

Preserving balance in the EU digital single market: How *Like Company* could reframe copyright and innovation in the generative AI era

Enrico Bonadio¹ | Giancarlo Frosio² | Christophe Geiger³ |
Andres Guadamuz⁴  | Stavroula Karapapa⁵ | Irimi Stamatoudi⁶

¹City St. George's, University of London, London, United Kingdom of Great Britain and Northern Ireland

²Queen's University Belfast, Belfast, United Kingdom of Great Britain and Northern Ireland

³Luiss Guido Carli University, Rome, Italy

⁴University of Sussex, Brighton, United Kingdom of Great Britain and Northern Ireland

⁵University of Essex, Colchester, United Kingdom of Great Britain and Northern Ireland

⁶University of Nicosia, Nicosia, Cyprus

Correspondence

Andres Guadamuz.

Email: a.guadamuz@sussex.ac.uk

Funding information

Google

Abstract

This study analyses the Court of Justice of the European Union's first referral on generative AI and copyright, that is, Case C-250/25 *Like Company v. Google*, which interrogates whether the generation of chatbot summaries of press articles infringes the press publishers' right under Article 15 CDSM Directive and the reproduction and communication rights under Articles 2 and 3 InfoSoc Directive. The analysis shows that large language model training cannot, as such, be characterised as reproduction because it encodes statistical patterns rather than fixing protected expression, while recognising that further empirical evidence would be valuable to clarify the precise technical steps involved in the training pipeline. RAG-generated summaries that paraphrase content within 'very short extracts' fall outside Article 15(1) and do not communicate publications to a new public. Alternatively, even if training were characterised as engaging temporary reproductions, the text-and-data mining exception under Article 4 CDSM Directive lawfully authorises such uses absent valid opt-outs. Regarding liability, the study concludes that

This is an open access article under the terms of the [Creative Commons Attribution](#) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *The Journal of World Intellectual Property* published by John Wiley & Sons Ltd.

individual users, not platform providers, bear responsibility for outputs generated through their models' learned parameters and retrieval mechanisms. This comprehensive analysis ensures coherence across EU copyright doctrine and sustains legal certainty for AI innovation within the EU digital single market.

1 | INTRODUCTION

The intersection of generative artificial intelligence ('GenAI') and copyright law has rapidly evolved from a theoretical horizon into a pressing arena of doctrinal conflict, reshaping long-standing assumptions about infringement and lawful uses of protected works in the digital environment. Against this backdrop, this study offers an analysis of *Like Company v Google Ireland Limited*,¹ which represents an important moment for the EU digital single market, marking the first instance where the Court of Justice of the European Union ('CJEU' or 'Court') is directly petitioned to interpret the 2019 Copyright in the Digital Single Market Directive ('CDSM')² in the specific context of Large Language Models ('LLMs'). The dispute we examine here arises from the operation of Google's GenAI chatbot, Gemini (formerly Bard), which is made available to users in the European Economic Area.

The primary aim is to critically analyse the legal issues of the *Like Company* referral. We also seek to move beyond the immediate dispute to address the systemic frictions between the economic rights offered by copyright (namely, the reproduction and communication to the public right) and the press publisher's right (PPR) under Article 15 CDSM, and the technological imperatives of AI training and deployment. By situating the case within the CDSM's broader tapestry as well as the InfoSoc Directive ('InfoSoc'),³ the study aims to provide a theoretical framework for understanding how, or indeed if, existing copyright doctrines can accommodate the non-linear, probabilistic nature of GenAI without stifling innovation or eviscerating the economic incentives of content creators.

1.1 | The facts of the case and the Hungarian national proceedings

The claimant, Like Company, is a Hungarian media publisher operating several news portals, including the regional news site *Balaton Környéke*, and the defendant is Google Ireland Limited (Google). The specific conflict arose from a 579-word article published by *Balaton Környéke* on 21 July 2023.⁴ The piece, titled 'Kozsó refuses to give up: the popular singer still wants to introduce dolphins to Lake Balaton', detailed the eccentric plans of a Hungarian celebrity to release dolphins into the freshwater lake. This article was unique to the claimant's site, containing specific details not widely reported elsewhere.⁵

According to the facts established by the Budapest Környéki Törvényszék (the Budapest District Court), the claimant tested the defendant's system by inputting a highly specific prompt into the Gemini interface: *Can you provide a summary in Hungarian of the online press publication that appeared on balatonkornyeke.hu regarding Kozsó's plan to introduce dolphins into the lake?*⁶ Crucially, the prompt did not merely ask for information about the singer; it explicitly requested a summary of the specific press publication found at a specific URL. In response, the chatbot generated a detailed summary that accurately recapitulated the key information contained in the original article.⁷

Like Company subsequently brought proceedings in Hungary. It asked the referring court to find that Google infringed its rights under national and EU law. In particular, it alleged that, between 13 June 2023 and 7 February 2024, the defendant made electronic copies of Like Company's online news articles so that they could be offered to

users.⁸ It further alleged that, over the same period, the defendant made those articles available online in a manner enabling members of the public to access them on demand, at a time and place of their choosing, without the publisher's authorisation.⁹

