THE NEW STATESMAN

Spotlight

Thought leadership and policy

AI Governance

Chi Onwurah MP

Josh Simons MP

Gina Neff

Jonathan Camrose

ASA

FORTINET

Netcompany

THE NEW STATESMAN

New Times, New Thinking

Just £8.99 a month*

Subscriber benefits include:

Unlimited access to newstatesman.com

Weekly magazine delivered direct to your door**

Early-bird and discounted tickets to *NS* events

Ad-free listening to the *NS* podcast

www.newstatesman.com/magazine

*Offer is for a digital subscription for £8.99 a month or a bundled print & digital subscription for £13.99 a month and is limited to subscribers within the UK. Your payment will auto-renew monthly at £8.99 (digital) or £13.99 (print & digital). **Both website access and magazine deliveries will only be available to bundled print & digital subscribers. Overseas prices available at newstatesman.com/subscribe or call +44 (0)808 284 9422 for further assistance.

Intelligence test

The government has set out plans not only to embed artificial intelligence across public services – from NHS diagnostics to tax compliance – but also to grow a world-leading domestic ecosystem.

The latest AI Opportunities Action Plan outlines ambitions to support homegrown innovation, strengthen compute infrastructure and build the skills base needed to make the UK "an AI superpower". But regulation is catching up slowly. Ongoing debates around copyright protections and transparency have drawn sharp criticism from the creative industries, for example, raising questions about how Britain will balance openness with safeguards.

As global powers race to shape the norms of AI governance, the UK must define its position – with clarity, ambition and care. This edition of *Spotlight* seeks to explore exactly what that position could and should look like.

Labour MP Chi Onwurah, chair of the Science, Innovation and Technology Committee, argues the country boasts unique strengths to leverage in becoming a trusted, innovative and ethical leader in AI development and governance, praising the government's commitment to employ AI for the public good (page 4). But the shadow science minister, Jonathan Camrose, believes it should be more ambitious solving the AI-copyright challenge (page 5).

His suggested solutions align more with those who believe a light-touch regulatory model akin to that employed in the US is the best framework for success. The EU has moved in the other direction, with a common regulatory and legal framework for AI coming into force last year. On page 8, we ask our panel of experts whether the UK needs its own AI act, while on page 11, co-chair of the All-Party Parliamentary Group on Artificial Intelligence Lord Clement-Jones makes the case for why regulations could be a catalyst for greater AI success.

Away from Westminster, we investigate the growing use of AI within local government (page 14); Josh Simons MP outlines efforts to establish Wigan as a hotbed for constructive AI intervention (page 16); and we look at how one solves the paradox of placing both AI and job creation at the heart of industrial strategy (page 20).

As AI reshapes policy, services and industry alike, the challenge – and opportunity – for Britain is clear: to lead with purpose, govern with foresight and grow with the public firmly in mind.

Contents

4 / View from parliament
Chi Onwurah: Defining the UK's role in global AI

5 / View from opposition
Jonathan Camrose calls for a sense of urgency

8 / Symposium
Does the UK require an AI Act?

11 / View from the Lords
Tim Clement-Jones: Regulations as AI catalyst

14 / Feature
The role of responsible AI in local government

16 / View from elsewhere
How Wigan is harnessing AI at a local level

20 / Feature
The AI, jobs, industrial strategy paradox

22 / Q&A
Ada Lovelace director Gaia Marcus

Spotlight

40-42 Hatton Garden
London
EC1N 8EB

THE NEW STATESMAN
Subscription inquiries:
help@subscribe.newstatesman.co.uk

Publisher
Peter Coombs

Sales Manager
Dominic Rae
Account Managers
Jugal Lalsodagar
Lucy Sloan
Nia Ayoade
Kalpesh Vadher

Design
Erica Weathers
Cover Illustration
Klawe Rzeczy

Editor
Phin Foster

Associate Editors
Jonny Ball
Megan Kenyon

Editorial
Harry Clarke-Ezzidio
Samir Jeraj

BRITISH SOCIETY OF MAGAZINE EDITORS AWARDS 2017 WINNER

First published as a supplement to the New Statesman of 13 June 2025.

© New Statesman Ltd. All rights reserved.

Registered as a newspaper in the UK and US.
The paper in this magazine is sourced from sustainable forests, responsibly managed to strict environmental, social and economic standards. The manufacturing mills have both FSC and PEFC certification and also ISO9001 and ISO14001 accreditation.

This supplement can be downloaded from:
newstatesman.com/spotlight/reports

The view from parliament

Chi Onwurah MP
House of Commons Science, Innovation and Technology Select Committee chair

"The UK must leverage its unique strengths to become a trusted leader in AI development and governance"

As Select Committee chair, I hear witness after witness reference the role of AI in everything from recruitment processes to CDP figures. Britain's role in the global AI ecosystem isn't just about growing our tech sector – already the world's third largest – but about laying the foundations of our economy for a new age.

Unlike its predecessors, this government tells a positive story about how AI could be shaped to serve the public good and boost Britain's growth as a centre for responsible AI use. I was delighted to see the Prime Minister throw the full force of government behind the AI Action Plan, pledging concrete steps to drive adoption in the public and private sectors; we need to go hard and fast to close the gap with international competitors.

Open-source AI, which I've long been committed to championing, represents the opportunity to build trust and democratise technology. The benefits are plain to see. It provides a mechanism for scrutiny, encourages international cooperation on a matter of huge general interest, allows the UK to become a diplomatic and technological bridge and stimulates entrepreneurship. Open-source initiatives could catalyse an entire generation of SMEs and start-ups.

Public sector productivity is one of the most enduring challenges in Britain, and the public directly faces the consequences. There is enormous potential for AI to improve such outcomes. I have long encouraged the parliamentary authorities to roll out secure AI tools that could save my staff time – time that could then be spent helping constituents in Newcastle. The government claims potential savings of up to £45bn through the digitisation of services.

In our NHS, there is huge scope for AI to support management and boost the productivity of hospitals and GP surgeries, saving precious time and money.

In climate tech, AI is being used to optimise energy systems and improve climate modelling. The Alan Turing Institute is using AI to help understand the effects of climate change, protect communities and natural habitats and develop ways of reducing greenhouse gas emissions.

Programmes such as the High Value Manufacturing Catapult, which investigates innovative technologies and scales up new products and processes, is integrating AI for predictive maintenance, process optimisation and supply chain resilience – boosting productivity and innovation across UK industry.

The Committee has heard from local government representatives who are using digitisation to enhance performance and efficiency. Sutton Council uses sensors to monitor the routines of elderly citizens. Not getting up at the usual time may signal an emergency, and then first-responders are on the scene much faster.

Those who say AI is irrelevant to the challenges of economic growth in the public sector fundamentally misunderstand the nature of the opportunities available, its presence everywhere in our lives already, the frenetic pace of its delivery and its ability to drive growth.

We are exceptionally gifted by our access to the infrastructure and institutions needed to conduct research and take advantage of the opportunities posed in adopting AI. I was lucky enough to attend the opening of Northumbria University's Centre for Responsible AI, which aims to enhance the UK as a leader in the ethical implementation of AI.

Britain's optimal role in the global AI ecosystem must be as a hub for great tech, a magnet for great talent and a centre for ethical and responsible deployment. We must leverage our unique strengths to become a trusted, innovative and ethical leader in AI development and governance.

MARTA SIGNORI

The view from the opposition

Jonathan Camrose
Shadow minister for science, innovation and technology

“The government is moving too slowly. It’s time for boldness”

When I first started grappling with the problem of AI and copyright as a minister in 2023, I grossly underestimated the difficulty: the two sides could surely be brought together to find compromises and workarounds. Sadly, it became ever clearer over months of talks that there was no mutually agreeable landing zone – we could satisfy the content creators or the AI labs, but not both.

As the government approaches the first anniversary of its own forlorn attempts to find the elusive middle ground and temperatures rise on both sides, I’ve been urging them to take a different approach.

