Institute for Public Policy Research

ACCELERATION IS NOT A STRATEGY

A FRAMEWORK FOR DIRECTING AI TOWARDS PUBLIC VALUE BEFORE IT'S TOO LATE

Roa Powell and
Carsten Jung

April 2026

ABOUT IPPR

IPPR, the Institute for Public Policy Research, is an independent charity working towards a fairer, greener, and more prosperous society. We are researchers, communicators, and policy experts creating tangible progressive change, and turning bold ideas into common sense realities. Working across the UK, IPPR, IPPR North, and IPPR Scotland are deeply connected to the people of our nations and regions, and the issues our communities face.

We have helped shape national conversations and progressive policy change for more than 30 years. From making the early case for the minimum wage and tackling regional inequality, to proposing a windfall tax on energy companies, IPPR's research and policy work has put forward practical solutions for the crises facing society.

IPPR
4th floor,
8 Storey's Gate
London
SW1P 3AY

E: info@ippr.org
www.ippr.org
Registered charity no: 800065 (England and Wales),
SC046557 (Scotland)
This paper was first published in April 2026.

The progressive policy think tank

CONTENTS

ABOUT THE AUTHORS

Roa Powell is a senior research fellow at IPPR.

Carsten Jung is interim associate director for economic policy and AI at IPPR.

ACKNOWLEDGEMENTS

We would like to thank Rainer Kattel, Matt Davies, Cosmina Dorobantu and Harry Quilter-Pinner for their advice and comments. We would also like to thank Abi Hynes, Richard Maclean and Rowena Mayhew for copyediting and producing the report to the highest standard.

Download
This document is available to download as a free PDF and in other formats at:
https://www.ippr.org/articles/acceleration-is-not-a-strategy

Citation
If you are using this document in your own writing, our preferred citation is:
Powell R and Jung C (2026) *Acceleration is not a strategy: A framework for directing AI towards public good before it's too late*, IPPR. https://www.ippr.org/articles/acceleration-is-not-a-strategy

SUMMARY

The politics of artificial intelligence (AI) is set to drastically change in 2026 as recent technical breakthroughs get implemented across the economy. This comes in the context of an AI ecosystem fundamentally misaligned with the public interest.

Over the past two years, IPPR has conducted a programme of work that both accepts that AI will be utterly transformative *and* asserts that governments have the tools to mould this technology to create public value. But time is running out. So far, **governments have struggled to articulate a clear vision for what it would mean for AI to go well**.

The public are understandably worried about AI (Ada Lovelace Institute 2025; Ipsos 2025), and they are becoming more worried. AI is now perceived to be one of the top three risks in the world, ranked only behind climate change and risk of war – a dramatic rise from 10th place in 2022 (Centre for Long-term Resilience 2025). This is already turning into political fights in the US as a growing group set themselves against AI entirely. And with signs of **'techlash' emerging in Europe, the potential for AI to become politically divisive should not be ignored** (Harding 2026).

Governments must stand ready to both protect people from the risks of AI and **deliberately steer this transformation towards public value**.

But policy has, so far, been too timid to do so. Experiments to harness AI in areas such as healthcare and education are welcome, and the UK is ahead of many countries on this, with initiatives including in-house government AI engineers and a series of AI big bets to improve public services. Yet, at the moment, efforts to show the public what AI is for do not think big enough given the pace of change. Meanwhile, efforts to rein in big-tech power have been modest and efforts to redistribute the benefits, for example by giving the public a stake in AI's economic upside, have been lacking entirely.

If the benefits of AI are to match up to the obvious risks, and if the concessions made to attract US tech investment are to prove worthwhile, governments will need to be far more assertive going forward.

IPPR has, over the past two years, conducted a number of case studies on directing AI deployment towards public value.[1] In this piece we draw reflections from this work and highlight next steps, with **recommendations for European governments seeking to demonstrate that they are intervening ambitiously in their citizens' interests**.

We also introduce a how-to guide for directing AI to public value, identifying priority policies for the near term (summarised in table S1 but set out in more detail in chapter 3).

1 *Transformed by AI: How generative artificial intelligence could affect work in the UK – and how to manage it* (Jung and Srinivasa Desikan 2024); *The direction of AI innovation in the UK* (Jung and Srinivasa Desikan 2025), *The new politics of AI* (Jung 2025), *AI's got news for you* (Powell and Jung 2025), *A pandemic preparedness approach to AI labour market disruption* (Jung, forthcoming 2026) and *AI and health creation* (Jung et al, forthcoming 2026).

TABLE S1

A how-to guide for directing AI to public value

Policy area	Near-term priorities
1. Steer AI towards the highest-impact uses.	Scale **bold bets on AI for public value** by embedding AI experts within schools, hospitals and local government. Use **procurement to pull the innovation ecosystem** towards transformative use cases. Use **market creation** to shape the AI ecosystem (eg, mandating AI-enabled flexibility for data centres).
2. Equip priority sectors to adopt AI well.	**Build AI adoption infrastructure**, focussing not just on big-ticket enablers such as data centres, but also on simple, sector-specific enablers. **Set standards** that provide clarity on AI liability to enable adoption in high-risk, high-reward sectors.
3. Rebalance power in the AI economy.	**Rein in incumbents** by enforcing competition rules and introducing **conduct requirements** that make big-tech products work for the wider economy. Invest in **multiple AI paradigms** beyond frontier models (eg, open-source AI) to avoid locking Europe into a narrow model of innovation.
4. Ensure the gains from AI are broadly shared.	Promote **pro-worker AI adoption** (eg, incentivising augmentation over full automation). **Give people a stake in the positive windfalls** from AI (eg, via sovereign AI funds).

Source: Authors' analysis

In 2025, countries across Europe made substantive progress towards accelerating the growth of their AI sectors. In 2026, the focus needs to shift to shaping these sectors in the public interest. **If governments can pull this off quickly, they can show the public what AI is for – and whose side they are on – before techlash deepens.**

1. INTRODUCTION: THE AI TIPPING POINT

AS AI BACKLASH GOES MAINSTREAM, GOVERNMENTS RISK BEING BOXED IN

AI is quickly becoming a mainstream issue for the public. In the US, AI's salience has increased more than any other issue in the past year (Blue Rose 2025). In the UK, it is now perceived to be one of the top three risks for humanity (34 per cent), ranked only behind climate change (49 per cent) and risk of war (48 per cent) – a dramatic rise from 10th place in 2022 (Centre for Long-term Resilience 2025). Meanwhile, just 16 per cent of people around the world are more excited than concerned about AI (Poushter et al 2025).

