

Policy Signals Briefing #1

Future-oriented policy debates: AI in the Creative Industries

A **ForesightLab** policy briefing

Dr Ekaterina Tarnovskaya, Dr Vicki Williams
and Prof Graham Hitchen

July 2026




About the Policy Signals briefing series

This is the first in a series of Policy Signals briefings examining emerging trends, early indicators of change, and longer-term policy trajectories shaping the creative technology landscape. Rather than focusing solely on recent announcements, the series takes a foresight-oriented approach to identify how government activity, institutional priorities, technological developments, and regulatory interventions may evolve over time. Published quarterly, the briefings assess emerging ideas and signals of change through a long-term policy and strategic lens: horizon-scanning across policy debates, think tank research, and wider grey literature to assess the longer-term potential and implications of emerging ideas and to identify the policy themes and strategic priorities likely to shape the future direction of the sector.

The series builds on the methodologies developed through the CoSTAR Foresight Lab programme, including the findings of the [*Creative Technology Futures: Policy Pathways to 2035 and Beyond*](#) report and associated resources, which establish baseline trajectories for the sector. The series places particular emphasis on four core thematic areas, which were developed in the *Creative Technology Futures* report: **AI, Creative Education, Growth, and International & Regional Collaboration**.

Suggested citation:

Tarnoskaya, E; Williams, V; Hitchen, G. (2026) "Policy Signals Briefing #1. Future-oriented policy debates: AI in the Creative Industries". CoSTAR Foresight Lab. DOI: 10.5281/zenodo.21129748



Introduction


To set the scene, it is worth noting a series of recent developments that point to a rapidly evolving AI landscape in the UK. Government initiatives such as the AI Growth Zones, the Creative Content Exchange, and the Sovereign AI Fund signal a strong commitment to accelerating AI adoption, strengthening domestic capabilities, and fostering place-based innovation.

A series of [updates on the North East's AI Growth Zone](#) were published – with some targets pertaining to skills and opportunities for local communities. This includes ambitions for 30,000 primary school children to benefit from AI and digital skills initiatives, 80,000 students across the region to benefit from training, 1,000 teachers equipped to teach AI and 150 local work placements. Alongside the update, Mayor Kim McGuinness published the AI Growth Zone prospectus for consultation – focussed on skills, adoption and innovation across the region.

A first iteration of the [Creative Content Exchange](#) (CCE), a UK Government-backed marketplace with the ambition to enable businesses to license digitised creative assets, first announced in the Creative Industries Sector Plan, has launched in beta. Funded as part of UK Research and Innovation's (UKRI) Mission Accelerator Programme, the CCE has launched a pilot with the National History Museum to test the market.


The [UK Sovereign AI Fund](#) officially launched, a £500m fund supported by the Government to support the growth of UK grown AI companies. The Fund's ambition is to establish a targeted approach, providing early-stage investments, fully funded access to UK super compute capability and provide visas for global R&D talent. Since the launch, the fund has announced backing for companies including Isomorphic Labs and Callosum.

Against this backdrop, this briefing identifies several major drivers of change in the AI landscape: these include growing power imbalances between creators and large AI firms, declining trust and social legitimacy around AI deployment, disruption to labour markets and creative talent pipelines, and the environmental pressures associated with high-compute AI infrastructure. These drivers suggest that policy around AI must address governance, fairness, skills, sustainability and public confidence.



The policy signals show momentum toward more accountable AI regulation. Weak signals include emerging ideas such as certified human-made content, new personality rights, collective data trusts, dynamic legislation and cultural institutions acting as data governance testbeds. Stronger signals include the rejection of copyright opt-out models, growing support for licensing and transparency, the influence of the EU AI Act, mandatory labelling and disclosure, structural employment changes, implementation gaps in business adoption, and the risk that AI benefits will deepen regional inequality.

The trajectories outlined suggest two possible futures. An undesirable trajectory would see reduced entry-level employment, unsustainable AI infrastructure, SME marginalisation and capital flight caused by policy uncertainty. A more desirable trajectory would involve stronger public investment, collective licensing and data governance models, sustainability standards, transparent and ethical AI adoption, and a more active state role in shaping a sovereign AI ecosystem that supports both competitiveness and public value.



