions themselves must also be affordable.

Open-source tools offer a natural fit toward this end. In the private sector, cost savings are consistently cited as the leading reason for choosing open-source tools (AI models and software) over proprietary alternatives.^{33,34} There are no licensing fees or vendor overheads.

Another key benefit is the ability to avoid vendor lock-in. Organisations that build on open models and employ open-source hardware retain control over their tech stack and reduce long-term dependence on single providers. Open-source also facilitates greater interoperability between platforms providing transparent standards, customisable code, and fostering a community-driven approach to solving integration challenges. Control and interoperability are especially important for public infrastructure. The UK's public sector is not known for being particularly agile in its procurement processes, and the integration of closed models from large foreign providers would entrench expensive dependencies for years to come.

Concerns about safety and security often deter some organisations from adopting open-source AI. However safety and security are also cited as a reason for choosing to adopt open tools.³⁵ The perceived benefit of procuring either open or closed likely depends on internal technical expertise within the organisation. Testing, maintaining, using and contributing to open-source software requires significant human skill and capital, which may deter some organisations from using AI openness and software solutions.³⁶ Addressing digital skills shortages will therefore be key to ensuring UK businesses and public sector alike can fully benefit from absorbing open-source innovations from all over the world.³⁷

2.4 ACTING AS AN ECONOMIC MULTIPLIER

Open-source software (OSS) has long served as a powerful catalyst for economic growth and innovation. It enables innovation, reduces duplication, and lowers the barriers to entry for new firms. Without any open-source contributions, the average country would lose around 2.2% of its GDP even after accounting for knowledge 'spillover' - the phenomenon where the benefits of one country's OSS contributions extend beyond its borders and are freely used by others,

31 Nasrullah, J. & Ramamurthy, A. (August 26, 2024). Announcing the Renaming of ODSA to "Open Chiplet Economy": A New Era in Open Silicon Innovation. https://www.opencompute.org/blog/announcing-the-renaming-of-odsa-to-open-chiplet-economy-a-new-era-in-open-silicon-innovation?utm_source=chatgpt.com

32 Grossner, C. (April 29, 2025). The Open Compute Project Accelerating Deployment of Next Gen AI Clusters. <https://www.opencompute.org/blog/the-open-compute-project-accelerating-deployment-of-next-gen-ai-clusters>

33 Bisht, A, et al. (2025). Open source technology in the age of AI. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/open-source-technology-in-the-age-of-ai>

34 Van Slyke, D. (March 14, 2019). Baidu, Facebook and Microsoft work together to define the OCP Accelerator Module specification. Open Compute Project. https://www.opencompute.org/blog/baidu-facebook-and-microsoft-work-together-to-define-the-ocp-accelerator-module-specification?utm_source=chatgpt.com

35 Bisht, A, et al. (2025). Open source technology in the age of AI. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/open-source-technology-in-the-age-of-ai>

36 Nagle, F., Wheeler, D. A., Lifshitz-Assaf, H., Ham, H., & Hoffman, J. L. (2020). Report on the 2020 FOSS Contributor Survey. The Linux Foundation & The Laboratory for Innovation Science at Harvard. <https://www.linuxfoundation.org/resources/publications/foss-contributor-2020>

37 Blind, K., Schubert, T. (2024). Estimating the GDP effect of Open Source Software and its complementarities with R&D and patents: evidence and policy implications. *J Technol Transf* 49, 466–491. <https://doi.org/10.1007/s10961-023-09993-x>

without direct compensation.³⁸ While knowledge spillover can dilute the domestic return on investment, the net global effect is overwhelmingly positive. More so, countries with strong R&D ecosystems, like the UK, capitalise most strongly.

A 2021 study commissioned by the European Commission illustrates this dynamic well.³⁹ It found that roughly €1 billion in OSS investment by EU-based firms resulted in an economic impact of €65–€95 billion. A 10% increase in OSS contributions was projected to increase EU GDP by 0.4%–0.6% annually and generate more than 600 additional Information and Communication Technology (ICT) start-ups. The European Commission subsequently cited the economic importance of OSS when introducing new rules to streamline the process for open-sourcing government software.⁴⁰ Building upon the established economic benefits of open-source software (OSS), it is reasonable to anticipate that open-source AI will yield similar advantages.⁴¹

2.5 DRIVING PUBLIC BENEFIT THROUGH PUBLIC AI

AI development efforts are rapidly transitioning from R&D to consumer application, and in this transition the direction of proprietary model development is motivated by economic interest serving commercial priorities. While this model drives rapid progress, it does not always align with the broader public interest, and vital applications that support inclusion, equity, or long-term social benefit can be overlooked. For example, large AI firms have less incentive to support minority and under-resourced languages like Cornish, Gaelic and Irish in the development of large language models. This stems from the relatively low number of speakers and the difficulty of acquiring training data-sets, even though these languages are important for cultural preservation, linguistic diversity, and digital access.⁴² Open-source AI enables communities, researchers, and public bodies to develop and sustain such tools, ensuring these needs are met.^{43,44,45}

AI models and supporting infrastructure including datasets, compute, and safety tooling can also be developed and maintained as a public good - or 'public AI' - that is accessible to the public and accountable to the public for their function and impact. A shift towards public AI investments is reflected in recent initiatives like Current AI, announced at the 2025 Paris AI Summit working to create an open, global infrastructure layer for public-interest AI projects.⁴⁶ It is backed by private and public investment across 10 governments and major philanthropies. Similarly, ROOST (Responsible Open-Source Tooling), also launched at the Paris Summit, is building shared safety tools and governance protocols to support the development of trustworthy open models.⁴⁷

Supporting open-source AI and building out public AI infrastructure is an opportunity to ensure that AI works for a broad range of public needs, enables more inclusive digital tools that fill gaps in service provision, and supports the development of AI systems that are aligned with public values. This approach also helps underpin public trust in the long-term accountability and public orientation of AI systems.

38 Blind, K., Schubert, T. (2024). Estimating the GDP effect of Open Source Software and its complementarities with R&D and patents: evidence and policy implications. *J Technol Transf* 49, 466–491. <https://doi.org/10.1007/s10961-023-09993-x>

39 European Commission. (September 2, 2021). Study about the impact of open source software and hardware on technological independence, competitiveness and innovation in the EU economy. Accessed 23 May, 2025. <https://digital-strategy.ec.europa.eu/en/library/study-about-impact-open-source-software-and-hardware-technological-independence-competitiveness-and>

40 https://ec.europa.eu/commission/presscorner/detail/en/ip_21_6649

41 (December 10, 2024). Statement from Economists on the Importance of Open Source AI. Accessed 24 May, 2025. <https://open.mozilla.org/economists/>

42 AI initiative gives Gaelic a foothold in the digital age. The University of Edinburgh. Accessed May 27, 2025. <https://cahss.ed.ac.uk/news-events/news/current-news/ai-initiative-gives-gaelic-a-foothold-in-the-digit>

43 Indian Ministry of Electronics and Information Technology (MeitY; 2025). 'About Bhashini'. <https://bhashini.gov.in/about-bhashini>

44 Typhoon (2025). 'About'. <https://opentypoon.ai/about>; Pipatanakul, K. et al. (2023). Typhoon: Thai Large Language Models. <https://arxiv.org/abs/2312.13951>

45 Tran, K., O'Sullivan, B. and Nguyen, H. (2024). UCCIX: Irish-eXcellence Large Language Model. <https://arxiv.org/html/2405.13010v1>

46 Current AI. Accessed 24 May, 2025. <https://www.currentai.org/>

47 ROOST. Accessed 24 May, 2025. <https://roost.tools/>

Finally, investment towards international collaboration on public AI across the AI stack is an opportunity to bolster the open-source counterweight to proprietary tech. The larger the collaboration, the stronger and more viable the open alternative is to relying on a handful of leading AI developers.

2.6 FACILITATING GREATER TECH SOVEREIGNTY AND INFLUENCE OVER AI FUTURES

All of the benefits outlined above ultimately contribute to building greater national AI sovereignty and resilience, ensuring that a country is not wholly dependent on US AI developers and service providers.

Attempting to build AI and all associated support infrastructure domestically is not a realistic option for many countries outside the US and China given the financial constraints and the size of investment needed. Instead, these countries can look outward toward embracing collaboration, open ecosystems, and shared resources both to build strength and resilience through collective innovation and to provide foundations for nurturing domestic AI champions. Countries like France, India, and Singapore are already following such AI openness strategies with varying degrees of success (see Section 3 and Appendix 1).

Importantly, embracing an AI openness strategy does not necessarily mean rejecting partnerships or investment opportunities from proprietary big tech firms where they arise, but, at a minimum, we need to complement them by establishing a strong, open counterweight to foreign proprietary dominance.

SECTION 3

WHY NOW?

The UK can harness the above benefits of open and open-source AI toward the delivery of the AI Action plan goals - section 5 provides a more detailed breakdown of a UK specific AI openness strategy. But the benefits of openness in software and AI development have been evident for decades, so why the hard push for the UK to commit to an AI openness strategy now?

We make three key arguments for acting now: **(1) The AI game is changing, (2) the geopolitical landscape is shifting, and (3) the UK needs a practical and resilient AI openness strategy to match.**

3.1 THE AI GAME IS CHANGING

Since the open release of the DeepSeek V3 and R1 models out of China, it has become increasingly apparent that the AI development game is changing. As discussed in section 2.1, the capability gap between proprietary frontier models and open-source models is quickly narrowing. Meanwhile advances in model efficiency and integration stemming from open-source innovation have provided the means for a shift in market focus toward smaller, specialised models instead of giant monoliths.

Enterprises and investors are starting to show preference for efficient, specialised models over large general purpose models because they are cheaper, faster, and easier to customise.^{48,49} China's own AI startup scene went through a major overhaul after DeepSeek's success.⁵⁰ Leading startups like Zhipu, 01.ai, Baichuan, and Moonshot have pivoted from foundational model training to specialised applications, cost-cutting, and niche markets to stay afloat as ballooning foundation model training costs without reliable revenue raise concerns among investors. DeepSeek experienced a surge in popularity in Europe and North America as well, with companies such as Microsoft and Perplexity quickly including the model in their services.⁵¹

48 DeBiase, D. (2024). Why Small Language Models Are The Next Big Thing In AI. Forbes. <https://www.forbes.com/sites/deandebiase/2024/11/25/why-small-language-models-are-the-next-big-thing-in-ai/>

49 Metz, R. (2024). In AI, Smaller, Cheaper Models Are Getting Big Attention. Bloomberg UK. <https://www.bloomberg.com/news/articles/2024-08-08/move-over-llms-small-ai-models-are-the-next-big-thing>

50 Olcott, E., McMorrow, R. & Wu, Z.. (2025). Chinese AI start-ups overhaul business models after DeepSeek's success. <https://www.ft.com/content/c19f3988-45d7-4a81-854d-9ba0d71812fe>

51 Kan (2025). 'DeepSeek Sees Huge Surge to Become Second Most Popular AI Chatbot'. PCMag. <https://uk.pcmag.com/ai/156531/deepseek-sees-huge-surge-to-become-second-most-popular-ai-chatbot>; Forlini (2025). 'DeepSeek Is Here to Stay as Microsoft, Perplexity Integrate Its Model'. PCMag. <https://uk.pcmag.com/ai/156495/deepseek-is-here-to-stay-as-microsoft-perplexity-integrate-its-model>

While scaling laws generally still hold - whereby highest performance is still achieved through brute force model size - building bigger is increasingly looking less economical. OpenAI, for example, offers a ChatGPT Pro subscription for \$200 per month.⁵² At such costs, paying for access to the biggest models is not an option for many individuals and organisations, especially in the cash-strapped public sector, or where the full range of a general purpose model's capabilities may not be needed by the adopter. Such would be the case, for example, where a model is to be employed specifically in the medical sector or to be employed primarily for sifting and analysing legal texts (on May 6 2025, the Solicitors Regulation Authority (SRA) approved the first purely AI based law firm, Garfield.Law⁵³). Meanwhile, OpenAI is losing money on its ChatGPT Pro subscription plan, and the company as a whole has not yet been profitable despite raising \$20 billion since its inception.⁵⁴ OpenAI's CEO Sam Altman has said that OpenAI was "on the wrong side of history" following DeepSeek's success and they may begin to open-source older models to help it compete.⁵⁵

The shift from single, high-capable models to simple, good-enough models is also being facilitated by the development of open standards for protocols that allow AI models to interface with external systems including data sources, wider software environments, and other AI models. Most notably, in November 2024 Anthropic open-sourced its Model Context Protocol (MCP)⁵⁶ which the company describes as a "USB-C port for AI applications".⁵⁷ It lets different AI models interface with apps, files, and data in a standard way so that developers do not need to build custom software connections in each case.⁵⁸ MCP has been adopted by OpenAI, Google DeepMind, and Microsoft. In a similar vein, Google released the Agent2Agent (A2A) open protocol in 2025.⁵⁹ APA compliments MCP⁶⁰ and enables developers to build "large-scale, multi-agent systems"⁶¹ by standardising model-to-model interface. It has buy-in from Accenture, Capgemini, Cohere, and Oracle, among others.⁶²

Together, these open protocols allow developers to interact with wider software ecosystems and combine models which might be individually less capable into powerful workflows. Moreso, by letting developers link models with external data sources, it is possible to move beyond reliance on single large models that are instilled with large swathes of 'knowledge' from training data, to smaller models, or networks of smaller specialised models, and can search and query in real time.

None of this to say that proprietary frontier AI development will come to a halt. Pushing the bleeding edge of AI capability currently remains the purview of brute force scaling,⁶³ and in the long term it may still be economical for a select few AI actors who are able to keep operating in this space, capitalising on benefits of being the first mover toward new advances such as agentic AI.

However, there are several reasons to be skeptical about the degree of first mover advantage in frontier development. First, the closing window between proprietary frontier models and open-

52 ChatGPT Pricing. Accessed May 5, 2025. <https://openai.com/chatgpt/pricing/>

53 Solicitors Regulation Authority (May 6, 2025). SRA approves first AI-driven law firm. <https://www.sra.org.uk/sra/news/press/garfield-ai-authorised/>

54 Wiggers, K. (2025). OpenAI is losing money on its pricey ChatGPT Pro plan, CEO Sam Altman says TechCrunch. <https://techcrunch.com/2025/01/05/openai-is-losing-money-on-its-pricey-chatgpt-pro-plan-ceo-sam-altman-says/>

55 Kahn (2025). 'DeepSeek has tilted the balance towards open source AI, but big security issues remain'. Fortune. <https://fortune.com/2025/02/04/sam-altman-openai-wrong-side-of-history-open-source-deepseek/>

56 Anthropic (2024). 'Introducing the Model Context Protocol'. <https://www.anthropic.com/news/model-context-protocol>

57 Anthropic (2025). 'Developer guide: Model Context Protocol (MCP)'. <https://docs.anthropic.com/en/docs/agents-and-tools/mcp>

58 Model Context Protocol (2025). Github. <https://github.com/modelcontextprotocol>

59 Surapaneni et al. (2025). 'Announcing the Agent2Agent Protocol (A2A)'. Google. <https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/>

60 Google (2025). 'Agent2Agent (A2A) Protocol'. GitHub. <https://github.com/google/A2A?tab=readme-ov-file>

61 Surapaneni et al. (2025). 'Announcing the Agent2Agent Protocol (A2A)'. Google. <https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/>

62 Surapaneni et al. (2025). 'Announcing the Agent2Agent Protocol (A2A)'. Google. <https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/>

63 Maslej, N. et al. (April 2025). The AI Index 2025 Annual Report. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University. <https://hai.stanford.edu/ai-index/2025-ai-index-report>

source innovation raises questions about the duration of expected return for being first mover. If open-source quickly democratises competitive model capability beyond the market leader, then the opportunity for the market leader to offer its product as a uniquely advanced capability is time limited. Second, the size of investment needed to drive up model capability via scaling is immense (e.g. training OpenAI's GPT-4 cost an estimated \$100million, Google's Gemini 1.0 Ultra an estimated \$192million, and Meta's Llama 3.1-405B an estimated \$170million⁶⁴) meaning there may be added economic advantage to being a second mover that waits to innovate and capitalise on open-source efficiency breakthroughs. Finally, as illustrated by OpenAI's case above, timely return on investment for large frontier AI model development is not a given.

In countries such as the US that have the investment backing to take a chance on first mover advantage,⁶⁵ it is possible the move could pay off in the long run. However, for the purpose of economic and AI industry growth in the rest of the world, leaning into the rich world of open-source innovation, collaborative resource sharing, and specialised application development is a stronger bet. It is also the bet already being made by countries in similar situations to the UK including France, India, Spain, Singapore, and Thailand (see Appendix 1 for case studies and lessons for the UK).

This does not mean the UK is out of the frontier AI game. In section 5.3 we discuss how the UK retains a strong position for influencing future directions of frontier AI development through academic and frontier safety contributions, both of which will be amplified through open science and resource sharing.

3.2 THE GEOPOLITICAL LANDSCAPE IS SHIFTING

While the AI industry is increasingly leaning towards specialised and efficient AI applications, shifting geopolitical pressures are also reinforcing the need for open and resilient AI infrastructure in the UK.

Heightening US-China tech rivalry combined with volatile US trade negotiations with traditional adversaries and allies alike are highlighting the unreliability of global AI infrastructure supply chains.

For example, in January 2025 the US Biden administration updated its export controls on the high performance computer chips needed to train advanced AI systems.⁶⁶ The new "AI Diffusion Rule" applies to compute hardware as well as trained model weights for proprietary models over 10^{26} FLOPs.⁶⁷ It also introduces a three-tiered export control system. Top tier countries including Australia, Canada, France, Japan, and the UK face limited export controls, while bottom tier countries including China, Cuba, Iran, North Korea, and Russia, face strict export bans. However, the Trump administration is changing changing up the rules again, doing away with the tiered approach,⁶⁸ and is reportedly looking instead to implement a global licencing

64 Maslej, N. et al. (April 2025). The AI Index 2025 Annual Report. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University. <https://hai.stanford.edu/ai-index/2025-ai-index-report>

65 For context The US recently secured a £20 billion investment from Saudi Arabian company DataVolt to build AI data centers and supporting energy infrastructure in the US. Meanwhile Google, DataVolt, Oracle, Salesforce, AMD, and Uber have committed to investing \$80 billion building cutting edge transformative technologies across the US and UAE. The UK cannot hope to compete. (May 13, 2025). Fact Sheet: President Donald J. Trump Secures Historic \$600 Billion Investment Commitment in Saudi Arabia. <https://www.whitehouse.gov/fact-sheets/2025/05/fact-sheet-president-donald-j-trump-secures-historic-600-billion-investment-commitment-in-saudi-arabia/>

66 Federal Register (January 2025). Framework for Artificial Intelligence Diffusion. <https://www.federalregister.gov/documents/2025/01/15/2025-00636/framework-for-artificial-intelligence-diffusion>

67 Note that there are currently no restrictions on the open release of model weights but the reason given is that no known open model has been trained that is over 10^{26} and the most capable models are still proprietary. These statements were made in January 2025 and the scene is quickly changing. Currently the most capable non-reasoning model is DeepSeek V3-0324 - an open-source model. When DeepSeek R2 - an open-source reasoning model - drops, it will likely be one of the best models in the world. This is to say, US export policy may well change in response.

