FEBRUARY 2026

GUY WARD-JACKSON
KEEGAN MCBRIDE
ELIZABETH SEGER
LAURA GILBERT
OTT VELSBERG
JAKOB MÖKANDER

Open Source: How Middle Powers Can Build Influence in the Age of AI

TONY BLAIR
INSTITUTE FOR
GLOBAL CHANGE

Contents

Foreword

Over the past six years, Ukraine, starting from close to zero, has made a massive digital leap, turning digitalisation – and now AI – into a national superpower to serve people, strengthen governance and defend the country.

We have built one of the world's leading digital states, vaulting from 102nd to fifth globally in digital government services. Ukraine became the first country in the world with digital passports. Our citizens can register a business online faster than almost anywhere else. We also launched the world's first fully online marriage process in our state super-app Diia – an innovation recognised by *TIME* in 2024 as one of the best inventions of the year. This is a true digital revolution, with Diia at its core. But Diia is more than an app – it is a full ecosystem that transforms how the government interacts with people. The Diia app serves over 23 million users, offering more than 30 digital documents and more than 70 digital services, including a fully integrated digital signature. The Diia portal adds more than 165 additional services, making Ukraine one of the most digitally advanced countries in citizen services.

Our digital state and defence innovation are not the results of top-down planning, but the fruits of open cooperation. We placed a bet on the private sector, on entrepreneurs and innovators who could move more quickly than the government.

In the weeks that followed Russia's full-scale invasion in 2022, digital transformation became a lifeline for Ukraine. We deployed critical services in days, not months: registration of internally displaced people, financial aid, eRecovery for rebuilding homes, War Bonds and eEnemy for secure reporting. We added Diia.TV and Diia.Radio to deliver verified information in shelters. Public-private cooperation delivered life-saving solutions such as Air Alert – built in days and scaled nationwide in weeks, replacing outdated sirens and reaching millions instantly.

In four years of defending our country, drones and robotics have transformed the battlefield. This progress is powered by Brave1 – our defence-tech cluster, launched to meet the frontline's evolving challenges. AI-enabled autonomous systems, data-driven fire control and real-time battlefield awareness save the lives of soldiers and civilians every day. More technology means fewer deaths.

Now, our next big thing is building the Agentic State – where AI does not wait for citizens to navigate bureaucracy but proactively delivers services. Ukraine's National AI Strategy is coming soon. We're already fine-tuning a national large language model (LLM) and have launched Diia.AI – the world's first AI assistant performing real government services, not just functioning as a chatbot. It will soon be in the Diia app. At the same time, the Digital Agenda 2030 will guide Ukraine's digital transformation, supporting recovery, EU integration and institutional resilience under ongoing Russian aggression.

Openness is the technical foundation of our cooperation model. Open source, open data and open standards are not nice-to-haves – they are the infrastructure of innovation in the AI era. When AI is used in defence, health care and government, you need to see inside it; open source makes that possible. And open data creates accountability.

Ukraine does not just advocate openness – we practise it. We have open-sourced Diia. We score 97 per cent on EU Open Data Maturity, placing us third in Europe. We are building strategic data sets as public goods: court documents, educational content, wartime archives and other data. And for our sovereign LLM, we chose Google's open-weight Gemma – because sovereignty means owning the adaptation layer, not paying for a black box.

This report from the Tony Blair Institute makes a timely and valuable contribution to a debate that matters for every country navigating the AI transition. It provides a clear-eyed framework for how middle powers can build genuine AI capability – not by chasing frontier model development, but by building full-stack ecosystems on open foundations. The report argues convincingly that open source is not merely virtuous, but effective: it accelerates diffusion, strengthens sovereign capability and creates the competitive dynamics that drive innovation. Its practical recommendations

on strategic data sets, procurement reform, national benchmarks and ecosystem governance validate the path Ukraine is already walking. We invite every like-minded nation to learn from our experience and build together. Openness isn't optional. It's survival. It's leadership.

Oleksandr Bornyakov

Acting Minister of Digital Transformation of Ukraine

Executive Summary

A new geopolitical and economic order is emerging, driven largely by advances in artificial intelligence. Middle powers are being forced to confront a new reality in which the development of the largest frontier AI models will remain concentrated among a small number of firms in the United States and China.

But the characterisation of an AI race in which middle powers are locked out or unable to compete is misleading. Being locked out of frontier model development does not mean being locked out of the long-term AI race. Too much attention is focused on who builds the largest and most advanced large language models (LLMs), while the more important question for most governments is how nations outside the US and China can capture the economic value of AI, build public-sector capacity and develop strategic leverage in areas of comparative advantage.

The bulk of AI's economic and strategic value will be captured downstream, through deployment, adaptation and integration into existing sectors and continuous innovation across the economy on top of AI.[1] This is where open source (meaning open models themselves, as well as the open software, tools and infrastructure surrounding them) becomes strategically important.

Open source matters not simply because open models are increasingly catching up to proprietary capabilities.[2] It is important because AI will only be as powerful and as effectively adopted as the open foundations on which it runs and the applications it is used to build.

This paper therefore focuses not only on "open-source AI" at the model level but also on the deliberate use of "openness" across the broader stack – from open tools to data sets and benchmarks (standardised capability tests). Leveraging openness in this way offers middle powers a rare lever to build national capabilities, retain agency in a concentrated global AI landscape and capture value without attempting to own the frontier.[3,4]

Despite its promise, however, open source has not attracted sufficient attention. So far, many government strategies have focused on investing in AI hardware infrastructure, from access to compute to preparing grid capacity. While an essential first step, the underlying infrastructure alone will not be sufficient to capture the value of AI in the long term. This is because the economic and strategic gains from automation tend to flow not solely to the countries that invent the technology, but to those that build applications, own complementary assets, and create innovation ecosystems able to absorb and then reallocate its value.[5,6]

Even where policy attention has been paid to open source, it has often been incomplete. Open source – and open-source models in particular – is not a silver bullet for total "AI sovereignty". National open-source AI strategies that amount to middle powers building their own open-source models from scratch are often an inefficient use of resources. This is because these models quickly lag behind the frontier and become obsolete, and because the approach falsely assumes that the entire value of AI lies in the models themselves.

An effective open-source strategy cannot be model-centric. Rather, it must work towards building an open-source ecosystem that leverages openness across the AI stack.

Doing this will lead to five tangible outcomes:

1. **Creation of sovereign capability.** This is not about building one-off models but supporting a domestic open ecosystem across the stack, with the second-order effects of increased national talent and innovation capacity.
2. **Reimagination of the state.** Governments must do more with less in an environment that is increasingly complex and integrated. Developing openness across the AI stack – and working with allied partners to do so – can help transform the state into an agile, adaptive platform capable of driving innovation at scale and pace.

3. **Generation of economic value through diffusion and adoption.** Open-source AI lowers barriers to experimentation and deployment, accelerating real-world adoption and enabling domestic firms and public institutions to capture more of the technology's economic value.
4. **Improvement of security, resilience and defence.** The widespread adoption of open tools broadens the base of users who can scrutinise, evaluate and improve or secure these systems. Government adoption of open source also reduces the risk of vendor lock-in and increases strategic optionality.
5. **Strengthening soft power and global influence.** By developing national open assets such as data sets, harnessing a strong domestic open ecosystem and contributing to international open-source projects, countries can influence the direction of AI globally. Furthermore, coordinated interoperable standards can pool purchasing power, enabling blocs of middle powers to exert greater collective influence over the global AI market.

Finally, and crucially, the case for open source holds even amid uncertainty about how the AI frontier will evolve. Whether the current capital-intensive scaling paradigm for LLMs persists or the gap between proprietary and open models widens, AI adoption will always hinge, to some degree, on efficiency, interoperability and adaptability. Openness is a catalyst for these qualities. Therefore, just as investing in data centres and compute is a robust long-term bet regardless of the form AI ultimately takes, investment in open source is a durable strategic choice because the advantages it brings underpin all plausible AI trajectories.

Recommendations

To develop a leading open-source ecosystem, policymakers should take five key actions.

1. **Establish a flagship open-source programme.** Instead of investing scarce public resources in building national LLMs from scratch, governments should establish flagship national or regional open-source

AI programmes to build the capacity to adapt, deploy and govern open models. These programmes should focus on developing the talent, infrastructure and ecosystems needed to turn any AI model into value and growth.

2. **Treat open-source tooling and maintenance as critical AI infrastructure.** Governments should focus on three things: building cost-effective public-sector tools in-house, maintaining critical open software infrastructure, and incentivising researchers to develop and sustain tools that accelerate AI-driven science. Across all three, international collaboration is essential to build and maintain shared open infrastructure.
3. **Curate strategic data sets and unlock value with pro-innovation regulations.** Governments should build the capability to curate and govern high-quality national and international data sets in priority sectors to drive downstream innovation in open models. This must be paired with pro-innovation data and copyright rules that give startups and researchers the legal certainty to fine-tune and deploy AI at scale.
4. **Shape the open-source market with government procurement.** Public purchasing power should be used to support open standards, drive startup dynamism, ensure the availability of reusable components and build a competitive open-source ecosystem that reduces dependence on single vendors.
5. **Leverage open benchmarks to drive innovation and adoption.** Governments should treat open, sector-specific benchmarks as strategic infrastructure for trust, adoption and interoperability. Developed by existing regulators and aligned internationally, shared benchmarks can open markets and strengthen collective purchasing power.