Early signs of AI becoming politicised are emerging. These range from opposition to data centres and the stalemate with creatives over copyright, to increasingly heated debates about children's safety, growing activism on AI's extreme risks and boycott movements such as 'QuitGPT' (QuitGPT 2026). In the past year alone there have been a growing number of protests (Key 2026), a resistance book by 10,000 authors (Spanoudi 2026) and guides published on opposing data centres (Blackmon 2025). **The coalition of people with strong anti-AI sentiment is growing.**

The risk is that justified concern hardens into blanket opposition that makes nuanced policy harder. In the US, AI has already become a political headache as a growing group are opposed to AI in principle and suspicious of leaders who appear to say that everything is ok (Blue Rose 2025). This has left politicians struggling to navigate opposition without simply "setting themselves against the future", especially as their economies increasingly rely ever more heavily on these same technology giants (Harding 2026).

Europe is far from immune. If governments do not prioritise public value, AI might 'rip through the economy' – being adopted widely and quickly without clear direction or guardrails and with net negative impacts for most people. In that context, public pressure for blunt, populist measures, from data-centre moratoriums to outright bans on AI deployment, is likely to grow and get in the way of a progressive, positive vision on AI.

THE STATE OF AI IN 2026 OFFERS LITTLE REASSURANCE

For now, anti-AI coalitions remain fragmented. But the state of AI in 2026 means this could change quickly, for example in response to a significant AI safety incident (Egan 2025) or a visible labour market shock (Leicht 2025).

Recent months have seen rapid progress in AI capabilities. Models can now autonomously complete complex tasks that would take humans hours. Claude Opus 4.6 is estimated to have a time horizon of nearly 12 hours in software tasks, meaning it can complete a range of coding, cybersecurity and machine learning tasks in one go that would take a skilled engineer 12 hours with 50 per cent success (METR 2026).

Many risks and opportunities that recently seemed science fiction, have become reality. AI agents can interact at scale in a 'social network for AI agents', including

in ways humans cannot interpret, make elaborate plans and even blackmail humans if doing so serves their goals (Janjeva et al 2026).

The biggest risks still lie ahead, and quick progress means we may not have time to avoid them all (Brundage 2026). What would have been a boosterish timeline to human-level AI a few years ago no longer appears to be so as even "long timelines to advanced AI have gotten crazy short" (Toner 2025).

As capabilities advance, politicians face three major risks:

1. **Capability progress is outpacing safety.** Safety testing is getting harder as AI models can identify if they are being tested and change their behaviour to *avoid* being caught (International AI Safety Report 2026). Standard evaluation set-ups can now underestimate capability ceilings (AI Security Institute 2026) and frontier AI companies[2] have removed hard commitments to not deploy models beyond certain capability thresholds unless safety measures are already proven to work. In other words, we cannot with certainty ensure that the latest AI models can be made safe (Häggström 2026).
2. **Labour market impacts are beginning to show, piquing anxiety in the face of high uncertainty.** Companies have attributed significant staff layoffs to AI (Bhaimiya 2026; Davenport and Srinivasan 2026). Coders have seen large parts of their jobs become automatable (Scheiber 2025; Lee 2026). And rather than 'augmenting' work, large swathes of new start-ups aim at replacing humans (Bort 2025; Business Wire 2025). The next generation of models, trained in even more gigantic data centres, will come online this year (Lee 2025). Yet economic predictions vary wildly, leaving anxiety to grow (Magistro et al 2025).
3. **Deployment remains patchy and gains are uneven.** There are pockets of genuinely transformative, positive AI deployment, but generic applications, deployment gaps and a lack of clear pull-through to public benefits could dwarf them (Jung and Srinivasa Desikan 2025). While AI's risks become obvious quickly, the benefits for most people are taking longer to materialise.

In this context, commitment on AI safety and governance is essential as it shows people that AI will not harm them. But this is not enough. There will always be risks, and people will only accept this if they can see what AI is actually for.

By articulating tangible benefits from AI and showing willingness to intervene to redistribute benefits, **politicians can begin to justify an agenda that embraces rather than resists new technology**.

2 Frontier AI companies are the small number of companies that develop highly capable, general-purpose AI models (eg, ChatGPT, Claude).

2. THE UK AND EU'S APPROACH TO AI IS MISSING DIRECTION

EUROPE'S AI AGENDA IS IMBALANCED

Policymakers are caught between two modes, neither of which aims at comprehensively steering AI's societal impact. On one side, there is AI acceleration – meaning growing the AI sector and accelerating deployment. On the other side, there is AI safety and governance – meaning avoiding clearly defined risks, for example through regulation and safety testing.

In recent years, Europe has pivoted dramatically towards acceleration. After a period where the UK and EU were strongly associated with AI governance, both are now racing to attract investment, expand computing power capacity, grow domestic AI champions and accelerate AI adoption (figure 2.1 shows this shifting emphasis over time).

FIGURE 2.1

European AI policy is shifting from prioritising AI governance towards a focus on AI acceleration

European AI policy timelines, 2023–25

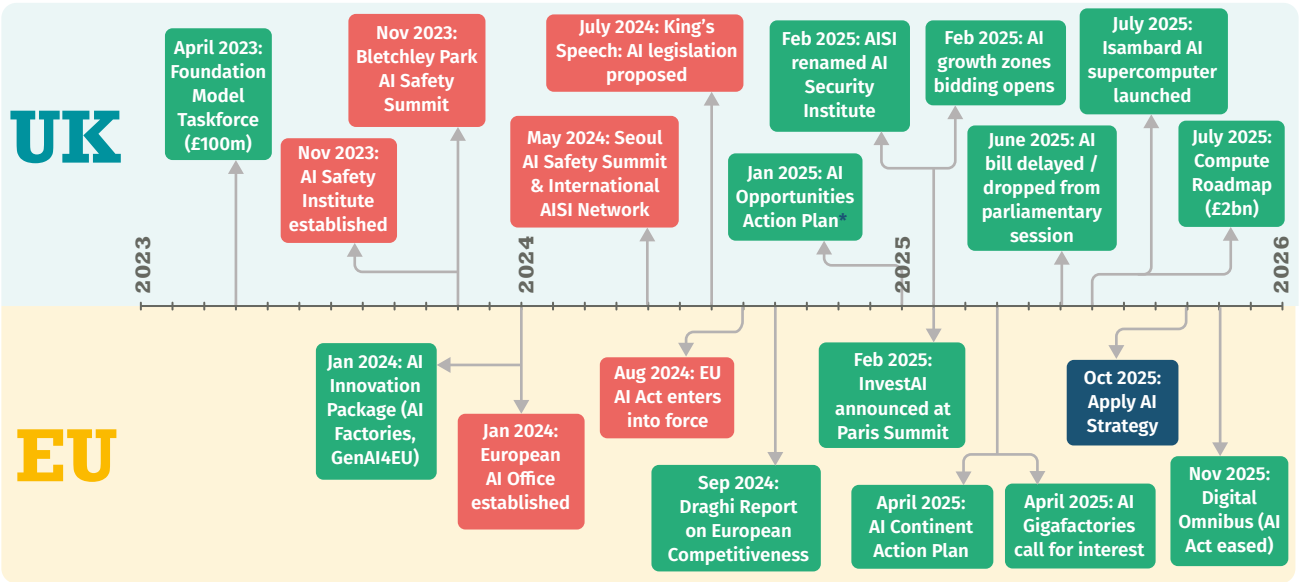

*The AI Opportunities Action Plan has some directional elements – contains specific missions and public service transformation plans

On directionism: A third orientation – directing AI towards specific public good outcomes through sector missions, institutional preparation and the redistribution of benefits. Where present, directional policies appear in blue. Currently, these appear as an afterthought compared to the substantive effort behind AI acceleration.