68 Bureau of Industry & Security (May 13, 2025). Department of Commerce Rescinds Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls. Accessed 24 May, 2025. <https://www.bis.gov/press-release/department-commerce-rescinds-biden-era-artificial-intelligence-diffusion-rule-strengthens-chip-related>

regime operating on the basis of individual government-to-government agreements.⁶⁹ This mechanism would feed into Trump's broader trade strategy and allow both chip and model weight controls to serve as another lever in trade and tariff negotiations. The UK currently seems to be in Trump's good books, but will it last?

The widening rift between the US and Europe, coupled with China's increasing influence, has prompted some to voice concerns that reliance on foreign AI could pose risks to national- and cyber-security. These concerns are two-fold.

First, some have expressed fears that AI services based in the US or China will store data in those countries – data which could, in theory, then be accessed by their respective governments.⁷⁰ Such data processing and access could lead to citizens' sensitive personal data being available to governments with malign agendas, and could potentially violate local data protection laws.⁷¹ In particular, these discussions have highlighted the potential for data gathered by AI services to be used for surveillance and espionage purposes by foreign governments.⁷² While these concerns have been raised regarding Chinese AI apps like DeepSeek,⁷³ they are also pertinent to US-based services.

Second, there have been concerns raised about economic risks due to infrastructural dependency on systems controlled by foreign powers. These concerns include the general risk of economic dependence on tech from unreliable or hostile states⁷⁴ and the security risk of foreign states exerting control over the functioning of critical infrastructure which includes AI.⁷⁵ There have also been fears about the potential for foreign governments to impose censorship via AI models,⁷⁶ as it appears China is doing.⁷⁷ In the EU, these debates are part of a desire for greater leadership and competitiveness in AI, which the EU believes it will not achieve by relying on American or Chinese products.⁷⁸ As a result, some governments have started to explore building sovereign tech stacks which minimise reliance on external actors. The EU's debates about creating a 'EuroStack' are perhaps the most prominent example and encompass calls for AI sovereignty.⁷⁹

Against this backdrop of rising tensions, tech-rivalry, unpredictable trade wars, and security concerns, countries are increasingly focusing on **AI sovereignty** to reduce reliance on foreign technology.

China has proceeded most successfully on this front.

Despite US export controls, DeepSeek was able to train its highly performant V3 model on Nvidia H800 chips (a degraded version of the blocked Nvidia H100 chips) and reportedly at a fraction of the cost to comparable US models.⁸⁰ There are some doubts about the cost

69 Freifeld, K. (2025). Exclusive: Trump officials eye changes to Biden's AI chip export rule, sources say. <https://www.reuters.com/world/china/trump-officials-eye-changes-bidens-ai-chip-export-rule-sources-say-2025-04-29/>

70 Fleming (2025). 'What is digital sovereignty and how are countries approaching it?'. World Economic Forum. <https://www.weforum.org/stories/2025/01/europe-digital-sovereignty/>

71 Fleming (2025). 'What is digital sovereignty and how are countries approaching it?'. World Economic Forum. <https://www.weforum.org/stories/2025/01/europe-digital-sovereignty/>

72 Marchand (2025). 'US House Select Committee Report Accuses DeepSeek of Spying and Circumventing Export Controls on Chips'. TechPolicy.press. <https://www.techpolicy.press/us-house-select-committee-report-accuses-deepseek-of-spying-and-circumventing-export-controls-on-chips/>

73 E.g., Watson (2025). 'Dangers of DeepSeek's privacy policy: Data risks in the age of AI'. Security. <https://www.securitymagazine.com/articles/101374-dangers-of-deepseeks-privacy-policy-data-risks-in-the-age-of-ai>

74 E.g., France24 (2025). 'Europe seeks tech independence amid strained ties with Trump's America'. France24. <https://www.france24.com/en/live-news/20250416-europe-seeks-to-break-its-us-tech-addiction>

75 European Parliament (2024). 'Security and defence implications of China's influence on critical infrastructure in the European Union (C/2024/5719)'. European Union. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52024IP0028>

76 Funk et al. (2023). 'Freedom on the Net 2023: The Repressive Power of Artificial Intelligence'. Freedom House. <https://freedomhouse.org/report/freedom-net/2023/repressive-power-artificial-intelligence>

77 Booth & Milmo (2025). 'Chinese AI chatbot DeepSeek censors itself in realtime, users report'. The Guardian. <https://www.theguardian.com/technology/2025/jan/28/chinese-ai-chatbot-deepseek-censors-itself-in-realtime-users-report>

78 European Commission (2025). 'Shaping Europe's leadership in artificial intelligence with the AI continent action plan'. https://commission.europa.eu/topics/eu-competitiveness/ai-continent_en

79 The EuroStack Industry Group (2025). 'European Digital Industry Ecosystem Calls for Deployment of the EuroStack in the EU and Member States by 2030'. <https://euro-stack.eu/the-white-paper/>; Schaake (2025). 'Europe's dependence on US tech is a critical weakness'. <https://www.ft.com/content/30d6f79f-d1ee-49dc-bff5-719f18c1a9e5>;

80 Liu, A. et al. (2025). DeepSeek-V3 Technical Report. <https://arxiv.org/abs/2412.19437>

reporting,⁸¹ but the efficiency gains they were able to achieve remain notable, and that the models were openly released served a blow to the US AI market.⁸²

Nvidia H800 and A800 chips have now also been banned for export to China by the AI Diffusion Rule, but the long term effects this will have on China's AI capability are uncertain and could backfire. Driven by export control developments, China has been gearing up its own domestic AI hardware supply chain.⁸³ In competition with Nvidia in chip design, China has Huawei; and while Nvidia's chips are fabricated by TSMC (Taiwan Semiconductor Manufacturing Corporation), China boasts SMIC (Semiconductor Manufacturing International Corporation).

Furthermore, China is committed to an open-source AI development strategy, seeing it as an opportunity to reduce research and development costs and to bolster its domestic compute market.⁸⁴ Rousing open-source community enthusiasm for models such as DeepSeek V3 and R1 will likely help increase the global competitiveness of China's Huawei chips by disseminating base models built for Huawei's software ecosystem, CANN (Compute Architecture for Neural Networks).⁸⁵ Currently Nvidia's ecosystem, CUDA (Compute Unified Device Architecture) is more widely used by AI developers making Nvidia chips more attractive (See appendix 1 for more on China's open-source strategy). Orchestrating a global shift from CANN to CUDA could take years, but China is playing the long game.

However, for any country outside the US or China, the mass investment in AI infrastructure needed to reach a similar level of national AI sovereignty is infeasible. Building a competitive sovereign AI industry robust to the whims of the world's technological and geopolitical giants requires looking outward toward sovereignty through collaboration.

India has perhaps been most successful to date from the perspective of a lower-resourced country at employing a strong commitment to open-source software and AI openness as a means towards greater national tech sovereignty.

Since at least 2018, India has sought to promote AI openness as part of a wider emphasis on open-source software.⁸⁶ India's approach to achieve these goals has centred on AI projects developed by the government.

India's government has a long, successful history of promoting open-source software⁸⁷ and open-access data.⁸⁸ In 2015, India mandated that all software used at a federal level had to be open source,⁸⁹ with state and regional governments also following suit.⁹⁰ As of 2024, it was estimated that the adoption of an open-source computer operating system in schools in the

81 Sheehan, M. & Winter-Levy, S. (2025). Chips, China, and a Lot of Money: The Factors Driving the DeepSeek AI Turmoil. <https://carnegieendowment.org/emissary/2025/01/deepseek-ai-china-chips-explainer?lang=en>

82 Milmo, D., Hawkins, A., Booth, R. & Kolwe, J. (2025). 'Sputnik moment': \$1tn wiped off US stocks after Chinese firm unveils AI chatbot. The Guardian. <https://www.theguardian.com/business/2025/jan/27/tech-shares-asia-europe-fall-china-ai-deepseek>

83 Allan, G. C. (April 2025). DeepSeek: A Deep Dive. Centre for Strategic & International Studies. <https://www.csis.org/analysis/deepseek-deep-dive>

84 Webster et al. (2017). 'Full Translation: China's 'New Generation Artificial Intelligence Development Plan''. Stanford University. <https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>

85 Allan, G. C. (April 2025). DeepSeek: A Deep Dive. Centre for Strategic & International Studies. <https://www.csis.org/analysis/deepseek-deep-dive>

86 National Institution for Transforming India (NITI) Aayog (2018). National Strategy for Artificial Intelligence #AIFORALL. <https://www.niti.gov.in/sites/default/files/2023-03/National-Strategy-for-Artificial-Intelligence.pdf>; also see IndiaAI (2025). 'India's vision for AI: Prime Minister's address at the AI Action Summit, Paris'. <https://indiaai.gov.in/article/india-s-vision-for-ai-prime-minister-s-address-at-the-ai-action-summit-paris>

87 E.g., OpenForge (2025). 'About'. <https://openforge.gov.in/openforge/about.php>; Global Digital Public Infrastructure Repository (2025). 'India'. https://www.dpi.global/globaldpi/india_list; see also Section 16 of India's 2023-24 national budget. Government of India (2023). Budget 2023-24. https://www.indiabudget.gov.in/doc/bspeech/bs2023_24.pdf; Government of India (2014). 'Policy on Adoption of Open Source Software for Government of India'. https://www.meity.gov.in/static/uploads/2024/02/policy_on_adoption_of_oss.pdf; Government of India (2015). Framework for Adoption of Open Source Software in e-Governance Systems. <https://egovstandards.gov.in/sites/default/files/2021-07/Framework%20for%20Adoption%20of%20Open%20Source%20Software%20in%20e-Governance%20Systems.pdf>

88 E.g., Open Government Data Platform (2025). 'About'. <https://www.data.gov.in/about>; Piriha (2015). 'How is open data changing India?'. World Economic Forum. <https://www.weforum.org/stories/2015/02/how-is-open-data-changing-india/>

89 Government of India (2015). 'Framework For Adoption of Open Source Software In e-Governance Systems'. <https://egovstandards.gov.in/sites/default/files/2021-07/Framework%20for%20Adoption%20of%20Open%20Source%20Software%20in%20e-Governance%20Systems.pdf>

90 De et al. (2015). 'Economic Impact of Free and Open Source Software Usage in Government Final Report'. International Centre for Free and Open Source Software (ICFOSS). [https://icfoss.in/doc/ICFOSS_economic-impact-free\(v3\).pdf](https://icfoss.in/doc/ICFOSS_economic-impact-free(v3).pdf)

state of Kerala had saved nearly ₹30,000,000,000 – approximately £265 million⁹¹ – compared to using proprietary software like Microsoft Windows.⁹²

Within this context, India's government has a stated goal of promoting AI openness through both investments and in-house development.⁹³ The government has led or supported several open-source AI projects via its centralised national mission for AI development and investment called IndiaAI.⁹⁴ IndiaAI provides AI researchers with access to compute,⁹⁵ promotes Indian open-source AI projects,⁹⁶ and is involved in the development of a public platform for data and model sharing similar to Hugging Face.⁹⁷

Closer to home, **France** specifically and the **EU** more widely are looking to follow-suit, though it is too early for realised benefits to be clearly identified.

France has chosen to promote open-source AI development primarily through government investments. The country's aim has been to develop AI national champions and support France's existing open-source AI developers, while avoiding dependence on monopolies for access to AI capabilities.⁹⁸

The country is home to an ecosystem of open-source AI developers and platforms for supporting open-source AI development.⁹⁹ These include Mistral AI,¹⁰⁰ an open-source frontier AI developer, and the French-American open-source AI model repository Hugging Face.¹⁰¹

In recent years, the French government has made a public commitment to supporting its open-source AI ecosystem and providing funding for more AI openness initiatives.¹⁰² Since 2023, France has announced several investments in open-source AI through grants, public-private partnerships, and co-investments made alongside venture capital. These include a contribution to Current AI¹⁰³ announced at the Paris AI Summit.

However, it is difficult to identify how much of France's success in AI openness can be attributed to the government's strategy. The French government's decision to promote AI openness as a national strategy was relatively recent, compared to countries like China, while projects invested in by the French government such as Current AI are still in development.

Meanwhile, the EU has taken a two-stranded approach to AI openness. In one strand, it has pursued innovative AI regulations with specific provisions for open-source AI via the AI Act (2024), which treats open-source AI models and systems more lightly than their closed-source

91 Based on an exchange rate of 1 Pound to 113.164 Indian Rupees, using exchange rates from 7/5/2025.

92 The Hindu Bureau (2024). 'KITE set to launch updated FOSS-based OS for public schools in Kerala'. The Hindu. <https://www.thehindu.com/news/national/kerala/kite-set-to-launch-free-updated-os-for-public-school-computers-in-kerala/article68553871.ece>

93 National Institution for Transforming India (NITI) Aayog (2018). National Strategy for Artificial Intelligence #AIFORALL. <https://www.niti.gov.in/sites/default/files/2023-03/National-Strategy-for-Artificial-Intelligence.pdf>; also see IndiaAI (2025). 'India's vision for AI: Prime Minister's address at the AI Action Summit, Paris'. <https://indiaai.gov.in/article/india-s-vision-for-ai-prime-minister-s-address-at-the-ai-action-summit-paris>

94 IndiaAI (2025). <https://indiaai.gov.in/>

95 ETech (2025). 'Explained: IndiaAI compute portal, AIKosha and other initiatives under the IndiaAI Mission'. The Economic Times. <https://economictimes.indiatimes.com/tech/technology/explained-indiaai-compute-portal-aikosha-and-other-initiatives-under-the-indiaai-mission/articleshow/118780355.cms>

96 E.g., Jeevanandam (2022). 'Eight interesting open-source Indian projects that can support AI research'. IndiaAI. <https://indiaai.gov.in/article/eight-interesting-open-source-indian-projects-that-can-support-ai-research>; Jeevanandam (2022). 'Sarvam AI launches open-source foundational models in 10 Indian languages'. IndiaAI. <https://indiaai.gov.in/article/sarvam-ai-launches-open-source-foundational-models-in-10-indian-languages>

97 Suri (2025). 'The Missing Pieces in India's AI Puzzle: Talent, Data, and R&D'. Carnegie Endowment for International Peace. <https://carnegieendowment.org/research/2025/02/the-missing-pieces-in-indias-ai-puzzle-talent-data-and-randd?lang=en>

98 Chatterjee & Volpicelli (2023). 'France bets big on open-source AI'. Politico. <https://www.politico.eu/article/open-source-artificial-intelligence-france-bets-big/>

99 Office of the President of France (2025). Make France an AI Powerhouse. <https://www.elysee.fr/admin/upload/default/0001/17/d9c1462e7337d353f918aac7d654b896b77c5349.pdf>

100 Mistral AI (2025). <https://mistral.ai/>

101 Hugging Face is a French-American company founded in New York by French AI developers. See Hugging Face (2025). <https://huggingface.co/huggingface>; Cai (2022). 'The \$2 Billion Emoji: Hugging Face Wants To Be Launchpad For A Machine Learning Revolution'. Forbes. <https://www.forbes.com/sites/kenrickcai/2022/05/09/the-2-billion-emoji-hugging-face-wants-to-be-launchpad-for-a-machine-learning-revolution/>

102 Chatterjee & Volpicelli (2023). 'France bets big on open-source AI'. Politico. <https://www.politico.eu/article/open-source-artificial-intelligence-france-bets-big/>

103 Current AI (2025). 'Current AI Launch Press Release'. <https://www.currentai.org/latest-updates/launchpressrelease>

or commercial counterparts.¹⁰⁴ In the other strand, the EU has emphasised investment in open-source AI through initiatives such as its Digital Europe Programme.¹⁰⁵ The EU has provided €20 million funding to the OpenEuroLLM project,¹⁰⁶ and is funding a project to make a European high performing open-source foundation model available for downstream finetuning.¹⁰⁷

The UK can take specific lessons from these examples and others outlined in Appendix 1. But at a higher level, the lesson here is that the UK must follow suit in looking to secure some degree of resilience in its AI dependencies and marketplace in the current geopolitical context.

3.3 THE UK NEEDS A PRACTICAL AND RESILIENT AI STRATEGY

With the market shifting towards smaller specialised AI models, combined with mounting urgency to reduce reliance on foreign proprietary AI supply chains, the time is ripe for the UK to lean more heavily into an AI openness strategy toward achieving the nation's AI goals.

It is a deeply pragmatic move - one that takes seriously the UK's strong but nonetheless lagging position in frontier AI development and that is driven by the well-documented benefits of open knowledge-sharing for tech innovation and industry growth.

By some measures the UK is third in the world for frontier AI development, but it still lags significantly behind the US and China.¹⁰⁸ The UK lacks the financial resources to build itself into an AI superpower by playing the US and China at their own game through sheer force of infrastructure and industry investment.

But what the UK does have to its advantage is (a) talent (leading AI expertise and scientific expertise) thanks to a trailblazing technology research sector driven by world-leading universities and research institutes,¹⁰⁹ and (b) large, high-quality data resources held by the NHS and Government Digital Service invaluable to domain specific AI application development.

Given the UK's limitations, liberty from heavy reliance on proprietary US big tech (both with respect to the AI tools and stack infrastructure they provide) must come through collaboration, knowledge-sharing, and shared resources. The future of AI can be fully dominated by big tech, or there can be a thriving open public resource counterpoint ranging from collective compute infrastructure and open data resource to open-source research and development, that nations worldwide can build and capitalise on their own terms. The larger the collaboration, the stronger the open-source counterpoint to proprietary big tech.

More excitingly, given the UK's strengths, the UK is well positioned to extract maximum value from the open ecosystems that it helps build and maintain. The UK can leverage its deep bench of AI and scientific expertise to develop robust, open resources and drive the development and adoption of AI solutions across both public and private sectors.

With a concentration of AI talent uniquely embedded within the public sector seated with AISI, the UK also has an opportunity to support industry by open-sourcing safety tooling, setting clear standards, and sharing implementation guidance. This same in-house expertise could also be leveraged to help facilitate more widespread public sector adoption of open tools. Section 5 will explore these levers in greater detail.

104 See European Union (2024). 'Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act)'. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>; hereafter EU AI Act 2024.