01 Drivers of change


Drivers of change are the underlying causal forces shaping long-term policy trajectories. Drawing on the CoSTAR Foresight Lab's Certain Uncertainties methodology, drivers are understood as systemic, multi-dimensional trends that have been shaping an era over time and are expected to continue influencing the next decade, even if their precise trajectories remain uncertain. In this context, drivers help answer the question: 'what is driving change?' and provide a framework for assessing whether current policy approaches adequately recognise and respond to these emerging forces.

Structural power imbalance and inequality

Advances in machine learning, including generative adversarial networks (GANs), transformers, and stable diffusion, have accelerated the concentration of AI capabilities within a small number of global firms. These companies are able to appropriate creative IP as training data for commercial models, often without clear systems of consent or remuneration. This has prompted widespread legal challenges against AI developers (e.g. OpenAI, Midjourney, Runway), which may set important precedents for future IP regimes.

The dominance of these firms also shapes the distribution of benefits. Evidence from industry bodies, such as the report by the Association of Illustrators and partner organisations, indicates that AI-driven economic gains are unevenly allocated. Because AI is dominated by multinational firms like Google, OpenAI, and Anthropic, much of the economic gain may flow overseas rather than to UK workers. Meanwhile, costs, such as job insecurity, reduced incomes for creatives, pressure on infrastructure, largely fall on the UK. Weak copyright protections and unlicensed data use could further harm UK creators, shifting future revenues to foreign firms and weakening the UK's role in the AI licensing market.

Growing tensions between individual creators and large multinational technology firms point to a deepening imbalance of power, reinforcing the need for collective mechanisms such as sector-specific data trusts and coordinated bargaining frameworks. Recent research from the KCL Centre for Data Futures suggests that individual opt-outs cannot fix the power imbalance between creators and AI developers, as enforcement burdens fall on rightsholders. Instead, collective bargaining offers a stronger solution, improving negotiating power and enabling more participatory AI governance. Governments should support collective representation through sector-specific bargaining mechanisms, data trusts, legal clarity on data rights, and inclusive access for all creators.



Social legitimacy and trust are constraints on growth

Public trust has become the limiting factor in AI-driven growth and AI adoption. While adoption of AI is advancing rapidly, large sections of the population remain disengaged, cautious or mistrustful, particularly where skills confidence is low and safeguards are unclear. Broader societal shifts, including changing expectations around work, wellbeing, and institutional accountability, have intensified scrutiny over how AI is deployed and governed. Evolving working patterns, such as remote and hybrid employment, the decline of shared workplaces, and shifts in attitudes towards 'good work', have weakened traditional institutional anchors of trust, further complicating adoption.

This reframes AI policy from innovation acceleration to trust infrastructure building. A report by the Tony Blair Institute argues that the limiting factor for AI-driven growth is no longer technical capability or supply, but social legitimacy, trust and confidence. This suggests that the government's AI agenda should be paired with visible, experience-based trust-building measures, such as use-case-led deployment, real-world evaluation, inclusive skills training and sustained public engagement. Without these measures, distrust and disengagement among population groups risk slowing adoption and economic benefits of AI.

Shifts in the labour market and creative talent pipeline


Labour markets are undergoing significant structural transformation, driven by automation, demographic change, evolving models of work and wider dynamics of workforce instability. The rise of gig-based and flexible working arrangements has heightened precarity, and cycles of hiring and layoffs have normalised the treatment of labour as a financial variable. Demographic pressures, particularly ageing populations, further complicate this landscape. As the proportion of retirees rises, pressures on productivity, welfare systems, and skills replenishment intensify, raising questions around immigration, labour supply, and long-term economic resilience.

The increasing adoption of AI is reshaping organisational structures towards a diamond-shaped model, with fewer entry-level roles and greater concentration in mid- and high-skill positions. Automation is particularly affecting junior and graduate pathways, contributing to a growing entry-level drought and undermining long-term skill development pipelines. This is accompanied by the rise of hybridisation of creative and technological skills, where employers within CIs increasingly demand a combination of technical and creative skills, as the recent Creative PEC report suggested.