Source: IPPR analysis

There has been significant churn as earlier policies are adapted to fit this new agenda, with the EU introducing a 'Digital Omnibus' to simplify its digital rulebook, and annual AI safety summits pivoting to focus not just on AI risk but also on a wide range of opportunities (European Commission 2025a; Neuille 2025).

Yet in shifting agendas, governments have skipped a crucial step. They have largely bypassed what we call 'AI directionism', in other words, the idea that shaping AI towards public value deserves as much attention as simply growing the AI sector.

ACCELERATION ALONE IS NOT A STRATEGY

The shift to acceleration is understandable. To have any stake in AI's future, governments have to keep up. Faced with intense geopolitical competition and dependence on foreign companies for core infrastructure, European governments are right that AI acceleration is part of the solution.

Yet it is often forgotten that it is just that – *part* of the solution. **Growing Europe's AI sector is a means to an end rather than an end in itself.**

Sovereign AI strategies often have three stated goals: economic growth, national security and strategic leverage. While the first two flow more straightforwardly from a strengthened European AI sector, the third takes more work. The goal is not just to build leverage over the AI stack for the sake of it, but also to use this leverage to shape the AI ecosystem in line with European values (Csernatoni 2024; Bria et al 2025).

This part of the agenda remains thin. Governments have been far more ambitious about growing their AI ecosystems than shaping them. In the context of AI backlash, this is a problem.

There have been some efforts to build public-value AI. In the UK this has included the government launching an incubator for AI, which sees AI engineers pioneer positive-use cases within government and through individual 'exemplars' on AI for good in sectors such as health and education (DSIT 2025a, 2026a, 2026b). In Europe, there have been similar initiatives, most notably the Apply AI Strategy, which sets out goals for AI use across sectors such as health and education (European Commission 2025b).

But these efforts remain modest compared to the far greater ambition behind AI acceleration. The headline metrics against which governments measure their success remain investment and AI sector growth, while public-value uses appear more as case studies than as part of the overall strategy. Many of the projects held up as showing AI's positive potential remain small scale, from assisting planning decisions to automating consultation responses and converting text to British spelling (Colbourn et al 2025).

If these efforts are meant to help justify an AI acceleration agenda, they alone are unconvincing given the scale of risk we must accept in return for these benefits. The ambition behind these two agendas needs to be closely matched, especially given that businesses can, by and large, be relied on to drive AI progress – but not to shape it in the public interest.

3. HOW TO IMPLEMENT AI DIRECTIONISM IN PRACTICE

Below we set out a how-to guide for AI directionism, highlighting policies that will help governments show people **what AI is for** and demonstrate **whose side they are on** (as summarised in table 3.1).

TABLE 3.1

Summary of policies within the AI directionism toolkit

Policy area	Priority policies	Governments' agenda so far
1. Steer AI towards the highest-impact uses.	Scale **bold bets on AI for public value** by embedding AI experts within schools, hospitals and local government. Use **procurement to pull innovation** towards transformative use cases. Use **market creation** to shape the AI ecosystem (eg, mandating AI-enabled flexibility for data centres).	More developed ↓ Less developed
2. Equip priority sectors to adopt AI well.	**Build AI adoption infrastructure**, focussing not just on big-ticket enablers such as data centres, but also on simple, sector-specific fixes. **Set standards** that provide clarity on AI liability to enable adoption in high-risk, high-reward sectors.	
3. Rebalance power in the AI economy.	**Rein in incumbents** by enforcing competition rules and introducing **conduct requirements** that make big-tech products work for the wider economy. Invest in **multiple AI paradigms** beyond frontier models (eg, open-source AI), to avoid locking Europe into a narrow model of innovation.	
4. Ensure the gains from AI are broadly shared.	Promote **pro-worker AI adoption** (eg, incentivising augmentation over full automation). **Give people a stake in the positive windfalls** from AI (eg, via sovereign AI funds).	

Source: Authors' analysis

1. STEER AI TOWARDS THE HIGHEST-IMPACT USES

Governments need policies that steer AI innovation towards specific public outcomes so that people can clearly see what AI is for.

Market forces alone will not do this. We found that 85 per cent of AI applications target general process improvements, with just 15 per cent targeting specific problems in specific sectors (Jung and Srinivasa Desikan 2025). The US AI company, Anthropic, similarly identified a capability overhang, with several sectors using AI far less

than they could (Anthropic 2026). Meanwhile the UK's Department for Science, Innovation and Technology found that AI adoption is concentrated on easy-use cases rather than on transformative, high-impact challenges (DSIT 2026c).

While market forces will eventually fill some gaps, many of the highest-impact AI use cases will be slow to emerge without government incentives. This will be true both in the public sector, and in parts of the private sector where public value is high but economic returns are low.

Identifying where intervention is needed will not be as simple as picking priority sectors and introducing innovation funding. Instead, close collaboration with sector experts is needed to develop sufficiently specific AI missions.

In the case studies we explored, we found that AI innovation in high public-value areas is lagging:

- In healthcare, AI innovation is narrowly indexed on use cases for treatment and administrative support, while incentives for prevention and health creation have been missing (Jung and Srinivasa Desikan 2025) (see figure 3.1).
- In news, AI could have positive impacts on our information environment, for example, AI for fact checking or AI-enhanced newsrooms. But such uses are insufficiently incentivised (Powell and Jung 2026).