105 European Commission (2025). 'Digital Europe Programme'. https://commission.europa.eu/funding-tenders/find-funding/eu-funding-programmes/digital-europe-programme_en; Sawers (2025). 'Open source LLMs hit Europe's digital sovereignty roadmap'. TechCrunch. <https://techcrunch.com/2025/02/16/open-source-llms-hit-europes-digital-sovereignty-roadmap/>

106 OpenEuroLLM (2025). <https://openeurollm.eu/>; Sawers (2025). 'Open source LLMs hit Europe's digital sovereignty roadmap'. TechCrunch. <https://techcrunch.com/2025/02/16/open-source-llms-hit-europes-digital-sovereignty-roadmap/>

107 European Commission (2024). 'Making available a high performing open-source European foundation model for fine-tuning (DIGITAL-2024-AI-06-FINETUNE)'. <https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/digital-2024-ai-06-finetune>

108 Stanford University Human-Centred Artificial Intelligence (2024). 'Global AI Power Rankings: Stanford HAI Tool Ranks 36 Countries in AI'. <https://hai.stanford.edu/news/global-ai-power-rankings-stanford-hai-tool-ranks-36-countries-ai>

109 Stanford University Human-Centred Artificial Intelligence (2024). 'Global AI Power Rankings: Stanford HAI Tool Ranks 36 Countries in AI'. <https://hai.stanford.edu/news/global-ai-power-rankings-stanford-hai-tool-ranks-36-countries-ai>

SECTION 4

SQUARING AI OPENNESS WITH THE UK'S AI SAFETY COMMITMENTS

Before diving into the specific pathways toward an AI Openness strategy for the UK, there is a pressing question about how leaning into open-source would square with the UK's existing AI safety commitments.

The UK has positioned itself as a global leader in AI safety, beginning with the AI Safety Summit series and continuing with the establishment of the world-leading AI Safety Institute, now AI Security Institute (AISi). The Labour government has also committed to passing legislation on frontier AI safety through the forthcoming AI Bill. In all of these contexts - as well as in parallel discussions held around the EU AI Act, California SB 1047, and the US NTIA's consultation on open-weight models¹¹⁰ - serious concerns have been expressed about model misuse and the dissemination of dangerous AI capabilities facilitated by open-source model sharing. Subsequent discussion and research pertains to if, how, and when restrictions should be placed on the open sharing of model weights and code to mitigate these risks.¹¹¹ For example, where might an application only be made available for use or fine-tuning via API to prevent malicious actors from bypassing misuse safeguards?

So, at first glance, AI safety considerations may appear at odds with adopting an AI strategy advocating for greater openness. Here we attempt to make some progress in squaring that circle and explain why AI openness and a serious commitment to AI safety can sit hand in hand.

To begin, AI is an unhelpfully large umbrella term and the risks of AI openness described above should not be taken as a convincing argument against open-source AI writ large. AI describes a massive net of technology ranging from: single predictive algorithms, to image recognition software used in self-driving cars, to generative AI systems like ChatGPT, to Agentic AI systems that can autonomously make decisions and take actions to achieve goals with little to no human

¹¹⁰ NTIA Report (July 2024). Dual-Use Foundation Models with Widely Available Model Weights. Retrieved May 5, 2025, from <https://www.ntia.gov/sites/default/files/publications/ntia-ai-open-model-report.pdf>

¹¹¹ See Seger et al. (2023) Section 3 for an overview of risk categories from open sharing of highly capable frontier AI models; Seger, E. et al. (2023). Open-Sourcing Highly Capable Foundation Models: An evaluation of risks, benefits, and alternative methods for pursuing open-source objectives. Retrieved May 20, 2025, from <https://www.governance.ai/research-paper/open-sourcing-highly-capablefoundation-models>

intervention (like booking a nice table at a restaurant or autonomously executing military operations). Some experts predict we will have Artificial General Intelligence (AGI) which would be capable of human-level reasoning, learning, and problem-solving across numerous domains within 10 years.¹¹²

This range is significant because it shows that safety concerns related to openness should not be applied equally to all AI systems. For the vast majority of AI use cases featuring narrower AI applications and models behind the frontier of AI development, potential harm from misuse is limited. They do not provide malicious actors with a capability uplift beyond what they might achieve utilising other available technologies like internet search.¹¹³ This is to say that current AI capabilities pose low “marginal risk”. Marginal Risk describes the risk a technology poses through intentional misuse relative to (a) pre-existing technologies or (b) closed-source versions.¹¹⁴ Attending to marginal risk is important to prevent fear-mongering and to ensure recommended interventions are proportional to the threat posed. For example, some analysis shows that the risk of open-source AI for disinformation, biorisk (developing biology weapons), and cybersecurity is currently marginal.^{115,116} This is because either the system output is of relatively limited use or because there are other key drivers of bottlenecks to harm. For instance, while generative AI can be used to generate disinformation, the key driver to harm is how it is disseminated online. With biorisk, capability uplift is uncertain, but there is a strong bottleneck on harm given the physical resources and expertise needed to develop pathogens in a lab.

The clear exception is with respect to the production of child sexual abuse material (CSAM) and non-consensual intimate imagery (NCII).¹¹⁷ Dissemination of open-source image generation models has enabled a proliferation of AI tools like “nudification apps” that bypass safeguards that would have been built into the original model’s inference code. The challenge, however, is that some proposed model release restrictions that might centre on large frontier models would not apply to these models which tend to be smaller, and lowering the bar on model release restrictions would encapsulate far too much. It is therefore necessary to focus on risk mitigation measures that extend throughout the AI life-cycle and do not just focus on the point of model release. For example, model hosting platforms like GitHub and HuggingFace could work to moderate and remove these applications. The UK’s Online Safety Act has also taken an important step in criminalising production of CSAM and NCII and mandating the removal of the illegal content (AI generated or otherwise) from online platforms.¹¹⁸

But moving on, aside from the (mostly) limited marginal risk posed by current AI systems, greater openness has a proven benefit for improving safety and security. The dynamic has been clearly demonstrated for decades with open-source software (OSS) allowing for transparency, reproducibility, interoperability, peer review, and faster identification and remediation of vulnerabilities. Even though OSS allows malicious actors a clear view of the system and the opportunity to scrutinise the code for vulnerabilities to leverage, a much larger community of neutral and ‘white hat’ actors are working to identify and fix vulnerability to defend the system from attack. This is to say, in the “offence-defence balance” of OSS - a term describing the relative ease of carrying out and defending against attack¹¹⁹ - the scales tip in favour of defence.

112 Continuously running prediction initiated 2020; Barnnet, M. (2020). When will the first general AI system be devised, tested, and publicly announced? Accessed 24 May, 2025. <https://www.metaculus.com/questions/5121/date-of-artificial-general-intelligence/>

113 Bommasani, R. et al. (December 13, 2023). Considerations for Governing Open Foundation Models. Accessed 24 May, 2025. <https://hai.stanford.edu/policy/issue-brief-considerations-governing-open-foundation-models>

114 Bommasani, R. et al. (December 13, 2023). Considerations for Governing Open Foundation Models. Accessed 24 May, 2025. <https://hai.stanford.edu/policy/issue-brief-considerations-governing-open-foundation-models>

115 NTIA Report (July 2024). Dual-Use Foundation Models with Widely Available Model Weights. Retrieved May 5, 2025, from <https://www.ntia.gov/sites/default/files/publications/ntia-ai-open-model-report.pdf>

116 Bommasani, R. et al. (December 13, 2023). Considerations for Governing Open Foundation Models. Accessed 24 May, 2025. <https://hai.stanford.edu/policy/issue-brief-considerations-governing-open-foundation-models>

117 Thiel, D., Stroebe, M., and Portnoff, R. (2023). Generative ML and CSAM: Implications and Mitigations. <https://purl.stanford.edu/jv206yg3793>.

118 Department for Science, Innovation & Technology (24 April, 2025). Online Safety Act: Explainer. Accessed 24 May, 2025. https://www.gov.uk/government/publications/online-safety-act-explainer/online-safety-act-explainer?utm_source=chatgpt.com

119 Garfinkel, B. and Dafoe, A. (2019). How does the offense-defense balance scale? *Journal of Strategic Studies*, 42(6):736–763. DOI: 10.1080/01402390.2019.1631810.

The question is whether the same benefits of OSS translate to AI. For smaller, narrower AI systems the answer is likely that they do. The offence-defence balance hinges on a variety of factors such as how difficult it would be for a malicious actor to find or manipulate safety vulnerabilities without access to the open-source model and, for defenders, how easy it is to develop and disseminate fixes.¹²⁰ For smaller, narrower models the relative simplicity and reduced capability surface make it more feasible for defenders to anticipate threats, rapidly iterate on mitigations, and benefit from community scrutiny. The case becomes less certain as we move towards larger, more complex and highly-capable general purpose AI.¹²¹ This is due to several challenges: (i) given our current lack of understanding of how advanced AI systems work internally, it may be difficult to identify the source of risk or failure; (ii) some harms, like bias or discrimination, may be baked into the training data and hard to eliminate entirely; (iii) addressing misuse may require broader social change beyond technical fixes; and (iv) the structure of AI systems introduces new failure modes specific to AI that are resistant to quick fixes. For example the stochastic nature of large language models may make it difficult to eliminate all negative outputs, and the inability to distinguish prompt injections from “regular” inputs may make it difficult to defend against such attacks. Whether these factors mean the dial actually tips in favor of offence is yet unknown.

What is clear, however, is that the benefits of openness for AI safety and security are dependent on being as open as possible. Unlike OSS which only concerns the release of code, AI openness is more nuanced. Model weights, code, documentation, and data can all be shared or withheld independently of each other. However, while a semi-open model can lay bare vulnerabilities for malicious actors (e.g. by sharing only model weights), it does not provide raw materials needed to fuel a crowdsourced safety benefit. It is only by passing on a model (its code and weights) along with all life-cycle documentation that downstream developers can meaningfully evaluate the model’s suitability for a particular task and test accordingly. As Alex Engler writes, there is “simply too much at stake for downstream developers to use AI systems they do not fully understand”.¹²²

In other words, AI openness poses risks that can often be reduced by being more open about more things. For example, better documentation (e.g. technical reports, model cards, and data cards with information model characteristics, training, and evaluation processes) and standardised documentation practices across platforms can enhance transparency and facilitate easier understanding and responsible use of AI artefacts. Open audits allow for independent verification of model safety and performance claims, fostering trust and accountability in the AI ecosystem. The availability of open datasets enables researchers and developers to train and test models on well understood, ethically sourced data. Open benchmarks provide standardised ways to evaluate and compare different models, promoting fair competition and progress tracking.

So what does this all mean for the question at hand: Is AI openness strategy for the UK compatible with the UK’s existing safety commitments? The high-level answer is, yes. Open-source has been a net benefit for technological safety and security for decades, and we should expect the same benefits to translate to AI where high-standards of openness are maintained.

Where concerns persist, they sit at the cutting edge of frontier AI development where the extent of potential future harms are unknown and increasing model size and complexity may require more difficult interventions to mend vulnerabilities and ensure performance safety. Especially as we watch the gap between proprietary frontier AI and open-source AI narrow, a question

120 Shevlane, T. and Dafoe, A. (2020). The Offense-Defense Balance of Scientific Knowledge: Does Publishing AI Research Reduce Misuse? <https://arxiv.org/abs/2001.00463>

121 Seger, E. et al. (2023). Open-Sourcing Highly Capable Foundation Models: An evaluation of risks, benefits, and alternative methods for pursuing open-source objectives. Retrieved May 20, 2025, from <https://www.governance.ai/research-paper/open-sourcing-highly-capablefoundation-models>

122 Engler, Al. (November 10, 2022). To Regulate General Purpose AI, Make the Model Move. Tech Policy Press. <https://techpolicy.press/to-regulate-general-purpose-ai-make-the-model-move/>

remains around if and when restrictions should be placed on model-sharing for models over a certain size or capability.

The concern about frontier AI safety is legitimate, however we recommend caution in laying restrictions on open model sharing without a clear view of the risks pathways being guarded against. A thorough analysis of the potential negative downstream consequences is also needed to avoid hampering safety and societal benefits of AI openness.

Finally, treading carefully with restrictions on model-sharing does not mean letting go of AI safety commitments in the least. Controlling model access is but one mechanism for mitigating risk of misuse and we caution against getting distracted by over indexing on its utility. Even where model-sharing restrictions are in place, they are not foolproof; model evaluation techniques are imperfect and the occasional model leak should be expected. More so, as the gap between open and closed model development continues to close, risk mitigation measures beyond model release will become increasingly important.¹²³ It is therefore essential that the UK’s AI safety research and intervention activities attend seriously to risk mitigation measures that can be implemented throughout the AI lifecycle by different actors in order to prevent, detect, and respond to risks. Table 1, adapted from Demos’s previous paper, *Open Horizons*, provides an overview of possible lifecycle intervention points.¹²⁴

TABLE 1
KEY PLAYERS AND RISK MITIGATION THROUGHOUT THE AI VALUE CHAIN¹²⁵

PLAYERS	RISK MITIGATION
(i) Preventing Risk	
Model Providers	<ul style="list-style-type: none">Develop and implement durable model-level interventions (see ‘technical solutions’ above for examples).
Model Providers & Model Adapters	<ul style="list-style-type: none">Responsibly source and filter training data to reduce bias and remove harmful content.Conduct internal safety and misuse evaluations to inform model release decisions.Provide clear user guidance documentation.

123 Currently leading open-source models are still developed in-house by companies like Meta and Deepseek with clear release decision-points that can be the target of regulation. But fully open and distributed development processes like the BigScience Initiative that yielded the 176 billion parameter model BLOOM are not so clearly contained. The development of BLOOM involved over 1000 researchers from over 250 institutions across more than 70 countries. (Introducing The World’s Largest Open Multilingual Language Model: BLOOM. Accessed 23 May, 2025. <https://bigscience.huggingface.co/blog/bloom>)

124 Seger & O’Dell (2024). Open Horizons: Exploring Nuanced Technical and Policy Approaches to Openness in AI. https://demos.co.uk/wp-content/uploads/2024/08/Mozilla-Report_2024.pdf

125 Also see Partnership on AI’s work on this subject. Srikuman, M., Chang, J. & Chmielinski, K. (2024). Risk Mitigation Strategies for the Open Foundation Model Value Chain. <https://partnershiponai.org/resource/risk-mitigation-strategies-for-the-open-foundation-model-value-chain/>

Model Hosting Services	<ul style="list-style-type: none"> Establish consistent structures for content moderation on their platforms. Assess whether hosted models meet the platform's standards for responsible model development and deployment including, for example, evidence of adequate safety testing and risk analysis, clear and complete documentation and model use guidance. More closely monitor and focus evaluations on the most frequently downloaded models. While the open-source ecosystem is vast, 70% of hosted models have 0 downloads while 1% account for 99% of downloads thus narrowing down "widely used models" to a more manageable range.¹²⁶
(ii) Detecting Risk	
Model Providers & Model Hosting Services	<ul style="list-style-type: none"> Implement and support incident reporting channels to allow external stakeholders to report safety concerns, vulnerabilities and AI incidents. Establish external audit and evaluation programs to facilitate access for auditors to critical components for detecting risk.
(iii) Responding to Risk	
Model Providers, Model Hosting Services, & App Providers	<ul style="list-style-type: none"> Establish decommissioning and incident response policies outlining the conditions under which a model is recalled and no longer hosted, or changes to licence are implemented to limit or prohibit certain uses.

A NOTE ON CYBERSECURITY AND DATA PROTECTION

AI safety is not purely about model misuse and dangerous capability dissemination. There are other often overlooked concerns about open model-sharing pertaining to cybersecurity and data protection.¹²⁷ These risks include training-time attacks like code injection and data poisoning where an attacker inserts unwanted code or data into a model to alter its behaviour. Post model deployment, attackers can try to get models to disclose sensitive data such as previous users' prompts or any personal data that was included in training data through clever prompting.¹²⁸ This is called "prompt extraction".¹²⁹

¹²⁶ Osborne, C., Ding, J., & Kirk, H. R. (2024). The AI Community Building the Future? A Quantitative Analysis of Development Activity on Hugging Face Hub. Retrieved July 22, 2024, from <https://arxiv.org/abs/2405.13058>

¹²⁷ A recent McKinsey report found that 62% of technologists and AI developers were worried about cybersecurity regarding open-source AI. Bisht et al. (2025). Open source technology in the age of AI. McKinsey. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/open-source-technology-in-the-age-of-ai>

¹²⁸ ICO (2025). 'How should we assess security and data minimisation in AI?'. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/how-should-we-assess-security-and-data-minimisation-in-ai/>

¹²⁹ UK Department for Science, Innovation & Technology (2024). 'Cyber security risks to artificial intelligence'. <https://www.gov.uk/government/publications/research-on-the-cyber-security-of-ai/cyber-security-risks-to-artificial-intelligence>

These risks exist for open and closed models, but the concern is that AI openness may heighten these risks by granting would-be attackers access to open-source models' key components – such as code, weights, or training data – that can use this access to identify vulnerabilities more effectively. The decentralised nature of some open-source AI development efforts may pose a challenge for vulnerability detection and accountability in the event of an incident.¹³⁰

On the flipside, some argue that open-source AI offers better security, transparency, and trust than closed-source alternatives.¹³¹ Like open-source software, AI openness models benefit from public scrutiny, enabling crowdsourced vulnerability detection and more robust security testing.¹³² Making model components accessible allows developers to identify and fix issues more efficiently. Additionally, open models are often easier to run locally, giving users greater control over data flows and reducing reliance on third-party cloud services—thereby minimising security risks.¹³³

Businesses looking to integrate AI tools into their internal or customer facing operations understandably cite security and data protection as reasons both for and against choosing to adopt open-source tools.¹³⁴ Which option is preferable will likely depend on the expertise of those looking to adopt open-source AI solutions and their ability to implement and maintain them properly.¹³⁵

In the meantime, building on the experience of OSS, there are a variety of measures that can be implemented throughout the AI lifecycle to help mitigate the security and privacy concerns surrounding open-source AI. These include strengthening the security protocols and hardening the defences in place for software used around AI models, implementing controls over software supply chains, regularly evaluating models for security and privacy risks, and implementing KYC checks on model downloads. The ICO recommends that AI developers take “appropriate steps” to identify privacy risks during the development lifecycle and to implement data protection best practices to prevent the exposure of sensitive data.¹³⁶

130 Wong (2025). 'Mapping the Open-Source AI Debate: Cybersecurity Implications and Policy Priorities'. RStreet. <https://www.rstreet.org/research/mapping-the-open-source-ai-debate-cybersecurity-implications-and-policy-priorities/#the-ldquo-open-rdquo-approach-to-ai-development>

131 E.g. Cable & Black (2024). 'With Open Source Artificial Intelligence, Don't Forget the Lessons of Open Source Software'. US Cybersecurity & Infrastructure Security Agency (CISA). <https://www.cisa.gov/news-events/news/open-source-artificial-intelligence-dont-forget-lessons-open-source-software>

132 E.g. Richardson (2025). 'Why open source is critical to the future of AI'. Red Hat. <https://www.redhat.com/en/blog/why-open-source-critical-future-ai>

133 Saran (2025). 'AI models explained: The benefits of open source AI models'. Computer Weekly. <https://www.computerweekly.com/feature/AI-models-explained-The-benefits-of-open-source-AI-models>

134 Bisht, A, et al. (2025). Open source technology in the age of AI. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/open-source-technology-in-the-age-of-ai>

135 Nagle, F., Wheeler, D. A., Lifshitz-Assaf, H., Ham, H., & Hoffman, J. L. (2020). Report on the 2020 FOSS Contributor Survey. The Linux Foundation & The Laboratory for Innovation Science at Harvard. <https://www.linuxfoundation.org/resources/publications/foss-contributor-2020>

136 ICO (2025). 'How should we assess security and data minimisation in AI?'. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/how-should-we-assess-security-and-data-minimisation-in-ai/>

SECTION 5

A PURPOSE BUILT AI OPENNESS STRATEGY FOR THE UK

So far, this paper has set out what an AI openness strategy is, outlined the national benefits of leaning into openness, examined examples from other countries, and addressed considerations regarding AI safety. This section dives into what a bespoke AI openness strategy for the UK might look like.