First, let’s think more clearly about where we want to get to. The eventual answer is surely a trusted, efficient marketplace for the use of copyrighted materials, where rights-holders can freely choose to license, sell or withdraw their property and developers can make rapid and affordable commercial choices. Prices will be set by supply and demand, with smaller rights-holders represented by collectives who process transactions on their behalf and distribute revenues.

Second, what are the barriers that prevent us building such a marketplace? Among a great many, two stand out: offshoring and transparency. The UK can make any laws it likes, but any AI developer who doesn’t like them can offshore training activity to a jurisdiction where they can legally conduct the same training with the same content. Not only would the problem remain unsolved for our rights-holders, but we would be pushing AI activity out of the country.

Which brings us to transparency: there is a view that we could require any AI model used in the UK to declare all of the content used in its creation, thus solving the offshoring problem. I’m afraid this is wishful thinking.

How would this vast quantity of information be sought, verified and audited? How would we enforce it? How would we establish a direct link between the expressed output and the suspected input? What if the training material is not copyrighted but an imitation of an imitation (and so on) of copyrighted material?

I don’t believe there is a form of words that can be made into a law that fixes AI and copyright in a way that satisfies all sides. I do, however, believe that a combination of technologies, standards and law can help us build the marketplace for copyright.

The key technology is going to be machine-readable digital watermarking that can be indelibly embedded in any file and would contain licence information.

Crucially, rights-holders would need to easily (or automatically) apply watermarks to all of their material. This starts to address the transparency problem, but it only works if everyone agrees to use the same design – or at least to choose from a limited number of designs. It would depend, in short, on the existence of globally agreed technical standards – hardly a novelty in the internet age. Armed with such standards, governments would be far better placed to make the laws to create a fair and trusted marketplace for copyrighted materials.

Globally standardised digital watermarks solve the transparency problem and transparency solves the offshoring problem.

The government is moving too slowly. Technology is driving change more quickly now – it’s time for boldness and agility. No country has solved this problem, and if we can then we have a chance to be a global leader in AI again. ●

MARTA SIGNORI

Advertorial

Digital sovereignty should sit at the core of the UK's AI strategy

The public and private sectors must collaborate to create homegrown tools built on transparent supply chains

By Richard Davies

In association with Netcompany

With the Prime Minister's AI Opportunities Action Plan, the UK government pledged to turbocharge the economy by infusing AI throughout the public sector. From hospitals leveraging AI for faster diagnoses to public sector teams freed from administrative drudgery, the goal is to use AI as the engine of British progress.

But as the government throws its weight behind this technological revolution, several crucial questions arise: who owns and controls the digital foundations upon which our AI-powered future will be built? What tools, platforms and companies make up the digital supply chains of public and private sector services? And how can we ensure that homegrown innovations in AI are scalable?

Ultimately, the challenge lies in establishing "digital sovereignty" – ensuring the UK can secure and govern the foundations of its AI-driven future.

In times of global unrest and economic uncertainty, digital sovereignty is a necessity, not a luxury. It means the UK retaining control over its critical technological infrastructure, data and algorithms. It's about ensuring that the tools underpinning our public services and industries are not black boxes managed from afar but transparent, accountable systems shaped by our values.

The risks of dependency are real. Over-reliance on foreign-owned platforms can expose our institutions to security vulnerabilities, regulatory misalignment and loss of control over sensitive data.

And yet, pragmatism will need to be practised. Technological supply chains will undoubtedly cross international lines. Achieving digital sovereignty, therefore, requires a balanced approach: ensuring transparency so the public can understand these supply chains, prioritising domestic and European technology solutions, and working with a carefully vetted group of international partners.

This approach will also help the UK tackle one of its biggest challenges with AI: scaling projects from proofs of concept to delivering value more quickly and widely. Digital sovereignty empowers the UK to set its own standards, foster innovation within a trusted ecosystem

and maintain control over the process of moving AI projects from concept to widespread implementation.

With this in mind, consider Humphrey, the UK government's new AI assistant, which is being trialled in 25 local authorities to streamline administrative tasks such as planning, archive searches and transcription.

Early results are promising. Government pilots found that Humphrey's "Minute" notetaking tool saved officials an hour of admin for every 60-minute meeting, freeing staff to focus on higher-value work and improving morale.

Other components, like "Consult", can analyse thousands of public consultation responses far faster than human teams, with comparable accuracy and significant projected cost savings across the civil service.

If the platform can continue to deliver such results as its adoption scales, Humphrey may serve as a valuable case study for public-sector AI implementation. At the same time, with increasing attention on how governments manage and govern AI tools, providing clear information about the platform's technical underpinnings – from the large language models powering it to its hosting setup – will help build confidence and set standards for future initiatives.

Digital sovereignty fits into a wider framework of responsible digitalisation – a guiding principle for Netcompany. It means deploying technology in ways that are ethical, transparent and aligned with societal needs.

Our experience delivering large-scale digital projects across the UK and European public sector has shown that responsible digitalisation is not only possible but essential for building trust and ensuring long-term impact. Whether deploying a digital patient registration service used by 98 per cent of English GPs or developing an AI-powered delay prediction tool for rail networks across Europe, we let our customers take control of their processes and data, foster collaboration and commit to re-using technologies, never developing the same tools twice.

The same goes for the EASLEY AI platform. Developed by Netcompany, EASLEY is a secure, model-independent generative AI solution for both public and private sector organisations. Unlike many off-the-shelf AI products, EASLEY puts data privacy and organisational control at its core. It integrates seamlessly with existing systems, allowing clients to switch between AI models as technology evolves – without relinquishing control over their data or processes. In practice, this means a local authority can automate document processing or improve citizen services with confidence, knowing their data never leaves UK or European jurisdiction.

SHUTTERSTOCK

Achieving digital sovereignty requires a pragmatic, balanced approach

Legacy IT systems are silent saboteurs of digital progress. Across the UK and Europe, outdated infrastructure drains budgets and stifles innovation, with up to 80 per cent of IT budgets spent just keeping these obsolete systems running – resources that could otherwise fund better digital services, innovation and security.

Earlier this month, we announced Feniks AI, a pioneering tool that accelerates the transition from legacy systems to modern, open architectures – cutting delivery times by up to 60 per cent. In short, what once took years can now be completed in months.

The tool has already delivered promising results in three large-scale public sector projects in Denmark, and we look forward to bringing it to the UK.

Feniks AI is built on Netcompany's unique methodology and platforms, developed through 25 years of experience delivering large-scale, business-critical projects across the public and private sectors in Europe. By embracing such solutions, we can help our customers break free from decades of digital debt and lay the foundations for a more innovative and secure future.

As the UK charts its course towards an AI-powered future, cross-sector collaboration is key to delivering digital transformation at scale. Partnerships focused on transparency, scalability and pragmatic digital sovereignty will best position the UK to become a leader in the development and deployment of AI.

In doing so, we can shape a digital landscape that is not only world-leading but also serves the needs and aspirations of our citizens. ●

Richard Davies is country managing partner at Netcompany

Does the UK need an AI Act?

Policy experts on whether there is a need for greater government intervention

Britain finds itself at a crossroads with AI. The stakes are heightened by the fact that out closest allies appear to be on diverging paths.

Last year, the EU passed its own AI act, seeking controlled consensus on how to regulate new technologies. The US, meanwhile, is pursuing a lighter-touch approach to AI – perhaps reflecting the potential financial rewards its Big Tech companies could lose if stifled by regulation. Prime Minister Keir Starmer and Science Secretary Peter Kyle seem to be mirroring the US strategy.

In the January launch of the government's AI Opportunities Action Plan, Kyle wants Britain to "shape the AI revolution rather than wait to see how it shapes us". Many have called for the government to bring forward an AI act, to lay the foundation for such leadership. Does Britain need one, and if so, how stringent should it be? *Spotlight* reached out to sectoral experts to give their views.

Gina Neff

Professor of responsible AI at Queen Mary University of London

This government is betting big on AI, making promises about turbo-charging innovation and investment. But regulatory safeguards are fragmented, public trust remains uncertain, and real accountability is unclear. Charging forward without a clear plan means AI will be parachuted into industries, workplaces, and public services with little assurance that it will serve the people who rely on it.