01

Capturing Value: From Models to Ecosystems

The rapid advance of artificial intelligence is forcing governments to confront a simple reality: the most powerful AI systems are increasingly central to economic growth and state capacity, yet control over them is concentrated in a small number of firms and countries. However, for most states, building a frontier model of their own is neither realistic nor necessary. The real challenge is securing reliable access to AI on terms that preserve agency: the ability to adapt and deploy AI across the economy without being locked into external platforms or losing control over where value is created. Our recent paper *Sovereignty in the Age of AI: Strategic Choices, Structural Dependencies and the Long Game Ahead* sets out this approach in full.

This paper builds on that analysis by examining how open-source AI can help middle powers – states with strong technological ecosystems, but who cannot compete at the very frontier of the model race – secure agency and capture the value of AI. It argues that open source should not be understood primarily as a means to build alternative national models, but rather as a way to shape the wider AI stack around them and to build an innovative open-source ecosystem domestically.

Terminology in This Paper

For the purposes of this paper, it is useful to distinguish between open-source AI models and openness across the wider AI stack. Open-source AI refers to open-weight models themselves. Openness across the stack refers to the software, tools, data, standards, infrastructure and documentation that make those models usable, governable and adaptable in practice. This paper is concerned not only with open models, but with the deliberate use of openness across these surrounding layers. When we refer to an "open-source strategy", we mean this broader, ecosystem-wide approach. For a more detailed discussion, see Annex A.

The case for shifting attention beyond models and towards ecosystems rests on three structural changes in how value is being captured in AI. First, a new intermediary layer between frontier models and deployment is rapidly emerging as a place where advantage is won: distillation (compressing large models into smaller ones), fine-tuning, evaluation, inference optimisation (making models faster and cheaper to run) and systems integration.[7] This intermediary layer is also where downstream innovation happens – turning models into low-latency, lower-cost systems that can run on-device or on-site in real-world settings.[8] These functions depend on shared tooling, standards and reusable components.

Second, small and specialised models are becoming central for many high-value tasks, shifting advantage from sheer scale to the ability to adapt models cheaply and close to the point of use.[9] These capabilities will diffuse fastest through open models and open infrastructure.[10]

Third, the rise of agentic systems will make interoperability and modularity decisive. As systems route across multiple models and tools, open interfaces, shared benchmarks and reusable components will become key to the agentic economy.[11]

It is therefore useful to think in terms of the AI technology "stack". While frontier models and compute infrastructure are highly concentrated, many other layers of the stack – including software frameworks, developer tools, data pipelines and deployment infrastructure – are already built on or shaped by open source. These layers play a disproportionate role in determining how quickly AI capabilities diffuse economically, how easily they can be adapted and who captures downstream value.[12]

FIGURE 1

The AI technology stack and the role of openness within it

Layer	What this layer does	Relevance of open source	Notes for policymakers
Physical compute (chips, fabrication plants)	Provides raw computational capacity	Low	Capital-intensive, geopolitically concentrated, largely closed
Cloud infrastructure	Scales compute, storage and networking	Low	Hyperscale cloud is concentrated; open standards affect portability
Hardware-acceleration software	Enables efficient use of GPUs/accelerators	Medium	Often proprietary, but open alternatives matter for resilience
Data	Training, fine-tuning and evaluation inputs	High	Open data sets and access regimes drive diffusion and advantage
Machine learning frameworks	Define, train and optimise models	High	Dominated by open source (e.g. PyTorch)
Models	Foundation and specialised model weights and code	Medium	Frontier models often closed; open models enable adaptation
Benchmarks & evaluation	Measure performance, safety, reliability	High	Critical for trust and interoperability
Inference & deployment tooling	Runs models efficiently in systems	High	Open runtimes reduce cost and lock-in
Adaptation & modification tools	Fine-tuning, distillation, specialisation	High	Key to downstream innovation
APIs & orchestration	Integrates models into workflows	High	Open interfaces reduce vendor dependence
Applications	End-user products and services	Variable	Often proprietary, built on open foundations
Safety, governance & documentation	Auditability and risk management	High	Open tools essential for scrutiny and trust

Source: TBI analysis, inspired by *Towards a Framework for Openness in Foundation Models* (Adrien Basdevant et al, 2024)[13]

A full-stack, ecosystem-wide approach has a few key strategic implications. First, because AI value increasingly sits beyond the models themselves, governments can retain real agency elsewhere in the stack: over data access, open tool-building and maintenance, procurement rules, and deployment in public and regulated sectors.

Second, investing in open-source ecosystem development and maintenance is a resilience issue, not just an innovation choice. The Log4j incident in 2021 illustrated the scale of this interdependence: a vulnerability in a small open-source logging library created cascading risks across governments, financial institutions and critical infrastructure worldwide.[14] Thus even without an explicit open-source strategy, every country's AI ambitions depend on the security and sustainability of global open-source infrastructure, leaving no government able to opt out.

CASE STUDY

Anthropic's Claude Code

A concrete illustration of this full-stack, ecosystem-wide logic can be seen in Anthropic's success with Claude Code.[15,16] Claude's growing strength as a coding assistant has been strongly supported by sustained investment in the layers around it: developer tooling, evaluation harnesses (systems for testing model performance), safety scaffolding (technical guardrails that constrain model behaviour), workflow integration and clear interfaces that enable the model to operate reliably within real software systems. Anthropic has focused on making Claude easy to adapt, test, deploy and plug into complex coding environments, so its capabilities compound through better tooling and tighter feedback loops rather than relying solely on raw model scale. The result is that Claude Code feels "smarter" in use because the surrounding ecosystem does more of the work – illustrating why value increasingly sits in the adaptable layers around models, not just in the models themselves.

For middle powers, the strategic task is therefore not to build a model, but to build a powerful national ecosystem that can leverage and build value on top of AI across the stack – and openness will be critical to that mission.

02

The Geopolitics of Open-Source AI

The fact that open-source AI has become attractive to middle powers is also what has made it a geopolitical issue. As open models continue to close the capability gap with proprietary systems – and more importantly as value creation shifts towards deployment, and government and startups around the world look to leverage open source – the US and China are increasingly viewing openness as a strategic lever to diffuse their models and supporting infrastructure globally.[17]

To understand why open source has become a geopolitical matter, it is important to recognise its long-standing role in shaping the digital economy. From Linux servers powering the cloud to widely used open-source tools such as web servers (Apache), databases (MySQL), programming languages (Python) and AI frameworks (PyTorch), the infrastructure and software supporting today's digital economy have been built on components that are collaboratively developed, transparently governed and globally reused.[18] For 40 years, open source has been the hidden foundation that has enabled the digital world.

As AI too becomes a foundational general-purpose technology, the same dynamic now applies. Shaping open-source AI is a way of influencing the foundational infrastructure on which the next phase of the global digital economy will be built – and therefore a central object of strategic competition between the US and China.

China has embraced openness as a strategic instrument, with the January 2025 release of DeepSeek-R1 disrupting the assumption that capital-intensive, proprietary models alone would shape AI's trajectory. China is now leading the open-source AI model race.[19] In recent months, Chinese models accounted for around 50 to 60 per cent of new open-model adoption, overtaking the US, which has fallen to roughly 25 to 30 per cent.[20]

Beyond AI, China has also been a historic leader in the development of and support for open-source software. Its domestic developer community, one of the world's largest, underpins major projects such as OpenHarmony, Kubernetes deployments and even open chip architectures such as RISC-V.[21]

The US, by contrast, has until recently been less supportive of open-source AI. The Biden administration's "small yard, high fence" approach to AI – emphasising AI safety regulations backed by export controls – went alongside an AI strategy focused on closed, proprietary models.[22] But the Trump administration's 2025 AI action plan marked a pivot, framing open-source AI as a strategic asset for diffusion and soft power.[23]

The calculation by both the US and China is that by exporting open models, they can align the world's middle powers with their AI systems and hardware stacks – winning global market share. There are two clear ways this alignment could emerge (which are not mutually exclusive).

First, *vertical lock-in.* This would be where alignment is achieved by exporting not just open models, but the wider technological stacks that support them. When a country adopts a new open model built on specific frameworks, hosted on specific clouds or optimised for certain chips under a government's influence, it often results in the adoption of the surrounding infrastructure, standards and regulations as well. For example, Huawei has marketed DeepSeek integration as part of its cloud offerings in countries such as South Africa and other parts of sub-Saharan Africa.[24]

Second, *horizontal lock-in.* This is where alignment is achieved by shaping the application and economic layers that sit atop AI models. As open systems spread, countries and firms may come to depend on particular model families or interoperability standards, embedding their digital economies and future agentic systems within the exporter's ecosystem logic.[25]

FIGURE 2

How vertical and horizontal lock-in could emerge

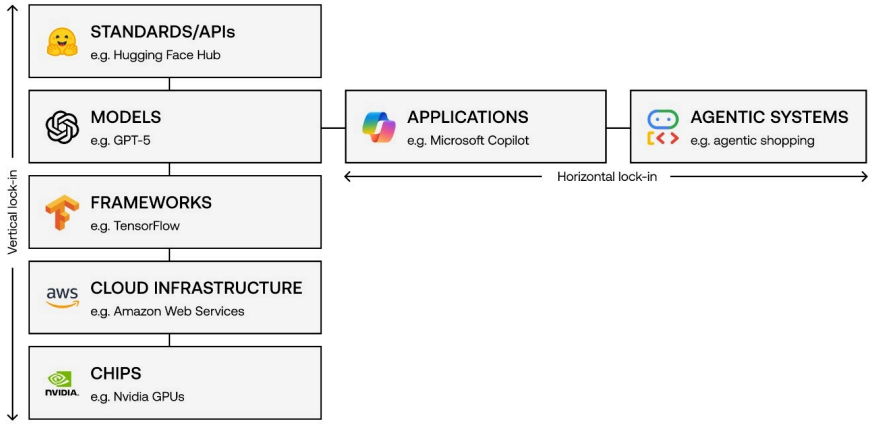

Source: TBI analysis

For policymakers, the US–China geopolitical context underlines a key point: while openness is an important strategic lever for middle powers, it is not a silver bullet for sovereignty. The diffusion of open-source models does not, by itself, de-risk countries from structural dependencies elsewhere in the AI stack – most notably, the continued concentration of hardware, cloud infrastructure and low-level software in a small number of firms and jurisdictions. Simply adopting open models, therefore, does not prevent lock-in or guarantee strategic autonomy.

