

Q&A: The UK's Copyright Report - A Gift to Creators, a Problem for AI

March 26, 2026

The UK Government has released its long-awaited copyright report, framed as an attempt to reconcile the competing interests of creators, technology companies and the wider innovation ecosystem. Rightsholders will welcome it, while the UK's AI sector will find less comfort.

Two core policy decisions (on training data and on the ownership of AI-generated outputs) mark a shift away from earlier, more developer-friendly proposals. Both decisions leave significant questions unanswered: how AI developers can lawfully assemble training data at scale, what happens to content produced with minimal human input, and whether the UK's current posture is sustainable in a world where capital and training runs are increasingly mobile.

In this Q&A, Oliver Howley, partner in Proskauer's TMT Group and one of The Lawyer's 2026 Hot 100, unpacks what the report says on these two decisions, what it leaves open, and what it means for developers, investors and rightsholders navigating the uncertainty ahead.

Q: What is the overall direction of the UK Government's new copyright report?

A: The report takes a creator-friendly approach, stepping back from earlier proposals that would have given AI developers a clearer legal framework for training and output ownership. That will be welcomed by rightsholders (including musicians, publishers, writers and visual artists) who have lobbied hard against a permissive text and data mining exception. However, for the UK's AI sector, the report leaves two foundational questions unresolved: how do you lawfully assemble training data at scale and who owns what an AI system produces? Until those questions have clear statutory answers, developers and investors are operating in a legally uncertain environment and in some cases are already factoring that into where they deploy and invest.

Q: What was the opt-out TDM exception and why did it matter?

A: Text and data mining exceptions allow AI developers to process large volumes of content (things like web pages, articles, books, images) to train their models, without needing individual licences for every piece of material. The previously proposed UK model would have made this lawful by default under a broad exception, while giving rightsholders the ability to opt out by reserving their rights. That opt-out mechanism matters in practice: it preserves rightsholder control while providing a workable legal baseline for developers. Article 4 of the EU Copyright Directive is a useful comparator, which implements a broadly similar structure and shows that a balanced opt-out model can be legislated coherently.

Q: What has replaced it?

A: The report steps back from that model and instead emphasises licensing, transparency and further evidence-gathering. The practical implication is a licensing-first environment: developers who want to use third-party content to train their models need to clear rights in advance. That shifts rightsholders into a gatekeeping position and assumes that rights can be assembled at scale through bilateral or collective agreements.

Q: Why is that approach considered unworkable in practice?

A: Because at scale, the arithmetic doesn't really work. A modern large language model may be trained on hundreds of billions of tokens of text drawn from across the open web, from news articles to academic papers, to forum posts, to digitised books, to code repos. Clearing rights for each of those sources through individual licences is not a legal strategy, because it's unworkable in practice. I've advised clients who have tried to construct licensing programmes along these lines and the practical barriers (things like identifying rightsholders, negotiating terms, handling jurisdictional variation, managing ongoing compliance) are immense even for a fraction of a typical training corpus.

Q: What about collective licensing? Isn't that the solution?

A: Collective licensing is often raised as the bridge between the impossible (clearing every right individually) and the unacceptable (training without permission). The theory is certainly attractive: collecting societies aggregate rights across large catalogues and AI developers pay a fee for access. The report acknowledges that the licensing market is developing and industry bodies such as the CLA have publicly indicated they are developing collective licence schemes for generative AI training. However, the UK market for AI training remains nascent and largely unproven at the scale required and no collective licence is operational yet. It's also unclear whether the schemes will cover the breadth of content that modern model training requires or whether they will be priced at a level the market can sustain. As a result, endorsing a licensing-first model before that infrastructure exists is a significant bet.

Q: How does the report deal with copyright in AI-generated outputs?

A: The report proposes removing copyright protection for wholly computer-generated works (content produced without any human creative input) while confirming that AI-assisted works (where a human author exercises sufficient creativity in the process) remain protectable under existing copyright principles. That proposal apparently reflects broad consultation support: the report notes that a substantial majority of online survey respondents were not in favour of maintaining the current protection for computer-generated works.

Q: Isn't that a reasonable distinction?

A: It's a legally coherent one, but it doesn't advance the position for developers as much as the framing suggests. The reassurance that AI-assisted works remain protectable largely restates what practitioners already knew - there was no serious dispute that a human author who uses AI as a tool retains copyright in the output, provided their creative contribution is sufficient. The live question is what happens at the other end of the spectrum: autonomous agents generating outputs without meaningful human direction, bulk content generation with minimal oversight, vibe-coded applications where the human input is primarily a prompt. That is really where the ownership question bites, commercially and legally, and the report doesn't provide us with answers.

Q: What was the existing statutory position, and does removing it make things better or worse?

A: The UK has a relatively unusual provision (section 9(3) of the Copyright, Designs and Patents Act 1988) which attributes authorship of computer-generated works to the person who makes the arrangements necessary for their creation. It is imperfect: it was drafted before generative AI existed and applying it to modern agentic systems requires a bit of interpretive gymnastics. However, it at least provided a starting point for allocating ownership of AI outputs to someone in the development chain. On the current proposal, removing it without a replacement allocation rule does not simplify the landscape - it leaves the ownership question more open in practice. Developers who were relying on section 9(3) as a foundation for their IP strategy now have less to work with in terms of a default allocation rule, even if practical reliance on that provision was already limited.

Q: What would a workable solution actually look like?

A: The underlying policy debate is genuinely contested, the tension between creator rights and AI development is real, and reasonable people disagree. However, what the industry is asking for legislatively is relatively specific. On inputs: the AI industry wants a workable opt-out exception that allows developers to train on publicly available content without licensing every page of the internet first – perhaps modelled on, or going further than, Article 4 of the EU Copyright Directive. On outputs: the industry wants a clear statutory allocation rule that gives developers a real baseline for ownership, rather than leaving them to rely on operational workarounds (and there are workarounds, but they are exactly that). Neither requires abandoning creator rights and both would give developers and investors the legal certainty they need to make long-term decisions.

Q: What is the competitive risk if the UK doesn't act?

A: There are early signs of it materialising. Training location decisions depend in part on copyright regimes, and the absence of a broad exception places the UK among the more restrictive jurisdictions internationally. In our experience, clients are already factoring this into siting and investment decisions, looking at jurisdictions through the lens of regulatory clarity and not just talent or compute availability. The UK needs to move, and it needs to move soon, if it wants to be in the conversation.

Q: What should we expect next?

A: The report signals ongoing work on transparency, on monitoring the development of licensing markets and on gathering evidence before any legislative intervention. The collective licensing market is explicitly described as market-led and the government's posture is to enable it to develop rather than legislate around it. That is a cautious and defensible approach, but it means continued ambiguity for developers operating now. I think the more interesting question is whether the lobbying dynamics shift - rightsholders have had the better of this debate so far, but the economic case for a workable TDM framework is strong and the competitive pressure from other jurisdictions will intensify. The UK Government has a long tradition of revisiting policy positions when the evidence shifts. On TDM in particular, a well-reasoned course correction would be welcomed - and on the current trajectory, it may well be necessary.

[Related Professionals](#)

- **Oliver R. Howley**
Partner