FIGURE 3.1

AI innovation in health is still skewed towards clinical applications

Business model focus of AI firms (percentage of the total number of firms active in the health space)

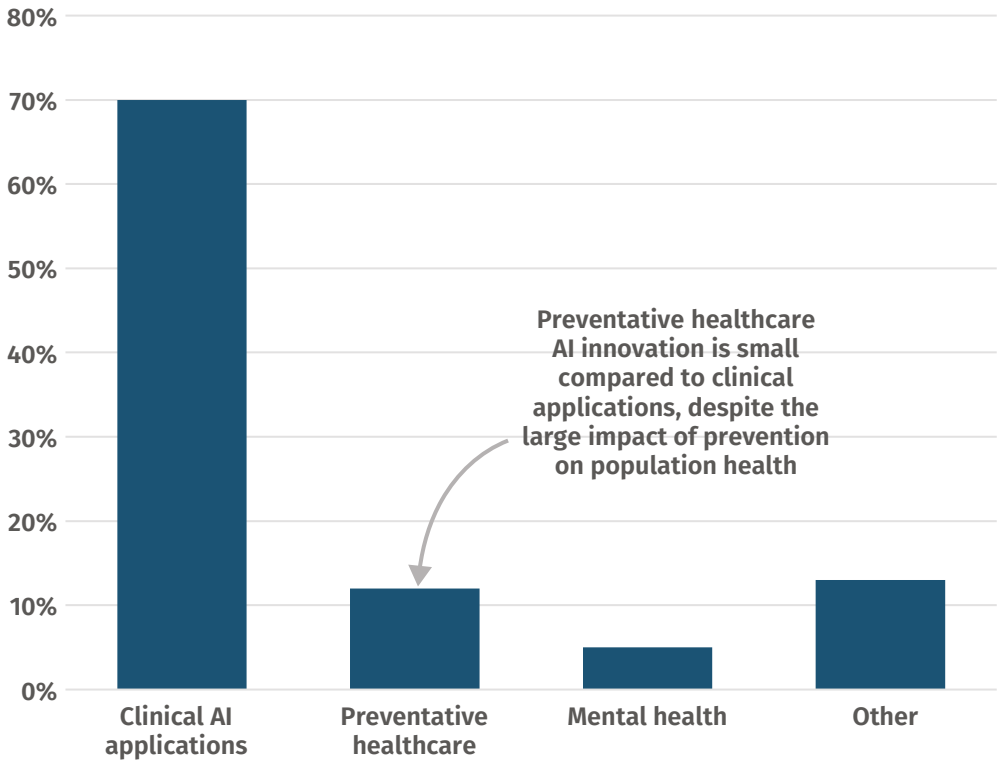

Source: This is based on Jung and Srinivasa Desikan (2025), where we conduct quantitative analysis of 3,200 organisations that have AI as part of their value proposition.

We propose three policies to push AI innovation towards public value.

Make bold bets on AI for public-value creation, embedding engineers within frontline teams

This is where governments' agendas are already most advanced in pilot form. Initiatives such as the UK's AI Exemplars scheme and team of in-house engineers at the Incubator for Artificial Intelligence (i.AI), and Europe's Apply AI Strategy, have begun to identify positive public-sector use cases.

These efforts should be scaled up as advanced AI reaches maturity. But, for this to succeed, priority use cases need to be selected in collaboration with sector experts, not imposed from the top down or shaped too much by the often-exaggerated promises of tech chief executives. In health, for example (Jung et al, forthcoming 2026), there is a stark mismatch between:

- what tech chief executives promise (that is, AI to cure most diseases) versus what is realistic
- the problems that AI innovators tackle (that is, diagnosis and treatment) versus the bottlenecks that health practitioners face (that is, how can we improve community care and health creation?).

Embedding forward-deployed engineers in schools, hospitals and local government can help solve this, identifying where AI will genuinely add value before feeding back to central strategy teams. These engineers should be given freedom to experiment and iterate applications – and allow public deliberation about them – before big bets are scaled, but they should be guided by a central 'constitution' that ensures deployment is focussed on delivering 'missions' (Mazzucato 2025).

This would see governments take a leaf from frontier AI labs' playbooks, where this has already been scaled in the private sector (Bousquette 2026). The Ministry of Justice in the UK has also trialled this, providing a model that can be copied (James and L'Huillier 2025).

Use outcomes-based procurement to pull innovation towards transformative outcomes

For ambitious AI use, governments will need private sector help. Yet at the moment, procurement is not set up to deliver AI innovation at scale *or* to actively shape the market. Instead, procurement is fragmented, risk-averse and often focussed on buying off-the-shelf products (for example, hardware or software licenses) rather than on outcomes (for example, reduced hospital admissions or measurable gains in teacher retention). Despite some recent pilots in the UK and the EU that change this, government capacity to do sophisticated outcomes-based procurement on new technologies is still highly limited (House of Commons Committee of Public Accounts 2025).

Procurement should be treated as a market-shaping tool. Governments can articulate clear outcomes they want to achieve and use procurement to create credible demand for the solutions. Doing this effectively will require strong central coordination, for example via a centralised AI procurement function where transformative use cases can be collated, risks tolerance remains high and an aggregate, cross-government demand signal can be generated.

Create markets through regulation and standards

Finally, governments can go further by creating a market for positive AI use cases. This has already been done successfully in green industrial policy. For example, the UK's zero emission vehicle mandate requires manufacturers to ensure a rising share of new car sales is zero emission, compelling the private sector to create the market.

This could be equally transformative in AI, incentivising European firms to develop positive use cases *and* giving European firms a head start in new markets.

Examples that governments should consider include:

- creating a market for flexible data centres (that is, those that shift demand away from peak hours where needed) by mandating that a rising share of data centres uses AI to help spread their load – this could simultaneously reduce the energy pressure of AI and create a market for innovative AI solutions in Europe
- creating a market for AI assurance by mandating that independent experts audit high-risk AI products, thus creating a market for assurance technologies, giving European companies a first-mover advantage and advancing AI safety goals (Kaprayoon and Anderson-Samways 2024).

STATE CAPACITY FOR MISSION-BASED POLICY MAKING

The above is effectively an outline of the need to do more 'mission-based' policy making that coordinates across departments. This is harder than siloed work within departments and requires some new capacities. But these are achievable and indeed already exist in embryonic form in G7 economies.

It will require:

- capacity for cross-government coordination (where multiple departments are relevant)
- increased technical expertise
- iterative learning by doing
- the ability to use procurement, regulation and finance towards outcomes.

The OECD has long made the case for mission-based policy making. Missions can work through dedicated teams, experimentation and feedback – they do not require a single grand plan, nor a huge societal mission, nor centralised decision making (OECD 2024).

The UK has the state capacity to do such mission-based policy in a narrowly focussed and disciplined way. There are useful building blocks in place – for instance, the prime minister's AI Exemplars programme, which is testing a small portfolio of practical AI applications in areas such as health, education and justice. But many of the use cases are focussed on operational improvement. This should be scaled and gradually aimed at more ambitious challenges.

Germany can call on one of the world's most advanced innovation systems in dealing with these challenges, but a new more agile and experimental approach to science, technology and innovation policy is needed (OECD 2022). France appears somewhat closer to a usable mission-state model. Its France 2030 investment programme combines central strategic capacity with distributed delivery across specialised operators. It sets concrete medium-term goals in specific domains and these are delivered through multiple operators and ecosystems rather than a single ministry.

2. EQUIP PRIORITY SECTORS TO ADOPT AI WELL

Often what limits AI adoption is **not technology readiness but infrastructural readiness**. If this is not addressed, even well-designed AI tools will fail to deliver public benefit at scale.