An AI act would signal that Britain is serious about making AI work for people, investing in the places that matter for the country, and harnessing the power of AI for good. An AI act would create oversight where there is ambiguity, insisting on transparency and accountability. An AI act could provide the foundation to unlock innovation for public benefit by answering key questions: who is liable when AI fails? When AI systems discriminate? When AI is weaponised?

Starmer's government borrows from Silicon Valley's logic, positioning AI regulation as the opposite of innovation.

Keir Starmer with Peter Kyle, who wants the government to help "shape the AI revolution"

Such logic ignores a crucial fact: the transition to AI will require a major leap for workers, communities and societies. Government must step in where markets won't or can't: levelling the playing field so powerful companies do not dominate our future, investing in education and skills so more people can benefit from opportunities, ensuring today's laws and regulations continue to be fit for purpose, and building digital futures with companies and civil society.

Under Conservative governments, the UK took a "proportionate", "pro-innovation" approach outlined in the AI White Paper, suggesting responsibility for safe and trustworthy AI rests with the country's existing 90 regulators. That was always envisioned to be a wait-and-see stop-gap before new measures. The AI Opportunities Action Plan sketches out support for the UK's AI industry, but does not go far enough on how to manage the social, cultural and economic transitions that we face.

With worries about the impact on entry-level jobs, on our children, on information integrity, on the environment, on the UK's creative sector, on growing inequality, on fair yet efficient public services: there is a long list of jobs now for government to do. Lack of action will only create confusion for businesses and uncertainty about rights and protections for workers, consumers and citizens. Without an AI act to help shore it up, the good work that is already happening in the UK won't be able to fully power benefits for everyone.

An AI act must go beyond data protections to establish transparency requirements and accountability provisions, outline safeguards for intellectual property, set clearer rules around and recourse for automated decision-making. These are responsibilities that tech companies are largely evading. Who can blame them? They have cornered global markets and will gain handsomely with our new investments in AI. A UK AI act could empower regulators with stronger enforcement tools to right the imbalance of power between British society and the world's biggest players in this sector.

An AI act would give real structure to this country's ambitions for AI. The UK needs clarity on what AI can and cannot do, and that won't come from piecemeal guidance – it will come from leaders with vision helping us build the society that we all so rightly deserve.

Marina Jirotka, Keri Grieman

Professor of human-centred computing at the University of Oxford; research associate, RoboTIPS project

The EU AI act entered into force not even a year ago, and there is already serious discussion on whether to reduce enforcement and simplify requirements on small and medium enterprises in order to reduce burdens on companies in a competitive international marketplace. The US House of Representatives has narrowly approved a bill that blocks states from enforcing AI regulations ▸

CHARLES MCQUILLAN / GETTY IMAGES

◂ for ten years, while forwarding one bipartisan federal act that criminalises AI deepfakes but does not address AI on a broader level. Large language model updates are rolled out faster than the speed of subscription model billing. AI is invading every corner of our lives, from messaging apps to autonomous vehicles – some used to excellent effect, others to endless annoyance.

The British government has chosen a policy of investment in AI – investing in the industry itself, in skill-building education and in inducing foreign talent. Its hesitancy to regulate seems borne out of the fear of hobbling a potential cash cow. However, this leaves the regulatory burden on individual sectors: piecemeal, often siloed and without enough regulatory AI experts to go around, with calls coming from inside the house – the companies themselves – for a liability system.

The UK needs clarity: for industry, for public trust and for the prevention of harm. There are problems that transcend individual industries: bias,discrimination, over-hype, environmental impact, intellectual property and privacy concerns, to name a few. A regulator is one way to tackle these issues, but can have varying levels of impact depending on structure: coordinating between industry bodies or taking a more direct role; working directly with companies or at arm's length; cooperative investigation or more bare-bones enforcement.

But whatever the UK is to do, it needs to provide regulatory clarity sooner rather than later: the longer the wait, the more we fail to address potential harms, but we also fall behind in market share as companies choose not to bet the bank on a smaller market with an unclear regulatory regime.

Baroness Beeban Kidron

House of Lords member and digital rights activist

All new technology ends up being regulated. On arrival greeted with awe. Claims made for its transformative nature and exceptionality. Early proponents build empires and make fortunes. But sooner or later, those with responsibilities for our collective good have a say. So here we are again with AI.

Of course we will regulate, but it seems that the political will has been captured.

Those with their hands on the technology are dictating the terms – terms that waver between nothing meaningful to almost nothing at all. While government valorises growth and efficiency without asking: growth for whom? Efficiency to what end?

In practical terms, an AI act should not seek to regulate AI as a technology but rather regulate its use across domains: in health (where it shows enormous benefit); in education (where its claims outweigh its delivery by an unacceptable margin); in transport (where insurers are calling the shots); and in information distribution (where its deliberate manipulation, unintended hallucination and careless spread damages more than it explains).

If we want AI to be a positive tool for humanity then it must be subject to the requirements of common goods. But in a world of excess capital restlessly seeking the next big thing, governments bent over to do the bidding of the already-too-powerful, and lobbyists who simultaneously claim it is too soon and too late, we see the waning of political will.

Regulation can be good or bad, but we are in troubling times where the limit of our ambition is to do what we can, not what we should – which gives it a bad name. And governments – including our own – legislate to hardwire the benefits of AI into the ever-increasing concentration of power and wealth of Silicon Valley.

Tech companies, AI or otherwise, are businesses. Why not subject them to corporate liability, consumer rights, product safety, anti-trust laws, human and children's rights? Why exempt them from tax, or the full whack for their cost to planet and society?

It's not soon and it is not too late – but it needs independence and imagination to make AI a public good, not wilful blindness to an old-school playbook of obfuscation and denial while power and money accumulate.

Yes, we need regulation, but we also need political will.

Michael Birtwistle

Associate director, Ada Lovelace Institute

AI is everywhere: our workplaces, public services, search engines, our social media and messaging apps. The risks of these systems are made clear in the government's *International AI Safety Report*.

Alongside long-standing harms like discrimination and 'hallucination' (where AI confidently generates false information), systemic harms such as job displacement, environmental costs and the capacity of newer 'AI agents' to misinform and manipulate are rapidly coming to the fore.

But there is currently no holistic body of law governing AI in the UK. Instead, developers, deployers and users must comply with a fragmented patchwork of rules, with many risks going unmanaged.

Crucially, our current approach disincentivises those building AI systems from taking responsibility for harms they are best placed to address; regulation tends to only look at downstream users.

Our recent national survey showed 88 per cent of people believe it's important that the government or regulators have powers to stop the use of a harmful AI product. Yet more than two years on from the Bletchley summit and its commitments, it's AI developers deciding whether to release unsafe models, according to criteria they set themselves.

The government's own market research has said this "wild west" is lowering business confidence to adopt.

These challenges can only be addressed by legislation, and now is a crucial time act.

The government has announced an AI bill, but its stated ambition (regulating "tomorrow's models not today's") is extremely narrow.

For those providing scrutiny in parliament, press and beyond, the real test of a bill will be whether it credibly responds to the growing list of everyday harms we see today — such as bias, misinformation, fraud and malicious content — and whether it equips government to manage them upstream at source. ●

The view from the Lords

Tim Clement-Jones
Liberal Democrat peer and spokesperson for the digital economy

"The idea that AI regulation stifles innovation needs turning on its head"

Ever since co-founding the All-Party Parliamentary Group on AI nine years ago, still ably administered by the Big Innovation Centre, I've been deeply involved in debating and advising on the implications of artificial intelligence. My optimism about AI's potential remains strong – from helping identify new Parkinson's treatments to DeepMind's protein structure predictions that could transform drug discovery and personalised medicine.

Yet this technology is unlike anything we've seen before. It's potentially more autonomous, with greater impact on human creativity and employment, and more opaque in its decision-making processes.