However, this should not undermine an open-source strategy. Instead, it clarifies what such a strategy should prioritise. The strategic value of openness lies not in the passive adoption of open models, but in creating a national open ecosystem that can adapt and build around them. In the context of geopolitical competition around open source, only a full-stack, ecosystem-wide approach allows middle powers to capture economic value from openness without locking themselves into a single model, company or country.

Leveraging openness means being a shaper of the broader open ecosystem, not a consumer of open-source models.

The Open-Source AI Landscape in Graphs

1. Open models – where the parameters learned during training (known as weights) are publicly available – are still behind, but are catching up to the capabilities of closed models.

FIGURE 3

Open models are closing the capability gap

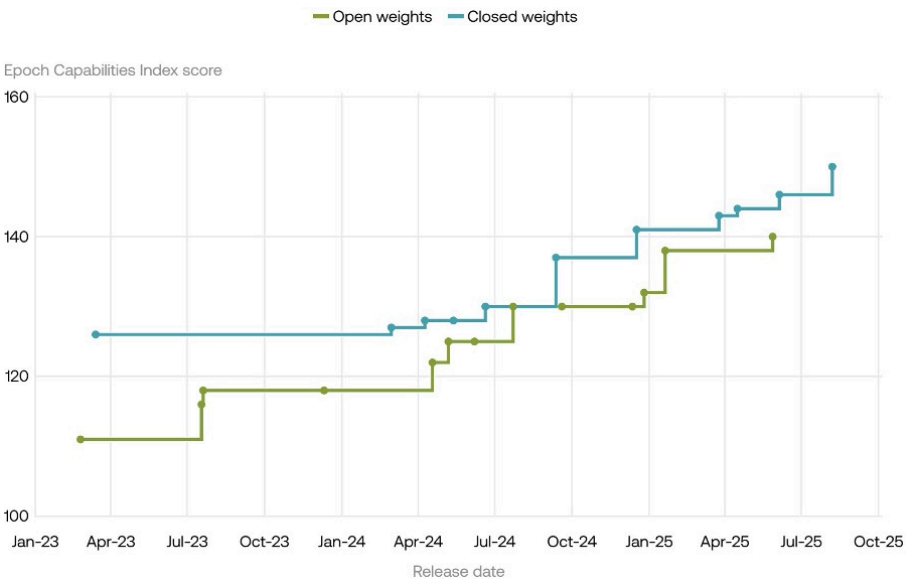

Source: Epoch AI[26]

Note: Based on the Epoch Capabilities Index score, which combines scores from many different AI benchmarks into a single "general capability" scale, allowing comparisons between models even over timespans long enough for single benchmarks to reach saturation.

2. China is increasingly taking the lead in the open-source AI race against the US.

FIGURE 4

China has overtaken the US in cumulative AI model downloads on Hugging Face

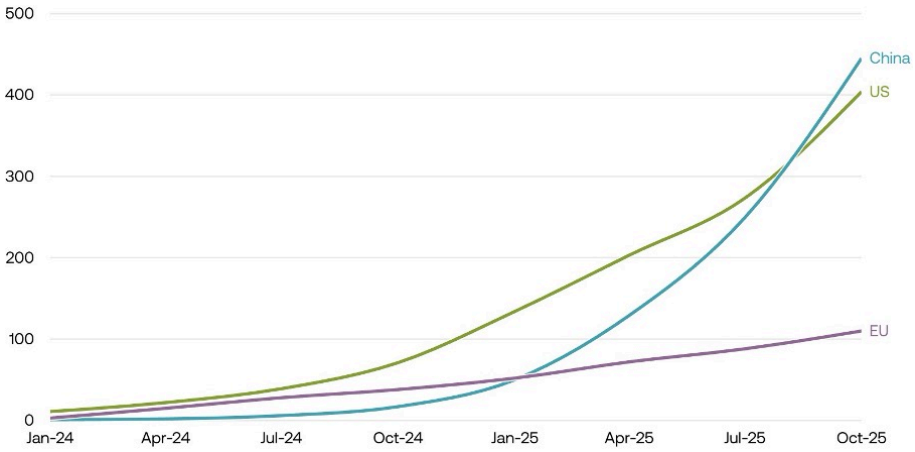

Source: The ATOM Project[27]

Note: Hugging Face is a global AI development platform centred on open-source tools and collaborative model sharing, supporting both open-weight and restricted machine-learning models.

FIGURE 5

China's Qwen has overtaken Meta's Llama on cumulative downloads of open models

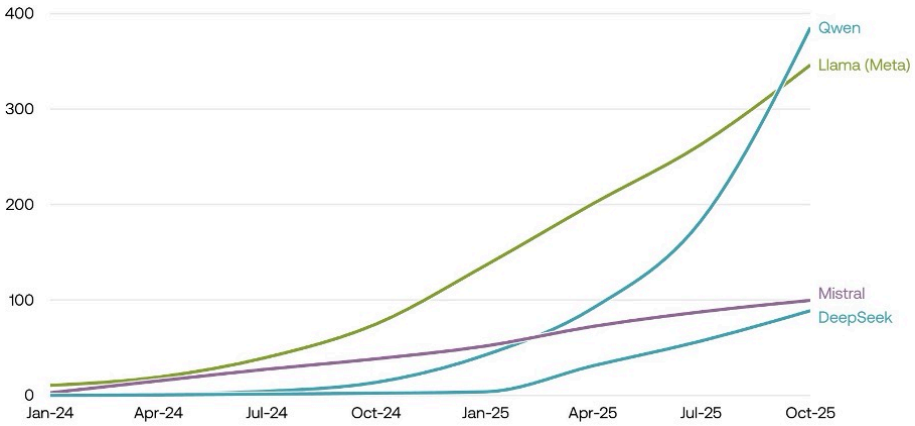

Source: The ATOM Project[28]

FIGURE 6

Qwen is also dominating new fine-tunes and derivatives on open-source models

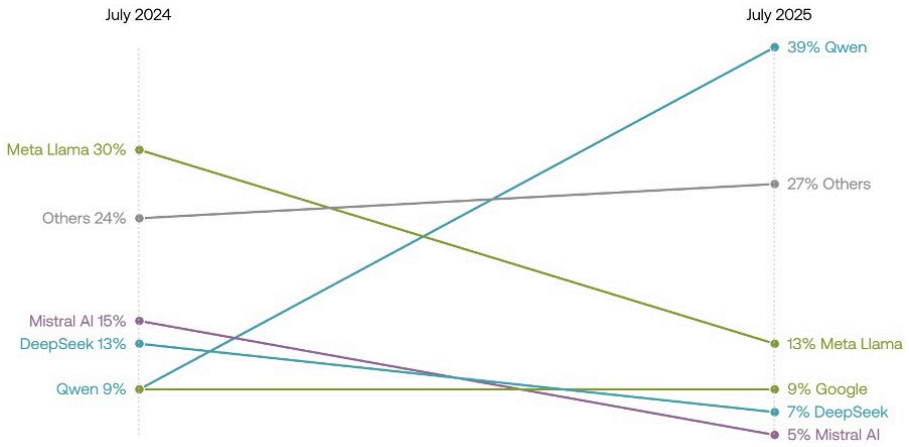

Source: The ATOM Project[29]

3. Outside the development of open-source AI models themselves, which is clearly dominated by the US and China, the picture becomes more dynamic when we look at broader open software communities and open-source model adoption. The US and China still dominate, but there is evidence of other dynamic open-source ecosystems around the world.