Governments need to develop sector-readiness plans, especially in sectors set to be most imminently transformed by AI, focussing on the following two priorities.

Build the institutional infrastructure that sectors need to adopt AI well

Simply 'dropping AI into' a system that is not currently working well will fail to deliver good outcomes. For example, in the case of public health, a key bottleneck for genuinely health-creating AI is the lack of underlying community health infrastructure (Jung et al, forthcoming 2026). Health systems are already failing to reach disadvantaged parts of the population, and layering AI on top will simply widen that gap. Building underlying institutional infrastructure is therefore an essential precondition for AI to deliver widespread public benefit. But there is no one-size-fits-all approach. Good AI policy thus requires sector-specific knowledge and state capacity to act on it.

Some infrastructural blockers to AI adoption are more universal (for example, a lack of computing power, poor data quality or, more recently, a lack of the permissions infrastructure for AI agents to safely execute multi-step tasks). But it is often the sector-specific bottlenecks that go unaddressed and they would be simplest for European governments to rectify.

Set clear standards on responsible AI to enable high-risk adoption

In many sectors, a key bottleneck to positive AI use is uncertainty about how to use it without undermining trust or incurring liability. This is especially true where AI use is high risk and high reward, such as in health or education. In the case studies we explored, this came through strongly:

- In news, there is no consensus on how far AI should extend into the journalistic workflow, on how its use should be disclosed to audiences or on what checks journalists should conduct before accepting AI outputs (Ajdinovic 2026; Powell and Jung 2026).
- In health, some AI applications sit within clinical applications and are highly regulated. But other AI applications sit within the wellbeing space where standards, let alone regulation, are minimal. Here, a lack of clear standards from government is preventing people from being able to easily access the most effective AI applications for them (De Freitas and Cohen 2025).

The arrival of AI agents makes this issue even more important. While agents are already roaming the internet, there is no standard for an auditable trail of their behaviour in case something goes wrong, and no standards for ensuring the good behaviour of networks of agents. Without rules of the road for agents, many businesses will be too nervous to adopt them (Chan et al 2025).

Clear assurance and standards for AI are not just governance tools, but also enablers of AI adoption. Where governments recognise this dual function, they can use these tools to enable high-impact AI adoption, with the UK government's work on AI assurance offering a positive example of this sort of initiative (DSIT 2025b).

3. REBALANCE POWER IN THE AI ECONOMY

New incentives and adoption infrastructure can help steer AI towards positive use cases, but governments will struggle to truly direct AI towards public benefits if they do not address the extraordinary concentration of power in the hands of a few technology companies.

AI's supply chain is dominated by a single chip designer (Nvidia holds roughly 80 per cent of the AI accelerator market; Hajdari 2025) and by three cloud providers ('hyperscalers'). And vertical integration means that the same companies are also increasingly shaping the application layer and determining which AI products reach consumers.

Already, this is causing harm. The UK's Competition and Markets Authority has found that concentration in the cloud market is leading to higher costs, less choice and lower-quality service for UK businesses (CMA 2025). Meanwhile, Google's integration of AI into its search engine has led to a collapse in traffic to publishers and to concern around the quality of sourcing for the public (Powell and Jung 2026).

Without intervention, these firms will continue to be the ones that shape AI, risking a market with fewer choices, higher prices, less privacy and ever-larger models with more severe environmental impacts (Radsch et al 2025).

Below we propose two realistic ways in which governments can begin to rebalance power. However, more research is urgently needed in order to predict just how severe concentration in AI markets might become – and what the likely consequences will be for European innovation.

Enforce competition policy to curb incumbent power

To start with, rebalancing power simply requires governments to be bolder in enforcing existing rules quickly and visibly. In the UK, the Digital Markets, Competition and Consumers Act 2024 already empowers the Competition and Markets Authority to designate firms with entrenched power *and* to apply conduct requirements before harm occurs. In the EU, the Digital Markets Act 2022 similarly allows the EU to introduce obligations for gatekeeper firms.

Yet, so far, governments have repeatedly not capitalised on this legislation. In the UK, commitment to robust competition enforcement is under active threat:

- The government has issued a strategic steer emphasising pro-investment priorities (DBT 2025).
- The chair of the Competition and Markets Authority has been removed and replaced with an ex-Amazon executive.
- The chair of the cloud inquiry, which looks into concentrated power within cloud markets, resigned, citing concerns about the glacial pace of change.
- The business department is now consulting on further reform to the competition regime.

The big threat to addressing market power is not a lack of good regulation, but a row-back in political will – often in the face of strong pressure from big tech. To address this, governments across Europe could do more to coordinate enforcement. When Europe acts together, exit threats from technology giants become hollow as these markets are simply too large to abandon. Greater alignment in the sequencing of action across the AI stack could allow competition authorities to move faster and harder.

In practice, this can work on two fronts. First, action from competition authorities can directly make AI products better for the public, for example by requiring Google to negotiate on fair terms with publishers and ensuring it attributes diverse sources in its AI answers. Second, open and contestable markets give European start-ups a genuine chance to compete. For this to happen, governments need to resource their competition authorities to act quickly where market power is already entrenched and ensure they are given a strategic steer to respond proactively should concentration spread to other parts of the AI stack.

Use industrial policy to build plural, public-interest alternatives

Beyond competition enforcement, governments can do more to shape their domestic AI industries. Both public procurement and investment can help, from advanced market commitments for domestic start-ups to public equity stakes in strategically important firms. Often, existing initiatives designed to accelerate growth in domestic AI sectors can be tweaked by giving them an explicit mission to also shape AI's impacts. Sovereign AI funds – including in the UK and France, for example – could be used not just to scale up domestic AI firms but also to grow companies building publicly beneficial AI (BpiFrance 2025; Sovereign AI 2026).

In designing these market-shaping initiatives, governments must remember that favouring frontier models inevitably concentrates power among the handful of companies with the resources to train them. A more diverse ecosystem would distribute power more widely, while still enabling UK growth.

Open-source AI,[3] for example, should be supported to provide value for money, strategic autonomy and transparency (Robinson 2025), alongside significant financial returns. For example, Public First (2026) suggests that open-source AI adoption has the potential to deliver £40 billion in the UK over a decade. Alongside this, the UK might consider support for edge computing (that is, small models that operate on local devices) and more specialised machine learning for science companies.

A fast follower model where "we wait and see what innovations yield the most value" before finding a "cost-effective way of doing it ourselves" is likely to have significant merit in European countries where competing at the frontier seems infeasible. But to do this, governments must hedge their bets, betting on multiple AI paradigms at once (Shearer and Davies 2026).