The conventional wisdom that regulation stifles innovation needs turning on its head. As AI becomes more powerful and pervasive, appropriate regulation isn't just about restricting harmful practices – it's key to driving widespread adoption and sustainable growth. Many potential AI adopters are hesitating not due to technological limitations but uncertainties about liability, ethical boundaries and public acceptance. Clear regulatory frameworks addressing algorithmic bias, data privacy and decision transparency can actually accelerate adoption by providing clarity and confidence.

Different jurisdictions are adopting varied approaches. The European Union's AI Act, with its risk-based framework, started coming into effect this year. Singapore has established comprehensive AI governance through its model AI governance framework. Even China regulates public-facing generative AI models with fairly heavy inspection regimes.

The UK's approach has been more cautious. The previous government held the AI Safety Summit at Bletchley Park and established the AI Safety Institute (now inexplicably renamed the AI Security Institute), but with no regulatory teeth. The current government has committed to binding regulation for companies developing the most powerful AI models, though progress remains slower than hoped. Notably, 60 countries – including Saudi Arabia and the UAE, but not Britain or the US – signed the Paris AI Action Summit declaration in February this year, committing to ensuring AI is "open, inclusive, transparent, ethical, safe, secure and trustworthy".

Several critical issues demand urgent attention. Intellectual property: the use of copyrighted material for training large language models without licensing has sparked substantial litigation and, in the UK, unprecedented parliamentary debate. Governments need to act decisively to ensure creative works aren't ingested into generative AI models without return to rights-holders, with transparency duties on developers.

Digital citizenship: we must equip citizens for the AI age, ensuring they understand how their data is used and AI's ethical implications. Beyond the UAE, Finland and Estonia, few governments are taking this seriously enough.

International convergence: despite differing regulatory regimes, we need developers to collaborate and commercialise innovations globally while ensuring consumer trust in common international ethical and safety standards.

Well-designed regulation can be a catalyst for AI adoption and innovation. Just as environmental regulations spurred cleaner technologies, AI regulations focusing on explainability and fairness could push developers toward more sophisticated, responsible systems.

The goal isn't whether to regulate AI, but how to regulate it promoting both innovation and responsibility. We need principles-based rather than overly prescriptive regulation, assessing risk and emphasising transparency and accountability without stifling creativity.

Achieving the balance between human potential and machine innovation isn't just possible – it's necessary as we step into an increasingly AI-driven world. That's what we must make a reality.

MARTA SIGNORI

Advertorial

Resilience starts with design and disclosure

Cybersecurity vendors are key to embedding security across product development and placing responsible disclosure at the core of operations

By Richard Woolfrey

In association with FORTINET

From the emergence of AI to the rise of 'ransomware as a service', the methods, tools and campaigns used by cybercriminals continue to evolve. Organisations must stay one step ahead, with cybersecurity critical to doing so – and essential for business survival.

Adopting a proactive rather than reactive approach towards cybersecurity – identifying and addressing gaps before an attack takes place – is crucial to minimising the risk of an attack and improving overall protection. But organisations can't do this alone. Cybersecurity vendors have a responsibility to lead, making it easier for customers to adopt resilient strategies.

So how can this best be done? By executing 'secure by design', which incorporates cybersecurity into every stage of the development and design process, vendors can build security into the products customers rely on, making their systems secure from the very beginning. Also, integrating responsible disclosure, where vulnerabilities are identified and addressed before exploitation or public disclosure, increases trust and transparency while reducing the likelihood of an attack.

It's essential that every vendor commits to these two approaches. Doing so allows them to support customers not just in protecting their data but also in safeguarding overall business operations and building long-term cyber resilience.

It's clear that cyber attacks are continuing to grow in scale and complexity. Fortinet's *2025 Global Threat Landscape Report* found that tools like FraudGPT, BlackmailerV3 and ElevenLabs are automating the generation of malware, phishing websites and deepfake videos, leading to increasingly effective threat campaigns while lowering barriers to entry. The underground economy for stolen credentials and direct corporate access has exploded in the past year, while cloud environments – on which most organisations rely – remain a top target for cybercriminals.

Ultimately, cybercriminals are getting smarter, and the rate of innovation is outpacing the speed that vendors and customers can protect themselves. Cyber resilience now demands more than traditional firewalls and anti-virus

Cybersecurity measures need to be incorporated from the very first stage and throughout the SDLC

DIYAJYOTI, SHUTTERSTOCK

software. Instead, it requires a proactive, multi-layered approach that makes cybersecurity an integral part of business strategy from the very beginning – an approach vendors must be accountable for.

One way this can be achieved is by ensuring products are secure by design. This means embedding cybersecurity measures into every stage of the software development cycle (SDLC), from planning and design to deployment, maintenance and decommissioning. While the benefits of this to customers include helping them adopt a more proactive approach towards cybersecurity, the responsibility lies with vendors to ensure products and solutions are secure by design.

Fortinet is committed to a secure-by-design approach and is actively driving progress towards industry-wide standards and innovations. This includes signing the Cybersecurity and Infrastructure Security Agency's Secure by Design Pledge in 2024.

The pledge outlines ways organisations can implement "secure by design", including security patches, multi-factor authentication (MFA) and default passwords.

Alongside these recommendations, implementing secure coding standards, including input validation, proper error handling and avoiding unsafe functions or outdated libraries, will play a role by minimising vulnerabilities being introduced during development and ensuring every line of code is written with security in mind. Using a memory-safe programming language can also prevent common vulnerabilities related to unsafe memory handling.

Employing hardware-backed cryptographic key management provides another way to protect sensitive data by storing encryption keys within dedicated hardware modules, such as 'hardware security modules' (HSMs) or 'trusted platform modules' (TPMs). This helps to prevent keys from being exposed in software memory, reducing the risk of key theft or compromise and enhancing system security.

In parallel, organisations should also invest in rigorous testing frameworks that include static and dynamic code analysis and vulnerability scanning throughout the development process. Testing supports continuous improvement by making sure security measures are implemented correctly and fed back into design.

Embedding security from the onset starts with cybersecurity vendors. Encouraging customers to adopt solutions that are secure by default means potential vulnerabilities can be identified, contained and mitigated before impacting the wider business, supporting organisations to stay ahead of evolving threats. ●

Richard Woolfrey is regional director, UK & Ireland, at Fortinet

Computer says maybe

Local government is adopting AI, but careful thought is needed

By Samir Jeraj

Artificial intelligence is already being widely used across local government to deliver services. These involve everything from customer service chatbots to writing paperwork to preventing damp.

Central government is currently piloting a set of AI tools for local government. Other councils are working with commercial AI providers. Somerset Council was one of the first local governments in the UK to join Microsoft's Early Access Program for AI.

Somerset started with Copilot, a generative AI chatbot, and has now moved on to Magic Notes, which turns meeting recordings into a first draft of paperwork and forms and is now used by around half of social care teams in the UK as a way to improve the quality of paperwork while reducing the time spent on it by 50 per cent. Somerset Council says this has allowed frontline staff to free up time to be spent on "person-centred work" and improved their work-life balance.

AI has great potential to improve democratic processes and decision-making and from there flows improved services for citizens, according to Alexander Iosad, director, government innovation policy at the Tony Blair Institute. AI can be used to engage with citizens and get data from them to accelerate the back-office process and to improve decision making, he explains.

One problem of the traditional process for policymaking is having "more data than you can turn into actionable insights", Iosad says. AI is not just able to process all of that information, and to do it faster so the information is still current or in real time, but is able to cross-reference it across services.

However, there are still significant hurdles to integrating AI tools into council services and realising its potential. A survey of county councils and unitary councils by the County Councils Network (CCN) found that while 85 per cent were using AI (and the remaining 15 per cent were considering using it), local governments felt that a lack of staff capacity, funding and training were holding them back. The Tony Blair Institute report says councils "lack the confidence, capabilities and infrastructure required to unleash this innovation".

The Tony Blair Institute worked with

a local council in the UK and estimated that AI could be applied to 26 per cent of tasks, resulting in a saving of one million work hours, or £30m in financial costs, each year. Scaled up nationally, this would mean around £8bn in cost savings.