FIGURE 7

Several middle powers make outsized contributions to open-source software (GitHub contributions per million population, 2025)

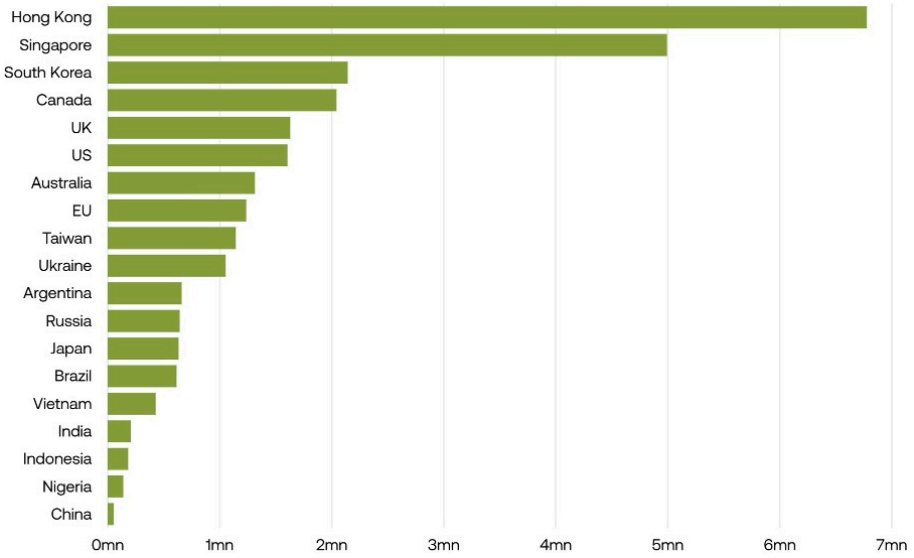

Source: GitHub[30]

FIGURE 8

For their population size, many middle powers overperform on open-model adoption

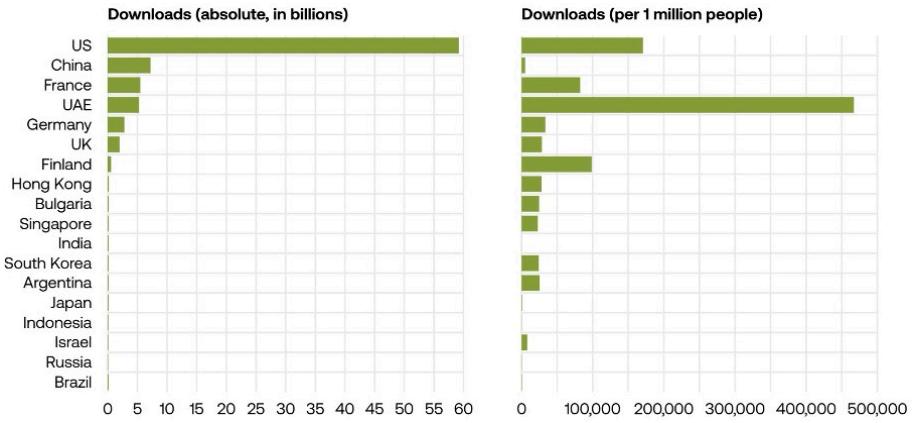

Source: Hugging Face[31]

FIGURE 9

Developer communities are also sizeable in some middle powers (Europe and Asia, developers per million, 2024)

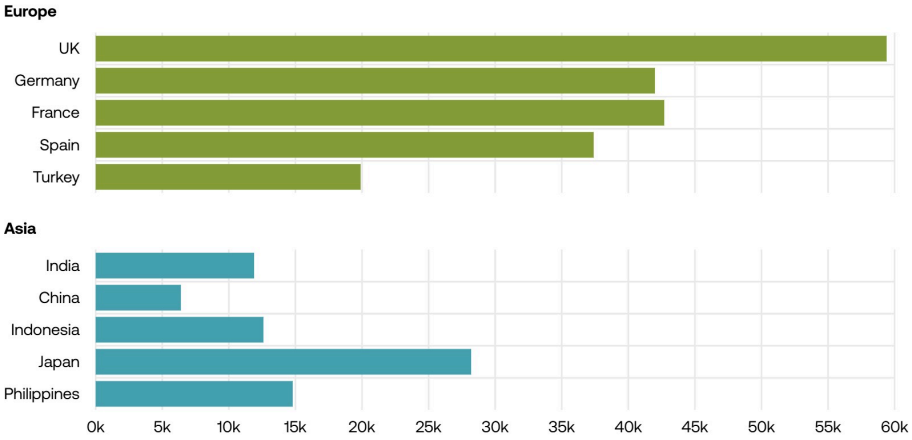

Source: GitHub[32]

03

The Advantages of an Open Ecosystem

When taken from an ecosystem-wide, full-stack perspective, a national open-source strategy delivers a set of distinct advantages for states. It does this not by competing at the frontier of AI model development, but by strengthening capability across the layers where AI is actually built and deployed.

See Annex B for our thought experiment on what a successful open strategy for a middle power looks like.

Advantage 1: Openness as a Foundation Sovereign Capability

Openness cannot deliver full AI sovereignty in a world where compute, cloud and software infrastructure remain so globally concentrated. Yet it can deliver something more realistic: *sovereign capability*.

Sovereign capability is a country's practical ability to reliably access, adapt and deploy AI systems on its own terms, even when upstream technology (chips, frontier models, hyperscale clouds) is controlled elsewhere. Most importantly, it means the ability to capture the economic value of AI domestically. It is achieved not by owning the frontier, but by sustaining full-stack domestic capability – the talent, institutions and tooling required to build, understand and improve AI across the pipeline.

Crucially, a common mistake is to assume that sovereign capability necessarily means building a national LLM. True sovereign capability is about having an ecosystem – the talent and know-how – that has the *capacity* to build on top of models and build from scratch if necessary.[33] Building a national open ecosystem of builders, rather than mere adopters, will be foundational to creating this sovereign capability and know-how.

Advantage 2: Openness as the Enabler of a Reimagined State

Governments are being asked to do more with less. They must deliver better services under tight fiscal constraints. This means being technologically innovative and adaptable, which an open approach accelerates in the following ways:

- **Adaptability**. An open approach enables governments to adapt technology to their own policy and service contexts not only by modifying tools, but by choosing between them. Open models and modular architectures give departments the capacity to evaluate, compare and decommission AI applications that do not deliver value, rather than being locked into monolithic systems. Teams can trial tools for correspondence handling, benefits processing or policy analysis, replace components as evidence accumulates, and maintain control over integration and exit over time.
- **Capability and control**. When combined with investment in technical in-house talent, openness brings technical understanding back into government. This builds internal expertise, supports better procurement decisions and reduces lock-in to a handful of external vendors.
- **The state as a platform for innovation**. Government can be an enabler, not only a user – for example, in releasing well-governed public data sets, tools and benchmarks that others can build on. The classic example is GPS: when the US opened precision signals and documentation in the early 2000s, what was originally a closed government defence asset became a general-purpose platform that catalysed entire sectors – logistics, navigation, location-based services – far beyond government itself.[34]

Beyond these specific enablers of innovation, however, there are two other second-order impacts of an open-source-first approach to government. First, a focus on building open-source AI tools offers governments the opportunity to collaborate with other like-minded governments – primarily where similar government structures and public services exist. For example,

there is a serious opportunity for European countries facing similar challenges of high debt and declining public services to collaborate on building open tools.[35]

Second, building public trust. Public trust in AI is a significant barrier to adoption in industry and in government.[36] By building open-source tools and sharing their impact and use cases, governments can bring the public on board for AI-driven public services – moving beyond abstract fears of AI towards actual applications that improve citizens' interactions with the state.

Advantage 3: Openness as a Multiplier of Diffusion and Downstream Innovation

AI creates economic value not when models are trained, but when they are widely deployed and used. Open-source AI and open software play a central role in this process by lowering barriers to access, reducing costs, and enabling experimentation and adaptation.

The economic value of openness drives value through two distinct but complementary channels:

1. **Adoption and use**: lowering the cost and friction of deploying AI across the economy.
2. **Value creation through building**: enabling firms, researchers and governments to use open models, tools and data as the baseline infrastructure from which to build new, sector- or task-specific applications and innovations.

Open-source software is already fundamental to economic growth. If no country contributed to open-source software development, GDP for the average country would be 2.2 per cent lower in the long-run.[37] In the UK, for example, open-source software contributed an estimated direct and indirect value of £46.5 billion to UK business in 2020 alone.[38]

Open-source AI models are also diffusing fast. Adoption rates are accelerating and nearly all developers have experimented with open models: 63 per cent of companies already use one and 89 per cent of AI-enabled firms rely on some open components in their stacks.[39] Small and medium-sized enterprises (SMEs), especially, are adopting open-source models faster than large corporations.

However, treating open source primarily as a cost-saving adoption strategy misses its deeper economic significance: the value lies not just in using open tools, but in building, customising and shaping them. The more transformative impact comes not from using open models as substitutes for proprietary ones, but from leveraging openness to build new applications.

AI for science sits squarely within this second channel. As outlined in our paper, *A New National Purpose: Accelerating UK Science in the Age of AI*, AI has enormous potential to accelerate R&D and drive productivity, but a network of open-source tools and data, as well as international collaboration, will be critical to unlocking the real benefits. The open release of AlphaFold 2's source code, for example, catalysed global biomedical and drug-discovery research by allowing scientists to adapt and extend DeepMind's protein-predicting tool.[40] More broadly, as AI methods grow more complex, open software stacks will be essential to avoid duplication in research and accelerate cumulative rather than fragmented progress.

This kind of downstream innovation is harder to quantify, especially where it has global, not just national, spillover effects, but it is where openness delivers its greatest economic value.

Advantage 4: Openness Driving Security and Optionality

Open-source AI is often portrayed as a security risk,[41] yet for governments it is often the opposite.[42] In many applications, openness has net security benefits.[43] Openness allows defence, intelligence and critical-infrastructure agencies to inspect, fine-tune and operate models within localised and

secure environments for sensitive areas. Access to model weights enables engineers to retrain and verify systems locally, diagnose errors and adapt models to mission-specific needs without relying on external vendors.

This is not hypothetical: open software already underpins 96 per cent of US civil and military codebases, from navy warships and Space Force satellites running on Linux-derived systems to AI-enabled F-16s powered by open orchestration frameworks like Kubernetes.[44]

More fundamentally, open ecosystems provide strategic optionality by ensuring the flexibility to diversify suppliers, architectures and standards.[45] For example, emerging interface standards such as the Model Context Protocol (MCP) show how value in open AI increasingly sits in shared interfaces and tooling that allow models to be swapped over time, rather than in ownership of any single model.[46] This widens the supplier pool, increases competition and reduces the risk of single-vendor dependence.

