University of Southampton

WEB SCIENCE INSTITUTE

Creative Sovereignty as National Security: A Governance Framework for Generative AI

WSI Position Paper 2026-02
March 2026

Matt Javanshir, Alistair Sackley, Thomas Irvine

About the WSI

The Web Science Institute (WSI) brings together the University of Southampton's world-class, interdisciplinary, socio-technical expertise in web science, data science and artificial intelligence (AI) to leverage the unique role of online technologies in tackling global challenges. We work to create collaborations within the University and with industry, governments and third sector organisations that bring interdisciplinary socio-technical insights and impacts to the world's most pressing problems.

https://www.southampton.ac.uk/research/institutes-centres/web-science-institute

Web Science Institute, Building 32, Highfield Campus, University of Southampton, SO17 1BJ wsi@soton.ac.uk (DOI:10.5258/SOTON/WSI-WP016)

About the Authors

Matt Javanshir is a Postgraduate Researcher at the Web Science Institute, researching governance mechanisms to support the UK's creative industries amid emerging developments in generative AI. He has previously served as a data analytics product owner and governance specialist in the Investment Banking sector, bringing over 15 years of industry experience, as well as more than a decade in the creative industries as a music composer, sound designer, and game developer.

Alistair Sackley is a Specialist Policy Officer at the University of Southampton, working with the Web Science Institute to support researchers, communities, and organisations in achieving meaningful policy impact. He is also a technology policy consultant, advising the U.S. Department of Justice on projects relating to serious organised crime and national security

Professor Thomas Irvine is Deputy Director of the Web Science Institute. From 2019 to 2023, he was a Fellow of the Alan Turing Institute, the UK’s national centre for AI and data science, where he led the project "Jazz as Social Machine."

Executive Summary

This paper argues that the UK currently lacks enforceable leverage over how its creative output is used in AI systems. As generative AI becomes core infrastructure for economic value and cultural production, this creates a structural asymmetry: the UK supplies creative and informational input, but does not control the conditions under which it is absorbed, transformed, and monetised. This weakens the UK's creative sovereignty and, over time, its economic and institutional resilience. It proposes a governance framework for the interaction between generative AI systems and copyrighted content in the UK. It proceeds from a simple premise: this is not only a question of copyright enforcement or licensing design, but of market governance - how value generated from UK-originating content is controlled, measured, and redistributed.

The central proposition is structured permission: AI training and deployment are permitted, but only under enforceable conditions - compulsory payment, auditable transparency, and market access requirements. In this model, a levy is not a general tax on AI, nor a relaxation of copyright. It is a maintenance budget for governability: a mechanism that ensures the system remains lawful, auditable, and capable of sustaining the creative industries on which it depends. This is not a free-training regime, nor is it a revenue-raising measure detached from sectoral outcomes; it is conditional participation in the UK AI market.

The framework has five components:

#	Description
1	**Conditional permission to train:** explicitly classify text and data mining (TDM) for commercial AI training as copyright infringement if performed without a compulsory licence. This licence can be agreed bilaterally or through remunerating rightsholders through participation in a creative sovereignty levy.
2	**A creative sovereignty levy:** Applied to major training events and high-volume inference, with proceeds ringfenced to compensate rightsholders and sustain the UK creative economy.
3	**Provenance and transparency requirements:** Aggregated, auditable disclosure of data sources sufficient to support fair remuneration without requiring perfect attribution.
4	**A purpose-based creative sovereignty trust:** An independent body with fiduciary duties to rightsholders, responsible for distributing levy proceeds and supporting sector sustainability.
5	**Independent audit and verification:** A dedicated audit utility to assess compliance with disclosure and remuneration obligations, with public reporting.

Together, these measures help transform copyright from a reactive legal constraint into a functional system of governance, enabling innovation while preserving accountability, economic value, and long-term creative capacity.