Organisations of all types, but particularly government and public services, are looking to AI to save them money, but less than half of those surveyed by the CCN said AI had resulted in savings so far. This compared to over 90 per cent reporting improved staff productivity and three quarters that said it improved services. Transformation of services usually requires significant investment in order to realise improvements to the quality of services and their savings. Big digital projects have historically been difficult to deliver and to realise their benefits, particularly in the public sector. In 2021, Birmingham City Council invested £19m in a new IT system with the aim of reducing costs, but spending on getting the system up and running has since spiralled to £90m.

There are also questions over some of the ways AI has already been deployed. Several councils have bought and use AI systems to identify families in need of support and young people who may be at risk of becoming involved in gangs, while others use them to highlight potential council tax fraud. These have been mired in allegations of bias, consent and lack of transparency.

In Rotterdam, an algorithm was used to assign a fraud risk score to residents based on data from 12,707 previous investigations. But in practice, women, young people, parents and migrants got higher scores and were more often flagged for investigation. Subjective assessments by caseworkers, such as whether a person was "flexible" enough to deal with the challenges they were facing, were incorporated into the data, giving it the sheen of neutrality. The city suspended its use in 2021 following an external ethics review. In 2015, Hackney Council paid for an Early Help Profiling System provided by the company Xantura that used data collected by the council to identify families that were "at risk" – the families were never informed their data would be used in this way.

These challenges mean the governance, transparency and privacy policies around AI in local government need to be robust. The Local Government Association says "councils must also be mindful of the risks and challenges associated with safe and responsible AI adoption", listing strong data foundations, data protection, bias and privacy as key issues to address. The Tony Blair Institute advises councillors to champion AI innovation, but to de-risk it by looking to tools that are already in use, collaborating, using existing standards and training the workforce to improve AI literacy.

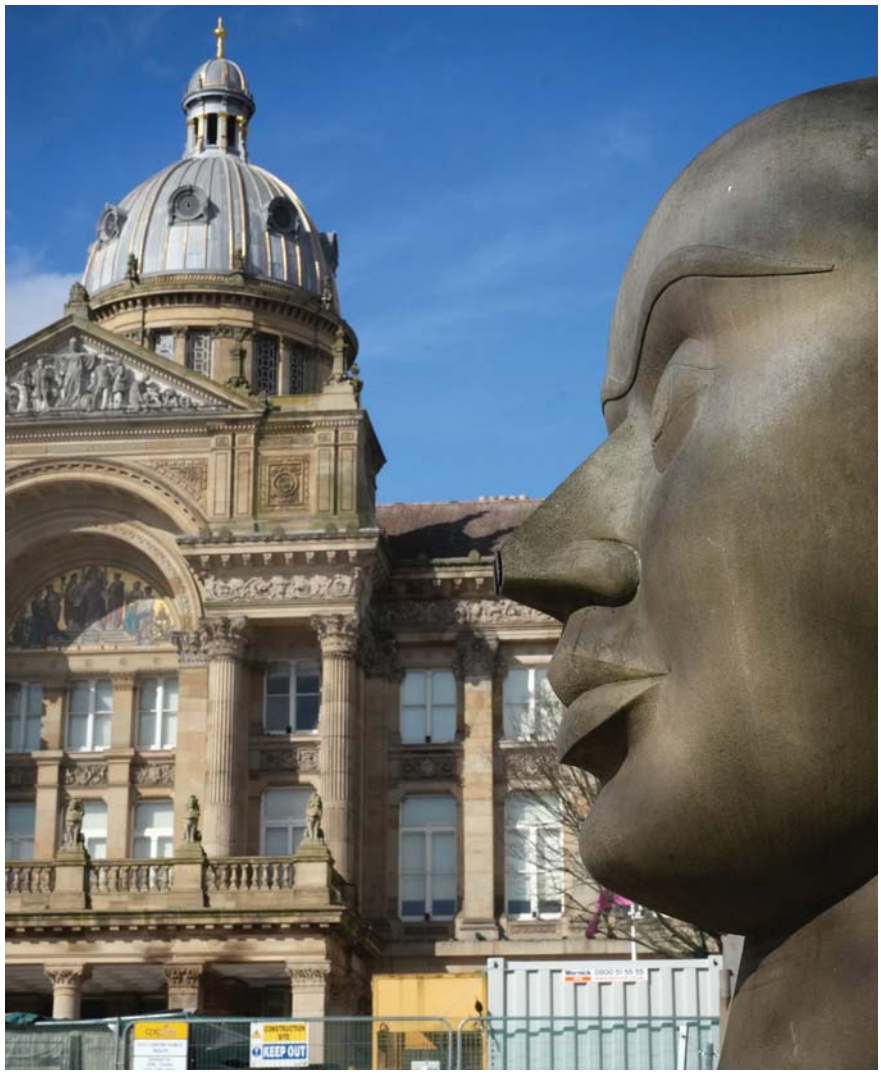

Birmingham Council is having to cut services severely to meet its costs

"The biggest challenge is data," says Iosad. "It's getting the right data and then setting up the systems allowing you to work with data effectively." Where the UK is arguably weakest, he says, is having data sets collected by government that can be linked and used together in order to really personalise services and to make effective decisions at local, regional and national levels.

"The *blueprint for modern digital government* has a very encouraging line on this, not just recognising that there is a challenge, but mostly setting out the ambition that every data set that is used in the public sector has an API," Iosad says. He adds this will also put in place the safeguards required for data.

What is clear is that AI tools and technologies will be taken and used by local government in the way that new technologies have been in the past. Some of the risks and challenges around technical literacy, training, choosing good suppliers and managing those contracts are familiar. However, the challenges of transparency, scrutiny and identifying bias in complex systems are more novel and require local government not just to look for an easy answer but the best solution to create a better place and life for residents.

The view from elsewhere

Josh Simons MP and David Molyneux

"We must be more granular and specific to harness AI growth"

Artificial Intelligence is often used as a shorthand for the future of our economy. But to say we want a future shaped by AI says little about what that future looks like. AI is a tool, a way of generating answers to questions from massive quantities of data. An AI-powered future is a future that uses data to achieve things we care about. What those things are is up to us.

Beyond Whitehall, there is little serious grappling with what this future should look like for places like Wigan. As the Member of Parliament for Makerfield and the leader of Wigan Council, we have been thinking carefully about how AI can be harnessed to give working people in our area a better future. While it's a work-in-progress, we wanted to share it, because it will matter for leaders across the country in the coming years.

The government's approach is centred around AI Growth Zones. These are effectively energy-intensive data centres, which are important for the future of AI in the UK but do not, in themselves, generate good jobs or productivity growth in local areas. Data processing is not like coal mining: there is no relationship between the people who do it and the physical place the processing happens. We must be clear that AI will not herald a return of place-based industrial capitalism.

To actually harness AI to generate growth and employment in a local area, we must be much more granular and specific. It starts by looking at the real-world businesses that generate data and could use it better to drive efficiencies and productivity. In Wigan, we have UK-wide strengths in sports, food manufacturing and other process manufacturing that is generally high-volume and low-margin. For instance, Heinz has one of the largest food factories in Europe in Wigan, producing three million tins of beans a day, and the largest wet-wipe factory in Europe is in the Makerfield constituency. Local business Uncle Joe's Mint Balls has grown from selling sweets at a Wigan market stall in 1898 to being exported across the world today. It is a particular favourite in Japan.

These factories work on fine margins, producing hundreds of thousands of cheap outputs. On a daily basis, they generate millions of datapoints. Unlocking that data, using it to make improvements in the process of packaging, or bringing products to market faster and more cheaply, can be transformative to the productivity and profitability of these companies, and in turn, the wages and work of local people.

Or take sport. Wigan Warriors, which is owned by the *New Statesman's* owner, Mike Danson, has been one of the best rugby league teams in the world for most of the game's history, and this season, Wigan Athletic and Wigan Warriors women's teams have been almost undefeated. Each game generates hundreds of datapoints about each player, team tactics, fan attendance and spending, and a host of other pieces of information that clubs can use to professionalise and secure their finances. Wigan can lead the way in this kind of analysis.

