

ARTIFICIAL INTELLIGENCE AND INTELLECTUAL PROPERTY RECENT GLOBAL DEVELOPMENTS AND MALAYSIA'S PROPOSED COPYRIGHT ACT 1987 REFORMS

by Michael Soo & Wong Qi Shu

Artificial intelligence (AI) is moving faster than the law and its enforcement. Across the United States (US) and Europe Union (EU), courts are being asked whether training an AI on copyrighted works requires the owner's consent and compensation, and whether the works an AI generates deserve protection at all. Each ruling settles only a fragment, and the fragments do not fit together. These are questions on which billions of dollars and the future of the creative economy now turn, yet the answers are emerging case by case and jurisdiction by jurisdiction, often in open conflict.

This article surveys the key global developments since the second quarter (Q2) of 2026: the Anthropic settlement in the US, *GEMA v. Suno Inc.* in Germany, *CNN v. Perplexity* in New York, and the first generative-AI reference before the Court of Justice of the European Union (CJEU). It then turns to Malaysia, which has taken a different course. Rather than await a domestic test case, Malaysia is answering these questions by legislation, through the proposed amendments to the Copyright Act 1987 (CA 1987) and the forthcoming AI Governance Bill. The second part examines those reforms and offers a view on the shape the framework ought to take.

Part 1 — Key Global Developments Since Q2 2026

Bartz et al. v. Anthropic PBC — A Billion-Dollar Price Tag, But No Precedent

On 20 July 2026, District Judge Araceli Martínez-Olguín of the Northern District of California granted final approval to the USD 1.5 billion copyright settlement under which Anthropic will compensate authors whose pirated books were used to train its Claude chatbot. The settlement resolves the claims arising from Judge William Alsup's summary judgment ruling of June 2025, which held that training on copyrighted books was not itself unlawful, but that Anthropic's acquisition of pirated copies from shadow libraries was. The settlement covers more than 482,000 works, with each yielding approximately USD 3,000 before fees and costs.

The case establishes the first concrete pricing benchmark for the unauthorised use of pirated content in AI training and is the largest known copyright recovery in history. However, its precedential reach ought not to be overstated. The matter was disposed of by settlement rather than a final adjudication on the merits: the damages for piracy were never tested at trial, and no appellate court has pronounced on the question. The first-instance ruling did characterise the training as fair use, but that holding remains a district court determination, untested on appeal and confined to US law. It offers persuasive guidance at most and does not settle whether using copyrighted material to train AI models falls within the fair use defence.

GEMA v. Suno Inc. — Extraterritorial Reach

On 31 July 2026, the Munich Regional Court ruled in favour of GEMA, the German music collecting society, against the US-based AI music generator Suno (case no. 42 O 763/25). The court held that both the training conducted in the US and the musical output generated in Germany infringed copyright, granting injunctive relief, disclosure and a declaration of entitlement to damages. Its significance lies in being the first EU ruling to hold that AI training carried out entirely outside the EU can still give rise to German copyright liability, the opposite of the territoriality approach that defeated the training claims in *Getty Images v. Stability AI* [2025] EWHC 2863 (Ch) before the English High Court.

This ruling should put international AI developers on urgent notice. An AI developer might have hoped to train its AI model in a country with permissive copyright rules, then rely on the lawfulness of that training as a shield against liability elsewhere. This decision closes that route at two levels, jurisdiction and applicable law.

First, on jurisdiction, the court reached the US training in two steps. Section 131(1) of the German Collecting Societies Act 2016 (CSA), read with Section 32 of the German Code of Civil Procedure (CCP), gave the court international jurisdiction over the domestic acts (ie, the memorisation in the model and the outputs in Germany), under the principle of double functionality. The concentration rule in Section 131(2) CSA then drew the US training into the same court, since a collecting society may bring all its claims against one infringer before a single competent court. However, it should be noted that this gateway is available only to collecting societies, and it requires at least one place of infringement to lie in Germany.

Second, on applicable law, the court applied US copyright law to the US training, in accordance with the country-of-protection principle, and still found no fair use protection, because fundamentally here simple prompts produced outputs substantially similar to the originals. Separately, under German law for the outputs in Germany, the text and data mining (TDM) exception could principally cover the training, but not the memorisation of the works inside the model, which could be reproduced on request. On that footing, the court held that the AI model itself was an infringing copy. The exception failed for a further reason as there was no lawful access, because the works had been stream-ripped in breach of a technical protection measure. A user who circumvents such a measure cannot rely on the exception.

The AI developer thus found no refuge, whether under the law of the country where it chose to train, or under German law governing the acts that took place in Germany. The jurisdictional principle is firm: a court may take extraterritorial jurisdiction over training conducted abroad, where its domestic law permits this under the principle of double functionality. Whether that training in fact amounts to infringement, however, turns on the law of the country where the training took place, and separately, whether the output amounts to infringement turns on the law of the country where the infringing act occurs.