Sustainability challenge

AI and related high-compute systems rely on complex global supply chains with substantial ecological footprints, spanning raw material extraction, manufacturing, energy consumption, and end-of-life disposal. From e-waste systems to resource-intensive data infrastructures, digital technologies are embedded in forms of environmental strain that are often obscured from end users.

The CIs are increasingly forced to account for the environmental and social externalities of the high-compute technologies they adopt. The CIs must develop new ways to measure the full carbon footprint of adopted technologies, including equipment, supply chains, and end-of-life emissions. For example, AI environmental impact is growing rapidly, with computing power increasing tenfold from 2018–2022 and inference now accounting for 80–90% of usage, according to the LSE blog. Data centres consumed 1.5% of global electricity in 2024, projected to double by 2030, potentially driving up emissions (already at ~1% of global totals) and adding up to 1.6 gigatonnes of CO₂ equivalent (GtCO₂e) emissions by 2035. They also place heavy pressure on resources, with AI expected to use 4.2–6.6 billion cubic metres of water by 2027, raising concerns about sustainability, transparency, and local environmental impacts.



02 Policy signals

Signals of change identify emerging events that may alter the policy landscape, ranging from ambiguous early signs to established directions. They are activities that alert about potential change, identified through horizon scanning processes where multiple signals are collated to explore possible futures. Signals capture very early, inherently emergent indications of change that may not yet be fully realised but are judged to be significant. Unlike trends, which reflect continuity and extend existing patterns over time, signals point to developments that are currently weak but have the potential to become more impactful. In a policy context, signals of change can be used to assess and validate the plausibility of future scenarios.

2.1 Weak signals

Weak signals are subtle indicators of emerging trends that are not yet mainstream policy but suggest a change in direction. They are ambiguous and uncertain but revealing small shifts or the early beginnings of something new. As the earliest signs of emerging change, their future impact remains unclear and often difficult to assess. Within policy tracking, weak signals are frequently embedded within more prominent announcements, lack a coherent narrative, or have not yet developed into widely recognised trends or areas of action. Despite this, they are critical for foresight work, as they can point to nascent developments that may evolve into significant policy directions over time.

Certified human standards and personality rights

Human-made content initiatives aim to establish clear market recognition and protection for works created by people. This includes developing human-made labels to differentiate and add value to authentic creative output, alongside the use of digital watermarks and metadata to trace and prevent unauthorised reuse in AI-generated content. Additional safeguards, such as anti-mimicry tools, opt-out mechanisms, and restrictions on replicating artists' distinctive styles, help protect intellectual property and individual expression.

Moreover, the KCL report argues that training AI on AI-generated content only without human data inputs leads to a decline in quality, making the protection of human creativity a technical necessity for the AI industry itself. Research suggests that AI models can degrade in quality when trained predominantly on AI-generated content, reinforcing the continued importance of human-created data and input in maintaining model performance and originality. Creativity is a 'social good' that sustains cultural diversity, but it is currently being treated as a mere raw material for commercial AI.


The [Human Made Mark initiative](#) is a UK-based campaign launched in April 2026 to protect and promote the value of human creativity in filmmaking in the face of growing AI-generated content. It introduces a certification that recognises films and commercials made with real people both in front of and behind the camera, aiming to create a premium for human-made work. By influencing audience choices, the initiative seeks to support filmmakers, studios, and artists, ensuring that the collaborative, human-driven nature of film continues to thrive and inspire future generations. As part of a wider movement, [Proudly Human](#) provides the overarching trust mark that enables audiences to choose human-created work, operating as a trusted, independent, international certification body that verifies and labels human-created content, services, and products across industries.

The report published by the [Society of Authors, Equity and other creative organisations](#), bringing together evidence from more than 10,000 creators across writing, illustration, music, photography and performance, is calling for new legal protections for a person's name, likeness, and distinctive creative style to close the gap where copyright alone is insufficient to prevent digital cloning. Copyright protects works, but not the individuals behind them. As digital replicas become more common, there is a growing need for personality rights to prevent unauthorised commercial use of a person's identity.