AI is a tool. How you decide to use that tool, and how an area develops distinctive expertise in improving and deploying AI, will depend on that area's existing strengths. For us in Wigan, it's not about building cutting-edge large language models or generative AI, it is about

potential for global businesses building and doing things in the physical world to use data and AI to drive up productivity and boost our economy.

Local authorities can make a real difference here. They must collect, store and process data much more effectively to understand how best to serve their populations and inform better decisions. To make preventative healthcare a reality, target financial assistance or educational programmes, or get people back to work, Wigan is working to improve how it gathers and processes data.

A more data-driven council can be a more community-led council, more tailored to the needs of specific, hyper-local places and people. This approach has seen success in London, where the London Office for Technology and Innovation has pooled data from a variety of boroughs to give better data sets, to intervene on issues from community health to tackling widespread mould problems.

This summer, we will launch Scale Space North – an AI and digital innovation hub working with Wigan Council's digital partner, Agilisys. Wigan's Civic Centre has been repurposed to act as a catalyst for local start-ups and new business propositions. Such partnerships and digital innovation are vital for the public sector to radically transform to meet community needs and to create skills and jobs for the future.

Wigan powered Britain's industrial revolution. Mining, for example, was hard, physical work, but it gave people jobs, skills and, importantly, purpose. Wigan helped Britain win wars, become the world's wealthiest nation and to protect our country from threats.

Despite its promise, the internet revolution has provided little for northern towns like Wigan. Our public realm has been battered, while too many people, from kids to pensioners, have retreated into the addictive haze of social media and fragmented digital worlds.

People are angry, and they are right to be. Levelling up did not provide what our towns really need: a huge boost of productivity, capital investment and real attention to the potentials and the problems of specific places, from the ground up. The job of people like us is to ensure the AI and data revolution actually benefits working people. Often, this won't always be shiny or glamorous, it will be about understanding the distinctive strengths of an area, and how they can be built on to generate, store and analyse data better, then integrate data-driven tools across businesses and public organisations.

AI is neither a silver bullet nor a destructive disaster. It is a tool that should be part of the future we choose to build. It is we, the people, who give meaning to family, work and communities. The shape of our AI-powered future is up to us.

Josh Simons is the Labour MP for Makerfield and David Molyneux is the leader of Wigan Council

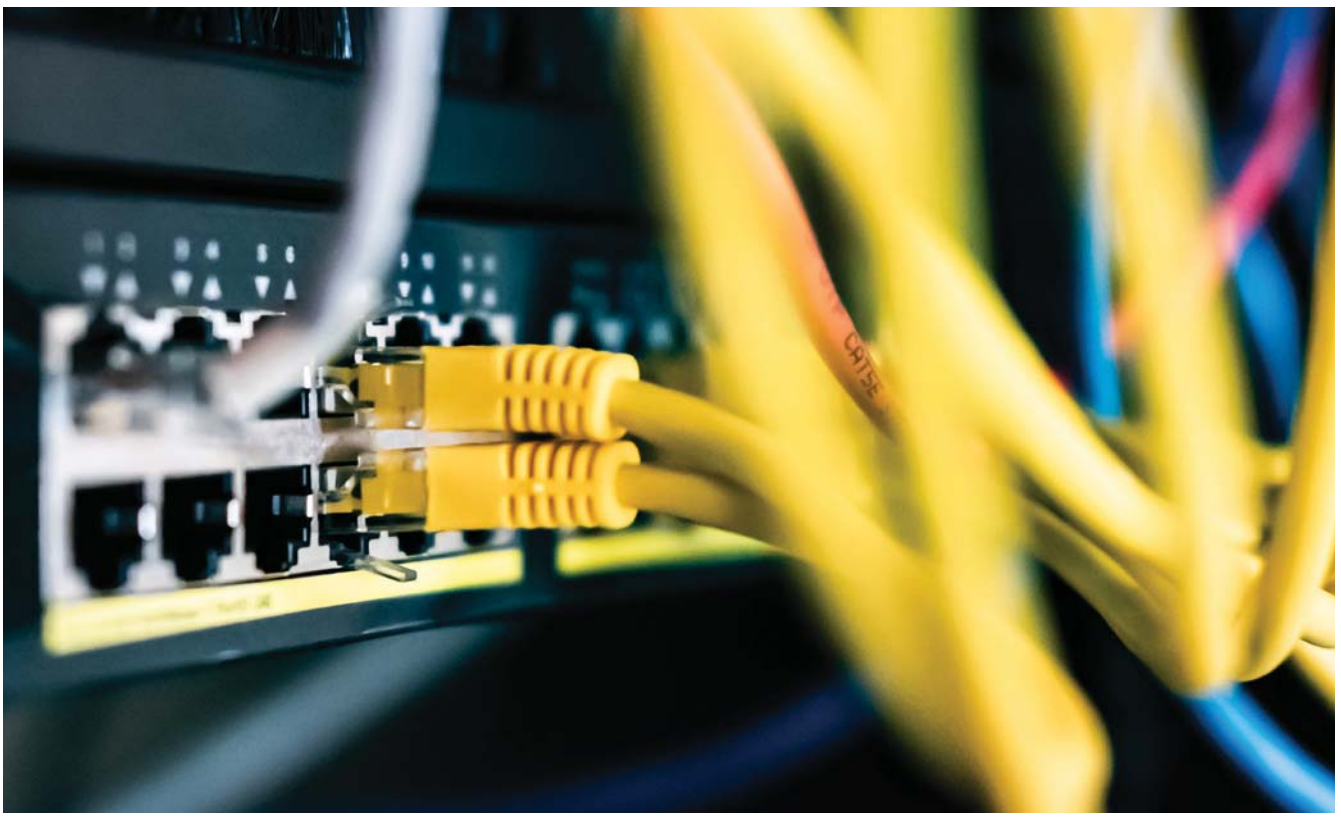

NICK BEER/SHUTTERSTOCK

Advertorial

A power for good?

AI is often seen as a threat, but it can help advertising regulation

By Guy Parker

In association with

For some, the advent of the worldwide web is still fresh in the memory. But technological leaps seem to happen with ever-increasing frequency, and we now all find ourselves blinking in the brilliant light at the dawn of the age of AI. At the Advertising Standards Authority (ASA), we've donned the sunglasses and rolled up our sleeves, and AI is already proving a game-changer in how we regulate.

The lightning speed with which AI has developed and integrated into our everyday lives inevitably raises legitimate concerns. What does it mean for jobs, data protection, originality, creativity, copyright, plagiarism, truth, bias, mis- and disinformation and what we think is fake vs real?

These are undoubtedly important issues to grapple with. But the technology also brings multiple benefits. As was the case in the mid-1990s with the launch of search, web browsers and online shops, there were innovators, early adopters, cautious sceptics and technology resisters. AI is no different. The ASA is firmly in the "early adopter" category. Four years ago, we appointed a head of data science and began building our AI capability; AI is now central to our transformation into a preventative and proactive regulator. Around 94 per cent of the 33,903 ads we had amended or withdrawn last year came from our proactive work using our AI-based Active Ad Monitoring system. The ability to be front-foot and take quick and effective action is crucial when regulating the vast online ecosystem. AI gives us much greater visibility of online ads.

Last year, our system scanned 28 million ads with machine learning and, increasingly, large language models finding the likely non-compliant ads we're interested in. That was a tenfold increase on 2023. Our target is to scan 50 million ads this year. AI-based tools are embedded in our work to help us monitor and tackle ads in high-priority areas and are now used in most of our projects, including our work on climate change and the environment, influencer marketing, financial advertising, prescription-only medicines, gambling and e-cigarettes. It's enabling us to carry out world-leading regulation – monitoring, identifying and tackling potential problem ads at pace and scale. Take one example: our ongoing climate change and environment project.

Following high-profile and precedent-setting rulings against major players in various industries, we're now seeing businesses adapting and evolving to make better evidenced, more precise green claims. Monthly sweeps using AI show high levels of compliance. Following our 2023 airline rulings on misleading "sustainable" and "eco-friendly" claims, of the circa 140,000 ads we monitored, we found just five that were clearly non-compliant.