4. DISTRIBUTE THE GAINS FROM ADVANCED AI

Finally, if governments want AI to deliver public benefit, they need to act now to ensure the benefits of AI are broadly distributed. Previous periods of rapid technological change resulted in highly concentrated gains and, without government steering, AI is likely to be a repeat of this, only at greater scale (Johnson and Acemoglu 2023).

Even as the speed and extent of AI's impact are still so uncertain, two priority areas stand out as having high payoffs in the immediate term, as set out below.

Use fiscal and industrial policy to maintain employment over the medium term, including in less exposed jobs

Many countries have scaled up tax breaks on software in recent years, meaning there is now a fiscal incentive to automate rather than augment workers (Dabla-Norris and de Mooij 2024; Acemoglu et al 2026). Yet, evidence suggests that people would prefer policy to focus on maintaining employment rather than merely replacing incomes (Bruenig 2026). It is also important to incentivise high employment rates, to ensure workers maintain leverage in the economy (Leicht 2025).

As a first step, governments should begin to rebalance tax and subsidy schemes so firms are rewarded for raising worker productivity rather than automating roles (Jung and Srinivasa Desikan 2024). In the longer term, they should also consider strengthening taxes on capital income as value in the economy shifts away from labour and towards capital (Dabla Norris and de Mooij 2024).

3 Open-source AI involves AI models that can be shared, modified or reproduced with little restriction on users, for example thanks to openly available weights or data. However, the degree of openness associated with models labelled as 'open source' varies significantly.

Even where work is augmented rather than automated, there is no guarantee that the benefits will flow to workers, with evidence suggesting that work can instead intensify while productivity benefits do not flow through to workers (Ranganathan and Maggie Ye 2026). Giving workers a meaningful voice in shaping AI deployment can reduce anxiety and ensure adoption increases job quality (Giuntella et al 2024; Evans and Kinton, forthcoming 2026).

Finally, and crucially, this will mean not just defending existing jobs but also growing sectors that are less exposed to automation risk. This points not only to green industries but also to 'social occupations' such as care and mental health, with serious retraining and transition support so workers can move into new, better jobs (Jung and Srinivasa Desikan 2024; Acemoglu et al 2026).

Give the public a direct stake in AI's upside

If we accept that AI is going to have huge economic upsides but are unsure about where exactly these will fall (for example, will it be to frontier AI model developers or to large AI deployers?), governments should put in place mechanisms to make sure citizens capture the upsides.

One way to do this is through using sovereign wealth funds that allow governments to take a stake in and then redistribute windfalls from AI (Anthropic 2025), where equity stakes – either direct or indirect via public schemes – provide some insurance against the radical uncertainty about the degree of AI's future impact, and ensure broad participation in the prosperity that AI might bring (Korinek and Lockwood 2026).

Building on IPPR's longstanding case for spreading capital ownership more widely through citizens' wealth funds and employee ownership trusts, we argue that sovereign AI funds could help (Roberts and Lawrence 2017). Already, the UK has advanced the idea that the state should take a strategic stake in AI by introducing a Sovereign AI Fund. From here, one step further would be to ensure that sovereign AI units do not just help create national champions but also ensure that some of that upside is captured and redistributed – for example, channelling returns into dividends and transition support for citizens most exposed to displacement.

4. CONCLUSION

European governments need a convincing positive vision for AI. Otherwise, it will not be long before the increasingly apparent risks feed anti-AI sentiment to the point where resistance builds and any nuanced agenda on AI becomes hard to deliver. As AI capabilities continue to advance at pace, the political space for an ambitious AI agenda is narrow.

We therefore argue that 2026 must be the year when European governments adjust their approach, recognising that simply growing the AI sector and hoping for spillover benefits is not a complete strategy. Alongside ramping up efforts on AI safety, governments must become far more interventionist in:

- steering AI towards clear public value missions
- preparing sectors for the radical transformation that AI will bring
- confronting the extreme concentration of power in the AI economy
- ensuring the benefits are broadly felt.

This how-to guide for implementing AI directionism in practice provides governments with a starting point for doing exactly this.

REFERENCES

Acemoglu D, Autor D and Johnson S (2026) 'Building pro-worker artificial intelligence', working paper 34854, National Bureau of Economic Research. https://economics.mit.edu/sites/default/files/2026-03/Building%20Pro-Worker%20Artificial%20Intelligence.pdf

Ada Lovelace Institute (2025) 'Nearly 9 in 10 people in the UK support independent regulation of AI', Ada Lovelace Institute website, 4 December 2025. https://www.adalovelaceinstitute.org/press-release/nearly-9-in-10-people-in-the-uk-support-independent-regulation-of-ai

AI Security Institute (2026) 'Evidence for inference scaling in AI cyber tasks: increased evaluation budgets reveal higher success rates', AI Security Institute website, 5 March 2026. https://www.aisi.gov.uk/blog/evidence-for-inference-scaling-in-ai-cyber-tasks-increased-evaluation-budgets-reveal-higher-success-rates

Ajdinovic M (2026) *Editorial identity and AI*, LSE. https://static1.squarespace.com/static/64d60527c01ae7106f2646e9/t/6943bc00f4458c5d7bbbc1d5/1766046720329/Editorial+Identity+report.pdf

Anthropic (2025) 'Preparing for AI's economic impact: Exploring policy responses', Anthropic website, 14 October 2025. https://www.anthropic.com/research/economic-policy-responses

Anthropic (2026) 'Labor market impacts of AI adoption: a new measure and early evidence', Anthropic website, 5 March 2026. https://www.anthropic.com/research/labor-market-impacts

Bhaimiya S (2026) 'Meta stock climbs nearly 3% on report of planned layoffs to offset AI spending', CNBC website, 16 March 2026. https://www.cnbc.com/2026/03/16/meta-ai-costs-mass-layoffs-20percent-up-premarket.html

Blackmon S (2025) 'How to resist data centers: a guide for local communities in Europe', AlgorithmWatch website, 13 November 2025. https://algorithmwatch.org/en/a-guide-to-data-centers

Blue Rose (2025) *AI is colliding with America's affordability crisis*, Blue Rose Research. https://data.blueroseresearch.org/hubfs/%5BBRR%5D%20AI%20Is%20Colliding%20With%20America%E2%80%99s%20Affordability%20Crisis-1.pdf

Bort J (2025) 'Artisan, the "stop hiring humans" AI agent startup, raises $25m – and is still hiring humans', TechCrunch website, 9 April 2025. https://techcrunch.com/2025/04/09/artisan-the-stop-hiring-humans-ai-agent-startup-raises-25m-and-is-still-hiring-humans

Bousquette I (2026) 'The hottest job in tech isn't very glamorous', *Wall Street Journal*, 12 March 2026. https://www.wsj.com/cio-journal/the-hottest-job-in-tech-isnt-very-glamorous-dc29ab3e