As a first step, an AI openness strategy for the UK requires alignment with an overarching goal-based agenda to guide decision-making and action. What are the high level goals which should be pursued? How might different kinds of activities promoting greater openness in AI development and deployment strategically work to support those goals? Without a unifying goal-based framework, all you have is an unstructured pile of recommendations.

This paper utilises the AI Opportunities Action Plan as the framework for an AI Openness strategy for the UK. Our reason for utilising the Action Plan are threefold:

- The Action Plan already alludes to openness but there is opportunity to more strongly embed it.
- The government has already bought into the Action Plan, accepting all proposals.
- The Action Plan's high-level goals are well suited to enhancement by AI openness.

In what follows we describe each of the Action Plan's high level goals in more detail, and we discuss opportunities for different aspects of openness to help achieve those goals. The detailed policy recommendations discussed throughout this section are collated in Appendix 2. Section 6 distills five high-level recommendations for government.

OUR ANALYSIS YIELDS THE FOLLOWING HIGH-LEVEL INSIGHT

Overall the Action Plan goals are primed to benefit from integrating measures to promote greater AI openness. This is demonstrated by the high number of existing Action Plan proposals on which our recommendations directly build (See Appendix 2 for cross-reference).

However there are two critical points that the Action Plan overlooks where AI openness will be of great benefit. These are:

- **Public Orientation** - Building AI systems and infrastructure for public benefit and to be in alignment with public needs and values.
- **Sovereignty through collaboration** - Recognising the importance of knowledge sharing and international collaboration as a route to greater self-sufficiency and distributed influence over AI futures.

The discussion that follows builds on insights from a multistakeholder expert workshop convened by Demos on May 12, 2025.

5.1 LAYING THE FOUNDATIONS FOR AI

The first high-level goal of the AI Action Plan identifies fundamental requirements for the UK to build a thriving domestic AI industry: (1) access to world-class computing, (2) access to high-quality data sets and data infrastructure, (3) a robust pipeline of talent, and (4) pioneering legislation underpinning safe and trustworthy AI development. Its associated proposals cover investments in compute infrastructure, making more data available for AI developers, expanding education to cover AI skills, attracting international talent, creating legislation to regulate AI, and policies to grow an AI safety and assurance industry in the UK.

Each of these aims can be enhanced by AI openness. We describe how below. The key takeaway is that **alongside compute, data, and human capital, a thriving open-source ecosystem should itself be treated as a core component of national AI infrastructure.**

5.1.1 Openness and compute access

The Action Plan rightly emphasises the importance of planning the UK's compute infrastructure investment. There are strategic ways openness can bolster the UK's compute capacity and help deliver public AI benefits.

To begin, the design of any new data centres should prioritise open standards for interoperability and sustainability such as by aligning with OCP Ready™ requirements.¹³⁷ This will help future-proof UK infrastructure and facilitate compatibility with global systems (See section 2.2). A portion of public compute investment could also go toward open-source

¹³⁷ OCP Ready™ Data Center Recognition Program. Accessed 27 May, 2025. <https://www.opencompute.org/projects/ocp-readytm-data-center-recognition-program>

hardware to reduce dependence on increasingly specialised proprietary hardware.¹³⁸

Second, given the scale of global demand and the resource intensiveness of such investments, the UK cannot realistically meet all of its computing needs independently. We recommend heavily weighting Proposal 6, which gestures at building international compute collaboration as a strategic imperative. The UK should actively pursue partnerships that support shared infrastructure development and access, such as through the EuroHPC Joint Undertaking¹³⁹ or EuroStack.¹⁴⁰

Finally, as the UK builds compute resources or secures access through international collaboration, the UK should look at prioritising and subsidising model access for open-source AI developers working on public-interest AI projects. Public interest AI will need to be defined (see section 5.1.4) but generally speaking the idea is to help lower barriers for industry growth in a way that will also catalyse AI innovation toward addressing a set of critical public interest challenges. The same can be done through the provision of data resources which we attend to in the next section.

5.1.2 Openness and data access

The Action Plan identifies the National Data Library (NDL) as a key initiative to make high-quality datasets more available for AI development. We strongly support this ambition and recommend the NDL adopt a strategic openness framework from the outset, that helps lower barriers to innovation for UK-based developers, startups, and researchers.

Wherever possible, public datasets should be released under open-access licences to make data usable and accessible across the economy. The government might consider, for example, mandating that data be open by default, with data producers publishing reasons for not opening data.¹⁴¹

Of course, full openness may not be appropriate for all datasets, particularly for sensitive data such as from the NHS, or where there are copyright and remuneration concerns. However, there is significant room between fully open and fully closed data-sharing. Some datasets could be made “locally open” with free access granted to domestic developers. There are also tools such as controlled data-sharing frameworks, synthetic data generation, and structured transparency methods¹⁴² that can be used to help preserve security while promoting transparency and innovation.

Finally, international data collaboration is also a powerful lever for driving AI innovation and enhancing the UK’s global standing. This is because both general-purpose and specialised AI systems increasingly require access to diverse, large-scale datasets that no single country can generate alone, making collaboration essential to building effective, representative, and contextually relevant AI solutions.

One timely opportunity that would need urgent attention is to **reset the UK’s relationship with the EU around the Data Union Strategy** which is set for review in July.¹⁴³ We recommend the UK propose a flagship open data-sharing initiative with European partners, centering on an

138 For example, Spain is investing in open-source chip design; Barcelona Supercomputing Center (2023). ‘BSC presents Sargantana, the new generation of the first open-source chips designed in Spain’. <https://www.bsc.es/news/bsc-news/bsc-presents-sargantana-the-new-generation-the-first-open-source-chips-designed-spain>

139 The European High Performance Computing Joint Undertaking (EuroHPC JU). Accessed 25 May, 2025. https://eurohpc-ju.europa.eu/index_en

140 EuroStack (2025). Deploying the EuroStack: What’s Needed Now. Accessed 25 May, 2025 <https://euro-stack.eu/the-white-paper/>

141 See principle 2 of the ODI Policy Manifesto (2024). Retrieved May 22, 2025, from https://theodi.cdn.ngo/media/documents/ODI_Policy_Manifesto.pdf

142 OpenMined (2021). Structured Transparency: Ensuring Input and Output Privacy. <https://blog.openmined.org/structured-transparency-input-output-privacy/>

143 A European Strategy for Data. Accessed 27 May, 2025. <https://digital-strategy.ec.europa.eu/en/policies/strategy-data>; European Commission (May 23, 2025). Commission seeks views on the use of data to develop Artificial Intelligence. Accessed 27 May, 2025. <https://digital-strategy.ec.europa.eu/en/news/commission-seeks-views-use-data-develop-artificial-intelligence>

initial public interest project of mutual public benefit. For example, the UK could offer access to the UK Biobank dataset for the purpose of a collaborative project joining expertise and resources to tackle a pressing medical challenge. This would signal renewed commitment to knowledge exchange, collaborative AI development, shared digital sovereignty, and building AI in the public interest.

More generally, **the UK should consider joining Current AI**, an international funding body and convener announced at the Paris AI Summit working to catalyse public investment in public AI.¹⁴⁴ It works to coordinate action across governments, philanthropic funders, and research communities with a particular focus on shifting norms and building infrastructure to facilitate data-sharing for public interest AI projects. Partnering with Current AI would be a potentially high-impact and relatively low-cost step in showing the government's commitment to developing AI for public benefit and aligning with international efforts to collaboratively build an AI openness counterpoint to proprietary AI infrastructure.

5.1.3 Openness and talent pipelines

The UK is exceptionally strong in AI talent, anchored by world-leading universities such as Oxford, Cambridge, Durham, and Edinburgh, and bolstered by prominent research hubs like Google DeepMind and the Alan Turing Institute. This has created a positive feedback loop where top-tier talent attracts more of the same - the best minds seek out environments where they can collaborate with peers.

A thriving open-source ecosystem alone is unlikely to be a primary driver for attracting more AI talent to the UK; factors like availability of prestigious positions and quality of life may play a more significant role. Although it is notable that DeepSeek's founder does view their open-source approach as a core to attracting talent.¹⁴⁵ Perhaps a greater concern, however, is that overly restrictive regulation of open-source could have a negative effect on the UK's leading AI talent pool. Regulation that might limit the ability of UK researchers to engage in international open-source innovation or to commercialise and re-contribute open-source solutions could deter talent and inhibit the UK's competitiveness.

5.1.4 Openness and regulation for safe and trusted AI development and deployment

As a final topic under the Action Plan's high-level goal for "laying foundations for AI" we turn to regulation for safe and trusted AI development and deployment. There are very many recommendations that could feasibly fall in this category, but here we will focus on the forthcoming UK AI Bill and on methods of decision-making in AI governance.

With respect to the UK AI Bill and AISI

Ministers have indicated that the UK AI Bill will focus on frontier AI, tabling legislation to mitigate risks from the most highly capable AI models developed by the largest AI companies. Given the narrowing gap between frontier AI capability and leading open-source AI innovation there is some concern that the forthcoming Bill could impact open-source AI development in the UK. Some discussion on model-sharing restriction is to be expected given the safety concerns discussed in section 4, but we caution against over-indexing on model-sharing restrictions as a risk mitigation method.

As discussed in section 4, while model-sharing is one mechanism for attempting to mitigate harm from misuse and dissemination of dangerous capabilities, it is fallible. Capability evaluations are imperfect, and model leaks should be expected. More so, UK legislation posing strict requirement of model-sharing is unlikely to have a strong effect on influencing the

¹⁴⁴ Current AI. Accessed 24 May, 2025. <https://www.currentai.org/>

¹⁴⁵ Ottinger, L. & Schneider, J. (February 1, 2025). DeepSeek: What it means and what happens next. ChinaTalk. <https://www.chinatalk.com/media/p/deepseek-what-it-means-and-what-happens>

direction of global AI legislation given the economic advantages of open-sourcing, but it very well could have a negative impact on the UK's AI industry, disallowing researchers to work with cutting-edge models or downstream developers to innovate on the latest capabilities.

This is not to say that the AI Bill should throw caution to the wind, but we wish to emphasise the importance of attending to risk mitigation strategies throughout the AI stack, noting opportunities for intervention at the levels of preventing risk, detecting risk and responding to risk with responsibilities laid out for AI actors including AI developers, AI providers, downstream model adapters, and model-hosting platforms (see section 4). We recommend looking to Demos¹⁴⁶ and Partnership on AI's¹⁴⁷ reports on this subject for specific intervention recommendations.

Creating and implementing regulation for AI along the spectrum from fully-open to proprietary models is a complex task. It must be an objective of regulation not to accidentally disincentivise open development. Toward this end, we recommend government look to see where openness can be beneficially integrated into regulation to ensure the UK takes a balanced approach. For example, with respect to the UK AI Bill and model-sharing, the UK could follow a similar path to the EU AI Act in introducing exemptions from regulatory obligations for models that meet a clear standard of openness and transparency (excluding when those models are put to use in high-risk context).¹⁴⁸ This would both ensure smaller open-source developers and distributed communities are not overburdened by unnecessary compliance requirements and would help incentivise proprietary developers to embrace greater openness. On the flip side, the UK AI Bill should also take inspiration from the EU AI Act in mandating transparency requirements for proprietary models in order to provide developers with the necessary information to responsibly integrate and monitor potentially high-risk systems.¹⁴⁹ See Section 6 and Appendix 2 for more details.

The UK AI Bill will also likely have implications for the future of the AI Security Institute (AISi). We recommend that AISi is charged with investigating and laying out guidelines for risk mitigation throughout the AI value chain; more downstream aspects of the work pertaining to societal harms may lie with the Systemic Risks program. More so, AISi could have a profound impact on the direction of AI safety globally through its internal and world-leading AI safety research function - by developing AI benchmarks, evaluations, guardrail and mitigation methods and open-sourcing safety tooling for others to utilise. The tooling could be made accessible either directly by AISi in a similar manner to Singapore's AI Verify program¹⁵⁰ or by contributing to ROOST,¹⁵¹ a nonprofit announced at the Paris AI Summit that builds open-source AI safety tools as a contribution to public AI infrastructure. In so doing, AISi could wield significant soft power by implicitly setting a leading global standard for AI safety testing. More so, AISi would be contributing a key resource to drive the development of the UK's own AI assurance sector, primed by the UK's abundant AI talent and safety focus, to serve AI developers globally.

There is some concern that sharing evaluations could allow nefarious actors to figure out how to game the tests. It may therefore be prudent to instate a short delay in open-sourcing the latest version, while the latest version is shared with trusted AI assurance providers within the UK and collaborators internationally. However, by open-sourcing the safety tooling, a wider community

146 Seger, E. & O'Dell, B. (2024). Open Horizons: Exploring Nuanced Technical and Policy Approaches to Openness in AI. https://demos.co.uk/wp-content/uploads/2024/08/Mozilla-Report_2024.pdf

147 Srikuman, M., Chang, J. & Chmielinski, K. (2024). Risk Mitigation Strategies for the Open Foundation Model Value Chain. <https://partnershiponai.org/resource/risk-mitigation-strategies-for-the-open-foundation-model-value-chain/>

148 Article 2: Scope. EU Artificial Intelligence Act. <https://artificialintelligenceact.eu/article/2/>

149 Article 13: Transparency and Provision of Information to Deployers. EU Artificial Intelligence Act. <https://artificialintelligenceact.eu/article/13/>

150 AI Verify Foundation (2025). 'About AI Verify Foundation'. <https://aiverifyfoundation.sg/ai-verify-foundation/>; Infocomm Media Development Authority (2023). 'Singapore launches AI Verify Foundation to shape the future of international AI standards through collaboration'. <https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/press-releases/2023/singapore-launches-ai-verify-foundation>

151 Surman, M. & Bdeir, A. (February 10, 2025). Open source AI Safety for Everyone. <https://blog.mozilla.org/en/mozilla/ai/roost-launch-ai-safety-tools-nonprofit/>

of researchers can help improve and iterate on AISI's version yielding a greater diversity of high-quality AI safety tools globally. This tradeoff will need to be carefully evaluated in formulating release plans.

With respect to open decision-making in AI governance

AI regulation has implications for AI openness, but so too does AI openness have implications for how AI is governed. AI openness extends beyond model components, artifacts, and safety tooling. AI openness governance mechanisms that are transparent, participatory, and inclusive can help ground AI systems in public values and build societal trust by enabling people impacted by AI to feed into decision-making about the transformative technology. It's about making AI something that is done with and for people, not something that is done to them.

To achieve this, the UK should consider embedding deliberative and multi-stakeholder processes into aspects of its AI governance infrastructure. These mechanisms can guide high-level strategic decisions about AI policy: for instance, by employing citizens' assemblies, participatory processes facilitated by civic tech (e.g. platforms such as Pol.is and Remesh) to engage diverse **multi stakeholder** deliberation. This is not to say that the day-to-day work of AI engineers at leading firms is subject to democratic vote. But there is certainly room at higher levels of abstraction for determining what values AI should be aligned with (e.g. via alignment assemblies), where to draw boundaries around sensitive or unacceptable use cases, and how to distribute AI resources toward public interest applications. This could involve parameterising a concept of 'public interest AI' as a basis for prioritising resource access and regulatory incentives to drive AI development for public benefit.

5.2 DRIVING CROSS-ECONOMY AI ADOPTION (IN THE PUBLIC INTEREST)

The second high-level action goal calls for the UK to adopt AI at scale across the public and private sector - "delivering services, transforming citizens' experiences, and improving productivity" by embracing AI. The goal sets out proposals for the government to identify areas where AI can be used, to set out a framework for AI procurement, to build AI prototyping and scaling capabilities, and to share knowledge with the general public about the results.

We are concerned, however, with how this high-level goal is currently articulated. There is nothing inherently wrong with driving forward AI adoption, but the current plan can read like a hammer seeking a nail - a "solutions first" approach angled at driving AI innovation and adoption for the sake of productivity gains and AI industry growth instead of a "needs first" approach rooted in serving public interest. On this topic, Demos has written about the need to embed AI rollout in a broader public service reform agenda;¹⁵² public services are struggling to deliver for citizens but, this is not solely, or even primarily, from a lack of adequate technology, but rather from under-resourced services hamstrung by a top-down, command-and-control model in which providers struggle to respond to the unique complexity of local contexts and needs. AI can help improve public services, but only if we look to the fundamental challenges facing public services first.

All that said, AI does provide a significant opportunity to improve service quality and delivery. Opportunities include freeing up service provider time from basic administration to attend more complex tasks,¹⁵³ improving timely communication with citizens,¹⁵⁴ and improving service quality with tools that enhance human performance on tasks ranging from medical diagnostics to flood prediction.¹⁵⁵

¹⁵² Knight, S. & Seger, E. (2024). Tech that liberates: A new vision for embedding AI in public services reform. Demos. <https://demos.co.uk/research/tech-that-liberates-a-new-vision-for-embedding-ai-in-public-service-reform/>

¹⁵³ Iosad, A., Railton, D. & Westgarth, T. (2024). Governing in the Age of AI: A New Model to Transform the State. <https://institute.global/insights/politics-and-governance/governing-in-the-age-of-ai-a-new-model-to-transform-the-state>

¹⁵⁴ Deloitte Digital. (January 2025). Artificial Intelligence (AI) and the Future of Public Engagement. <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/public-sector/amc-ai-the-future-of-public-engagement.pdf>

¹⁵⁵ Google. (September 2024). Unlocking the UK's AI Potential. Accessed 27 May, 2025. https://static.googleusercontent.com/media/publicpolicy.google/en/resources/uk_ai_opportunity_agenda_en.pdf

To the extent that AI can help improve public service, openness is key both to lowering barriers to innovation and adoption and helping to ensure that AI development and adoption is in the public interest.

Section 2.1 and 2.2 speak in detail to the benefits of AI openness to driving innovation and supporting industry. In this case we are particularly interested in laying the foundations for UK research and industry at the AI application layer, building solutions for public procurement. We outline various recommendations to aid in scanning for, piloting, and scaling effective AI tools for public sector adoption in Appendix 2. A couple to highlight include creating *open-source* and data-rich AI experimentation environments so that developers working both with and without the public sector can benefit from the resource to innovate for public sector application (this builds on Action Plan Proposal 37). We also recommend prioritising the development of the AI Knowledge Hub for publishing best-practice guidance, results, case-studies and open-source solutions (referring to Action Plan Proposal 45).