Importantly, we're not removing humans from the equation. Our experts are and will remain central to our regulation. While our AI capability has dramatically improved the efficiency of our monitoring (weeding out the millions of ads that stick to the rules and aren't a problem), it filters and flags potential problem ads to our human specialists for their expert assessment. AI is assisting rather replacing our people. There are a lot of open questions about how AI will impact industries, positively and negatively. And that's certainly true of advertising, as ever at the forefront of technological change.

We know that the use of AI is already changing advertising. There are big efficiency and effectiveness gains in play. Lower-cost ad ideation and creation, hyper-personalisation and improved customer experience. Quicker and better media planning and buying. Get this right and ads will be cheaper to make and send, and be more engaging and relevant to receive. UK businesses and the British economy will be boosted. But in all of this, responsible ads must not be sacrificed at the altar of advances in technology.

We're well aware of the many potential benefits and problems AI poses for advertising. Think back to the story from Glasgow, where AI-generated ads promised a Willy Wonka-themed event that wasn't quite as advertised. The advertising of certain AI products and services certainly throws up broader ethical considerations. On our radar are ads for AI tech offering mental health support (substituting human therapists), essay-writing tools that pass work off as original, and chat boxes that act as a partner or friend. We don't regulate the products themselves, but in all these examples there is potential for ads to be misleading, irresponsible or harmful.

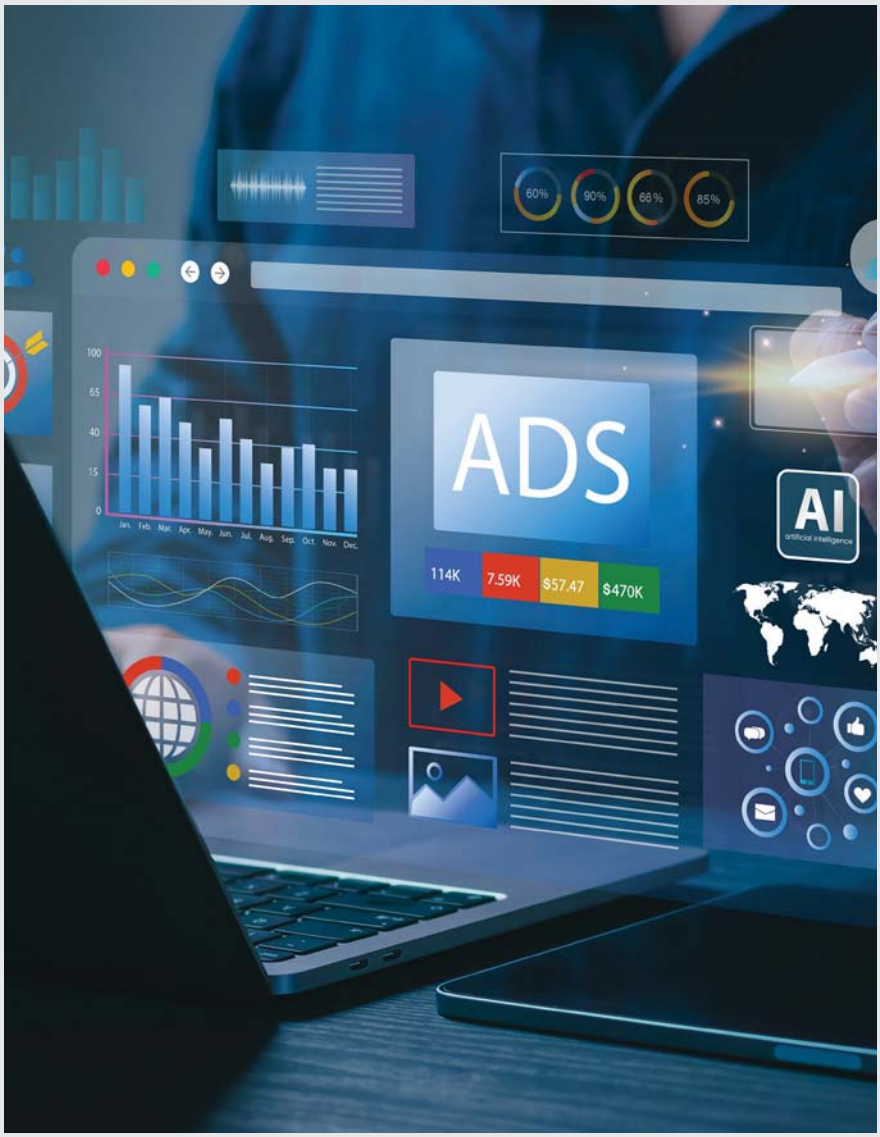

Advertising using AI throws up a number of ethical considerations

How can businesses use AI safely and responsibly? What does that mean for advertisers?

Our media and technology-neutral rules already cover most of the risks. Ads can't mislead, a principle as old as the hills. In the past, that might have been using photo-editing software; today, it might be through generative AI. Adverts must not be likely to cause harm or serious or widespread offence either. Generative AI might be an unsurpassed pattern-recogniser, but it's not a human and may well miss the nuance of judging prevailing standards in society when producing ad content.

Advertisers who harness AI can't abdicate responsibility for the creative content that it produces. That's why we urge businesses to be careful: use the good of AI, but avoid the bad. Put in place human checks and balances.

At the ASA, we're determined to take full advantage of technological advances, developing our Active Ad Monitoring system further and making even more use of large language models to speed up review of ads. Actively experimenting with how these tools can make our internal processes more efficient. And continuing to keep a close eye on how AI is used in advertising.

We are witnessing the next technological revolution that will change society in ways the internet did, perhaps even more. We can say with confidence that our use of AI is already delivering world-leading advertising regulation.

Guy Parker is chief executive of the ASA

BOY ANTHONY / SHUTTERSTOCK

Feature

AI, employment and the UK's industrial strategy

Invest 2035 aims to harness AI's potential while safeguarding jobs. Can the two ambitions coexist?

By Phin Foster

AI has emerged as both a panacea and the harbinger of a dystopian future. For policymakers crafting Britain's industrial future, the challenge is fully burnishing one side of that coin, while minimising exposure to the other.

Nowhere is this more evident than in the UK's Invest 2035 strategy, setting out a vision of economic renewal rooted in advanced technologies and regional growth. "Jobs will be at the heart of our modern industrial strategy," wrote Chancellor Rachel Reeves and Jonathan Reynolds, Secretary of State for Business and Trade, in their foreword to the strategy's draft consultation. At the same time, it promises "an ambitious approach to grow the AI sector and drive responsible adoption across the economy". The question is how, or whether, those two ambitions can fully succeed in tandem.

PwC's latest *Global AI Jobs Barometer*, published in early June, analysed close to a billion job postings across six continents. It found that demand for roles with high AI exposure expanded at a slower pace than those less affected by AI. The gap between these groups has widened since 2020, with jobs least exposed to AI experiencing a surge in listings.

Invest 2035 focuses on eight "growth-driving" sectors that are undoubtedly seeing greater AI penetration: advanced manufacturing; clean energy; the creative industries; defence; digital and technologies; financial services; life sciences; and professional and business services. PwC's report also found that roles with substantial AI exposure have undergone significantly more changes in skill requirements over the past five years. If jobs are to be at the heart of this new economy, one major challenge is how the UK equips its current and emerging workforce to fully engage.

The government's primary response to such questions lies with Skills England. Formally launched in 2024, the body is tasked with coordinating a fragmented post-16 skills system, and aligning it more closely with national economic priorities – including, crucially, the opportunities and risks posed by AI. With a mandate to "drive forward a skills system that meets the needs of employers, learners and the wider economy", it is a key lever in delivering not just more jobs, but better ones.

"Skills England... will ensure that our

workforce is equipped with the necessary skills to meet the demands of the modern economy," Phil Smith, Skills England chair, told MPs in a parliamentary debate in February.

The idea of lifelong learning – once a political platitude – has become a central pillar of this transition. As the economic landscape shifts faster than traditional education systems can keep pace with, workers increasingly need access to flexible, modular training that fits around existing jobs and responsibilities. Yet, as the IPPR has pointed out, the UK's investment in adult skills still lags behind many OECD peers.