Introduction and Background

Generative artificial intelligence is rapidly becoming core infrastructure for cultural production, economic value creation, and political communication. Systems that generate text, images, music, and video now shape how societies interpret reality, distribute attention, and assign meaning. This places generative AI not only within the domain of innovation policy, but also squarely within that of national resilience. The United Kingdom's cultural and creative sectors are among its most powerful sources of economic value and international influence. Yet the training and deployment of generative AI systems increasingly depend on opaque global supply chains of data, models, and platforms that are largely controlled outside the UK. This creates a strategic vulnerability: the UK benefits from creative output but does not control the conditions under which that output is absorbed, transformed, and monetised by AI systems.

We argue that the question of copyright and AI is therefore not only a matter of intellectual property or industrial competitiveness. It is a question of creative sovereignty: the UK's capacity to sustain, govern, and benefit from its own cultural production in an AI-mediated world. Creative sovereignty is not about cultural protectionism, nor about restricting technological development. It is about ensuring that the economic and informational systems built on UK creative labour remain auditable, accountable, and aligned with the public interest. Without such alignment, the UK risks becoming a passive supplier of raw cultural material to global AI infrastructures over which it has limited governance leverage. Current policy debates focus primarily on whether AI developers should be permitted to train on copyrighted works, and under what conditions. This framing is necessary but incomplete. The deeper challenge is how to design a system in which large-scale AI development can proceed while preserving incentives for creative production, maintaining public trust, and ensuring that the UK retains meaningful control over the cultural and informational environment that shapes its democracy.

This paper proposes that the most viable route forward is through structured permission: allowing AI training to proceed under a system of mandatory, auditable, and economically aligned obligations. It argues for a UK Creative Sovereignty Levy, paired with independent training-data provenance and audit capabilities, and governed through a creative sovereignty trust. Together, these mechanisms aim to transform copyright from a blunt legal constraint into a functional governance tool. The objective is not to slow innovation, but to stabilise it. A system that lacks legitimacy, transparency, and fair remuneration will ultimately fail, whether through litigation, public backlash, or political intervention. A system that aligns economic incentives with accountability, by contrast, can support both technological progress and cultural resilience. Crucially, part of that resilience comes from reinvesting levy proceeds into the infrastructure that makes the regime workable - provenance systems, rights metadata, audit capability, and the capacity of smaller creators to participate on fairer terms. The system therefore does more than distribute compensation; it pays for the systems that make compensation and accountability possible.

National Security and Concentration of the Generative AI Ecosystem

Generative AI is not a neutral technological layer; it is an integrated industrial system in which control over data acquisition, compute infrastructure, capital, model development, and platform deployment is concentrated in a small number of vertically integrated firms. For the United Kingdom, the question is not whether these systems are innovative, but whether their dominance leaves the UK with any meaningful governance leverage over a technology that shapes culture, markets, and political communication. While training data is sourced globally, it is aggregated through US-owned platforms; compute infrastructure is primarily provided by US hyperscalers; and model development and deployment are governed by US corporate priorities. UK creators, institutions, and users contribute substantial value to this ecosystem, but do so largely on terms defined elsewhere.

This configuration produces a structural asymmetry in which the UK supplies cultural and informational input, but does not control system design, auditability, or rule-setting. Current UK policy arrangements reflect this position: memoranda with US firms and infrastructure commitments under agreements such as the Tech Prosperity Deal provide access to capability, but not control over governance standards, transparency obligations, or long-term system design. At the same time, the generative AI ecosystem increasingly exhibits characteristics already recognised by government as systemic national risks: concentrated platform dominance, opaque supply chains, and growing influence over the information environment. Systems with these properties shape public understanding, economic value, and political legitimacy, yet remain structurally resistant to independent inspection or accountability.

This places copyright policy at the centre of a wider governance challenge. The dispute is not simply about legal exceptions for text and data mining, but about whether the UK can establish enforceable conditions over how its creative output is absorbed into AI systems and transformed into economic and informational power. If the current trajectory continues, the UK risks embedding a permanent strategic position as a cultural data exporter and technological services importer, weakening domestic creative capacity, institutional authority, and economic resilience. Creative sovereignty, in this context, is not about resisting foreign influence but retaining the capacity to govern how domestic cultural production participates in global AI systems as a matter of national security.