Turning to AI adoption, procuring open-source solutions significantly lowers costs of adoption, prevents vendor lock-in, makes maintenance easier, and lowers the risk of creating government contractor monopolies. In a rapid scope-prototype-scale environment, open models also allow for faster iteration and smoother integration with existing systems (see section 2.3). The Indian government provides an excellent case study for the benefits of procuring open-source. They have found prioritising open-source in the public sector to be helpful for lowering costs, increasing transparency, and ensuring that the government's unique needs are met (see Appendix 1). Building on Action Plan Proposals 34 and 43 we recommend government procure open-source solutions wherever possible, and that the scalable tech stack called for in Proposal 41 is subject to the infrastructure interoperability and open-source standards suggested in Proposal 42.

Finally, open-source AI procurement alongside greater knowledge-sharing among AI adopters and greater transparency about procurement decisions, can help ensure that AI adoption is oriented towards public needs. Adopting open-source tools allows more flexibility in tailoring applications to specific contexts and needs. Greater transparency around the evaluations and performance benchmark metrics can also help better assess AI tools prior to public sector deployment. This transparency is key to public sector accountability, enabling public oversight of AI use in high-risk contexts (e.g. health and housing) and underpinning greater public trust in AI-augmented service provision. We also recommend attending to the suggestions we presented in section 5.1.4 regarding AI openness governance, coordinating public participation in decision-making about priority public service AI applications as well as unacceptable use cases.

5.3 SECURING A FUTURE FOR HOMEGROWN AI (AI SOVEREIGNTY)

This third high-level Action Plan goal involves building up “true national champions at critical layers of the AI stack”. It comes with one core proposal: to “create a new unit, UK Sovereign AI, with the power to partner with the private sector to deliver the clear mandate of maximising the UK's stake in frontier AI”.

In the context of the AI Opportunities Action Plan, AI sovereignty could be conceptualised as aiming to ensure that the UK has reliable **access** to high-quality AI models that are not subject to external control, that the **value** of AI-led economic transformation is captured in the UK, and that the UK has **influence** over the global development and deployment of frontier AI.

Openness certainly helps with access, influence, and accruing value (we will expand below), but not necessarily through the mechanisms currently sought. The Action Plan is clear on its goal for seeing the UK achieve significant frontier AI ambitions, building on the 2021 AI strategy for

establishing the UK as an independently competitive frontier “AI superpower”¹⁵⁶ - repeated again in the government’s response to the Action Plan.¹⁵⁷

However, as discussed in Section 3, it is not clear that the UK should be holding on so strongly to its current frontier ambitions. The level of investment needed to play at proprietary frontier AI development on the same level with the US and China are far out of our range (consider the \$100 billion the Trump administration recently secured for AI infrastructure and development in deals with the UAE).¹⁵⁸ More so, the kinds of AI solutions that will help improve public service delivery and yield productivity gains for businesses are very often not developed on the back of massive general purpose models, but from narrower, specialised application development. Correspondingly, much of the opportunity for near and medium term economic growth from AI also sits behind the frontier with companies laser-focused on building AI for adoption - on filling gaps with AI solutions for real-world benefit.

So in what ways does openness help in pursuit of these sovereignty goals if not through helping establish the UK as a leading frontier AI powerhouse?

Value

Ultimately value will accrue to AI industries that succeed at putting AI tools in consumer hands. Given the UK’s incredible pool of AI talent, high quality data resources, and domain specific scientific expertise situated around world leading Universities, the UK is in a prime position to develop consumer-focused, solutions-oriented AI.

Laying foundations for a thriving AI openness ecosystem in the UK will help the UK capture value by lowering barriers to entry for new business (see section 2.2) and by acting as an economic multiplier across industries (see section 2.4). As articulated in section 5.1, nurturing homegrown AI industry in the UK would be facilitated by open or locally open compute and data resources, and, as we will speak to next, by ensuring reliable access to powerful open-source models on which downstream developers can innovate and build new AI solutions.

Access

There are two ways to guarantee access to highly-capable, safe, and aligned AI models that cannot be taxed or turned off by third parties: (1) build your own, or (2) collaborate.

The idea of building a sovereign AI model for the UK has been batted around quite a bit, featuring proposals for initiatives like “BritGPT”¹⁵⁹. Other countries have proceeded with building national models, notably India, Spain, and Thailand (see Appendix 1). However, training a competitive, general-purpose model domestically may not be the most effective path for the UK. Doing so would require significant investment and the model would quickly come to lag significantly behind frontier lab AI capabilities without continued investment and development efforts.

Instead we recommend the UK commit resources and researcher hours to open collaboration on AI development. The idea is to choose our partners from other liberal democracies similarly committed to AI safety, and to pool resources to build a thriving open-source counterpoint to proprietary big tech and cloud computing services that guarantees access for all. The more partners contributing to the collaboration, the stronger the open-source counterpoints. By contributing data resources to training (as discussed in 6.1.2) the UK can also build off the co-

156 (September 2021). National AI Strategy. Accessed May 27, 2025. https://assets.publishing.service.gov.uk/media/614db4d1e90e077a2cbdf3c4/National_AI_Strategy_-_PDF_version.pdf

157 (January 13, 2025). AI Opportunities Action Plan: government response. Accessed May 27, 2025. <https://www.gov.uk/government/publications/ai-opportunities-action-plan-government-response/ai-opportunities-action-plan-government-response>

158 (May 13, 2025). Fact Sheet: President Donald J. Trump Secures Historic \$600 Billion Investment Commitment in Saudi Arabia. <https://www.whitehouse.gov/fact-sheets/2025/05/fact-sheet-president-donald-j-trump-secures-historic-600-billion-investment-commitment-in-saudi-arabia/>

159 Belfield, H. (2023). Great British Cloud and BritGPT: the UK’s AI Industrial Strategy Must Play to Our Strengths. Labour for Long Term. <https://www.labourlongterm.org/briefings/great-british-cloud-and-britgpt-the-uks-ai-industrial-strategy-must-play-to-our-strengths>

developed open-source models to better serve the UK's unique needs - for example, preserving Cornish, Gaelic, Irish, and Scots linguistic culture throughout the AI transition.

The UK would not need to start from scratch in this endeavour. As discussed in 3.5 there are existing collaborative and public AI initiatives the UK could join such as Current AI. The UK could also enter into specific project-oriented collaborations such as that recommended in section 5.1.2 - resetting the UK's relationship with the EU around the Data Union Strategy which is set for review in July with flagship public interest AI development projects of mutual public benefit.¹⁶⁰

Of course securing AI model access is not just about securing access to the model itself. It all comes back around to securing access to the relevant compute, data, and talent resources needed to underpin development and continued use. On these topics we refer the reader back to Sections 5.1.1, 5.1.2, and 5.1.3 respectively.

Influence

When it comes to influencing the future frontier AI development, leaning into AI openness is the UK's ticket. Throughout this report we have emphasised the strength of the UK's AI expertise and data resource, and in particular the unique concentration of talent within the UK Government's AI Security Institute. The UK can influence the future direction of AI development by sharing this expertise widely and remaining a powerful contributor to the global AI research environment. And as discussed in section 5.1.4, the UK has a particular opportunity for influence and building global dependence on UK services through its AI safety research specialty. AISI can play a central function in developing high-quality safety tooling - benchmarks, evaluations, risk mitigations, and guardrails - that it shares openly to implicitly set a global standard for AI safety. In turn, the open-tooling also underpins the growth of the UK's own AI assurance ecosystem - capitalising on AISI's reputation - to provide gold-standing AI testing services globally as AI applications enter the market at an ever more rapid pace.

¹⁶⁰ A European Strategy for Data. Accessed 27 May, 2025. <https://digital-strategy.ec.europa.eu/en/policies/strategy-data>; European Commission (May 23, 2025). Commission seeks views on the use of data to develop Artificial Intelligence. Accessed 27 May, 2025. <https://digital-strategy.ec.europa.eu/en/news/commission-seeks-views-use-data-develop-artificial-intelligence>

SECTION 7

HIGH-LEVEL RECOMMENDATIONS

This report has executed on three aims: to make the case for why the UK should adopt an AI openness strategy to help realise its AI ambitions (sections 2-3); to help square the UK's AI safety commitments with the prospect of an AI openness strategy (section 4); and to present an overview of what it would look like to integrate an AI openness agenda as part of the AI Opportunities Action Plan (section 5).

The core message is that alongside compute, data, and human capital, a thriving open-source ecosystem should itself be treated as a core component of robust national AI infrastructure and key to ensuring AI development and deployment delivers for UK citizens.

For those interested in specific policy opportunities for integration into the AI Action Plan delivery, please see Appendix 2 which collates recommendations from section 5 and cross-references with the Action Plan.

In this final section we distill the discussion from across the report into five high-level recommendations for government.

1. Commit to an AI openness strategy for the UK and look to deploy it through integrations with the existing AI Opportunities Action Plan agenda

The benefits of AI openness to a country looking to grow its AI industry, drive AI adoption, and secure reliable access to AI resources are numerous. Thriving open-source ecosystems are known to drive innovation (Section 2.1), support industry growth (section 2.2), support flexible adoption (section 2.3), multiply economic benefits (section 2.4), and via those mechanisms, underpin greater technological self-reliance (section 2.6).

Now with the market shifting towards smaller specialised AI models (section 3.1), combined with mounting urgency to reduce reliance on foreign proprietary AI supply chains (section 3.2), the time is ripe for the UK to lean more heavily into an AI openness strategy toward achieving the nation's AI goals. The UK is primed to extract maximum value from the open ecosystems that it helps build and maintain, capitalising on our deep bench of scientific and AI expertise rooted in our world-leading universities and AI institutions (Section 3.3).

The existing AI Opportunities Action Plan is structured well to accommodate integration of AI openness objectives with modification of standing proposals and minor additions to place greater emphasis on international collaboration and public benefit (Section 5 and Appendix 2.)

Committing to an AI openness agenda does not mean bailing on plans to attract investment from international investors and big tech firms. It means investing in a third way, and a strong open-source counterpoint to proprietary AI to underpin national AI sovereignty and ensure UK AI stays oriented with public interest and values.

2. Use a commitment to AI openness to demonstrate dedication to building and deploying AI in the public interest

There is a notable division between the government's stated AI goals and apparent public interest. Government announcements centre on frontier AI ambitions, frontier AI safety, driving adoption, and reaping productivity gains by "mainlining AI into the veins of this enterprising nation".¹⁶¹ Meanwhile top of mind for civil society and the public are concerns about impacts on employment, data rights, and deepfakes. This is a damaging mismatch in narrative and communicated priorities that is eroding trust.

A public commitment to AI openness, and particularly to investment in collaborative public AI initiatives like Current AI would be a strong play toward demonstrating dedication to building, deploying, and governing AI for public benefit and in alignment with public interests and values (Section 2.5). 'Public AI' describes open-source AI and AI infrastructure built and maintained like a public good - accessible to the public and accountable to the public for its function and impact.¹⁶²

A similar opportunity exists in embedding open and participatory democratic governance into decision-making about public interest issues—such as prioritising AI application development projects, defining the parameters of 'public interest AI,' or identifying values for AI alignment (5.1.4).

Both of these opportunities reinforce the Action Plan's stated high-level goal of driving AI adoption to improve public service function and delivery for citizens.

3. Pursue AI Sovereignty through outward collaboration and resource sharing, and promoting open development guidelines

AI sovereignty is a core goal of the AI Action Plan with sub-goals of ensuring UK access to high quality AI systems and ensuring value accrues to UK AI industry through the AI transition. But the UK is not strongly positioned to execute on these goals in isolation (section 5.3). It is not feasible for the UK to match the kind of investment dedicated in the US and China. Consider, for instance, the \$20 billion secured by the US from the UAE for building AI compute infrastructure.¹⁶³ For the UK to ensure it maintains a leading role in AI research and development and to ensure reliable access to high-quality AI tools, the UK needs to look outward toward international collaboration and resource-sharing. For example, the UK could pursue partnerships through EuroHPC¹⁶⁴ (section 5.1.1) or aim to reset a data-sharing agreement with the EU with the review of the EU's Data Union Strategy (section 5.1.2).¹⁶⁵

¹⁶¹ (January 13, 2025). Prime Minister sets out blueprint to turbocharge AI. Press Release. <https://www.gov.uk/government/news/prime-minister-sets-out-blueprint-to-turbocharge-ai>

¹⁶² Surman, M., Marda, N. and Sun, J. (September 30, 2024). Public AI. Mozilla. <https://www.mozillafoundation.org/en/research/library/public-ai/>

¹⁶³ (May 13, 2025). Fact Sheet: President Donald J. Trump Secures Historic \$600 Billion Investment Commitment in Saudi Arabia. <https://www.whitehouse.gov/fact-sheets/2025/05/fact-sheet-president-donald-j-trump-secures-historic-600-billion-investment-commitment-in-saudi-arabia/>

¹⁶⁴ The European High Performance Computing Joint Undertaking (EuroHPC JU). Accessed 25 May, 2025. https://eurohpc-ju.europa.eu/index_en

¹⁶⁵ A European Strategy for Data. Accessed 27 May, 2025. <https://digital-strategy.ec.europa.eu/en/policies/strategy-data>;

But outward collaboration is not just about bolstering the pool of resources available to the UK - intergovernmental collaboration agreements alone can be reneged - it is about contributing to a thriving global open-source counterpoint to proprietary big tech. Open models, open data, open safety tooling, and open hardware designs can be used, developed and improved by anyone. The larger collective action, the larger and more viable the counterpoint.

4. Influence a positive future for AI globally by openly sharing AISI's AI safety research insights and tools

The UK's foremost advantage in AI arguably lies in its remarkable AI talent base anchored by world-leading universities such as Oxford, Cambridge, Durham, and Edinburgh, and bolstered by prominent research hubs like Google DeepMind and the Alan Turing Institute. These foundations coupled with a strong political will gave rise to the AI Security Institute (AISI) - now one of the world's foremost AI safety research centres uniquely seated within government.

Through AISI, the UK government has the opportunity to influence the direction of development for frontier AI globally. This effect will be more profound the more widely AISI shares its research insights. In line with Action Plan proposals 23 and 29, we recommend continued support for the growth of AISI's safety research and model evaluation capacity, and wherever possible the tools should be open-sourced and made accessible, either directly from AISI in a similar manner to Singapore's AI Verify program or by contributing to ROOST,¹⁶⁶ a nonprofit that builds open-source AI safety tools as a contribution to public AI infrastructure (Section 5.1.4). While AISI has open-sourced some of its existing tools, such as Inspect,¹⁶⁷ a much stronger emphasis could be placed on doing so. In this way, AISI could wield significant soft power by implicitly setting a leading global standard for AI safety testing. AISI would also benefit from the return contribution of a global research community improving, iterating, and innovating on AISI's work.

More so, AISI would be contributing a key resource to drive the growth of the UK's own AI assurance sector. As AI tools increasingly transition from research projects to real world application, AI auditing and assurance services will grow into a high-demand industry of its own. With its abundant AI talent and safety research focus, the UK is primed to house this industry to serve AI developers globally. Knowledge-sharing from AISI could help catalyse this growth.

5. Use the forthcoming UK AI Bill as an opportunity to promote greater transparency and openness in AI development across the board

The UK's forthcoming AI Bill offers a prime opportunity to enshrine the UK's commitment to AI safety while fostering a more transparent, open, and innovation-friendly development environment. Creating and implementing regulation for AI along the spectrum from fully-open to proprietary models is a complex task. Therefore, it must be an objective of regulation not to accidentally disincentivise open development. Toward this end, we recommend the government look to see where openness can be beneficially integrated into regulation to ensure the UK takes a balanced approach.

As the UK shapes its legislative response to the rapid evolution of AI technologies, embedding transparency requirements for highly capable proprietary models (similar to Article 13 of the EU AI Act)¹⁶⁸ would ensure that downstream developers have the necessary information to responsibly integrate and monitor potentially high-risk systems. This would enable more

¹⁶⁶ Surman, M. & Bdeir, A. (February 10, 2025). Open source AI Safety for Everyone. <https://blog.mozilla.org/en/mozilla/ai/roost-launch-ai-safety-tools-nonprofit/>

¹⁶⁷ AI Security Institute (2025). 'Inspect: An open-source framework for large language model evaluations'. <https://inspect.aisi.org.uk/>

¹⁶⁸ Article 13: Transparency and Provision of Information to Deployers. EU Artificial Intelligence Act. <https://artificialintelligenceact.eu/article/13/>

effective oversight and safety testing across the AI value chain, while acknowledging that open-source models already meet many of these transparency benchmarks.

In tandem, the Bill should introduce exemptions from regulatory obligations for models that meet a clear standard of openness and transparency, excluding when those models are put to use in high-risk context (e.g. integrating into a medical diagnostic tool or hiring system).¹⁶⁹ This mirrors Article 2(12)¹⁷⁰ of the EU AI Act and would serve a dual purpose: reducing unnecessary compliance burdens on smaller open-source developers and distributed communities, and incentivising proprietary developers to embrace greater openness. This exemption also reinforces the discussion in section 4 and in our previous *Open Horizons* paper¹⁷¹ that emphasises the importance of employing risk mitigation measures throughout the AI lifecycle. The EU AI act places responsibility on the entity deploying an open-source model into a high-risk context to meet various requirements like implementing continuous risk management frameworks¹⁷² and maintain records of the system's functioning for monitoring and incident investigation.¹⁷³ The UK AI Bill could go a step further in exploring other AI lifecycle interventions as well, such as requirements on model hosting platforms to moderate the most widely used models,¹⁷⁴ and for hosting platforms and application providers to establish decommissioning and incident response policies outlining the conditions under which a model would be recalled.

169 Article 6: Classification Rules for High-Risk AI Systems. EU Artificial Intelligence Act. <https://artificialintelligenceact.eu/article/6/>

170 Article 2: Scope. EU Artificial Intelligence Act. <https://artificialintelligenceact.eu/article/2/>

171 Seger, E. & B. O'Dell (2024). Open Horizons: Exploring nuanced technical and policy approaches to openness in AI. <https://demos.co.uk/research/open-horizons-exploring-nuanced-technical-and-policy-approaches-to-openness-in-ai/>

172 Article 9: Risk Management System. EU Artificial Intelligence Act. <https://artificialintelligenceact.eu/article/9/>

173 Article 12: Record Keeping. EU Artificial Intelligence Act. <https://artificialintelligenceact.eu/article/12/>

174 While the open-source ecosystem is vast, 70% of hosted models have 0 downloads while 1% account for 99% of downloads thus narrowing down "widely downloaded models" to a more manageable range; Osborne, C., Ding, J., & Kirk, H. R. (2024). The AI Community Building the Future? A Quantitative Analysis of Development Activity on Hugging Face Hub. Retrieved July 22, 2024, from <https://arxiv.org/abs/2405.13058>

CONCLUSION

The United Kingdom stands at a crossroads in its AI journey. As geopolitical shifts, evolving market dynamics, and accelerating technological advancements redefine the global AI landscape, the imperative for a distinct, resilient, and forward-thinking strategy is clear. This report has made the case for embracing AI openness through open-source models, collaborative infrastructure, and transparent governance. It promises an 'open dividend' from the AI-powered growth agenda.