Of course, the security benefits of open-source AI models and open-source software do not map perfectly onto one another – and must continue to be studied. Software code can be audited line by line, but AI models remain non-deterministic black boxes where inputs do not necessarily translate into the same outputs.[47] This does not mean that open-source AI models are inherently insecure; instead, leveraging them for security purposes will require government experimentation and adaptation.

Advantage 5: Openness Underpinning Soft Power and International Collaboration

By contributing to shared codebases, tools, benchmarks and governance frameworks, middle powers can do three things.[48] First, governments can collaborate on the open tools they are developing for public services – meaning no one country must shoulder the cost of tool development. Second, increasing interoperability reduces the likelihood that single countries or companies will dominate the future *agentic* economy, which will rely on a complex, often open, stack. Finally, in cases such as benchmarking,

collaboration can enable middle powers to agree common standards, pool demand and therefore increase their leverage in the global AI market (see Recommendation 5).

For middle powers, therefore, the open-source agenda must be understood not only as national capacity building but as an exercise in international influence.

The rest of the paper sets out the five policy levers governments should pull to achieve that agenda.

Recommendations

Realising the vision set out in the previous chapters means moving from a model-centric approach to an ecosystem-focused strategy. The aim is to build domestic capability to adapt and create value on top of open models. The alternative for middle powers is to remain passive adopters of AI systems and standards shaped elsewhere, with little opportunity to harness value or influence how AI develops.

Recommendation 1: Establish a Flagship Open-Source Programme

Recommendation: *Governments should establish flagship national or regional open-source AI programmes that build domestic capability to fine-tune, evaluate and deploy models – and stimulate a vibrant open ecosystem. These programmes could take various forms, such as a flagship distilled-model project or a coordinated regional initiative. The aim is not to build a national LLM from scratch or to compete at the frontier, but to use open-source projects to develop talent, ensure cultural alignment, strengthen infrastructure and foster a broader national open ecosystem.*

Countries are looking for ways to participate meaningfully in the AI race, but are confronted with the hard realities of fiscal constraints relative to leading private laboratories and the likelihood of quickly becoming obsolete as frontier AI advances rapidly. There are two approaches to national model-building that are commonly discussed:

- **Building nationally oriented "distilled" models – smaller, more compact AIs built off the back of existing open-source foundation models.** Countries such as India, the UAE and Sweden have built distilled national LLMs that are fine-tuned to fit their bespoke language and culture. [49,50,51]

- **Building a national LLM from scratch.** Some have called for a "US DeepSeek" project, which, in turn, has been echoed by calls for a "UK DeepSeek" equivalent.[52,53] The Netherlands has also decided to follow this approach.[54]

Attempts to build a full national LLM from scratch, although offering the illusion of AI sovereignty, are rarely a sound use of public resources. Training a frontier-scale model requires billions of dollars in compute, specialised engineering talent and a continuous R&D effort – costs that even large economies struggle to meet. Any model developed from scratch through a government-led initiative would almost certainly lag behind frontier capabilities within months, while also incurring significant opportunity costs that could be better invested in data infrastructure, tool development, research or support for domestic companies. Similar initiatives to build from-scratch technology capabilities in Europe, for example, do not have a strong track record.[55]

Further, creating a national model from scratch does little to reduce structural dependencies. Countries would still rely on foreign chips, cloud infrastructure and closed-source upstream innovations. For most governments, this approach risks producing an expensive, quickly obsolete artefact rather than a lasting national capability.

Nonetheless, there are still meaningful advantages to building a distilled and fine-tuned national LLM on top of an existing open-source foundation model. However, countries have to be clear-eyed about what those advantages are – and are not.

Distilled models are not a silver bullet to AI sovereignty, will almost certainly lag the frontier and will not de-risk middle powers from the geopolitical dynamics at play. But there are benefits. First, a distilled national LLM preserves linguistic and cultural uniqueness. Though it by no means resolves the underlying problems of technological dependence, this will be important as AI continues to diffuse across the economy, culture and literature.

Moreover, distilled models can be reoriented to fit national values – which may become an increasingly important consideration as AI becomes more pervasive.

Second, and more significantly, the real value lies not in having the national LLM but in the ecosystem- and capability-building that is a byproduct of having a national flagship open-source project. Having the expertise to repeatedly train, adapt and productise open models (and document how it was done) is valuable.

In a sense, the paradox is that what is important for middle powers is not to *build* models, but to have the *capacity* to build models.

Several precedents show how this can be done in practice. For example, AI Singapore's SEA-LION project, supported by the National Research Foundation, offers a model for *regionally* focused open development, building one of the largest multilingual data sets for South-East Asian languages.[56]

Whatever approach a country takes to model-building, each should be underpinned by the same core long-term aims to build:

- **Full-stack technical expertise:** the ability within national institutions to train, fine-tune, evaluate and deploy models, so that knowledge of architectures, data pipelines and optimisation is domestically held.
- **Ecosystem depth and continuity:** a self-reinforcing base of researchers, engineers and firms who can adapt and extend open models for local use cases, ensuring progress does not depend on a single flagship laboratory or contract.
- **Data and model assets as public goods:** curated national data sets and reusable model artefacts (weights, logs, benchmarks) that lower barriers for future research and private-sector adoption.
- **Global interoperability and visibility:** participation in the global open community to influence norms on safety, licensing and evaluation, while ensuring domestic contributions are widely adopted.

To achieve this, all credible national attempts to build open-source R&D capability will need the following:

- **Strategic public compute:** dedicated resources for both **training and inference** prioritised for startups and researchers.
- **Curated national or regional data sets:** national- or domain-specific data sets under open or controlled licences.
- **Reusable open tools and standards:** shared training code, evaluation frameworks and infrastructure that ensure interoperability and reproducibility.
- **Institutional anchors:** at least one public-interest laboratory or consortium that concentrates expertise, maintains continuity and sets open-science norms.
- **Talent pathways:** fellowships, secondments and startup partnerships that embed technical practitioners within the open-model-building process.
- **Sustained funding and mission alignment:** multi-year public investment rather than short-term prestige projects.

The goal is not to build *a* model (though this may be a byproduct), but instead to build the capacity and ecosystem needed to create, adapt and understand any model.

Recommendation 2: Treat Open-Source Tooling and Maintenance as Critical AI Infrastructure

Recommendation: *Governments should make building and maintaining open-source tools a cornerstone of their national AI strategies, recognising that both value and resilience lie not only in AI models themselves, but also in the software layers that underpin them, and the tools and applications that enable real-world usage.*

AI will only be as helpful as the software ecosystem it stands on. In a simplified sense, AI's value depends on two layers of tooling: the underlying software that makes models run reliably, and the higher-level tools that adapt and integrate them into real workflows. That is why open source matters well beyond the models themselves.

Governments therefore need to make open-source tool-building a core component of their AI capabilities – especially for three core areas:

1. Tool maintenance for critical infrastructure, which requires continual monitoring and funding.
2. Tool-building for government, which requires in-house technical capability.
3. Tool-building for science, which requires rewiring academic incentives to make the creation of reusable tools an integral part of research.

CASE STUDY

Scikit-Learn and the Power of Tooling

Scikit-learn illustrates how durable economic value in AI often emerges from open tooling rather than frontier models.[57] Originating in France's public-research ecosystem (the National Institute for Research in Digital Science and Technology) before becoming a global open-source project, it became foundational by standardising workflows for the training, evaluation and deployment of models that run efficiently on modest hardware. Its success lies in making machine learning usable in real-world settings – in SMEs, public services and regulated environments – without dependence on hyperscale compute or proprietary platforms. Crucially, scikit-learn's maintainers have successfully made the case to the French government that widely used "everyday" machine learning tools – downloaded over 3.5 billion times – are as strategically important to fund as newer, more headline-grabbing AI models, and that public support should cover long-term maintenance and security as well as new features. Today, with over 3.8 billion downloads in total and 150 million downloads per month, scikit-learn underpins applications across health care, finance and government worldwide. The scikit-learn library and wider ecosystem are stewarded by Probabl, the mission-driven company focused on developing and sustaining a complete suite of open-source tools for data science.[58]

TOOL-MAINTENANCE FOR CRITICAL INFRASTRUCTURE

All governments are exposed to critical open-source infrastructure whether they know it or not. Funding open-source maintenance must therefore be a core part of a nation's open-source strategy. Robust third-party risk

management is essential to map and monitor dependencies, set clear security and continuity standards for suppliers, and ensure that a disruption or compromise in one vendor's stack does not cascade into a systemic failure of critical digital services.[59]

There are various ways governments can mitigate critical open-source software risks. The three most feasible are:

1. **Establish a national open-source trust fund** to provide sustainable funding for software maintenance or, in some cases, for building brand new tools where this is deemed to be of security or strategic importance. Germany has already done this with the Sovereign Tech Agency. Governments could also consider pairing this with an investment strategy, whereby the government invests in national open-source companies and takes a share of the profit – either directly or through a national wealth fund or alternative vehicle. Generally, this should focus on maintenance rather than building brand new software, because higher returns are usually seen in investment in existing critical, but underfunded, tools.[60]
2. **Establish a national programme of red-teaming competitions to identify systemic software risks in critical digital infrastructure.** Governments should run recurring, structured exercises targeting deep layers of the software stack – especially open components underpinning public-sector systems – and publish the resulting vulnerability data. This should support contingency planning, provide ongoing assessments of the health of key open-source maintenance communities, and be coordinated with allied nations to surface cross-border risks.
3. **Work with international allies** to maintain shared software infrastructure that underpins major applications both in the government and the private sector – as has been outlined in the proposed European Sovereign Tech Agency.[61]

CASE STUDY

Maintaining Critical Open-Source Infrastructure: The German Sovereign Tech Agency

Germany's Sovereign Tech Agency provides a pragmatic model for how governments can support the open components that underpin their digital services.[62] One core pillar is the Sovereign Tech Fund, which directs long-term funding to foundational libraries, protocols and tools that are widely relied on but chronically undermaintained. Projects are selected based on systemic importance, security relevance and clear evidence of underinvestment.