Proposed Framework - Overview

To address the structural vulnerabilities identified above, this paper proposes a UK copyright regime designed explicitly as a governance mechanism for the generative AI supply chain. The framework comprises three interlocking components, supported by clear thresholds and sequencing for implementation.

For the purposes of administration, four threshold terms should be defined early in legislation or secondary regulation: "frontier AI developer" means an organisation

training or materially fine-tuning general-purpose generative models above a specified compute/capability threshold; "large-scale deployer" means an entity making such systems available in the UK market above a defined user, revenue, or usage threshold; "major training event" means a discrete training or fine-tuning cycle that materially improves capability or commercial performance; and "high-volume inference" means model output generated as part of a commercial or revenue-generating service, including advertising-supported or cross-subsidised services, above a threshold set to capture meaningful commercial extraction while excluding small-scale and genuinely non-commercial use.

1. **Explicitly classifying TDM for commercial training as copyright infringement if performed without a compulsory licence** - the regime would explicitly classify AI training as infringement under UK copyright law unless paired with compulsory collective licensing, mandatory payment, and auditable disclosure. This is not an opt-out exception or a free-training regime. It is conditional permission. Developers seeking to train or deploy covered systems in the UK market would be required to meet minimum transparency standards, submit aggregated provenance disclosures, and accept independent verification of the controls used to generate those disclosures. Together, these measures convert copyright from a blunt legal constraint into an enforceable interface for accountability, remuneration, and oversight.

2. **The Creative Sovereignty Levy** - the levy would apply by default to frontier AI developers and large-scale deployers who do not agree terms with rightsholders bilaterally, and would comprise two components: (a) a training-linked charge triggered by major training or fine-tuning events, which are discrete, identifiable instances when a model undergoes capability-relevant training on data; and (b) an inference-linked charge tied to high-volume output or revenue demonstrably derived from model inference. This dual structure recognises that AI development is iterative rather than one-time, while reducing incentives to shift value artificially between training, deployment, advertising, bundling, or cross-subsidised service layers. Proceeds would be ringfenced to compensate rightsholders, maintain the long-term sustainability of UK creative industries, and fund the provenance, rights-metadata, and audit infrastructure required to keep the regime governable over time. At minimum, aggregated provenance disclosure should distinguish between key categories of source material - for example licensed commercial datasets, web-crawled material, publisher-provided corpora, user-contributed content, and public-domain or openly licensed sources - and report them at a regular cadence for covered models and significant updates. The standard should be detailed enough to support equitable distribution and audit, while avoiding impossible claims of perfect attribution.

3. **Purpose-Based Creative Sovereignty Trust** - we propose that administration of the levy would rest with an independent governance vehicle acting as an umbrella collective management organisation (CMO), operating under a statutory mandate and responsible for negotiating rates and distributing

revenues through existing or adapted CMOs. Crucially, this body would be institutionally separate from the audit function, which would be responsible for verifying disclosures and compliance. This segregation of duty preserves integrity, reduces capture risk, and mirrors established best practice in financial and competition governance. To ensure robust governance of compensation and sectoral support, we propose that this vehicle adapts a data trust model that has been deployed in healthcare and other sectors. Data trusts are purpose-based non-charitable trusts: the purpose is incorporated into a trust instrument with defined rules, trustees are under a duty to fulfil that purpose, and they can be held to account if they fail to do so.

We propose to adapt this model to the question of intellectual property in the creative industries. Data trusts would be integrated into the governance structure alongside the hypothecated levy on training events and inference. The trust would provide a mechanism for content creators to determine how they are compensated: the purpose-based non-charitable trust would support artists and creative workers, either through individual compensation (following models such as PRS for Music) or through direct or indirect funding of the UK's creative and cultural ecosystem. This structure ensures that:

- Trustees have fiduciary duties to creators, not to AI companies or commercial intermediaries, and can be elected or selected from existing elected representatives within the creative sector or those of other rightsholders.
- Purpose (supporting creative sustainability) is legally enforceable and cannot be diluted
- Decisions about distribution (whether individual royalties, collective cultural investment, or mixed approaches) are made by trustees accountable to beneficiaries
- The governance mechanism is resilient to capture by well-resourced commercial actors.