An AI openness strategy offers the UK a practical and scalable pathway to achieving the ambitions laid out in the AI Opportunities Action Plan: laying the foundations for domestic AI capability, driving cross-economy adoption in the public interest, and securing a sovereign position as an AI maker. Openness enables faster innovation, supports flexible and cost-effective adoption, reduces dependencies on foreign technology providers, and amplifies the UK's influence in shaping global AI norms.

Global trends, from the rise of efficient open-source models to shifting supply chains, point to a future where open ecosystems are increasingly competitive and necessary to establish a viable counterpoint to proprietary giants. Countries like India, China, France, and Singapore are already making a move towards AI openness, and the UK has the expertise, infrastructure, and policy momentum to lead in this space.

The choice is not between openness and safety, or between sovereignty and collaboration. A purpose-built UK AI openness strategy can deliver all these objectives.

APPENDIX 1

OPEN-SOURCE AI STRATEGY IMPLEMENTATION AROUND THE WORLD

This appendix provides an overview of AI openness strategies other countries have been employing and the results they are yielding, with lessons outlined for the UK.

The countries analysed have pursued AI openness for a variety of overlapping reasons. Taken together, these are to:

- Attract international AI startups, talent, and investment
- Develop AI models as a public good
- Enable the diffusion of innovations between companies and labs
- Ensure AI models are developed in local/national languages
- Limit market concentration and monopoly power
- Lower barriers to accessing data and compute for domestic AI developers
- Lower barriers to upholding safety and reliability standards
- Lower costs and barriers for downstream AI use
- Lower costs and investment barriers for domestic AI developers
- Promote interoperability and modularity standards
- Promote national self-reliance in AI development
- Promote principles of transparency and openness
- Reduce dependence on foreign AI developers
- Reduce secondary development costs
- Reduce the impact of international trade wars
- Influence domestic and international AI standards
- Support existing open-source AI developers and the open-source ecosystem
- Widen access to AI tools and lower costs for end-users

Countries have pursued these goals in different ways. In our analysis, we have identified three types of AI openness strategy:

- **Focused investment:** A focus on using government investment and market mechanisms to spur open-source AI development beyond government. Examples: European Union, France.
- **Government led initiatives:** A focus on government developing open-source AI models, open-source development platforms, safety testing platforms, or other resources in-house or via collaborations. Examples: India, Singapore.

- **Standard setting policies:** A focus on policies which establish open-source as the *de facto* or *de jure* AI development standard to be followed by developers in the private and non-profit sectors. Example: China.

While countries can and do adopt policies associated with more than one type, they tend to lean more towards one cluster of policies associated with a larger strategy.

CHINA

China's approach to open-source AI can be considered the prime example of a standard-setting strategy.

Strategy Overview:

The Chinese government has advocated for open-source AI development in its national plans from at least 2017,¹⁷⁵ and has reaffirmed this aim several times in the years since.¹⁷⁶ China's reasons for doing so include reducing AI development costs, mitigating the impact of trade wars, and limiting China's dependence on foreign AI developers.

China's government views encouraging open-source AI development as a means of setting *de facto* standards for AI.¹⁷⁷ To this end, China has set up a consortium of 23 leading private sector AI developers — called the National AI Team for Governance (NAITG)¹⁷⁸ — which follows the government's lead and plays a standard-setting role for the sector. The government has been able to use the NAITG as a central point of intervention to direct the country's AI sector and establish openness as a standard for Chinese AI.¹⁷⁹ The Chinese government and NAITG have invested in AI openness development platforms to encourage innovation, which offer open access to AI datasets, toolkits, libraries and other resources.¹⁸⁰

Alongside China's standard-setting efforts and investment in the open-source AI ecosystem, Chinese AI companies have also invested heavily in ensuring they have the compute available to compete internationally. Driven by export controls on advanced AI chips, such as Nvidia's H100 processors, China has been gearing up its own domestic AI hardware supply chain.¹⁸¹ In competition with Nvidia in chip design, China has Huawei; and while Nvidia's chips are fabricated by TSMC (Taiwan Semiconductor Manufacturing Corporation), China boasts SMIC (Semiconductor Manufacturing International Corporation).

175 Webster et al. (2017). 'Full Translation: China's 'New Generation Artificial Intelligence Development Plan'. Stanford University. <https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>

176 E.g., Xinhua News Agency (2021). 'Outline of the People's Republic of China 14th Five-Year Plan for National Economic and Social Development and Long-Range Objectives for 2035 中华人民共和国国民经济和社会发展第十四个五年规划和2035年远景目标纲要'. Centre for Security and Emerging Technology, Georgetown University. <https://cset.georgetown.edu/publication/china-14th-five-year-plan/>

177 See this publication by China Electronic Standardisation Institute's Artificial Intelligence Standardisation White Paper, a subordinate of China's Ministry of Science and Technology (MIIT): China Electronic Standardisation Institute's Artificial Intelligence (2021). 'Artificial Intelligence Standardization White Paper (2021 Edition) 人工智能标准化白皮书 (2021版)'. Centre for Security and Emerging Technology, Georgetown University. https://cset.georgetown.edu/wp-content/uploads/t0393_AI_white_paper_EN.pdf

178 Zhu (2024). 'China's Approach to AI Standardisation'. Finnish Institute of International Affairs. https://fiia.fi/wp-content/uploads/2024/08/bp391_chinas-approach-to-ai-standardisation.pdf

179 Zhu (2024). 'China's Approach to AI Standardisation'. Finnish Institute of International Affairs. https://fiia.fi/wp-content/uploads/2024/08/bp391_chinas-approach-to-ai-standardisation.pdf; Bloom (2025). 'DeepSeek: how China's embrace of open-source AI caused a geopolitical earthquake'. The Conversation. <https://theconversation.com/deepseek-how-chinas-embrace-of-open-source-ai-caused-a-geopolitical-earthquake-249563>; Larsen (2019). 'Drafting China's National AI Team for Governance'. Stanford University. <https://digichina.stanford.edu/work/drafting-chinas-national-ai-team-for-governance/>

180 Ministry of Science and Technology (MIIT; 2019). '科技部关于印发《国家新一代人工智能开放创新平台建设工作指引》的通知'. https://www.gov.cn/xinwen/2019-08/04/content_5418542.htm; Larsen (2019). 'Drafting China's National AI Team for Governance'. Stanford University. <https://digichina.stanford.edu/work/drafting-chinas-national-ai-team-for-governance/>

181 Allan, G. C. (April 2025). DeepSeek: A Deep Dive. Centre for Strategic & International Studies. <https://www.csis.org/analysis/deepseek-deep-dive>

Realised Benefits:

China has several home-grown world-leading AI developers who have released open-source models with cutting-edge capabilities. These include Alibaba's Qwen¹⁸² and Tencent's Hunyan video generation model.¹⁸³ In particular, the release of DeepSeek's R1¹⁸⁴ sent shockwaves through the world of AI in January 2025 by showing it was possible to develop open-source models with leading capabilities at a significantly lower cost than closed-source competitors such as OpenAI's GPT-4.¹⁸⁵

Despite US export controls, DeepSeek was able to train its highly performant V3 model on Nvidia H800 chips (a degraded version of the blocked Nvidia H100 chips) and reportedly at a fraction of the cost to comparable US models.¹⁸⁶ There are some doubts about the cost reporting,¹⁸⁷ but the efficiency gains they were able to achieve remain notable, and that the models were openly released served a blow to the US AI market.¹⁸⁸ Rousing open-source community enthusiasm for models such as DeepSeek V3 and R1 will likely help increase the global competitiveness of China's Huawei chips by disseminating base models built for Huawei's software ecosystem, CANN (Compute Architecture for Neural Networks).¹⁸⁹

Notes for the UK:

China's strategy and policies point to the power of standard-setting. The UK could take lessons from how China has used tools like the NAITG to centrally set open-source as the standard way of developing AI. However, the UK government is unlikely to be able or want to exercise the same degree of centralised control over private AI developers as China has via its NAITG. Unlike the UK, China's government exercises an unusually high degree of top-down power over companies¹⁹⁰ – to the extent that it owns stakes in many companies – and influences corporate decision-making in ways that British politicians may consider illiberal. Nor does the UK have comparable levels of funding, compute, or talent available to pursue its aims at the scale China has, or deploy the resources at a comparable scale.

EUROPEAN UNION

The European Union (EU)'s approach to AI openness has combined innovation in AI regulation with an investment-led strategy for promoting open-source development.

Strategy Overview:

The EU's approach to its AI openness strategy is two-stranded. On one hand, it has pursued innovative AI regulations with specific provisions for open-source AI via the AI Act (2024), which treats open-source AI models and systems more lightly than their closed-source or commercial counterparts.¹⁹¹ For example, open-source AI systems are exempted entirely

182 Alibaba Cloud (2025). 'Qwen2.5'. GitHub. <https://github.com/QwenLM/Qwen2.5>

183 Tencent (2025). 'Hunyan Video'. <https://aivideo.hunyan.tencent.com/>; Tencent (2025). 'HunyanVideo: A Systematic Framework For Large Video Generation Model'. GitHub. <https://github.com/Tencent/HunyanVideo>

184 DeepSeek (2025). 'DeepSeek-R1'. GitHub. <https://github.com/deepseek-ai/DeepSeek-R1>

185 Ng et al. (2025). 'DeepSeek: The Chinese AI app that has the world talking'. BBC News. <https://www.bbc.co.uk/news/articles/c5yv5976z9po>

186 Liu, A. et al. (2025). DeepSeek-V3 Technical Report. <https://arxiv.org/abs/2412.19437>

187 Sheehan, M. & Winter-Levy, S. (2025). Chips, China, and a Lot of Money: The Factors Driving the DeepSeek AI Turmoil. <https://carnegieendowment.org/emissary/2025/01/deepseek-ai-china-chips-explainer?lang=en>

188 Milmo, D., Hawkins, A., Booth, R. & Kollewe, J. (2025). 'Sputnik moment': \$1tn wiped off US stocks after Chinese firm unveils AI chatbot. The Guardian. <https://www.theguardian.com/business/2025/jan/27/tech-shares-asia-europe-fall-china-ai-deepseek>

189 Allan, G. C. (April 2025). DeepSeek: A Deep Dive. Centre for Strategic & International Studies. <https://www.csis.org/analysis/deepseek-deep-dive>

190 Allen et al. (2022). 'Centralization or Decentralization? The Evolution of State-Ownership in China'. SSRN. <https://dx.doi.org/10.2139/ssrn.4283197>; Roberts (2021). 'Xi Jinping's politics in command economy'. Atlantic Council. <https://www.atlanticcouncil.org/in-depth-research-reports/issue-brief/issue-brief-xi-jinpings-politics-in-command-economy/>

191 See European Union (2024). 'Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act)'. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj; hereafter EU AI Act 2024>.

from the Act's regulations so long as they meet certain requirements:¹⁹² they must not be (a) considered 'general-purpose',¹⁹³ (b) sold commercially,¹⁹⁴ (c) considered 'high-risk',¹⁹⁵ (d) reliant on prohibited practices such as subliminal messaging,¹⁹⁶ or (e) intended to interact directly with individuals or expose individuals to AI-generated content.¹⁹⁷

On the other hand, the EU has emphasised investment in open-source AI through initiatives such as its Digital Europe Programme.¹⁹⁸ The EU has provided €20 million funding to the OpenEuroLLM project,¹⁹⁹ which is developing an open-source family of LLMs in European languages, and is funding a project to make a European high-performing open-source foundation model available for downstream finetuning.²⁰⁰ The EU's investments in AI openness appear designed to complement national-level investments and development efforts by member states such as France (see below).

Realised Benefits:

The EU's legislative efforts have positioned it as a regulatory innovator in AI. The AI Act is arguably the first regulation of its kind worldwide. However, the downstream effects of these regulations are still to be determined, including its effect on open-source development.

Meanwhile, the EU's investments in AI openness have led to notable projects such as OpenEuroLLM, which have brought together partners from industry and academia. As these projects are still in the development phase, it remains to be seen whether they will achieve the EU's aim of creating accessible high-performance open-source AI models.

Notes for the UK:

While the EU's position as a supranational body means it has a different role and capabilities, the UK can take lessons from its multi-stranded approach to AI openness. The AI Act illustrates how a strong legislative framework can create regulatory certainty and transparency for open-source AI developers, while building trust for users. As Demos has highlighted previously,²⁰¹ the UK may wish to consider adopting similar exemptions for open-source AI as in the Act. Meanwhile, in terms of its investments in AI openness, the UK could adopt the EU's emphasis on open-source development when it provides funding for projects.

FRANCE

France has chosen to promote open-source AI development through government investments. The country's aim has been to develop AI national champions and support France's existing open-source AI developers, while avoiding dependence on monopolies for access to AI capabilities.²⁰²

192 EU AI Act 2024, Article 2(12). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

193 See EU AI Act 2024, Article 3(63) and Article 3(66) for definitions. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

194 EU AI Act 2024, Article 2(12). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

195 See EU AI Act 2024, Article 6 for definitions. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

196 EU AI Act 2024, Article 5(1). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

197 EU AI Act 2024, Article 50. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

198 European Commission (2025). 'Digital Europe Programme'. https://commission.europa.eu/funding-tenders/find-funding/eu-funding-programmes/digital-europe-programme_en; Sawers (2025). 'Open source LLMs hit Europe's digital sovereignty roadmap'. TechCrunch. <https://techcrunch.com/2025/02/16/open-source-llms-hit-europes-digital-sovereignty-roadmap/>

199 OpenEuroLLM (2025). <https://openeurollm.eu/>; Sawers (2025). 'Open source LLMs hit Europe's digital sovereignty roadmap'. TechCrunch. <https://techcrunch.com/2025/02/16/open-source-llms-hit-europes-digital-sovereignty-roadmap/>

200 European Commission (2024). 'Making available a high performing open-source European foundation model for fine-tuning (DIGITAL-2024-AI-06-FINETUNE)'. <https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/digital-2024-ai-06-finetune>

201 Seger & O'Dell (2024). Open Horizons: Exploring Nuanced Technical and Policy Approaches to Openness in AI. https://demos.co.uk/wp-content/uploads/2024/08/Mozilla-Report_2024.pdf

202 Chatterjee & Volpicelli (2023). 'France bets big on open-source AI'. Politico. <https://www.politico.eu/article/open-source-artificial-intelligence-france-bets-big/>

Strategy Overview:

France is home to an ecosystem of open-source AI developers and platforms for supporting open-source AI development.²⁰³ These include Mistral AI,²⁰⁴ an open-source frontier AI developer, and the open-source AI model repository Hugging Face.²⁰⁵

In recent years, the French government has made a public commitment to supporting its open-source AI ecosystem and providing funding for more AI openness initiatives.²⁰⁶ Since 2023, France has announced several investments in open-source AI through grants, public-private partnerships, and co-investments made alongside venture capital. These include a \$400 million investment in Current AI²⁰⁷ announced at the Paris AI Summit. It is a partnership of countries and AI companies which seeks to develop AI projects in the public interest.

Realised Benefits:

France is home to several fast-growing AI startups and labs which focus on open-source, such as Mistral and Kyutai.²⁰⁸ However, it is difficult to identify how much of France's success in AI openness can be attributed to the government's strategy. The French government's decision to promote AI openness as a national strategy was relatively recent, compared to countries like China, while projects invested in by the French government such as Current AI are still in development.

Notes for the UK:

France is a useful comparison point for the UK: it is a close neighbour, trading partner, and has access to similar resources. Both countries share a policy goal of building AI 'national champions'. Given these similarities, the UK could consider adopting France's emphasis on AI openness as part of its investment strategy.

INDIA

Since at least 2018, India has sought to promote AI openness as part of a wider emphasis on open-source software, with the overarching goal of establishing national self-reliance within its tech ecosystem.²⁰⁹ India's approach to achieve these goals has centred on AI projects developed by the government.

Strategy Overview:

India's government has a history of promoting open-source software²¹⁰ and open access data.²¹¹ In 2015, India mandated that all software used at a federal level had to be open source,²¹² with

203 Office of the President of France (2025). Make France an AI Powerhouse. <https://www.elysee.fr/admin/upload/default/0001/17/d9c1462e7337d353f918aac7d654b896b77c5349.pdf>

204 Mistral AI (2025). <https://mistral.ai/>

205 Hugging Face is a French-American company founded in New York by French AI developers. See Hugging Face (2025). <https://huggingface.co/huggingface>; Cai (2022). 'The \$2 Billion Emoji: Hugging Face Wants To Be Launchpad For A Machine Learning Revolution'. Forbes. <https://www.forbes.com/sites/kenrickcai/2022/05/09/the-2-billion-emoji-hugging-face-wants-to-be-launchpad-for-a-machine-learning-revolution/>

206 Chatterjee & Volpicelli (2023). 'France bets big on open-source AI'. Politico. <https://www.politico.eu/article/open-source-artificial-intelligence-france-bets-big/>

207 Current AI (2025). 'Current AI Launch Press Release'. <https://www.currentai.org/latest-updates/launchpressrelease>

208 Kyutai (2025). <https://kyutai.org/>

209 National Institution for Transforming India (NITI) Aayog (2018). National Strategy for Artificial Intelligence #AIFORALL. <https://www.niti.gov.in/sites/default/files/2023-03/National-Strategy-for-Artificial-Intelligence.pdf>; also see IndiaAI (2025). 'India's vision for AI: Prime Minister's address at the AI Action Summit, Paris'. <https://indiaai.gov.in/article/india-s-vision-for-ai-prime-minister-s-address-at-the-ai-action-summit-paris> 210 E.g., OpenForge (2025). 'About'. <https://openforge.gov.in/openforge/about.php>; Global Digital Public Infrastructure Repository (2025).

'India'. https://www.dpi.global/globaldpi/india_list; see also Section 16 of India's 2023-24 national budget. Government of India (2023). Budget 2023-24. https://www.indiabudget.gov.in/doc/bspeech/bs2023_24.pdf; Government of India (2014). 'Policy on Adoption of Open Source Software for Government of India'. https://www.meity.gov.in/static/uploads/2024/02/policy_on_adoption_of_oss.pdf; Government of India (2015). Framework for Adoption of Open Source Software in e-Governance Systems. <https://egovstandards.gov.in/sites/default/files/2021-07/Framework%20for%20Adoption%20of%20Open%20Source%20Software%20in%20e-Governance%20Systems.pdf>

211 E.g., Open Government Data Platform (2025). 'About'. <https://www.data.gov.in/about>; Pirihaar (2015). 'How is open data changing India?'. World Economic Forum. <https://www.weforum.org/stories/2015/02/how-is-open-data-changing-india/>

212 Government of India (2015). 'Framework For Adoption of Open Source Software In e-Governance Systems'. <https://egovstandards.gov.in/sites/default/files/2021-07/Framework%20for%20Adoption%20of%20Open%20Source%20Software%20in%20e-Governance%20Systems.pdf>

state and regional governments also following suit.²¹³ As of 2024, it was estimated that the adoption of an open-source computer operating system in schools in the state of Kerala had saved nearly ₹30,000,000,000 – approximately £265 million²¹⁴ – compared to using proprietary software like Microsoft Windows.²¹⁵

Within this context, India's government has a stated goal of promoting AI openness through both investments and in-house development.²¹⁶ India's motives include promoting national independence and self-reliance, ensuring transparency, improving the inclusion of underrepresented groups and languages, lowering costs, and encouraging interoperability.