Cable News Network, Inc. (CNN) v. Perplexity AI, Inc. — AI Scraping the News

On 28 May 2026, CNN filed suit against Perplexity in the US District Court for the Southern District of New York (case no. 1:26-cv-04427), alleging that the AI search engine provider unlawfully copied thousands of CNN stories, videos and images to power its products and distributes identical or substantially similar competing content. The suit is CNN's first AI copyright action and advances a theory of infringement directed at the real-time scraping and redistribution of news content in AI-generated responses, distinct from the earlier wave of litigation.

The dispute is joined on the boundary between fact and expression. Perplexity's answer is a single proposition, that one cannot copyright facts, invoking the idea-expression dichotomy under which copyright protects original expression and not the underlying facts. CNN's response is that it asserts ownership not of the news itself but of the original expression through which it is reported: the selection of words, the editorial structure, the sourcing and the framing. The threshold question is therefore not whether facts are protected, since plainly they are not, but whether Perplexity reproduced CNN's protected expression.

Perplexity has yet to file its defence, but fair use is the defence every AI defendant in this wave has run, and it is expected here. It may prove difficult on these facts. Reproducing CNN's reporting in the responses served to readers arguably serves the same news-

reporting purpose as CNN itself, which is less likely to be transformative. The degree of exact reproduction weighs against Perplexity on the amount taken, and, most significantly, a response that spares the reader any need to visit CNN substitutes for the original and harms its market.

Like Company v Google Ireland Limited — Upcoming CJEU Decision

Like Company v Google Ireland Limited (case no. C-250/25) is the first preliminary reference before the CJEU to address generative AI and copyright. It arises from a dispute in which Like Company, a Hungarian press publisher, alleges that Google's chatbot Gemini reproduces and summarises substantial portions of its copyrighted articles without authorisation. Finding that the dispute turned on the interpretation of EU law, the Budapest court stayed the proceedings on 03.04.2025 and referred four questions to the CJEU.

The reference asks the CJEU to apply the InfoSoc Directive (2001/29/EC) and the Digital Single Market Directive (2019/790) (DSM Directive) to generative AI on the questions below: (1) whether a chatbot's answers that partially reproduce press content amount to a communication to the public; (2) whether training the model is an act of reproduction; (3) whether the TDM exception under article 4 of the DSM Directive covers it; and (4) whether output generated from a user's prompt is attributable to the service provider.

Advocate General Maciej Szpunar's opinion is scheduled for 3 September 2026, with the judgment expected some months later. It is highly significant as a CJEU ruling will bind all 27 Member States at once. As the first pronouncement by a supranational court on these questions, it is likely to shape courts and legislatures beyond the EU too, much as the DSM Directive's TDM regime has already become a reference point for jurisdictions, including Malaysia which is now designing their own exceptions.

The Common Thread, and Why Malaysia Is Not Waiting

The past AI-related legal developments share a common thread. Courts worldwide are grappling

with the same unresolved questions: whether AI training requires consent and compensation, and whether AI-generated outputs attract protection. Yet litigation is delivering answers piecemeal and often inconsistently. Malaysia has yet to record a court decision on AI and copyright. Rather than await its own test case, it is answering these questions by legislation.

Part 2 — Malaysia's Proposed Copyright Act Reforms: A Future-Ready Framework For AI

Malaysia's response is being assembled through two complementary instruments. The forthcoming AI Governance Bill governs accountability in the use of AI, while the proposed amendments to the CA 1987 regulate the rights of owners whose works are used in AI models, together with the rights arising from AI-generated works. Read together, they are intended to form a coherent legal framework for AI in Malaysia.

The Intellectual Property Corporation of Malaysia (MyIPO) describes the reform as a future-ready copyright framework. Unveiled through a public consultation that closed on 14 August 2026, it represents the most comprehensive revision of the CA 1987 since the digital-era amendments, signalling an intent to align Malaysia with global standards while enhancing local enforcement and commercialisation. The proposals now advance towards finalisation.

The Heart of the Reformation: Designing the TDM Exception

The central design question concerns the shape of the TDM exception. Its purpose is to give AI developers clear legal certainty at the outset: a defined statutory safe harbour within which computational analysis may proceed, without the uncertainty on whether the use would later be excused as fair dealing or fair use in litigation. In this regard, two international models, drawn from Singapore and Japan, are under consideration.

Japan's purpose-based model operates as a two-stage filter. A purpose test first distinguishes uses that enjoy a work's expression (such as reading or viewing) from information-analysis uses (such as data mining and AI training), where the work is exploited for its data rather than its expression. Only information-analysis uses qualify for the exception. A subsequent market-harm test then excludes any use that would unreasonably prejudice the legitimate interests of the rights holder, in particular by substituting for the original or causing economic harm. The Japanese model's strength is that it directly targets market substitution, however its weakness is that both stages turn on evaluative judgment which remains uncertain until interpretation accumulates.

On the other hand, Singapore's lawful-access model takes a simpler path. It permits computational data analysis for any purpose, provided one condition is met: the user must have lawful access to the work, by purchase, licence, subscription or a legitimately open source. The exception inherently excludes material acquired unlawfully, such as pirated copies or unauthorised downloads. Its strength is certainty, lawful access being a comparatively bright-line condition an AI developer can satisfy

and document in advance, however, unlike Japan, it carries no express market-substitution safeguard.