To make that anti-capture claim credible, the legislation should also specify governance guardrails: a balanced trustee selection model that reserves representation for smaller and long-tail creators alongside larger rightsholders; strict conflict-of-interest rules and disclosure requirements for trustees and advisers; publication norms for annual distributions, methodologies, and audit summaries; and an appeals and oversight pathway so both creators and developers can challenge decisions without collapsing the regime into continuous litigation.

Enforcement should attach to UK market deployment. Covered developers and deployers should be required to register systems made available at scale in the UK, comply with disclosure and levy obligations as a condition of deployment, and face civil penalties, procurement exclusion, and platform-level restrictions for non-compliance. This is where the UK's leverage sits: not in controlling where training occurs, but in conditioning lawful access to the UK market.

Proposed Framework – Illustrative Example

Consider a major AI company, Platform A, developing a new generative music model. The company gathers a training dataset that includes copyrighted recordings, compositions, and sheet music; much of it from UK artists and publishers. When Platform A begins the initial training run using this dataset, this constitutes a "major training event" and triggers the first component of the levy (the training-linked charge). Platform A must disclose, in aggregated form, the provenance of its training data to the independent audit utility, reporting the types and volumes of copyrighted material used. Based on the scale of the training run (measured by factors such as deployed compute resources, dataset size, and model capabilities), Platform A pays a calculated levy to the independent collecting body (for works that have not been licenced bilaterally with rightsholders). This payment is compulsory and non-negotiable. Six months later, Platform A fine-tunes the model on a new dataset of recent UK music releases to improve performance. This fine-tuning constitutes another training event, triggering a second levy payment and renewed disclosure obligations. The audit utility verifies that Platform A's reported data provenance aligns with available evidence regarding its data acquisition practices (and the robustness of associated controls), with its findings a matter of public record.

Platform A then deploys the music model commercially through a subscription service that allows users to generate custom tracks. This deployment triggers the second component of the levy (the inference-linked charge, for underlying works that have not been licenced bilaterally with rightsholders), calculated based on the volume and commercial value of outputs generated. As the service scales and generates millions of AI-created tracks for paying customers, Platform A continues to pay inference-linked charges tied to this high-volume output. Throughout this process, the proceeds from both training and inference levies flow to the creative sovereignty trust. Trustees, with fiduciary duties to UK creative workers and the creative ecosystem, determine distribution in a manner that is compliant with their mandate, domestic law, and any international conventions the UK is currently party to. This might include individual payments to rightsholders whose work contributed to training datasets (similar to PRS distributions), collective investment in creative infrastructure (venues, education, development programmes), or a combination of both approaches designed to sustain the creative sector that made the training data valuable in the first place.

The segregated audit utility independently verifies Platform A's compliance, ensuring that the system remains legitimate and free of conflicts of interest. Platform A benefits from legal certainty and a predictable operating environment; UK creators receive meaningful support; and the government retains governance leverage over how British creative output participates in global AI systems.

Proposed Framework – Rationale and Benefits

This framework is designed to address the structural governance deficit and the enforceable uncertainty faced by both creators and developers, enabling the UK to continue to flourish as a creative powerhouse and a nexus of technological innovation. Its purpose is not only to compensate creators after value has been extracted, but to maintain the institutional and technical conditions under which compensation, audit, and sectoral sustainability remain possible.