To achieve these goals, the government has led or supported several open-source AI projects via a centralised national mission for AI development and investment, IndiaAI.²¹⁷ IndiaAI provides AI researchers with access to compute²¹⁸ and promotes Indian open-source AI projects.²¹⁹ Moreover, since 2019, the Indian government has led the development of the Bhashini initiative – a family of open-source large-language models supporting 22 Indian languages and dialects.²²⁰ Looking ahead, IndiaAI is involved in the development of a public platform for data and model sharing similar to Hugging Face.²²¹

Realised Benefits:

India's AI openness strategy appears to have been a success. Projects such as Bhashini have led to the creation of open-source AI models as public goods: Bhashini provides an AI-powered platform for translation, chatbots, text summarisation and more which can be integrated into downstream applications by Indian developers.²²² Meanwhile, IndiaAI plays a key role in India's AI development ecosystem through its compute resources. Looking ahead, IndiaAI is involved in the development of a public platform for data and model sharing similar to Hugging Face.²²³

Notes for the UK:

The UK government can learn lessons from how India's government has developed open-source AI models as public goods. By taking this approach, India has created a low-cost resource for use by both public bodies and private software developers downstream. Likewise, IndiaAI indicates how the UK government could play a more direct role in providing open infrastructure which lowers barriers to entry for AI development.

213 De et al. (2015). 'Economic Impact of Free and Open Source Software Usage in Government Final Report'. International Centre for Free and Open Source Software (ICFOSS). [https://icfoss.in/doc/ICFOSS_economic-impact-free\(v3\).pdf](https://icfoss.in/doc/ICFOSS_economic-impact-free(v3).pdf)

214 Based on an exchange rate of 1 Pound to 113.164 Indian Rupees, using exchange rates from 7/5/2025.

215 The Hindu Bureau (2024). 'KITE set to launch updated FOSS-based OS for public schools in Kerala'. The Hindu. <https://www.thehindu.com/news/national/kerala/kite-set-to-launch-free-updated-os-for-public-school-computers-in-kerala/article68553871.ece>

216 National Institution for Transforming India (NITI) Aayog (2018). National Strategy for Artificial Intelligence #AIFORALL. <https://www.niti.gov.in/sites/default/files/2023-03/National-Strategy-for-Artificial-Intelligence.pdf>; also see IndiaAI (2025). 'India's vision for AI: Prime Minister's address at the AI Action Summit, Paris'. <https://indiaai.gov.in/article/india-s-vision-for-ai-prime-minister-s-address-at-the-ai-action-summit-paris>

217 IndiaAI (2025). <https://indiaai.gov.in/>

218 ETech (2025). 'Explained: IndiaAI compute portal, AIKosha and other initiatives under the IndiaAI Mission'. The Economic Times. <https://economictimes.indiatimes.com/tech/technology/explained-indiaai-compute-portal-aikosha-and-other-initiatives-under-the-indiaai-mission/articleshow/118780355.cms>

219 E.g., Jeevanandam (2022). 'Eight interesting open-source Indian projects that can support AI research'. IndiaAI. <https://indiaai.gov.in/article/eight-interesting-open-source-indian-projects-that-can-support-ai-research>; Jeevanandam (2022). 'Sarvam AI launches open-source foundational models in 10 Indian languages'. IndiaAI. <https://indiaai.gov.in/article/sarvam-ai-launches-open-source-foundational-models-in-10-indian-languages>

220 Indian Ministry of Electronics and Information Technology (MeitY; 2025). 'About Bhashini'. <https://bhashini.gov.in/about-bhashini>

221 Suri (2025). 'The Missing Pieces in India's AI Puzzle: Talent, Data, and R&D'. Carnegie Endowment for International Peace. <https://carnegieendowment.org/research/2025/02/the-missing-pieces-in-indias-ai-puzzle-talent-data-and-randd?lang=en>

222 Indian Ministry of Electronics and Information Technology (MeitY; 2025). 'About Bhashini'. <https://bhashini.gov.in/about-bhashini>

223 Suri (2025). 'The Missing Pieces in India's AI Puzzle: Talent, Data, and R&D'. Carnegie Endowment for International Peace. <https://carnegieendowment.org/research/2025/02/the-missing-pieces-in-indias-ai-puzzle-talent-data-and-randd?lang=en>

SINGAPORE

Singapore's AI strategy has an emphasis on supporting the open-source AI ecosystem. Beyond models and systems themselves, Singapore has pioneered an open-source approach to AI safety and reliability testing platforms.

Strategy Overview:

Singapore has incorporated AI openness into its national AI strategy – with a particular emphasis on open-source tools for AI safety and assurance.²²⁴

Like India and Spain, Singapore's government has led the development of a family of open-source LLMs in South-East Asian languages called SEA-LION.²²⁵ The stated goal behind SEA-LION has been to ensure that AI models in local languages are easily available to South-East Asian developers at a low cost. In April 2025, SEA-LION released two hybrid reasoning models based on Meta's Llama which they claim can outperform Deepseek R1 and GPT-4o-mini.²²⁶

Singapore has also invested heavily on open-source software to support AI capability and safety testing. For example, Singapore's national AI development initiative AI Singapore has released SEA-HELM, an open-source benchmarking platform aimed at South-East Asian AI models.²²⁷

In 2023, Singapore's government launched the AI Verify Foundation: an institute intended to set AI reliability and safety standards, develop assurance frameworks, and build open-source testing tools.²²⁸ The Foundation runs AI Verify, an open-source AI governance testing framework and software toolkit designed to be consistent with AI governance frameworks from the European Union, OECD, and Singapore.²²⁹ It has also launched Project Moonshot, an open-source Evaluation Toolkit for generative AI models and large-language models.²³⁰ Furthermore, Singapore's Cyber Security Agency has set out advice on addressing security risks when using open-source models.²³¹

Besides AI safety and assurance projects, the Singaporean government has partnered with AI developers such as Meta to create AI accelerator programs focused on open-source AI development.²³²

Realised Benefits:

Singapore has pioneered AI safety and assurance platforms through AI Verify. By making these tools open-source, Singapore has simultaneously lowered barriers to safety testing for domestic AI developers and has built a platform for other countries to build on. This means Singapore's AI safety and assurance standards are more likely to be adopted internationally. Meanwhile, Singapore has seen continued success in its development of its open-source family of LLMs, SEA-LION, which offers models based on languages used across the region.

224 Government of Singapore (2023). Singapore National AI Strategy 2.0: AI for the Public Good for Singapore and the World. <https://file.go.gov.sg/nais2023.pdf>; Sharon (2024). 'Singapore's AI Vision: Inclusivity, Innovation and Responsibility'. Open Gov. <https://opengovasia.com/2024/10/05/singapores-ai-vision-inclusivity-innovation-and-responsibility/>

225 SEA-LION.AI (2025). 'SEA-LION: South-East Asian Languages in One Network'. <https://sea-lion.ai/>

226 SEA-LION.AI (2025). 'SEA-LION v3.5 and Updated v3: Enhanced Language Models for Southeast Asia'. <https://sea-lion.ai/sea-lion-v3-5-and-updated-v3-enhanced-language-models-for-southeast-asia/>

227 AI Singapore (2025). 'SEA-HELM'. Github. <https://github.com/aisingapore/SEA-HELM>

228 AI Verify Foundation (2025). 'About AI Verify Foundation'. <https://aiverifyfoundation.sg/ai-verify-foundation/>; Infocomm Media Development Authority (2023). 'Singapore launches AI Verify Foundation to shape the future of international AI standards through collaboration'. <https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/press-releases/2023/singapore-launches-ai-verify-foundation>

229 AI Verify Foundation (2025). 'What is AI Verify?'. <https://aiverifyfoundation.sg/what-is-ai-verify/>; AI Verify Foundation (2025). 'AI Verify'. Github. <https://github.com/aiverify-foundation/aiverify>

230 AI Verify Foundation (2025). 'Project Moonshot'. <https://aiverifyfoundation.sg/project-moonshot/>; AI Verify Foundation (2025). 'Project Moonshot'. Github. <https://github.com/aiverify-foundation/moonshot>

231 For example, see Section 3.1 of Cyber Security Agency (2024). Companion Guide on Securing AI Systems. Government of Singapore. <https://isomer-user-content.by.gov.sg/36/3cfb3cd5-0228-4d27-a596-3860ef751708/Companion%20Guide%20on%20Securing%20AI%20Systems.pdf>

232 Singapore Economic Development Board (2024). 'Meta unveils AI Accelerator program to support open source AI solutions in Asia Pacific'. <https://www.edb.gov.sg/en/about-edb/media-releases-publications/meta-unveils-ai-accelerator-program-to-support-open-source-ai-solutions-in-asia-pacific.html>

Notes for the UK:

The UK could follow Singapore's emphasis on promoting an open-source approach to AI safety and reliability testing. Singapore's AI Verify Foundation is similar in some respects to the UK's AI Security Institute (AISi) in its focus on AI model safety testing, benchmarking, and compliance. Both institutions have released open-source AI safety testing toolkits to the public.²³³ AI Verify provides a helpful example of how to make the development of open-source AI testing platforms a core goal of a national AI safety body and illustrates a mechanism by which the UK's AISi could influence global AI safety through openly sharing high-quality safety tooling.

SPAIN

Spain's national AI strategy includes an emphasis on openness across its AI stack. The country's approach to achieving its goals has been primarily government-led.

Strategy Overview:

The Spanish government has stated that it sees AI openness as a way to promote principles of transparency and openness and to ensure Spanish-language AI models are developed.²³⁴ Spain has used government-led investment and AI development as a means of producing AI as a public good: subsidised by public funds and available to all for downstream use.

Like India, the Spanish government is working on developing a family of open-source Spanish-language foundation models called ALIA, which are to be "characterised by maximum transparency and openness".²³⁵ ALIA is intended to form a key component of "a public infrastructure of AI resources".²³⁶ It is intended to provide high-quality AI models in the official languages spoken in Spain – Spanish, Catalan, Valencian, Basque, and Galician – as an alternative to models trained primarily in foreign languages.

Spain is also pursuing an open-source approach to AI hardware. For example, Spain's 2024 National AI Strategy includes a goal to build AI compute within the open-source hardware (OSHW) paradigm. Open-source hardware principles require hardware designs, manufacturing processes, firmware, and other IP needed to produce and operate a device to be released publicly on an open licence.²³⁷ In Spain's case, this would mean that all chips used – as well as the system designs of the data centres themselves – must meet open-source hardware standards.

Spain's AI efforts are to be supervised and certified by a newly created public AI advisory agency, the Spanish Agency for the Supervision of Artificial Intelligence (AESIA). AESIA has the stated aim of promoting "ethics, innovation, and transparency".²³⁸

Realised Benefits:

Spain has made some progress in its goals. After an initial development period, the ALIA project has released its first models and datasets to the public.²³⁹ The first wave of releases includes four open-source models with different capabilities.²⁴⁰ Meanwhile, Spain has made its Lagarto family

233 E.g. AI Security Institute (2025). 'Inspect: An open-source framework for large language model evaluations'. <https://inspect.aisi.org.uk/>

234 Government of Spain (2024). 2024 Artificial Intelligence Strategy. https://digital.gob.es/dam/en/portalmtdfp/DigitalizacionIA/1_DOSSIER_AI_ENGLISH_15_JULIO.pdf

235 Government of Spain (2024). 2024 Artificial Intelligence Strategy. https://digital.gob.es/dam/en/portalmtdfp/DigitalizacionIA/1_DOSSIER_AI_ENGLISH_15_JULIO.pdf

236 ALIA (2025). 'ALIA: The public AI infrastructure in Spanish and co-official languages'. <https://alia.gob.es/eng/>

237 OSHWA Certification (2025). 'Open Source Hardware Basics'. <https://certification.oshwa.org/basics.html>

238 AESIA (2025). <https://aesia.digital.gob.es/en/es>

239 ALIA Kit (2025). <https://langtech-bsc.gitbook.io/alia-kit>

240 EU Open Source Observatory (2025). 'Spanish Government promotes open access to its ALIA AI models'. <https://interoperable-europe.ec.europa.eu/collection/open-source-observatory-osor/news/spanish-authorities-release-alia-ai-models>

high-performance computing chips available under an open-source hardware license.²⁴¹

Notes for the UK:

Spain provides a helpful comparison for the UK insofar as it is a medium-sized European country with an expressed interest in developing sovereign AI models. It offers an interesting example of how openness can be pursued across the AI technology stack.

The UK can learn lessons from how Spain has gone about developing its family of foundation models as an open-source public good. By building these models in this manner, Spain has created a low-cost resource for use by both public bodies and private downstream AI developers.

Meanwhile, Spain's exploration of open-source hardware within its efforts to expand its compute capacity provides an interesting example of how AI openness may go beyond models themselves. As the UK seeks to expand its own compute capacity and wider AI infrastructure, it may wish to consider whether supporting greater openness in hardware and software might be beneficial.

THAILAND

The Thai government has very recently set out its national strategy for AI (May 2, 2025), with the aim of making Thailand into a regional centre for AI development.²⁴² Thailand's strategy includes investing in the development of open-source platforms for AI development and supporting the development of open-source AI models. Thailand is already home to several open-source Thai-language AI development efforts, such as Typhoon.²⁴³

241 Barcelona Supercomputing Center (2023). 'BSC presents Sargantana, the new generation of the first open-source chips designed in Spain'. <https://www.bsc.es/news/bsc-news/bsc-presents-sargantana-the-new-generation-the-first-open-source-chips-designed-spain>

242 The Nation Thailand (2025). 'Thailand outlines ambitious AI strategy to become regional hub'. <https://www.nationthailand.com/business/tech/40049494>

243 Typhoon (2025). 'About'. <https://opentyphoon.ai/about>; Pipatanakul, K. et al. (2023). Typhoon: Thai Large Language Models. <https://arxiv.org/abs/2312.13951>

APPENDIX 2

COLLATED RECOMMENDATIONS CROSS-REFERENCED WITH AI OPPORTUNITIES ACTION PLAN PROPOSALS

FROM SECTION 5.1 - LAYING THE FOUNDATIONS FOR AI

TABLE 1

AI OPENNESS RECOMMENDATIONS FOR “LAYING FOUNDATIONS FOR AI”

RECOMMENDATION		DESCRIPTION
For building and securing reliable access to sustainable compute infrastructure		
1a	<u>Building on Action Plan (AP) proposals 2 and 3:</u> Continue plans to expand the UK’s AI Research Resource (AIRR) and prioritise subsidising access for entrepreneurs developing and open-sourcing ‘public interest AI’.	Proposal 3 notes the need for strategically prioritising AIRR compute allocation instead of spreading it thin. We recommend providing subsidised or free access to researchers undertaking projects of national strategic importance and to UK based AI entrepreneurs focused on developing and open-sourcing public interest AI (see recommendation 1i on defining ‘public interest AI’).
1b	<u>Building on AP proposals 1-4:</u> Follow OCP Ready™ requirements in building new data centres / invest in open-source compute hardware like in Spain and France.	Proposals 1-4 call for a long-term plan for expanding the UK’s domestic compute capacity and allocating it according to strategic goals. When investing in the UK’s compute capacity, the UK could follow Spain’s example and invest in open-source compute hardware. ²⁴⁴ France too is investing in open-source cloud compute infrastructure. ²⁴⁵ New Data Centres should follow OCP Ready™ requirements ²⁴⁶ and participate in OCP community projects. ²⁴⁷ Open-source hardware principles suggest hardware designs,

244 Barcelona Supercomputing Center (2023). ‘BSC presents Sargantana, the new generation of the first open-source chips designed in Spain’. <https://www.bsc.es/news/bsc-news/bsc-presents-sargantana-the-new-generation-the-first-open-source-chips-designed-spain>

245 French Government Plans to Invest €1.8 Billion to Support the French Cloud Industry. Research Connect. Accessed 27 May, 2025. <https://myresearchconnect.com/french-government-plans-to-invest-e1-8-billion-to-support-the-french-cloud-industry/>

246 OCP Ready™ Data Center Recognition Program. Accessed 27 May, 2025. <https://www.opencompute.org/projects/ocp-readytm-data-center-recognition-program>

247 OCP Projects. Accessed 27 May, 2025. <https://www.opencompute.org/projects>

RECOMMENDATION		DESCRIPTION
		<p>manufacturing processes, firmware, and other IP needed to produce and operate a device be released publicly on an open licence.²⁴⁸</p> <p>Government representatives could also consider attending the OCP Globals summit to learn more.²⁴⁹</p>
1c	Attend to (AP) proposal 6 - build international compute collaborations with like minded countries.	<p>Proposal 6 calls for the UK to agree international compute partnerships with like-minded countries. The UK should pursue these collaborations as part of its wider strategy. For example, the UK could take part in European efforts to build an independent tech stack for the region that includes AI compute such as with EuroHPC Joint Undertaking²⁵⁰ and the EuroStack proposal.²⁵¹</p>
For unlocking data assets in public and private sectors		
1d	Building on AP Proposals 7-9: Identify and develop high-impact datasets that can be made either fully or locally openly to UK researchers and innovators, as well as in international AI development partnerships for public interest AI.	<p>Proposals 7-9 set out ways for the UK to identify, curate and share access to high-quality datasets for AI. As part of these efforts, the UK should explore the full range of options available for data access – such as fully open to all, locally open to UK developers, and closed access based on partnerships. The UK should evaluate which access model is most appropriate for each dataset, depending on factors like the data's sensitivity and value, with a preference for greater openness. The UK should also draw on its experience to develop and publicly share guidance on best practices for data dissemination.</p> <p>Adopting a range of data access models, with an emphasis on open access, could provide the UK with flexibility in how it uses these strategic resources. By making data access fully or locally open, the UK can maximise the number of potential developers that could benefit and spur innovation.</p>
1e	Partner with Current AI as a commitment towards public AI and collaborative AI innovation.	<p>Current AI an international funding body and convener announced at the Paris AI Summit working to catalyse public investment in public AI. It works to coordinate action across governments, philanthropic funders, and research communities with a particular focus on shifting norms and building infrastructure to facilitate data sharing for public interest AI projects.</p>

248 OSHWA Certification (2025). 'Open Source Hardware Basics'. <https://certification.oshwa.org/basics.html>

249 2025 OP Global Summit. <https://www.opencompute.org/summit/global-summit>

250 The European High Performance Computing Joint Undertaking (EuroHPC JU). Accessed 25 May, 2025. https://eurohpc-ju.europa.eu/index_en

251 EuroStack (2025). Deploying the EuroStack: What's Needed Now. Accessed 25 May, 2025 <https://euro-stack.eu/the-white-paper/>