Collective bargaining

Collective bargaining is an important tool for managing the impact of AI, allowing creative workers to negotiate flexible, sector-specific agreements that can adapt more quickly than legislation to technological change, according to [AI for creative workers - TUC manifesto](#). It would enable creative workers to have a voice in how AI is developed and used, helping to protect jobs, rights, and working conditions. To be effective, [governments and regulators must support collective organisation](#) and ensure creative workers and their unions are properly consulted in shaping AI policy and regulation, recognising that they often lack the resources to advocate for themselves. Similarly, [mainstream outlets](#) advocate for an approach to AI and technology, which is aligned with the public's expectations and interests.

[Early discussions](#) point away from individual 'opt-outs' toward the creation of sector-specific 'data trusts' for managing collective consent and remuneration. Individual opt-out systems cannot overcome the power imbalance between creators and major AI developers. Rightsholders face practical challenges, including tracking where their work is used, navigating inconsistent opt-out processes, and covering enforcement costs. The [KCL report](#) argues that collective bargaining is a more effective solution. It strengthens creators' negotiating power with AI companies and supports more participatory governance of AI.

Four decorative orange lines of varying lengths and angles originate from the bottom left corner and extend towards the right side of the page, creating a dynamic graphic element.

Organisations like the Serpentine Galleries are piloting Choral Data Trusts Experiment to test how cultural sectors can set new standards for collective consent and attribution in AI datasets, according to the [Creative UK report](#). It created a collective AI dataset for vocal sound, introducing roles like data stewards and trusted intermediaries to ensure consent and fair attribution. The project tested new models of data governance, informing guidance on collective licensing and data trusts for creative content. The project meets key R&D criteria and shows how the cultural sector can support responsible AI development.

[Community Models](#) is an open-source framework developed by the [Collective Intelligence Project](#) to enable communities to directly shape AI system behaviour through collective decision-making. It uses a deliberative process to create a shared constitution of principles and then applies these principles to guide AI responses through structured evaluation and revision. Trials with communities across the world showed that non-experts can meaningfully define AI guidelines tailored to their contexts. The findings suggest that community governance is most effective in institutional settings where individuals lack control and highlight the importance of linking abstract values to concrete needs, ensuring practical usefulness, and prioritising real-world deployment to realise the potential of democratic, collectively governed AI systems.

Dynamic, reversible legislation

The [Social Market Foundation report](#) provides suggestions for ‘regulatory real options’ signalling a shift toward iterative, software-like lawmaking that can be rolled back or expanded based on real-time model breakthroughs. Challenges in assessing the impacts of AI regulation address the ‘valuation gap’ in AI governance arguing that because AI carries near-existential risks and moves with unprecedented velocity, the standard tools of government (like cost-benefit analysis) are ineffective. The report proposes a shift toward ‘real options’ and ‘precautionary’ frameworks, suggesting that regulation should be treated as a financial option with the choice to delay, expand, or abandon. [Policy proposals by Demos](#) are emerging to exempt AI models from certain regulatory obligations if they meet high standards of openness and transparency, provided they are not used in high-risk contexts. It would reduce compliance burdens on small developers and incentivise developers to embrace greater openness.

2.2 Weak signals

Strong signals are high confidence indicators of emerging trends suggesting a change in policy. While still carrying some degree of uncertainty, they are perceived as more impactful in terms of what they imply and the direction of change they point to. They are often supported by more visible, maturely developed evidence and may originate from authoritative or influential sources, lending them greater credibility. Within policy tracking, strong signals typically reflect established shifts or clear directions of travel, including areas where action is already underway. Even when building on prior developments, they provide a stronger sense of momentum, indicating that a particular policy pathway is becoming more defined and likely to shape future decisions.

Rising public demand for structured, accountable oversight and rejection of the 'opt-out' copyright recommendation.

The overwhelming rejection of active opt-out plans (with 95% of 10,000 respondents in a UK government consultation opposed) has forced a government reset of copyright policy. Mainstream outlets reported creative workers' public campaigns against weakening copyright protections, arguing that AI companies should license creative works rather than scrape them for free (see also an article by Copyright Alliance on the most notable AI-related campaigns launched by creators). CIs leaders warn that the government-generated uncertainty is hurting the economy by preventing creators from signing private licensing deals with AI firms.