In a widely discussed 2024 report on AI and labour market disruption, the IPPR argued that the UK faces a binary future. "Already existing generative AI could lead to big labour market disruption or it could hugely boost economic growth. Either way, it is set to be a gamechanger for millions of us," said senior economist Carsten Jung.

In that sense, industrial strategy cannot be separated from social policy. The government's proposed Advanced British Standard – a unified post-16 qualification intended to replace A-levels and T-levels – must prepare young people not only with subject knowledge but with the adaptability and analytical skills required in a rapidly evolving labour market. As automation touches roles from radiology to retail, the core employability question shifts from what you know to how quickly you can learn.

That shift is not evenly distributed across society. In towns with industrial legacies or fragile labour markets, AI is more likely to displace than create jobs – unless there is targeted, place-based intervention. Invest 2035 makes regional rebalancing a core ambition. But the delivery depends on ensuring that skills provision reaches not just growing tech clusters but also under-served communities. Community learning providers like the Workers' Educational Association (WEA) have a role not just in teaching but in building confidence and trust – especially for older or insecure workers who may feel alienated by the pace of technological change. WEA chief executive Simon Parkinson has called for long-term policy and funding stability so providers can scale up their work: taking training "to where people are, not where policy is most comfortable".

There are, however, lingering gaps between strategy and delivery. While government rhetoric supports "responsible adoption" of AI, it remains vague on how to mitigate job displacement in sectors most vulnerable to automation. Some analysts argue that Invest 2035, like its predecessors, risks overstating short-term innovation gains while underestimating longer-term disruption.

Rachel Reeves insists jobs will be at the heart of Labour's industrial strategy

Meanwhile, employer responsibility remains a crucial, under-addressed issue. If businesses are to adopt AI in a way that benefits workers, not just bottom lines, they must be incentivised to invest in staff retraining. At present, many treat upskilling as an externality – a public good best delivered by someone else.

Models such as Local Skills Improvement Plans (LSIPs), designed to give employers a greater role in shaping local post-16 training, have potential, but questions have been raised in regards to consistency of ambition and effectiveness. Without stronger national coordination, these initiatives may amount to well-meaning but fragmented efforts, rather than transformative change.

Indeed, the overarching challenge is neither technological nor economic – it is political. AI, like past waves of automation, will not distribute its rewards evenly or inevitably. The outcome will depend on the state's ability to shape markets and institutions in the public interest.

"AI... will transform jobs, destroy old ones, create new ones, trigger the development of new products and services and allow us to do things we could not do before," Jung writes in 2025's *The New Politics of AI: Why Fast Technological Change Requires Bold Policy Targets*. "But given its immense potential for change, it is important to steer it towards helping us solve big societal problems."

Invest 2035 boasts similarly grandiose ambitions. Its success will not rest on ambition alone, however, but on how convincingly it connects the dots between technology, training and trust. For AI and jobs to serve as dual engines of growth, the UK must resist the temptation to treat them as separate problems. They are, in reality, two faces of the same future.

Q&A

Gaia Marcus: "The deployment of AI often outpaces our ability to govern it effectively"

Gaia Marcus is the director of the Ada Lovelace Institute. Prior to joining, she was the deputy director (advanced analytics and local capabilities) at the Department for Levelling Up, Housing and Communities. The Ada Lovelace Institute is an independent research institute funded by the Nuffield Foundation.

What first attracted you to AI as a policy area?
I most enjoy roles that sit between technology and people, motivated by both an excitement for the new and a desire for fairness and opportunity for all. My first 'real job' was as social network analyst at the RSA, where I worked with local communities to translate quantitative and network research into projects boosting wellbeing and social connections. Having gone deeply into data policy as head of the UK's National Data Strategy, I strongly feel that AI policy is one of the biggest challenges of our time. How do we grasp the potential of these new technologies to build futures we all want to be part of?

When you look at developments in AI in the public sector, what most excites you?
I'm always excited by efforts to solve real societal problems using all the possible tools to hand. But tools alone aren't solutions.

Deploying AI in the public sector needs trust and legitimacy – and listening to the public is the only way to ensure these technologies work well, work for everyone and work in context. To me, the most exciting conversations are those involving the public's hopes, fears and aspirations about AI and data-driven technologies.

What frightens you?
"Frighten" is a big word, but there are things I worry about.

First, AI is a global value and supply chain and needs to be governed accordingly. Currently we are over-reliant on a few tech companies at most steps of this supply chain, who are largely marking their own homework and setting the narrative "weather" on policy.

Second, the deployment of AI often outpaces our ability to govern it effectively – from deepfakes to AI assistants and facial recognition. We need our safeguards to catch up, especially when it comes to use in the public sector. The public can't easily opt out of using public services and tends to hold them to a higher standard if things go wrong.

And finally, we're likely to see uneven advances in AI capabilities, with some progress in areas like quantitative reasoning but less in realising bold claims about productivity gains. We could see real disruption to our livelihoods, relationships and how we trust information.

What's the biggest misconception about plans to use AI in the public sector?
That it will magically solve entrenched, complex problems. AI influences and is influenced by the context it is used in, often with unintended consequences.

The effectiveness of an AI tool in the public sector depends on its interaction with existing social systems, values and trust. And often the underlying data just isn't good enough.

Can you give an example of where AI has been employed for public good?
It depends on the type of AI! Advancements in machine learning have improved weather forecasting, as well as speech-to-text and translation tools.

In 2024, AI also contributed to two Nobel Prize-winning achievements. The question is how we ensure the benefits are evenly distributed and contribute to the betterment of society.

What would you most like to see from government on the AI agenda?
To "be in the driving seat of AI", we need to know the brakes work.

Our research shows the UK public hold nuanced views on AI, but expect action from government on policy – 72 per cent say that laws and regulations would increase their comfort with AI. ●

MARTA SIGNORI

THE NEW STATESMAN
EVENTS

NS Future of Healthcare Conference

28 October 2025

The *New Statesman's* Future of Healthcare conference is an annual event addressing the most pressing policy issues in health and social care.

Through panels, live interviews, speeches and debates, you will hear insights from experts, politicians and industry leaders. Our 2025 conference will examine issues including prevention, innovation, the ageing population and the future of the NHS.

Long waiting lists and an overstretched workforce are challenges only too familiar in our health and social care systems. But how can we rethink these issues to improve patient outcomes, empower the workforce and build a resilient, forward-looking NHS? And how can a focus on prevention help to drive economic growth?

Topics for discussion include:

- How can we tackle the root causes of poor health?
- How can research and development unlock a new era of positive patient outcomes?
- How can we build a social care system that empowers independent living for older and disabled people?
- How can we build a healthcare service that meets the needs of people with long-term conditions?
- How can levelling up the gender health gap boost Britain's economy?
- How can the UK become a leader in the early adoption of new healthcare technologies?

Register for tickets
nsmg.live/event/future-of-healthcare/

Sponsor

novo nordisk®

Scan for more info

THE NEW STATESMAN EVENTS

NS Politics Live

19 November London

The *New Statesman*'s annual Politics Live conference takes on the most important topics in UK politics. Through panel discussions, live interviews and debates, leading lights from Westminster, industry and beyond come together to interrogate the most topical and urgent policy issues of the day.

Join us on **Wednesday 19th of November 2025**, to hear insights from experts, politicians and industry leaders; network with fellow influential guests and address the issues that are shaping the present and will determine the country's future.

Scan QR code to register

Key topics:

- Forging a green future: How can the UK lead the world in the net zero transition?
- Trade and investment: How can we drive economic growth?
- How do we ensure 1.5 million new homes are fit for the future?
- Britain's security in a volatile world
- Building a nation of innovators: How can technology drive growth and productivity?
- How can we reduce the economic burden of the greatest health challenges facing the UK?

Sponsors

DeHavilland Sellafield Ltd serco SIEMENS

Scan for more info

The above organisations have provided sponsorship funding to support this event and have had no influence over the content of the event or selection of speakers.