In authors' collective view, the soundest design adopts both models as two cumulative lines of defence. The first is lawful access: every work used to train a model must be lawfully acquired, by purchase, licence, subscription or a legitimately open source, with no cover for pirated or otherwise unlawfully obtained material. This bright-line gate at the input stage carries most of the weight, because it is here that the owner's consent is deemed to be given and some form of compensation flows from the lawful access. The second is a market-harm backstop: even where access is lawful, the exception should not extend to uses that substitute for or materially damage the market for the original work. This filter is deliberately narrow, a safeguard against the egregious case rather than a licence for judicial second-guessing, so the certainty conferred by the access gate is preserved. A use is excepted only if it passes both.

This dual structure would serve MyIPO's stated aims directly. The lawful-access requirement ensures that works are used with the owner's consent, and that owners are fairly compensated at the point of acquisition, honouring the creative sectors' demand for consent, compensation and transparency. The market-harm backstop protects the revenue streams on which continued creation depends. Together they secure certainty for AI developers and fair treatment for rights holders in a single exception.

Opt-In vs Opt-Out

Two further proposals refine the boundary of what counts as lawful material for AI training. The first concerns the mode of authorisation. In their consultation responses, creative-sector stakeholders have expressed a preference for an opt-in approach, meaning their works should be available for mining only where the owner has affirmatively consented, rather than swept in unless the owner opts out.

AI developers would unsurprisingly prefer the opposite default of opt-out, under which works are available for mining unless the owner has affirmatively withdrawn them, the burden falling on the owner to withdraw rather than on the AI developers to secure countless individual permissions.

The disagreement between the two sides comes down to one question: is the right to authorise TDM separate from the right to access to a work? Owners are likely to argue that it is separate, in the same way that different rights already attach to different uses of a copyright work, and that mining accordingly requires its own consent regardless of how the work was accessed. AI developers are likely to argue the opposite: if the exception already demands that every work be lawfully acquired, the owner has already granted access and been compensated at the point of acquisition, so requiring a separate consent for TDM imposes a further layer of authorisation on top of one already given, a double consent for a single body of works. This would be unworkable at the scale AI training demands, where a model may draw on millions of works, and would negate much of the certainty the exception is meant to confer.

The choice between these defaults mode of authorisation is the sharpest point of balance in the whole TDM exception design. Therefore, MyIPO will need to weigh this carefully, as its resolution will determine the practicality of AI training in Malaysia.

Orphan Works: Unlocking the Unfindable

A related proposal addresses a different problem: works that no one can license because their owners cannot be found. The reform explores a lawful mechanism for such orphan works, defined as works whose rights holders cannot be identified or located despite a diligent search. Where that search is conducted and documented, the work may lawfully be used. Such a mechanism would unlock archives and datasets that would otherwise lie idle. It does not weaken copyright protection, since access remains conditioned on a genuine and documented search for the owner.

Its significance for AI training is considerable. The provenance of training data is increasingly determinative of liability, and a diligent-search safe harbour offers AI developers a lawful route to material whose ownership is obscure, without harming any rights holder.

Human Creativity as the Threshold: Certainty by the Statute

For AI-generated works, the threshold question is one of subsistence before ownership: whether the output qualifies for copyright protection at all, before one asks who owns it. The reform's signal choice is to settle this by legislation rather than await judicial interpretation in Malaysia. Under the instant proposals, protection requires sufficient human creativity: it will continue to hinge on sufficient human authorship and creative contribution, consistent with the human-centred originality threshold the CA 1987 has always maintained, with the human author being the first owner of the work. It will be clear that wholly autonomous outputs produced without meaningful human contribution, such as, simple prompts and instructions, would not be eligible for protection at all. Fixing this in the statute, rather than leaving it to a test case Malaysia has yet to see, gives certainty on two points: which outputs are protected, and who owns them.

Conclusion

The global picture is one of hard questions answered in fragments. Across the jurisdictions surveyed above, courts are reaching for the same underlying principles. Yet each court reaches its own answer. What emerges is not a settled law of AI and copyright, but a patchwork of first-instance rulings, still open to appeal, still confined to their own jurisdictions, and still capable of pulling in different directions.

Malaysia's response is to supply by legislation the certainty that litigation elsewhere has not managed to deliver. The AI Governance Bill addresses conduct, governing how AI systems are built, deployed and used. The Copyright Act amendments address rights, governing what may lawfully be used to train those systems and what protection their output may claim. Read together, the two instruments assemble a coherent and deliberate legal framework for AI.

Whether that ambition will be realised will depend on the choices still to be made, particularly the shape of the TDM exception, the opt-in and opt-out question canvassed above, and how the balance is struck between the competing interests of AI developers and creators. If that balance is struck well, the reform will mark a decisive shift from soft-law governance towards enforceable obligations, and position Malaysia among the region's leaders in responsible AI adoption.

Michael Soo | Partner
Intellectual Property
michaelsoo@rdslawpartners.com

Wong Qi Shu | Associate
Intellectual Property
qishu@rdslawpartners.com