Several publications (e.g. the report by Institute for the Future of Work, Queen Mary University of London and The Alan Turing Institute) signal a strong preference for a licence-first regime, including mandatory licensing, transparency, and enforceable standards. Publications signal a move away from a broad TDM exception toward a highly granular, tiered regulatory system that treats different types of data differently.

A report published by the Association of Illustrators in coalition with the Society of Authors, ISM, Equity, and AOP argues that the UK CIs are facing a crisis due to unregulated Generative AI. The report calls on the UK government to urgently implement a new regulatory framework (CLEAR), signalling strong professional resistance. The CLEAR Framework shows specific legislative demands providing a blueprint for what good regulation looks like from the creator's perspective:


Consent: Clarify law so work cannot be taken without prior permission.

Licensing: Move from scraping to statutory licensing with fair pay.

Evidence/Ethics: Ensure transparency about what data was used.

Accountability: Make AI companies liable for the outputs they generate.

Remuneration: Guarantee that creators are paid and attributed for their contribution.

A series of five parallel orange lines of varying lengths, starting from the left edge of the page and extending towards the right, creating a sense of movement or a stylized graphic element at the bottom of the page.


The report shifts the debate to justice, arguing that the UK's soft power and culture are at stake. Without immediate policy intervention to protect human originality, AI will result in a digital hollowing out of the UK's CIs.

Similarly, the findings of the [Ada Lovelace Institute's report](#) highlight a significant gap between public expectations and government ambitions for AI, which could weaken trust and reduce support for future policies. While the government is investing heavily in AI, a lack of clear commitments to safe and responsible use risks undermining public confidence. Without trust, people are less likely to adopt new technologies or support government initiatives. The public prioritises fairness, safety, and positive social outcomes over economic growth or rapid innovation. Many feel excluded from decision-making and worry that government policies may favour large technology companies over the public interest. There is strong support for independent oversight, with most people backing the creation of a regulator with real enforcement powers. The public also expects robust monitoring systems, including transparency, independent standards, and clear accountability throughout the lifecycle of AI systems.

Likewise, [AI policies implemented by governments globally](#) are shifting from broad principles to concrete rules, focusing on disclosure, consent, watermarking, and accountability. Most countries continue to assign copyright and responsibility to human creators, while strengthening protections for performers' voice. AI use is increasingly embedded in compliance processes, with documentation and transparency becoming standard practice across the CIs.

Transparency and mandatory AI labelling and disclosure

There continues to be a significant level of debate regarding creative practitioners demanding mandatory disclosure of training data and clear labelling of AI-generated content. The [Fabian Society's opinion piece](#) argues that unchecked use of generative AI is accelerating the erosion of literary culture and threatening the economic viability, diversity and legal rights of authors and creators in CIs, unless government intervention restores transparency and enforces existing copyright law. It calls for mandatory disclosure of AI training data, auditable records, clear labelling of AI-generated content, and international regulatory alignment, arguing that without enforcement, AI risks hollowing out the CIs it depends on. It signals that AI development shifts toward licensed, auditable training practices.


Four orange lines of varying lengths and slopes originate from the bottom left corner and extend towards the right side of the page, creating a dynamic, abstract graphic element.

A [KCL Centre for Data Futures' report](#) argues for developing transparency standards along the AI-development and deployment pipelines and especially for AI outputs, working towards international coordination to prevent regulatory arbitrage and protect data ecosystem integrity. The report calls for transparency in training data and for AI-generated outputs to protect ecosystem integrity. Similarly, [case studies globally](#) demonstrate that AI transparency is becoming a key commercial and reputational issue, as platforms introduce labelling rules that affect content visibility, eligibility, and audience trust. Disclosure is shifting from a back-end requirement to a consumer-facing feature, influencing market access. AI tools that automatically alter content are raising new questions about creative control, requiring clearer tracking and governance across production and distribution.