The agency complements this with targeted resilience work such as engineering contributions, coordinated bug-fix efforts, and independent code audits for high-impact components. Its fellowship programme also provides structured support to the maintenance community that keeps this infrastructure running. The approach is simple but effective (and, crucially, flexible) – identify critical dependencies, fund their upkeep, and reinforce the small set of people (often "lone-wolf" maintainers) and projects that the wider system rests on.

TOOL-BUILDING FOR GOVERNMENT

An "open-first" approach to using AI to build capability in government and public services is essential for a few reasons. First, an open-source approach enables governments to experiment with and build AI-enabled applications that are more specific to their needs at much lower cost.[63] Second, because these tools are open-sourced, more individuals in the

ecosystem are likely to spot issues with the tools, increasing the likelihood of reliability. Estonia, for instance, opens its voting software every election to improve resilience and security.[64]

Third, critically, open building gives like-minded countries the opportunity to collaborate in tool-building, increasing the power-to-weight ratio of fiscally constrained governments. And finally, by sharing and being more open about the AI tools being used in government and public services, governments are more likely to build trust and secure buy-in from their own citizens – rather than AI experimentation being viewed as a black box.

The UK's Incubator for Artificial Intelligence (i.AI) offers a useful model for governments to build practical AI capability in-house.[65] It combines a small, technically skilled team with an open-first approach, producing reusable components that other departments can adopt and adapt. This reduces costs, avoids duplicated effort and creates transparency around how AI is being used in public services.

But i.AI's experience also shows that certain disciplines are essential for this model to work. Redbox, one of its early tools, demonstrated the need for continuous validation of user needs. Although designed for complex document analysis, most demand was for basic secure LLM access.[66] It also showed the value of prioritising "good enough" capability to reach users quickly, rather than over-optimising for niche features. As the wider market evolved and tools improved, Redbox has since been retired – with the Incubator focusing on other areas where there is real value add.

These lessons show that in-house, open-source capabilities can deliver real value, but only when paired with disciplined product management, strong technical skills and a willingness to iterate or pivot as conditions change.

As part of building in-house capability, it is also important that governments incentivise talent that would otherwise be drawn to the private sector. Of course, part of this talent diversion is due to salaries. But a frequently missed point is also *recognition*: companies such as Google and Meta enable open-source developers to take credit for their tools – which is a key pull factor for

those who wish to build a name for themselves. Governments should consider building similar recognition-focused incentives structures and agency to attract in-house tool-builders.

SOFTWARE TOOL-BUILDING FOR SCIENCE

The goal of an open-source strategy is about more than improving government capability and resilience; it is also about powering innovation within the broader economy. Nowhere is this potential more evident than in using fine-tuned open-source models to accelerate AI-driven research. The US recently highlighted this in a new executive order announcing the launch of the "Genesis Mission", which intends to automate scientific research and discovery.[67]

If governments want AI to accelerate their R&D bases, tool-building will be a key part of that. The right incentives need to be in place for researchers to build, share, maintain and scale reusable AI-enabling tools. This includes both the "first-mile tools" that, for example, are needed to curate new, domain-specific data sets and the "last-mile tools" that make AI useful for specific applications.

But in many cases these incentives do not exist. As outlined in *Accelerating UK Science in the Age of AI*, foundational AI tools originate in universities but are abandoned because maintenance and engineering work earn little credit in current academic structures.

Governments, therefore, need to treat tool development as a core research activity, incentivising it by funding long-term maintenance, recognising open-source contributions in grants and promotions, and embedding collaborative engineering roles in major research programmes.

This would turn publicly funded science into a continuous source of shared infrastructure that accelerates AI-enabled R&D – driving long-term productivity and potentially building national comparative advantage in AI-for-science domains.

Recommendation 3: Curate Strategic Data Sets and Unlock Value With Pro-Innovation Regulations

Recommendation: *Governments should invest in curating and utilising high-quality, publicly available national or international data sets for priority research areas and sectors to drive open downstream innovation and build sector-specific comparative advantage in AI. At the same time, governments should avoid imposing burdensome regulations on data use and access that could lock out researchers and startups seeking to fine-tune open-source models.*

For middle powers, combining their strategic data assets – both open and closed – with AI offers a route to compete in the longer-term AI game of downstream adoption rather than the short-term race for frontier models. As other parts of the AI stack, such as compute, remain concentrated in a handful of global firms, national or regional advantage can be built by building and applying high-quality open data sets in strategic sectors. There are two reasons for this:

1. Data quality – not model scale – is increasingly becoming a key bottleneck to AI progress.[68] With open web data largely exhausted, the next leap in capability will likely come from narrower models fine-tuned on high-quality, domain-specific data sets.
2. The best opportunity for middle powers to pursue comparative AI advantages lies downstream in applying AI to domain-specific sectors and problems, where well-curated data sets enable fine-tuning and innovation that general-purpose frontier models cannot match.

To ensure data advantage can be leveraged alongside an open ecosystem, there are two key things governments need to get right. First, governments should have an active strategy to build and maintain national public data sets that can be leveraged by the wider open ecosystem to drive new innovations, solutions and companies. Second, governments need to ensure that their broader regulations around data use and access – especially concerning copyright – do not stifle the open-source ecosystem.

BUILDING NATIONAL DATA SETS

To create new, open national data sets, governments should follow some core principles. As outlined in our paper *Governing in the Age of AI: Building Britain's National Data Library*, these core principles are:

- **Start with demand, not supply (for the most part).** Identify where real user demand exists across research institutions, public services or industry, and which specific data sets or linkages would unlock value. Understanding who needs the data and why is the foundation for legitimacy and potential co-investment if private-sector buy-in is to be obtained. That said, this should not preclude space for experimentation: some of the most valuable data assets only reveal their demand over time, just as the full economic value of GPS only emerged after it was opened up by the US government rather than being fully understood in advance.
- **Prioritise one or two strategic data sets.** Overly broad ambitions risk paralysis. Focus on sectors aligned with national missions such as health, energy, logistics, supply and demand for public services, building, transport or manufacturing, where early partnerships and visible results can demonstrate proof of concept and attract further investment. If fiscal or political capital is limited, concentrate on making one data domain work seamlessly end to end. A functioning, trusted exemplar will build confidence and set operational standards more effectively than a diffuse national project.
- **Adopt a federated, not centralised, model.** Link data where they sit, enabling secure and interoperable access rather than wholesale aggregation. The idea is not to create a giant data lake that centralises all government-held data in one place, but to ensure that the infrastructure, governance and access controls are in place for scalable data sharing.[69] A federated approach preserves data ownership within departments, maintains clear data lineage and auditability, reduces the risk of proliferating out-of-date or conflicting copies, and avoids the security and political challenges of creating a single monolithic repository.

- **Avoid "home bias" licensing regimes:** Restricting a data set to domestic users may seem attractive, but it ultimately limits its utility by shrinking the user base and discouraging the development of complementary tools and standards that make the data set valuable. A global-access approach is typically far more effective. First, it maximises scientific and economic impact by encouraging wider use and experimentation. Second, it attracts investment, partnerships and talent that help create a self-reinforcing domestic innovation ecosystem. Third, it still preserves national advantage, as domestic institutions retain natural benefits through proximity and collaboration networks. However, there may be instances (for example, health data) where it does make sense to have home-bias licensing regimes – either for privacy reasons or for purposeful national advantage.

The UK's recently announced OpenBind consortium represents a concrete example. It aims to generate a data set of more than 500,000 experimentally validated protein-ligand structures – about 20 times the amount available over the past 50 years – for AI-driven drug discovery.[70] Larger-scale examples include the UK Biobank, the world's largest genomic and health database, which has the potential to accelerate AI-driven health care.[71]

International collaboration on certain data sets should also be seen as a strategic opportunity, not a concession of national advantage. Some of the highest value AI data sets will not be producible by one country alone. Collaborative efforts such as the Protein Data Bank – the data set that underpinned Google DeepMind's major AlphaFold breakthrough – are cases in point.[72]

Yet not every data initiative needs to be a formal data set or "National Data Library". There are more systematic ways not just to create static data sets, but also to incentivise continuous, long-term national data creation.

In some areas, for example, the challenge may be to encourage the digitisation and structuring of existing data. The task here becomes more about incentives. For example, in scientific research, vast volumes of "dark data" remain buried across universities and research institutions:

unpublished experimental results, null findings and technical notes that stay in handwritten laboratory books.[73] This represents an enormous, untapped resource for AI. Creating incentives to capture and digitise this type of data within the national research base would help develop new training material for scientific AI models and catalyse domestic R&D ecosystems.

Such an approach would effectively make the national academic research base a continuous, strategic data asset – a far more radical step than building a singular one-off data set.

PRO-INNOVATION DATA REGULATION

Leveraging data as a strategic asset is not just about building things. Regulatory blockers must also be removed. Heated debates over AI and copyright regulation are a case in point. Although the AI and copyright debate is often framed as Big Tech versus creatives, especially in places such as the UK,[74] the reality is that the true losers of excessively strict data regulation are small players such as startups, researchers and entrepreneurs who leverage open source to create jobs and the next generation of leading global companies.