- **Legal clarity and reduced litigation risk**: By permitting training under clear, compulsory licensing conditions, the regime removes legal uncertainty while avoiding protracted court battles that would benefit neither creators nor developers.
- **Alignment of incentives**: Economic incentives are aligned with transparency and auditability. Companies that operate openly and contribute appropriately gain market access; those that do not face enforcement action.
- **Governance leverage**: The UK establishes enforceable conditions over how its creative output participates in AI systems, rather than relying on voluntary compliance or post-hoc litigation.
- **Cross-sector consistency:** Creative sovereignty impacts several sectors across the economy, such as literature, art, film, design, and music. This framework is designed to address the governance deficit and structural remuneration across all such sectors in a consistent manner.
- **Institutional resilience**: The separation of collection, audit, and trust governance functions, combined with the legal architecture of trusts, creates a system resistant to commercial capture and political interference.
- **Flexibility and accountability**: Data trust structures allow for adaptive governance. As the AI ecosystem evolves, trustees can adjust distribution mechanisms while remaining accountable for their fiduciary duties.
- **Economic sustainability:** By ringfencing proceeds for creative sector support, the framework ensures that the industries providing training value receive resources to sustain production, innovation, and employment. A defined share should also be reinvested into the substrate that makes the system workable: provenance infrastructure, rights metadata and reservation tooling, support for smaller creators to participate in licensing, and public-interest audit capacity. In that sense, the levy functions as a maintenance budget for governability as well as a compensation mechanism.
- **International precedent**: The UK would establish a model that other jurisdictions could adopt or coordinate with, reducing regulatory fragmentation and forum-shopping by AI developers.

Proposed Framework - Robustness and Edge Cases

This framework has been designed to remain workable under foreseeable edge cases and future developments. First, it distinguishes between frontier-scale commercial systems and small-scale or open-source activity by applying thresholds tied to market deployment and monetisation, ensuring that research, experimentation, and non-commercial use are not inadvertently constrained. Second, liability attaches to commercial deployment into the UK market, rather than to the geographic location of training, corporate domicile, or underlying infrastructure, mirroring established approaches in competition, consumer protection, and data protection law. Third, the framework does not depend on perfect visibility into training data or model internals. The audit utility is explicitly tasked with verifying controls, processes, and disclosures, rather than reconstructing ground truth or claiming omniscience. Finally, the separation between training-linked and inference-linked obligations allows the regime to adapt as AI development shifts toward more efficient training methods, continuous fine-tuning, or on-device inference, keeping governance aligned with economic extraction rather than any single technical architecture.

Because firms will optimise against whatever tests are introduced, secondary regulation should define the classification test for a major training event, specify anti-avoidance rules for inference-linked revenue attribution, and set penalties for misclassification or materially misleading provenance disclosures. That does not eliminate gaming, but it reduces the interpretive slack that would otherwise turn the regime into a paper exercise.

In the first 6-12 months, government could establish registration and minimum transparency requirements for large-scale deployment, designate or stand up the audit utility, and begin procurement alignment around compliant systems. Levy collection can then be phased in once reporting categories, audit processes, and appeals pathways are operational. This staged approach makes the proposal administratively credible while preserving the long-term objective of durable institution-building.

Conclusion

Artificial Intelligence presents significant opportunities for economic growth and productivity, but it also introduces new dependencies that are already reshaping the UK's cultural and informational environment. These systems are increasingly central to the production of meaning, value, and influence, yet their development and governance remain concentrated outside the UK. The central risk identified in this paper is a governance deficit. As generative AI systems scale, the UK faces diminishing visibility and control over how its creative output is incorporated into global models and monetised. Left unaddressed, this deficit will erode creative capacity, weaken soft power, and reduce the UK's ability to act independently in domains that affect public trust and democratic resilience. The proposed intervention - permitted training paired with compulsory licensing, minimum transparency standards, independent audit, and a purpose-based governance vehicle - addresses that risk directly by conditioning access to the UK market on lawful, auditable participation.

By transforming copyright into a functional governance mechanism, supported by robust institutional architecture, the UK can secure both technological progress and creative resilience, ensuring that the creative industries, which contribute substantially to national prosperity and influence, continue to thrive in an AI-transformed economy. Success should be measured not only in receipts distributed, but in whether the UK builds a stable licensing environment, lowers litigation risk, improves public trust, and becomes a standards-setter rather than a rule-taker. The choice is not between innovation and sovereignty, but between governed and ungoverned innovation. This proposal offers a route to retain control over critical cultural and informational systems while keeping Britain open for business.