	RECOMMENDATION	DESCRIPTION
		Partnering with Current AI would be a potentially high-impact and relatively low-cost step in showing the government's commitment to developing AI for public benefit and aligning with international efforts to collaboratively build an AI openness counterpoint to proprietary AI infrastructure.
1f	Urgently reset relationship with the EU on the Data Union Strategy ²⁵² with a collaborative project and open data-sharing.	Through open collaboration partnerships, the UK can benefit from data resources from beyond the country's borders. For example, the EU is leading the way in curating and enhancing access to high-quality data across the region. A UK-EU partnership on data curation and sharing could build on recent improvements in our bilateral relationship to enhance the UK's domestic AI efforts.
1g	<u>Building on AP Proposals 3 and 7</u> Build open-source AI innovation platforms for downstream developers.	<p>AP Proposal 3 calls for the UK to strategically allocate sovereign compute to AI developers, while Proposal 7 proposes to identify high-impact public datasets and make them available for private sector AI training. Both goals could be furthered by offering access through public AI openness innovation platforms.</p> <p>These platforms provide private sector developers with resources they might otherwise be unable to access by bundling together experimentation sandboxes, data, compute, foundation models for fine-tuning, and other useful tools. They leverage the accessibility, customisability, and lower costs that AI openness offers, and can be built as open-source software. China has successfully built more than 10 such platforms to spur AI innovations and support downstream developers who want to integrate AI into their products, with each platform dedicated to a specific subdomain of AI.²⁵³</p> <p>This proposal also aligns with the missions set out in the Action Plan for the UK Sovereign AI Unit and could be included as part of its portfolio (see section 5.3).</p>

252 A European Strategy for Data. Accessed 27 May, 2025. <https://digital-strategy.ec.europa.eu/en/policies/strategy-data>; European Commission (May 23, 2025). Commission seeks views on the use of data to develop Artificial Intelligence. Accessed 27 May, 2025. <https://digital-strategy.ec.europa.eu/en/news/commission-seeks-views-use-data-develop-artificial-intelligence>

253 Ding (2023). 'China's Uncharacteristic Approach to Artificial Intelligence (AI) Development'. University of California Institute on Global Conflict and Cooperation. <https://www.ucigcc.org/blog/chinas-uncharacteristic-approach-to-artificial-intelligence-ai-development/>; also Cricchio et al. (2025). 'China's new knowledge brokers. A patent citations network analysis of the artificial intelligence open innovation ecosystem'. Journal of Engineering and Technology Management. <https://www.sciencedirect.com/science/article/abs/pii/S0923474825000116>

RECOMMENDATION		DESCRIPTION
For enabling safe and trusted AI development and deployment through regulation and assurance		
1h	In the forthcoming UK AI Bill: Establish transparency requirements for highly capable proprietary models. <i>(As in Article 13 of the EU AI Act)</i> ²⁵⁴	This allows downstream developers to responsibly integrate potentially high-risk systems into new applications, and enables more thorough safety testing at different stages along the AI value chain. Open-source models will already meet this requirement.
1i	In the forthcoming UK AI Bill: Introduce exemptions from AI regulation for models that meet a certain standard of openness/transparency. <i>(as in the EU AI Act article 2(12)).</i> ²⁵⁵ <i>The exemption does not apply to models classified high-risk.)</i> ²⁵⁶	<p>This will incentivise greater transparency and information sharing among AI developers toward the facilitation of a thriving open-source ecosystem in the UK. Meanwhile, smaller open-source developers and dispersed open-source communities are not burdened by regulatory requirements that they do not have the resources or coordination to take action. Higher-risk systems can be bracketed off to not qualify for exemption irrespective of openness (e.g. 'high-risk' AI in the EU AI Act).</p> <p>There is some risk, however, that if regulatory exemptions are used to encourage greater openness but the transparency requirements for earning the incentive are too stringent, some providers who might otherwise have offered semi-open access to their models may be motivated to close model access. The perfect could become the enemy of the good.</p>
1j	<u>Building on AP proposals 23 and 29</u> : continue to support and grow AISI's safety research and model evaluation capacity. AISI should also invest heavily in new AI assurance and cybersecurity testing tool development, both internally and through its fast grants program. Wherever possible the tools should be open-sourced and made accessible, either directly from AISI in a similar manner to Singapore's AI Verify program or by contributing to ROOST.	<p>While AISI has open-sourced some of its existing tools, such as Inspect,²⁵⁷ a much stronger emphasis could be placed on doing so.</p> <p>There are several benefits for the UK in open-sourcing AI safety, security and assurance tooling. (1) It enables downstream developers to access high-quality testing tools without needing to invest significant resources. (2) It sets 'soft' domestic and international standards by sharing AI assurance and cybersecurity tools widely. (3) It helps the UK exert influence over the development of safe and secure AI globally. (4) It can enhance the UK's own safety research by enabling greater collaboration.</p>

254 EU AI Act. Article 13: Transparency and Provision of Information to Deployers. <https://artificialintelligenceact.eu/article/13/#:~:text=This%20article%20states%20that%20high,limitations%2C%20and%20any%20potential%20risks>

255 EU AI Act. Article 2: Scope. Retrieved 22 May, 2025, from <https://artificialintelligenceact.eu/article/2/>

256 EU AI Act. Article 6: Classification Rules for High-Risk AI Systems. Retrieved 22 May, 2025, from <https://artificialintelligenceact.eu/article/6/>

257 AI Security Institute (2025). 'Inspect: An open-source framework for large language model evaluations'. <https://inspect.aisi.org.uk/>; for a discussion, see OpenUK (2024). 'Case Study: UK AI Safety Institute's Inspect Testing Platform'. <https://openuk.uk/case-studies/ukaisafetyinstitute-from-phasethree/>

	RECOMMENDATION	DESCRIPTION
		<p>There may be some concerns about malicious actors using access to open safety tooling to figure out how to game tests. It is therefore important to carefully consider the prudence of opening all safety tools. Some versions might be reserved for time constrained periods for internal testing of particularly high-risk applications.</p>
1k	Incorporate democratic processes into government decisions around AI to help centre the government AI goals around public benefit and underpin public trust.	<p>Relevant decisions for public input may include, for example, decisions around spending, compute resource allocations, or AI integration into public services.</p> <p>In particular, we recommend using open deliberative processes to define a concept of “public interest AI” on which decisions about benefits for public interest AI developers (PIAI) can be based. For example, PIAI developers could be prioritised and subsidised for public compute resource allocation.</p> <p>Practical options for facilitating democratic input might include e.g., implementing democratically selected oversight boards, and employing participatory processes facilitated by civic tech (e.g. platforms such as Polis and Remesh) to engage diverse multi stakeholder deliberation.</p> <p>Many decisions about AI - e.g. individual coding decisions - do not lend themselves to wide public engagement. Research is needed to establish a taxonomy of AI governance decisions that would benefit from wider deliberative engagement.</p>

FROM SECTION 5.2 - DRIVING CROSS-ECONOMY AI ADOPTION

TABLE 2

AI OPENNESS RECOMMENDATIONS FOR “DRIVING AI ADOPTION (IN THE PUBLIC INTEREST)”

	RECOMMENDATION	DESCRIPTION
For scanning for, piloting, and scaling effective AI tools for public sector adoption		
2a	<u>Building on Action Plan (AP) Proposal 34:</u> Government should procure open-source AI solutions wherever possible.	Proposal 34 calls for the government to adopt a consistent framework for how to source AI and suggests it should “support open-source solutions”. The UK can expand on this proposal by making open-source AI the preferred option in its procurement framework wherever possible. This could lower procurement costs and avoid the government becoming locked-in to a given AI vendor.
2b	<u>Building on AP Proposal 35:</u> Develop a rapid AI prototyping capability within government, and make it open source.	Proposal 35 suggests the government develops a rapid AI prototyping capability. This capability should have expertise in and an emphasis on leveraging open-source AI solutions. Prioritising AI openness could allow the government to prototype faster, at a lower cost, with greater flexibility, and with better interoperability with other systems.
2c	<u>Building on AP Proposals 41 and 42:</u> In the development or procurement of a scalable AI tech stack, mandate infrastructure interoperability, code reusability and open-sourcing.	Proposal 41 calls for the development or procurement of a scalable AI tech stack, while Proposal 42 suggests the UK should mandate infrastructure interoperability, code reusability and open sourcing. The UK should ensure that the open-sourcing mandates set out in Proposal 42 are applied across the tech stack outlined in Proposal 41. Countries such as India have had great success in mandating use of open-source software and infrastructure across government, leading to significant cost savings and greater interoperability. The UK should learn from these examples and adopt an open-source first policy to receive similar benefits.
For enabling public and private sectors to reinforce one another		
2d	<u>Building on AP Proposal 37:</u> Create data-rich and open-source AI experimentation environments, for use by public sector AI.	The proposed experimentation environment could be built as an open-source platform. This would also allow developers outside the public sector to replicate the environment, build on its innovations, and benefit from the development effort involved.

	RECOMMENDATION	DESCRIPTION
2e	<u>Building on AP Proposal 43:</u> Use government procurement strategy to incentivise open model development.	Proposal 43 notes correctly that the UK government is likely to be both the “largest customer” for AI and a “market shaper”, and suggests using this power to influence the direction of UK AI development. However, the proposal is not specific about what standards and requirements to promote. AI openness could be adopted as one of these standards for the government to promote. A policy like this from the UK government could encourage AI developers to focus on AI openness development, which could lower procurement costs and make interoperability easier in the long-term.
2f	<u>Building on AP Proposal 44:</u> Adopt an Application Programme Interfaces (API) mandate akin to that rolled out at Amazon. This should be enabled through open-source protocols like MCP (see section 3.1)	Proposal 44 suggests that digital government infrastructure could be used to “create new opportunities for innovators” and mentions Amazon’s API mandate as an example of how this could be achieved. This required all teams’ data and functionality to be exposed through APIs (Application Programme Interfaces). All standard documentation interactions, like compliance or planning, could be done through APIs, to which companies could connect their own tools. The UK should adopt such a mandate, which could be assisted by requiring use of emerging AI openness protocols such as MCP and A2A. Such a policy would make it easier to combine products and services from different sources, and would enhance innovators’ ability to build on government development efforts. It could also benefit the government by encouraging greater interoperability between its products and the wider AI ecosystem.
For ensure AI development, scaling, and adoption is in the public interest		
2g	<u>Building on AP Proposal 45:</u> Publish best-practice guidance, results, case-studies and open-source solutions through a single “AI Knowledge Hub” .	The Knowledge Hub could also be formulated to facilitate knowledge-sharing between local government AI procurers to facilitate cross governmental knowledge sharing.
2h	Provide transparency regarding the evaluation and performance metrics used to assess public sector AI tools.	We see this as core to public sector accountability, enabling public oversight of AI use in high-risk contexts (e.g. health and housing) and underpinning greater public trust in AI-augmented service provision.
2i	Employ open and collaborative AI governance mechanisms	See recommendation 1k.

FROM SECTION 5.3 - SECURING A FUTURE FOR HOMEGROWN AI

Recommendations - Securing a future for homegrown AI

Value: For ensuring the UK accrues value from the global AI transition through its homegrown AI industry

Ultimately value will accrue to AI industries that succeed at putting AI tools in consumer hands. Laying foundations for a thriving AI openness ecosystem in the UK will help the UK capture value by lowering barriers to entry for new business and by acting as an economic multiplier across industries. Recommendations toward this end draw from Section 5.1 and Table 1 in this appendix. Key recommendations to attend to are:

- 1a-c on compute access
- 1d-e on data access
- 1g on building open-source AI innovation platforms to support downstream developers

Government can also support industry and help build a thriving open-source ecosystem by providing resources and setting standards through its role as a procurer. See recommendation 2d-f on helping public and private sector AI development activities to reinforce one another.

Access: For ensuring the UK has reliable access to powerful AI models that are safe and aligned with liberal democratic values

Through collaborative open-source and public AI initiatives, the UK can contribute to and benefit from shared AI resources that remain permanently accessible and that cannot be restricted, taxed, or withdrawn by other nations.

- 1c on compute collaboration
- 1d-e on data partnership
- 1f on international model development collaboration

Influence: For facilitating UK influence over the direction of AI futures globally

Through the UK's strong AI research environment:

Open science, and specifically early movers in open science, have outsized influence on the direction of open-source developments. UK will continue to have strong influence over frontier AI innovations by supporting strong AI research environments:

- Support academic researchers with compute access and data access (1a-g)
- Encourage international research collaborations (1f)
- Balanced AI safety regulation with the forthcoming AI bill (1h-i)

Through the UK's leading AI safety research within AISI:

- See proposal 1j

Licence to publish

Demos – Licence to Publish

The work (as defined below) is provided under the terms of this licence ('licence'). The work is protected by copyright and/or other applicable law. Any use of the work other than as authorized under this licence is prohibited. By exercising any rights to the work provided here, you accept and agree to be bound by the terms of this licence. Demos grants you the rights contained here in consideration of your acceptance of such terms and conditions.

1 Definitions

a 'Collective Work' means a work, such as a periodical issue, anthology or encyclopedia, in which the Work in its entirety in unmodified form, along with a number of other contributions, constituting separate and independent works in themselves, are assembled into a collective whole. A work that constitutes a Collective Work will not be considered a Derivative Work (as defined below) for the purposes of this Licence.

b 'Derivative Work' means a work based upon the Work or upon the Work and other pre-existing works, such as a musical arrangement, dramatization, fictionalization, motion picture version, sound recording, art reproduction, abridgment, condensation, or any other form in which the Work may be recast, transformed, or adapted, except that a work that constitutes a Collective Work or a translation from English into another language will not be considered a Derivative Work for the purpose of this Licence.

c 'Licensor' means the individual or entity that offers the Work under the terms of this Licence.

d 'Original Author' means the individual or entity who created the Work.

e 'Work' means the copyrightable work of authorship offered under the terms of this Licence.

f 'You' means an individual or entity exercising rights under this Licence who has not previously violated the terms of this Licence with respect to the Work, or who has received express permission from Demos to exercise rights under this Licence despite a previous violation.

2 Fair Use Rights

Nothing in this licence is intended to reduce, limit, or restrict any rights arising from fair use, first sale or other limitations on the exclusive rights of the copyright owner under copyright law or other applicable laws.

3 Licence Grant

Subject to the terms and conditions of this Licence, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) licence to exercise the rights in the Work as stated below:

a to reproduce the Work, to incorporate the Work into one or more Collective Works, and to reproduce the Work as incorporated in the Collective Works;

b to distribute copies or phono-records of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission the Work including as incorporated in Collective Works; The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary to exercise the rights in other media and formats. All rights not expressly granted by Licensor are hereby reserved.

4 Restrictions

The licence granted in Section 3 above is expressly made subject to and limited by the following restrictions:

a You may distribute, publicly display, publicly perform, or publicly digitally perform the Work only under the terms of this Licence, and You must include a copy of, or the Uniform Resource Identifier for, this Licence with every copy or phono-record of the Work You distribute, publicly display, publicly perform, or publicly digitally perform. You may not offer or impose any terms on the Work that alter or restrict the terms of this Licence or the recipients' exercise of the rights granted hereunder. You may not sublicense the Work. You must keep intact all notices that refer to this Licence and to the disclaimer of warranties. You may not distribute, publicly display, publicly perform, or publicly digitally perform the Work with any technological measures that control access or use of the Work in a manner inconsistent with the terms of this Licence Agreement. The above applies to the Work as incorporated in a Collective Work, but this does not require the Collective Work apart from the Work itself to be made subject to the terms of this Licence. If You create a Collective Work, upon notice from any Licensor You must, to the extent practicable, remove from the Collective Work any reference to such Licensor or the Original Author, as requested.

b You may not exercise any of the rights granted to You in Section 3 above in any manner that is primarily intended

for or directed toward commercial advantage or private monetary compensation. The exchange of the Work for other copyrighted works by means of digital file sharing or otherwise shall not be considered to be intended for or directed toward commercial advantage or private monetary compensation, provided there is no payment of any monetary compensation in connection with the exchange of copyrighted works.

c If you distribute, publicly display, publicly perform, or publicly digitally perform the Work or any Collective Works, you must keep intact all copyright notices for the Work and give the Original Author credit reasonable to the medium or means You are utilizing by conveying the name (or pseudonym if applicable) of the Original Author if supplied; the title of the Work if supplied. Such credit may be implemented in any reasonable manner; provided, however, that in the case of a Collective Work, at a minimum such credit will appear where any other comparable authorship credit appears and in a manner at least as prominent as such other comparable authorship credit.

5 Representations, Warranties and Disclaimer

a By offering the Work for public release under this Licence, Licensor represents and warrants that, to the best of Licensor's knowledge after reasonable inquiry:

i Licensor has secured all rights in the Work necessary to grant the licence rights hereunder and to permit the lawful exercise of the rights granted hereunder without You having any obligation to pay any royalties, compulsory licence fees, residuals or any other payments;

ii The Work does not infringe the copyright, trademark, publicity rights, common law rights or any other right of any third party or constitute defamation, invasion of privacy or other tortious injury to any third party.

b Except as expressly stated in this licence or otherwise agreed in writing or required by applicable law, the work is licenced on an 'as is' basis, without warranties of any kind, either express or implied including, without limitation, any warranties regarding the contents or accuracy of the work.

6 Limitation on Liability

Except to the extent required by applicable law, and except for damages arising from liability to a third party resulting from breach of the warranties in section 5, in no event will licensor be liable to you on any legal theory for any special, incidental, consequential, punitive or exemplary damages arising out of this licence or the use of the work, even if licensor has been advised of the possibility of such damages.

7 Termination

a This Licence and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this Licence. Individuals or entities who have received Collective Works from You under this Licence, however, will not have their licences terminated provided such individuals or entities remain in full compliance with those licences. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this Licence.

b Subject to the above terms and conditions, the licence granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different licence terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this Licence (or any other licence that has been, or is required to be, granted under the terms of this Licence), and this Licence will continue in full force and effect unless terminated as stated above.

8 Miscellaneous

a Each time You distribute or publicly digitally perform the Work or a Collective Work, Demos offers to the recipient a licence to the Work on the same terms and conditions as the licence granted to You under this Licence.

b If any provision of this Licence is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Licence, and without further action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

c No term or provision of this Licence shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.

d This Licence constitutes the entire agreement between the parties with respect to the Work licenced here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This Licence may not be modified without the mutual written agreement of Demos and You.

DEMOS

Demos is a champion of people, ideas and democracy. We bring people together. We bridge divides. We listen and we understand. We are practical about the problems we face, but endlessly optimistic and ambitious about our capacity, together, to overcome them.

At a crossroads in Britain's history, we need ideas for renewal, reconnection and the restoration of hope. Challenges from populism to climate change remain unsolved, and a technological revolution dawns, but the centre of politics has been intellectually paralysed. Demos will change that. We can counter the impossible promises of the political extremes, and challenge despair – by bringing to life an aspirational narrative about the future of Britain that is rooted in the hopes and ambitions of people from across our country.

Demos is an independent, educational charity, registered in England and Wales. (Charity Registration no. 1042046)

Find out more at www.demos.co.uk

DEMOS

PUBLISHED BY DEMOS JUNE 2025

© DEMOS. SOME RIGHTS RESERVED.

15 WHITEHALL, LONDON, SW1A 2DD

T: 020 3878 3955

HELLO@DEMOS.CO.UK

WWW.DEMOS.CO.UK