Well-intentioned liability, data-protection and copyright rules can unintentionally stifle the open-source ecosystem. This is for two reasons. First, because the very companies that try to take a transparent, open approach are most exposed to legal risk. Their openness makes them easier targets for copyright or data-use claims. Stability AI, for example, was sued by Getty Images and several US artists because, by releasing models trained on publicly available data, it made its methods and materials visible, and therefore litigable.[75]

Second, strict data and copyright regulations squeeze out the very people driving the open ecosystem: researchers and startups looking to challenge incumbents through downstream innovation who lack the financial or legal firepower to absorb licensing costs.

For open ecosystems to thrive, as our report *Rebooting Copyright: How the UK Can Be a Global Leader in the Arts and AI* argued, data-use laws must provide both legal certainty and technical feasibility: pragmatic opt-out tools that can actually be implemented rather than imposing URL-level reporting burdens on small developers.

The key is therefore for governments to build data infrastructure ambitious enough to drive comparative advantage and innovation, yet be sufficiently light touch on regulation to enable their national open ecosystem to thrive.

Recommendation 4: Shape the Open-Source Market With Government Procurement

Recommendation: *Governments should leverage procurement as an industrial-strategy tool to grow an open ecosystem, build state capacity through digitalisation and drive domestic local technology. Public purchasing power should be used to shape markets by creating incentives for open standards, driving startup dynamism and building a competitive open-source ecosystem that reduces dependence on single vendors.*

Procurement is one of the most powerful yet often overlooked levers in national AI strategies. The way in which governments buy technology shapes both market structure and their own technical capabilities. In the context of open-source AI, it serves two critical functions. First, as an industrial-strategy tool: directing public-sector demand towards open and interoperable systems can stimulate innovation, lower entry barriers and reduce dependency on a few proprietary incumbents. Second, as a capability-building tool: deploying and maintaining open systems in government enhances internal technical literacy, data infrastructure and engineering capacity, ensuring the state can use and govern AI (see Recommendation 2).

Underpinning governments' approach to promoting open-source software and AI should be four key tenets:

1. **Invest in capability rather than rigid targets.** Instead of setting arbitrary quotas for open-source procurement, focus on building technical expertise within departments. Appoint open-source AI technical champions in major ministries to identify viable use cases for open tools and software.
2. **Develop open-source AI procurement guidelines.** Provide practical evaluation frameworks that, first, explain the different tiers of "openness" (especially for AI models) and where having certain levels of openness is important for government applications (for example, does open data access matter or will APIs do?); and second, weigh the cost benefit of open versus proprietary systems across applications – acknowledging that open source is not always cheaper or better depending on different contexts.
3. **Map and regularly review the national AI and open-source industrial base.** Assess dependencies, vendor concentration and supply-chain resilience to identify where consolidation could threaten innovation or procurement agility – and speak to open-source players in industry.
4. **Prioritise interoperability and competition.** Procurement frameworks should prioritise modularity and open standards, ensuring contracts remain competitive and adaptable over time.

To do this, government procurement must be made suitable for SMEs, as these are the vendors driving open-source products. The key is to avoid having procurement models that lock SMEs out due to prescriptiveness or drawn-out contracts.

In particular, challenge-based procurement – which involves the government posing a problem statement to be solved – ensures that contracts are end-use driven rather than specification-led (the latter tends to shut out SMEs).[76,77] Similarly, modular contracts where, rather than using long contract timelines, teams break down each procurement exercise into separate stages, would also drive a more competitive, SME-led open-source procurement framework approach.

Crucially, however, openness is not a blanket solution. Evidence suggests that regulations or mandates requiring quotas for the adoption of open source in government tend not to work.[78] Instead, this must be undertaken on a case-by-case, experimental basis.

In each government use case, the trade-offs between open source and proprietary procurement can be boiled down to:

- **Performance gaps:** There is still a gap between open-source AI and proprietary model capabilities. Not all government applications will need cutting-edge capability, but some may.
- **Cost considerations:** While open models are free to access, they require significant technical resources for fine-tuning and inference to make them useful. For some generalised applications, off-the-shelf procurement will be more cost-effective.
- **Technical capacity requirements:** Successful implementation requires expertise in model evaluation, fine-tuning and ongoing maintenance. Not all government departments or specific use cases will have access to the in-house capability necessary to build open-source AI tools.

These caveats should not undermine the benefits of openness, but the point is that an approach promoting open-source procurement should not deter off-the-shelf purchasing for the sake of it. Procurement should be a strategic instrument, not a compliance exercise.

Recommendation 5: Leverage Open Benchmarks to Drive Innovation and Adoption

Recommendation: *Governments should prioritise the development of openly available benchmarks that build trust, enable adoption in high-stakes sectors and create interoperability across markets. To do this, governments should leverage existing sector-specific regulators and bodies – and work internationally to ensure global interoperability that opens markets. Aligned*

countries should also consider leveraging interoperable benchmarks to exercise collective purchasing power, giving them demand-side influence over the direction of global AI systems.

Rather than viewing benchmarks primarily as tools for safety, policymakers should recognise them as an important part of openness in the broader stack that enable:

- Widespread adoption of AI applications, especially in high-stakes sectors.
- The interoperability of global markets, enabling smaller companies building specific AI applications to scale.
- Middle powers to work together to increase their purchasing power for AI applications and therefore their leverage in the global AI system.

BENCHMARKS AS A FOUNDATION FOR ADOPTION

The economic benefits of AI will only be realised if organisations trust the models they are using. Yet as our recent polling in *What the UK Thinks About AI: Building Public Trust to Accelerate Adoption* highlighted, trust is currently a key barrier to AI adoption.

In high-stakes sectors, for example, generic performance benchmarks are not enough to reassure potential adopters. A health service assessing AI for cancer diagnostics or a financial regulator evaluating fraud-detection tools cannot rely on broad benchmarks that bear little resemblance to real operational environments.

This is why the development of sector-specific performance benchmarks is essential – and why countries should rely on existing regulators, rather than creating new, general-purpose "AI regulators". The expertise needed to design meaningful benchmarks already sits in sectoral bodies: medical colleges, financial conduct authorities, civil aviation regulators, environmental agencies and others.

Drawing on that expertise ensures benchmarks are:

- **Construct-valid**, meaning they measure the capabilities that actually matter in practice.
- **Aligned with existing regulatory regimes**, lowering compliance friction.
- **Openly shared**, so developers and regulators can evaluate models consistently.
- **Regularly updated**, as practices, risks and technologies evolve.

INTERNATIONAL COLLABORATION TO OPEN GLOBAL MARKETS AND UNLOCK COLLECTIVE PURCHASING POWER

Benchmarks determine not only whether AI is trusted and adopted domestically, but how easily AI applications can move across jurisdictions. Without shared or comparable benchmarks, each country risks defining its own bespoke evaluation requirements. This fragments markets and tilts the playing field towards the largest incumbents, who can afford to tailor systems to multiple regulatory regimes.

Shared, open benchmarks ensure that:

- Domestic firms can compete on transparent, open criteria rather than vendor-defined tests.
- New national champions can emerge in downstream applications, not only upstream model development.
- SMEs, public bodies and open-source developers can integrate AI without navigating fragmented compliance regimes – and therefore scale.
- Purchasing decisions in government and industry favour genuine performance and reliability.

For middle powers, interoperable benchmarks also create the foundations for collective purchasing power. When governments converge on shared standards for evaluating models and applications, they also align the criteria against which they buy them. This enables middle powers – individually small but collectively significant – to act in concert. Instead of each country procuring bespoke systems under different national tests, interoperable

benchmarks allow for joint procurement frameworks, whether at the scale of Europe, the Commonwealth or coalitions of aligned states. This dramatically increases buyer leverage in negotiations with large model providers and infrastructure vendors, ensuring that smaller states retain meaningful agency in shaping the AI ecosystem – but by buying, not selling.

If middle powers want to punch above their weight in the next decade, then using open, collective benchmarks is a powerful lever to build national champions and aggregate demand-side heft.

Conclusion

The capabilities of frontier AI systems will continue to advance. As they do, the resources required to sustain this progress will grow far beyond the means of most governments. Leaders must therefore work to secure access to frontier AI – but for most countries, competing directly at the frontier is neither effective nor realistic. As we have argued in previous reports, the more viable strategy is to build the ecosystems, infrastructure and capabilities required to deploy AI at scale.[79]

Yet investments in hardware and infrastructure alone will not amount to a comprehensive AI strategy nor capture long-term value. The existence of a strong open-source ecosystem is a prerequisite for any country that wishes to harness the opportunities offered by AI. Open source will play a key role in driving innovation by lowering costs and, more importantly, accelerating the downstream innovation which will be the true determinant of AI's value.

As this paper sets out, building this ecosystem will require middle powers to sustain investment across several concrete pillars: targeted open-source programmes that prioritise ecosystem development over model-building prestige; tool-building and maintenance for government, critical infrastructure and science; high-quality national data sets; procurement that rewards openness and interoperability; and interoperable sector-specific benchmarks that accelerate adoption and open markets.

In the long term, the AI race will not be determined by who builds the best model. It will be determined by who builds the best ecosystem.

Annex A: Defining Open Source

For more on how we are defining language used in this paper, please see Annex A.

Annex B: Thought Experiment

For our thought experiment on what a successful middle-power open-source strategy might look like, please see Annex B.

Acknowledgements

We would like to extend our thanks to the experts who offered their advice and guidance in the development of this report (while noting that contribution does not equal endorsement of all the points made and does not reflect the views of respective employers).