Regardless of domestic policy, UK firms are increasingly defaulting to EU AI Act standards for documentation and transparency to ensure continued access to European markets. For example, a [Conversation article](#) argues that the EU AI Act significantly affects UK businesses and consumers. The EU is treating AI as a compliance issue by introducing a legally binding, enforceable framework for AI governance, including fines, system withdrawal and mandatory design requirements. Its extraterritorial reach means UK firms whose AI systems are used in the EU must comply. While the UK has opted for a more flexible, principles-based regulatory model, the article argues that Europe's more prescriptive approach may ultimately shape global standards, forcing UK firms in practice to align with EU rules regardless of domestic policy choices. The publication signals that AI regulation is moving from ethics, guidelines and experimentation into enforceable legal obligations and regulatory borders matter less than markets.

AI implementation gap

30% of UK tech founders report that their businesses would not exist without AI, because the major AI model providers reshape the ecosystem, according to the [Tech Nation Report 2026](#). According to PwC's 29th UK CEO Survey, 81% of UK CEOs now rank AI, technology and data as their top investment priorities for 2026, yet only 9% say AI is deployed at scale in a way that consistently improves performance, according to the [TechUK report](#). The gap is not technological capability but organisational readiness. Successful AI adoption requires whole-organisation transformation, investment in workforce skills rather than just specialist hiring, and sustained leadership ownership. According to the [Tech Nation Report 2026](#), UK founders say unlocking AI's potential requires stronger government support, particularly in energy infrastructure, talent, and safety. Half believe tax incentives for AI investment would help most, while 40% call for more funding in energy and skills. Priorities also vary by location: London founders focus on compute and commercialisation, whereas those elsewhere emphasise AI safety and regional development.

Four orange lines of varying lengths and angles originate from the bottom left corner and extend towards the right side of the page, creating a dynamic, abstract graphic element.


A report by the [Tony Blair Institute for Global Change](#) similarly argues that public trust is the main limiting factor for AI-driven growth. Although generative AI adoption is increasing quickly, many people remain hesitant or disengaged, especially where they lack skills or clarity about safeguards. To address this, the government should complement its AI strategy with visible efforts to build trust in practice, such as deploying real-world use cases, evaluating impacts, improving inclusive skills training, and maintaining ongoing public engagement. [This is underscored by four recommendations for accelerating mission-aligned AI deployment:](#)

- establishing a dedicated unit to track AI deployment, assess gaps, and analyse impacts on jobs and skills;
- translating broad government missions into specific, cross-departmental problem statements with designated AI leads to drive targeted use cases;
- aligning innovation policy and funding with the missions to support priority challenges;
- stimulating demand through procurement, subsidies, and financing mechanisms that encourage widespread adoption, particularly among SMEs, while fostering collaboration between the public sector and AI vendors.

[A growing body of work](#) also shows that procurement is becoming a critical governance limitation: AI purchasing now carries significant risks related to data use, liability, vendor dependency and compliance, extending well beyond initial contracts. A responsible procurement model can be structured around four pillars:

- upskilling procurement teams to understand AI systems and governance implications;
- tailoring procurement approaches to different types of AI;
- aligning decisions with established governance frameworks and standards;
- embedding procurement within ongoing, organisation-wide AI governance processes through continuous review and oversight.

This functions as a strong signal in policy terms because it reflects a well-evidenced and recognised move beyond technological constraints towards institutional and governance challenges, with practical responses (such as skills development and governance frameworks) already taking shape. Supported by visible industry evidence and expert analysis, it points to a new direction of travel, indicating that future policy is likely to organise around the priorities of capability-building and governance reform.

Four orange lines of varying lengths and slopes originate from the left edge of the page and converge towards the bottom right corner, creating a dynamic graphic element.