BpiFrance (2025) 'Bpifrance deploys €10 billion to develop the AI ecosystem and facilitate the adoption of artificial intelligence by French companies', BpiFrance website, 27 March 2025. https://www.bpifrance.com/2025/03/27/bpifrance-deploys-e10-billion-to-develop-the-ai-ecosystem-and-facilitate-the-adoption-of-artificial-intelligence-by-french-companies

Bria F, Timmers P and Genone F (2025) *EuroStack – a European alternative for digital sovereignty*, CEPS. https://cdn.ceps.eu/wp-content/uploads/2025/02/EuroStack_2025.pdf

Bruenig M (2026) 'AI could destroy the labor market. We already know how to fix it', Substack website, 19 March 2026. https://substack.com/inbox/post/191419144

Brundage M (2026) 'We're in triage mode for AI policy', Substack website, 22 February 2026. https://milesbrundage.substack.com/p/were-in-triage-mode-for-ai-policy

Business Wire (2025) 'Decagon raises $131m at $15bn valuation to deliver concierge customer experience with AI agents', Business Wire website, 23 June 2025. https://www.businesswire.com/news/home/20250623894798/en/Decagon-Raises-%24131M-at-%241.5B-Valuation-to-Deliver-Concierge-Customer-Experience-with-AI-Agents

Centre for Long-term Resilience (2025) 'New poll shows high public demand for the government to address extreme risks from AI', Centre for Long-term Resilience website, 21 January 2025. https://www.longtermresilience.org/new-poll-shows-high-public-demand-for-the-government-to-address-extreme-risks-from-ai

Chan A, Wei K, Huang S, Rajkumar N, Perrie E, Lazar S, Hadfield K and Anderljung M (2025) 'Infrastructure for AI agents', research paper, GovAI. https://arxiv.org/pdf/2501.10114

Colbourn J, Messer S and Edwards G (2025) 'Extract: Using AI to unlock historic planning data', MHCLG Digital blog, 12 June 2025. https://mhclgdigital.blog.gov.uk/2025/06/12/extract-using-ai-to-unlock-historic-planning-data

Competition and Markets Authority [CMA] (2025) *Cloud infrastructure services: Final decision report*. https://assets.publishing.service.gov.uk/media/688b8891fdde2b8f73469544/final_decision_report.pdf

Csernatoni R (2024) 'Charting the geopolitics and European governance of artificial intelligence', Carnegie Endowment for International Peace website, 6 March 2024. https://carnegieendowment.org/russia-eurasia/research/2024/03/charting-the-geopolitics-and-european-governance-of-artificial-intelligence

Dabla-Norris E and de Mooij R (2024) 'Fiscal policy can help broaden the gains of AI to humanity', International Monetary Fund website, 17 June 2024. https://www.imf.org/en/blogs/articles/2024/06/17/fiscal-policy-can-help-broaden-the-gains-of-ai-to-humanity

Davenport TH and Srinivasan L (2026) 'Companies are laying off workers because of AI's potential, not its performance', *Harvard Business Review*, 29 January 2026. https://hbr.org/2026/01/companies-are-laying-off-workers-because-of-ais-potential-not-its-performance

De Freitas J and Cohen I (2025) 'Unregulated emotional risks of AI wellness apps', *Nature Machine Intelligence*, 7: 813–815. https://www.nature.com/articles/s42256-025-01051-5

Department for Business & Trade [DBT] (2025) 'Strategic steer to the Competition and Markets Authority', 15 May 2025. https://www.gov.uk/government/publications/strategic-steer-to-the-competition-and-markets-authority/strategic-steer-to-the-competition-and-markets-authority

Department for Science, Innovation and Technology [DSIT] (2025a) 'AI programme', GOV.UK website, 17 August 2025. https://www.gov.uk/guidance/ai-exemplars-programme

Department for Science, Innovation and Technology [DSIT] (2025b) 'Trusted third-party AI assurance roadmap', GOV.UK website, 3 September 2025. https://www.gov.uk/government/publications/trusted-third-party-ai-assurance-roadmap/trusted-third-party-ai-assurance-roadmap

Department for Science, Innovation and Technology [DSIT] (2026a) 'Bold bet on AI to keep UK at forefront of science and research breakthroughs from healthcare to better public services', press release, 19 February 2026. https://www.gov.uk/government/news/bold-bet-on-ai-to-keep-uk-at-forefront-of-science-and-research-breakthroughs-from-healthcare-to-better-public-services

Department for Science, Innovation and Technology [DSIT] (2026b) '450,000 disadvantaged pupils could benefit from AI tutoring tools', press release, 26 January 2026. https://www.gov.uk/government/news/450000-disadvantaged-pupils-could-benefit-from-ai-tutoring-tools

Department for Science, Innovation and Technology [DSIT] (2026c) 'AI adoption research', GOV.UK website, 13 February 2026. https://www.gov.uk/government/publications/ai-adoption-research/ai-adoption-research

Egan J (2025) 'The United States must avoid AI's Chernobyl moment', Center for a New American Security website, 10 March 2025. https://www.cnas.org/publications/commentary/the-united-states-must-avoid-ais-chernobyl-moment

European Commission (2025b) 'Digital Omnibus regulation proposal', European Commission website, 19 November 2025. https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-regulation-proposal

European Commission (2025c) 'Apply AI Strategy', European Commission website. https://digital-strategy.ec.europa.eu/en/policies/apply-ai

Evans J and Kinton A (forthcoming 2026) *Worker power in the age of AI*, IPPR.

Giuntella O, Konig J and Stella L (2024) 'Artificial intelligence and the wellbeing of workers', *Scientific Reports*, 15(1): 20087. https://www.researchgate.net/publication/392940780_Artificial_intelligence_and_the_wellbeing_of_workers

Häggström O (2026) 'Anthropic backpedals on their former responsible scaling policy', Substack website, 25 February 2026. https://haggstrom.substack.com/p/anthropic-backpedals-on-their-former

Hajdari U (2025) 'Nvidia shares slip as AI accelerator race shifts interest to Google chips', Euronews website, 25 November 2025. https://www.euronews.com/business/2025/11/25/nvidia-shares-slip-as-ai-accelerator-race-shifts-interest-to-google-chips

Harding V (2026) 'The coalition the Democrats didn't ask for', Substack website, 27 February 2026. https://verityharding.substack.com/p/the-coalition-the-democrats-didnt

House of Commons Committee of Public Accounts (2025) *Government's relationship with digital technology suppliers: Twenty-seventh report of session 2024–25* (HC 640), Parliament UK. https://publications.parliament.uk/pa/cm5901/cmselect/cmpubacc/640/report.html

International AI Safety Report (2026) *International AI safety report 2026*. https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026

Ipsos (2025) 'Ipsos global survey: the world is worried about AI', *Policy* magazine website, 23 October 2025. https://www.policymagazine.ca/ipsos-global-survey-the-world-is-worried-about-ai