- Arno Amabile – Advisor of the French President's Special Envoy for AI
- Dean Ball – Foundation for American Innovation
- Jennifer Barth – OpenUK
- Michael Boyce – US Digital Response
- Amanda Brock – OpenUK
- Ben Brooks – Berkman Klein Center
- Rory Byrne – University of Cambridge
- Hilary Carter – Linux Foundation
- Henry de Zoete – Saïd Business School
- David Eaves – Institute for Innovation and Public Purpose, UCL
- Max Fenkell – Scale AI
- Martin Goodson – Evolution AI
- Emmie Hine – Safe AI Forum
- Luukas Ilves – Ukraine Ministry of Digital Transformation
- Shayne Longpre – MIT and Data Provenance Initiative
- Emma Merrill – OpenTeams
- Cailean Osborne – Probabl AI
- Matthew Pickering – Well-Typed
- Aris Richardson – RAND School
- Maeve Ryan – Meta
- Imanol Schlag – ETH AI Center

- Niclas Stoltenberg – Amplyfi
- Paul Teather – Amplyfi
- Lauren Wagner – Abundance Institute
- Jonny Williams – Red Hat
- Minghui Zhou – Peking University

Contributor: Graham Drake, TBI

Endnotes

1 https://jeffreyjding.github.io/documents/Diffusion%20Deficit%20working%20paper%20August%202022.pdf

2 https://epoch.ai/data-insights/open-weights-vs-closed-weights-models.

3 https://demos.co.uk/research/the-open-dividend-building-an-ai-openness-strategy-to-unlock-the-uks-ai-potential/

4 We use the term "open-source strategy" to refer more broadly to openness across the stack.

5 https://digitaleconomy.stanford.edu/news/the-turing-trap-the-promise-peril-of-human-like-artificial-intelligence/

6 https://www.siliconcontinent.com/p/the-smart-second-mover

7 https://arxiv.org/abs/2512.03073

8 https://arxiv.org/abs/2108.07258

9 https://newsletter.semianalysis.com/p/on-device-ai-double-edged-sword

10 https://arxiv.org/pdf/2506.02153

11 https://market.us/report/agentic-ai-market/

12 For a more detailed breakdown of the AI "stack" and where openness fits in, see the large body of existing literature including: https://arxiv.org/pdf/2405.15802; https://demos.co.uk/research/the-open-dividend-building-an-ai-openness-strategy-to-unlock-the-uks-ai-potential/; https://arxiv.org/abs/2403.13784.

13 https://arxiv.org/pdf/2405.15802. This table is inspired partly by the original framework from Adrien Basdevant et al, *Towards a Framework for Openness in Foundation Models: Proceedings from the Columbia Convening on Openness in Artificial Intelligence* (2024). The "High / Medium / Low" categorisation indicates the relative potential for open-source approaches to shape capability, resilience and value creation at each layer of the AI stack. It is not a measure of how open or closed these layers currently are, which varies widely in practice.

14 https://builtin.com/articles/log4j-vulerability-explained

15 https://www.anthropic.com/engineering/advanced-tool-use

16 https://www.startuphub.ai/ai-news/ai-research/2025/anthropics-claude-code-plugins-open-the-floodgates

17 https://epoch.ai/data-insights/open-weights-vs-closed-weights-models

18 https://www.atlanticcouncil.org/in-depth-research-reports/report/open-source-software-as-infrastructure

19 https://www.bbc.co.uk/news/articles/c86v52gv726o

20 https://atomproject.ai/

21 https://www.economist.com/business/2025/06/17/why-china-is-giving-away-its-tech-for-free

22 https://www.cnas.org/publications/commentary/goodbye-to-small-yard-high-fence

23 https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf

24 https://hai.stanford.edu/assets/files/hai-digichina-issue-brief-beyond-deepseek-chinas-diverse-open-weight-ai-ecosystem-policy-implications.pdf

25 https://arxiv.org/pdf/2505.04345

26 https://epoch.ai/data-insights/open-weights-vs-closed-weights-models

27 https://atomproject.ai/

28 https://atomproject.ai/

29 https://atomproject.ai/

30 https://github.blog/news-insights/octoverse/octoverse-2024/

31 https://huggingface.co/spaces/economies-open-ai/open-model-evolution

32 https://github.blog/news-insights/octoverse/octoverse-2024/

33 https://press.airstreet.com/p/sovereign-ai-paradox

34 https://www.skyatnightmagazine.com/space-science/how-global-positioning-system-gps-changed-world

35 https://institute.global/insights/tech-and-digitalisation/europe-in-the-age-of-ai-how-technology-leadership-can-boost-competitiveness-and-security

36 https://institute.global/insights/tech-and-digitalisation/what-the-uk-thinks-about-ai-building-public-trust-to-accelerate-adoption

37 https://link.springer.com/article/10.1007/s10961-023-09993-x

38 https://openuk.uk/press-releases-posts/open-source-software-contributed-an-estimated-46-5bn-to-uk-business-in-2020-according-to-openuk/

39 https://www.linuxfoundation.org/hubfs/LF%20Research/lfr%5Fmarket%5Fimpact%5F052025a.pdf?hsLang=en

40 https://www.nature.com/articles/d41586-024-01555-x

41 https://cfg.eu/ai-governance-challenges-part-3-proliferation/

42 https://www.newamerica.org/oti/press-releases/open-source-ai-models-are-not-inherently-security-risks-but-they-are-integral-to-democracy-states-new-oti-report

43 https://www.rstreet.org/research/mapping-the-open-source-ai-debate-cybersecurity-implications-and-policy-priorities/

44 https://arxiv.org/pdf/2408.10026#:~:text=Preliminary%20evidence%20suggests%20that%20an%20open%20foundation%20model,impacts%20

45 https://www.kellton.com/kellton-tech-blog/why-vendor-lock-in-is-riskier-in-genai-era-and-how-to-avoid-it

46 https://modelcontextprotocol.io/docs/getting-started/intro

47 https://link.springer.com/content/pdf/10.1007/s12559-023-10179-8.pdf

48 https://demos.co.uk/wp-content/uploads/2025/06/The-Open-Dividend%5FReport%5F2025.ac-2.pdf

49 https://pinnacleiit.com/blogs/current%5Faffairs/bharatgpt/

50 https://www.tii.ae/news/uaes-technology-innovation-institute-launches-open-source-falcon-40b-large-language-model

51 https://www.ai.se/en/project/gpt-sw3

52 https://www.interconnects.ai/p/the-american-deepseek-project

53 https://www.spectator.co.uk/article/how-deepseek-can-help-britain/

54 https://www.tno.nl/en/newsroom/2023/11/netherlands-starts-realisation-gpt-nl/

55 https://www.politico.eu/article/chaos-and-infighting-are-killing-europes-grand-cloud-project/

56 https://aisingapore.org/aiproducts/southeast-asian-languages-in-one-network-data-seald/

57 https://dl.acm.org/doi/10.5555/1953048.2078195

58 https://sifted.eu/articles/probabl-open-source-13m-seed

59 https://www.oii.ox.ac.uk/news-events/reports/government-resilience-in-the-digital-age/

60 https://peekaboo-vision.blogspot.com/2020/01/dont-fund-software-that-doesnt-exist.html

61 https://www.thesixthfield.com/p/germany-sovereign-tech-agency

62 https://www.sovereign.tech/

63 https://digital-strategy.ec.europa.eu/en/news/thematic-roadmap-open-source-and-inputs-common-trust-principles

64 https://www.cambridge.org/core/journals/contemporary-european-history/article/free-to-choose-evoting-infrastructure-and-the-origins-of-estonias-digital-republic/4A59FC654A185D9B8C3D3BBB6B5AC793

65 https://ai.gov.uk/

66 https://ai.gov.uk/blogs/redbox-reflections-5-key-lessons-from-building-and-sunsetting-our-government-ai-chatbot/

67 https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission/

68 https://www.forbes.com/sites/kolawolesamueladebayo/2025/02/25/why-proprietary-data-is-the-new-gold-for-ai-companies/

69 https://institute.global/insights/tech-and-digitalisation/governing-in-the-age-of-ai-building-britains-national-data-library

70 https://www.gov.uk/government/news/uk-to-become-world-leader-in-drug-discovery-as-technology-secretary-heads-for-london-tech-week

71 https://www.gov.uk/government/news/uk-backed-ai-companies-to-transform-british-cancer-care-and-spark-new-drug-breakthroughs

72 https://www.rcsb.org/

73 https://institute.global/insights/tech-and-digitalisation/a-new-national-purpose-accelerating-uk-science-in-the-age-of-ai

74 https://www.leeandthompson.com/news/long-read-big-tech-vs-creators-whose-side-is-the-uk-government-on-when-it-comes-to-ai/

75 https://jipel.law.nyu.edu/andersen-v-stability-ai-the-landmark-case-unpacking-the-copyright-risks-of-ai-image-generators/

76 https://www.diu.mil/work-with-us

77 https://www.gov.ca.gov/wp-content/uploads/2019/01/1.8.19-EO-N-04-19.pdf

78 https://digital-strategy.ec.europa.eu/en/library/study-about-impact-open-source-software-and-hardware-technological-independence-competitiveness-and

79 https://institute.global/insights/tech-and-digitalisation/sovereignty-security-scale-a-uk-strategy-for-ai-infrastructure

TONY BLAIR
INSTITUTE FOR
GLOBAL CHANGE

Follow us

facebook.com/instituteglobal

x.com/instituteGC

instagram.com/institutegc

General enquiries

info@institute.global