Structural employment shifts caused by AI adoption

The 'great graduate job drought', as coined by [the FT article](#), describes a deepening 'graduate job drought' driven by economic uncertainty, cautious hiring and the rapid adoption of AI, which is disproportionately squeezing entry-level and junior white-collar roles in the UK. The article refers to the data provided by the [Labour Market Report Building a Future of Work That Works](#) (LinkedIn's report) indicating that UK hiring remains around 25% below pre-pandemic levels across all sectors; for example, hiring in the Technology, Information and Media industry is 27% below pre-pandemic levels, while the Entertainment Providers industry is 13% below its pre-pandemic level. This suggests the emergence of a 'diamond-shaped' organization, characterized by a narrow junior base and crowded mid-level roles, which threatens the long-term talent pipeline. Graduate employment is no longer the default outcome of higher education, signalling a greater emphasis on applied skills, credentials and AI literacy over degrees alone.

Another signal is the co-occurrence spike in job ads demanding both AI skills and creativity, as [the Creative PEC report](#) suggested. By analysing Adzuna job data from 2016 to 2024, the Creative PEC report found that GenAI has not reduced the demand for creativity, instead, it has redefined it. Employers are increasingly seeking hybrid 'createch' professionals, signalling that creativity requires technical competency. This is a strong signal that skill hybridization is the new market standard: employers are no longer looking for just an artist or just a coder, but a hybrid 'createch' professional.

Establishment of the AI Economics Institute: [A proposal for a formal body](#) to specifically study the long-term impact of AI on work signals that the government is preparing for significant structural labor market disruption.

Increasing inequality

The benefits of AI adoption remain unevenly distributed geographically. Evidence shows that highly-skilled creative clusters absorb AI benefits much faster than other regions, as the [Creative PEC report](#) has shown. AI and creativity skills are most concentrated in established creative clusters, positioning them to benefit most from AI-driven growth. This reinforces regional inequality, as already-advanced areas capture disproportionate gains while others fall behind.

This dynamic may be further intensified by the broader economic impacts of AI. [Geoffrey Hinton warns that AI could deepen inequality](#) by enabling wealthy individuals and corporations to replace workers at scale, driving widespread unemployment and rising profits. In such a scenario, the gains from AI would accrue to a small elite, leaving most people economically worse off - an outcome he attributes to the workings of capitalism rather than the technology itself.

Addressing this requires more inclusive, place-based approaches. This includes investing in skill training to support workers adapting to AI, and strengthening regional ecosystems through initiatives like microclusters, innovation hubs, and coworking spaces, helping spread the benefits of AI more evenly across the economy.

03 Trajectories

Trajectories describe plausible future paths based on the interaction of the signals and drivers. They represent the direction of travel of key drivers of change, showing how these forces may evolve over time and shape the future. Rather than presenting fully developed visions, trajectories outline potential pathways, often explored through underlying tensions. In foresight analysis, trajectories emerge where signals converge around specific drivers, with early indicators of change helping to reveal how these drivers may unfold.

As part of the foresight process, trajectories are designed to support exploration and discussion. They provide structured ways of considering how different forces of change may interact and what implications could emerge under different conditions. The trajectories presented here should not be interpreted as preferred, endorsed, or inevitable futures. Instead, they are analytical constructs intended to challenge assumptions, surface uncertainties, and help stakeholders consider a range of possible developments and responses.

The trajectories presented below are based on findings from the [*Creative Technology Futures: Policy Pathways to 2035 and Beyond*](#) report. They explore plausible futures for the adoption and governance of AI across the creative industries, highlighting how different technological, regulatory, and economic developments could shape creative work, business models, and cultural production.

Undesirable trajectories

One significant risk concerns the erosion of entry-level employment. Because AI systems increasingly perform routine creative and technical tasks, junior roles may be substantially reduced. A plausible outcome is [*the emergence of a 'diamond-shaped' workforce*](#), characterised by a narrowing base of entry-level roles and a concentration of mid- and senior-level positions. In this future scenario, synthetic workflows replace junior labour, undermining long-term skills development and limiting opportunities for new entrants to gain experience.

The second area of concern relates to the environmental sustainability of AI infrastructure. The growth in data centre construction risks proceeding without robust sustainability requirements. Without clear policy frameworks governing energy sourcing, water use, and land impact, the expansion could contribute to increased carbon emissions and ecological strain. This trajectory reflects a broader failure to integrate AI development with national climate objectives, exacerbating the misalignment between technological growth and environmental policy.