James D and L'Huillier S (2025) 'Announcing the first UK government forward deployed engineers', Ministry of Justice blog, 4 November 2025. https://ai.justice.gov.uk/blog/announcing-the-first-uk-government-forward-deployed-engineers

Janjeva A, Ashurst C and Hennessy R (2026) 'Agentic AI in the wild: lessons from Moltbook and OpenClaw', Centre for Emerging Technology and Security website, 12 February 2026. https://cetas.turing.ac.uk/publications/agentic-ai-wild-lessons-moltbook-and-openclaw

Johnson S and Acemoglu D (2023) *Power and progress: Our thousand-year struggle over technology and prosperity*, Public Affairs

Jung C (2025) 'The new politics of AI: Why fast technological change requires bold policy targets', IPPR website, 7 February 2025. https://www.ippr.org/articles/new-politics-of-ai

Jung C (forthcoming 2026) *A pandemic preparedness approach to AI labour market disruption*, IPPR

Jung C, Rees, S and Powell R (forthcoming 2026) *AI and health creation*, IPPR

Jung C and Srinivasa Desikan B (2024) *Transformed by AI: How generative artificial intelligence could affect work in the UK – and how to manage it*, IPPR. https://www.ippr.org/articles/transformed-by-ai

Jung C and Srinivasa Desikan B (2025) *The direction of AI innovation in the UK: Insights from a new database of AI firms*, IPPR. https://www.ippr.org/articles/the-direction-of-ai-innovation-in-the-uk

Kaprayoon J and Anderson-Samways B (2024) *Assuring growth: Making the UK a global leader in AI assurance technology*, Social Market Foundation. https://www.smf.co.uk/publications/assuring-growth-ai-uk-leader

Key A (2026) '"Scream if you want to move slower!" A nascent AI protest coalition comes together in London', Transformer website, 9 March 2026. https://www.transformernews.ai/p/scream-if-you-want-to-move-slower-pause-ai-pull-the-plug

Korinek A and Lockwood L (2026) 'Preparing for the economic impacts of advanced AI', working paper, Brookings Institution. http://www.brookings.edu/wp-content/uploads/2026/01/Korinek-Lockwood-FINAL-for-website.pdf

Lee J (2025) 'Today's largest data center can do more GPT-4-scale training runs each month', Epoch AI website, 4 November 2025. https://epoch.ai/data-insights/gpt-4s-trainable

Lee TB (2026) 'Sorry, skeptics: AI really is changing coding', Understanding AI website, 27 February 2026. https://www.understandingai.org/p/sorry-skeptics-ai-really-is-changing

Leicht A (2025) 'AI & jobs: Leverage without labor', Leicht Substack, 5 November 2025. https://writing.antonleicht.me/p/ai-and-jobs-leverage-without-labor

Magistro B, Borwein S, Alvarez M, Bonikowski B and Loewen PJ (2025) 'The coming AI backlash', *Foreign Affairs*, 13 October 2025. https://www.foreignaffairs.com/united-states/coming-ai-backlash

Mazzucato M (2025) 'AI for what? Public value creation versus extractive rents', *Mission Economics*, 3 February 2025. https://marianamazzucato.substack.com/p/ai-for-what

Model Evaluation & Threat Research [METR] (2026) 'Task-completion time horizons of frontier AI models', METR website, 3 March 2026. https://metr.org/time-horizons

Neuille H (2025) 'From safety to action: The upcoming French AI summit', Tech Policy Press website, 5 February 2025. https://www.techpolicy.press/from-safety-to-action-the-upcoming-french-ai-summit

Organisation for Economic Co-operation and Development [OECD] (2022) *OECD reviews of innovation policy: Germany 2022*. https://www.oecd.org/en/publications/oecd-reviews-of-innovation-policy-germany-2022_50b32331-en.html

Organisation for Economic Co-operation and Development [OECD] (2024) 'Designing effective governance to enable mission success', OECD website, 6 December 2024. https://www.oecd.org/en/publications/designing-effective-governance-to-enable-mission-success_898bca89-en.html

Poushter J, Fagan M and Corichi M (2025) 'How people around the world view AI', Pew Research Center website, 15 October 2025. https://www.pewresearch.org/global/2025/10/15/how-people-around-the-world-view-ai

Powell R and Jung C (2026) *AI's got news for you: Can AI improve our information environment?*, IPPR. https://www.ippr.org/articles/ais-got-news-for-you

Public First (2026) *Capability, control and growth: Achieving sovereignty through open-source AI*. https://ukopensourceai.publicfirst.co

QuitGPT (2026) 'ChatGPT takes Trump's killer robot deal. It's time to quit'. https://quitgpt.org

Radsch CC, von Thun M and Nie M (2025) 'AI and market concentration', Open Markets Institute website. https://www.openmarketsinstitute.org/publications/expert-brief-ai-and-market-concentration-courtney-radsch-max-vonthun

Ranganathan A and Maggie Ye X (2026) 'AI doesn't reduce work, it intensifies it', *Harvard Business Review*, 9 February 2026. https://hbr.org/2026/02/ai-doesnt-reduce-work-it-intensifies-it

Roberts C and Lawrence M (2017) *Wealth in the twenty-first century: Inequalities and drivers*, IPPR. https://ippr-org.files.svdcdn.com/production/Downloads/cej-wealth-in-the-21st-century-october-2017.pdf

Robinson S (2025) *Keeping the door open: A roadmap for integrating open-source AI in public services*, Social Market Foundation. https://www.smf.co.uk/publications/open-source-ai-public-services

Scheiber N (2025) 'At Amazon, some coders say their jobs have begun to resemble warehouse work', *New York Times*, 25 May 2025. https://www.nytimes.com/2025/05/25/business/amazon-ai-coders.html

Shearer E and Davies M (2026) 'Code dependency: Big tech market power and why progressives need an AI strategy', *Renewal*, 27 March 2026. https://renewal.org.uk/articles/code-dependency-big-tech-market-power-and-why-progressives-need-an-ai-strategy

Sovereign AI (2026) 'Sov/AI', Sovereign AI website. https://www.sovereignai.gov.uk

Spanoudi M (2026) 'Around 10k authors protest AI commercial research exception with publication of empty book', *The Bookseller*, 10 March 2026. https://www.thebookseller.com/news/around-10k-authors-protest-ai-commercial-research-exception-with-publication-of-empty-book

Toner H (2025) 'Long timelines to advanced AI have gotten crazy short', Substack website, 1 April 2025. https://helentoner.substack.com/p/long-timelines-to-advanced-ai-have

Institute for Public Policy Research

GET IN TOUCH

For more information about the Institute for Public Policy Research, please go to www.ippr.org

You can also e-mail info@ippr.org or find us on X @ippr

Institute for Public Policy Research
Registered Charity no. 800065 (England & Wales),
SC046557 (Scotland), Company no, 2292601 (England & Wales)

The progressive policy think tank