The third undesirable area involves marginalisation of SMEs. If government strategies continue to favour large multinational firms, smaller creative businesses may struggle to access capital and technological infrastructure. This imbalance risks facilitating an ecosystem in which well-resourced corporations consolidate their advantage, while SMEs face prolonged disadvantage and market exit. The results from the [*TechUK*](#)

conducted survey among 500 business leaders from both the technology sector and the wider economy showed that the UK tech sector has shifted from being optimistic to being frustrated with government rhetoric: 84% of tech businesses are unconvinced that economic growth is a genuine government priority and 45% of tech businesses have considered relocating investment or operations outside the UK. Hence, government-generated uncertainty lead to a capital escape, in which businesses relocate to more favourable markets in the US or EU.

Desired trajectories

In contrast, a more positive future can be foreseen through coordinated and strategic policy interventions. The expansion of funding mechanisms would align AI and the CIs. Increased public investment through grants, tax incentives, and cross-sector R&D programmes could enable creative organisations to integrate AI in ways that augment human creativity. This approach reframes the relationship between AI and creative practice as mutually reinforcing.

Closely linked to this is the emergence of new data governance models, in which high-quality human-generated data is treated as a ‘social good’ or ‘public utility’. One of the plausible futures could be a data commons model where governments fund digital infrastructures and creators remunerated through collective funds rather than individual licenses. The UK establishes a statutory licensing scheme, in which scraping becomes illegal, and AI developers must pay into a central fund. This creates a regulated market where human-authored is a premium, licensed data category.

The second area involves the establishment of robust sustainability policy frameworks for AI development. Prioritising environmental performance across the AI lifecycle would align technological innovation with climate goals. This could include requirements for renewable energy use in data centres, carbon reporting mechanisms, and support for energy-efficient R&D practices. Embedding sustainability at the core of AI policy would position the UK as a leader in responsible innovation.

The third area centres on redefining human and AI interaction within creative processes. AI is understood as a collaborative partner that is widely accessible, ethically trained, and integrated into creative workflows. Such a shift may be reinforced by norms around transparency and fair data sourcing, enabling practitioners to engage with AI systems as co-creators.


The UK could pursue an ‘activist state’ model of technological leadership, in which government acts as a primary investor, customer, and ecosystem architect. Recent commitment of £2.5 billion in AI and quantum investment, including £1 billion allocated to domestic procurement, illustrates the potential for this trajectory. If successfully implemented, it could foster a sovereign AI ecosystem by 2030, anchored by public procurement, growth labs, and strengthened ties with European markets. In this trajectory, the UK positions itself as a global hub for AI innovation, balancing economic competitiveness with regulatory oversight and public value creation.

04 Policy implications and conclusion

The briefing shows that AI policy is entering a complex and contested phase. While AI presents major opportunities for economic growth, innovation and creative productivity, these benefits are not guaranteed to be widely or fairly distributed. Without stronger governance, the expansion of AI could intensify existing inequalities, weaken creative labour markets, undermine trust, concentrate economic value in large multinational firms and place growing pressure on environmental resources.

The future of AI in the CIs will depend less on technical capability and more on the institutional frameworks built around it. Issues such as consent, licensing, transparency, remuneration, skills development, regional inclusion and sustainability are becoming core policy questions. The rejection of broad copyright opt-out models and the growing demand for licence-first, auditable and accountable systems suggest that creators, industry bodies and the public are pushing for a more rights-based approach to AI governance.

The UK therefore faces a strategic choice for the future. One path would allow AI development to remain shaped by market concentration, weak safeguards and uneven adoption. The alternative is a more deliberate policy model in which government acts as investor, regulator, convenor and ecosystem-builder. This would require coordinated action across copyright law, industrial strategy, skills policy, regional development, environmental regulation and public procurement. If managed well, AI could strengthen the UK's CIs by supporting new forms of human-machine collaboration, creating fair markets for high-quality human-generated data, and positioning the UK as a leader in responsible innovation.



CoSTAR Foresight Lab
Goldsmiths, University of London
New Cross
London SE14 6NW

costarforesightlab@gold.ac.uk
costarnetwork.co.uk/labs/foresightlab