Google's defense relied on technical and doctrinal counter-arguments standard in GenAI litigation but untested at the CJEU level. It argued that the generation of the chatbot's output did not involve reproduction or constituted an instance of making content available to the public.¹⁰ In any case, even if a reproduction occurred, Google pointed out that it was covered by the text and data mining (TDM) exception under Article 4 CDSM or the temporary reproduction exception under Article 5(1) InfoSoc.¹¹

1.2 | The questions referred to the CJEU

Recognising that the resolution of the dispute hinged on the interpretation of EU law rather than disputed facts, the Hungarian court stayed the proceedings on 3 April 2025. It first acknowledged that the application of copyright concepts like 'reproduction' and 'communication to the public' to the probabilistic outputs of neural networks was still untested under EU law. Consequently, it referred four preliminary questions to the CJEU, pausing the domestic litigation to seek guidance on how EU copyright law applies to the architecture of LLMs.

The first question focuses on the communication to the public right: 'Does the display of text in chatbot responses that partially matches press publisher web pages constitute 'communication to the public' under EU copyright law?' This question strikes at the heart of the PPR. The national court asks whether displaying the output of AI chatbot responses that mirrors press articles constitutes communication to the public under Article 15(1) CDSM and Article 3(2) InfoSoc ('Question 1').

The second question revolves around the reproduction right: *Does the process of training a LLM constitute 'reproduction' within the meaning of Article 2 InfoSoc?* Here, the focus shifts from the *output* (the summary) to the *input* (the training process). The referral asks whether the technical act of decomposing text into tokens, converting them into numerical vectors, and adjusting model weights constitutes a copyright-relevant reproduction. In other words, this challenges the court to decide whether copies made during training are merely functional steps in a data processing workflow or instead constitute infringing acts of copying protected expression ('Question 2').

The third question is about the CDSM's TDM exception: *If training constitutes reproduction, does it fall within the TDM exception provided in Article 4 CDSM?* Contingent on an affirmative answer to the second question, the CJEU is asked to clarify if training GenAI by reproducing lawfully accessible works falls within or outside such exception ('Question 3').

The fourth and final question tackles the issue of volition and causation: *Do user-prompted chatbot outputs that reproduce press publication content constitute reproduction by the chatbot service provider?* Thus, if a user explicitly asks the bot to summarise a specific article (as *Like Company* did), who is the infringer? Is it the user who elicited the copy, or the service provider whose system was capable of producing it? The question has implications for platform liability, potentially determining whether GenAI providers are direct infringers or merely intermediaries shielding behind safe harbors ('Question 4').

Although the reference is triggered by the specific PPR under Article 15 CDSM, the referring court frames the interpretative core as a problem of systemic fit: whether, in cases such as *Like Company*, the general communication to the public right under Article 3(2) InfoSoc "may exclude" the application of the specific PPR.¹² This formulation goes to the heart of the legislative architecture, raising the question of whether Article 15 CDSM functions as a *lex specialis*, that is a calibrated interface for press content with deliberate carve-outs for 'very short extracts', or whether those limits can be effectively bypassed by pleading general InfoSoc rights. Accordingly, while this study examines the broader doctrinal implications of GenAI for reproduction and communication to the public under Articles 2 and 3 InfoSoc, it anchors that analysis in the specific logic of Article 15 CDSM, treating the PPR not as a redundant add-on but as the governing regime for this category of online use, thereby ensuring the legislature's intended balance is not hollowed out.

The article proceeds as follows. Section 2 offers a technical analysis of how GenAI and LLMs like Gemini operate, clarifying key notions such as training and retrieval-augmented generation in order to ground the subsequent legal discussion. Section 3 then turns to the PPR, unpacking Article 15 CDSM, its legislative history, and its policy objectives. Section 4 examines the reproduction right, while Section 5 analyses the TDM exception and its significance for AI training workflows. Section 6 then addresses the communication to the public right, and Section 7 explores platform and user liability, while Section 8 draws together the main conclusions. This structure does not follow the sequential order of the four questions referred to the CJEU. Rather, the sections are logically positioned to build a coherent argument. Technical foundations must precede legal analysis; understanding how Gemini operates is essential before assessing whether its training or outputs infringe exclusive rights. The PPR and the reproduction right are examined before communication to the public, as the former concepts are foundational to evaluating the latter. The TDM exception is situated between reproduction and communication to the public because it operates as a defence to reproduction claims and thus requires prior explication of what constitutes infringing reproduction. Finally, platform and user's liability are addressed after substantive rights analysis, mirroring the CJEU referral's sequencing, as liability necessarily depends upon prior determination of whether infringement occurred.

Although the authors' analysis supports the view that GenAI models such as Gemini can operate within the existing copyright framework and do not inherently violate established legal principles, this technical compliance does not diminish the authors' fundamental conviction that content creators must receive adequate remuneration in the GenAI era. The development and deployment of sophisticated GenAI models depend substantially upon the creative works and journalistic content produced by authors, publishers, and journalists, whose intellectual labour underpins the data sets upon which GenAI systems are trained. Accordingly, the authors are cognisant of and supportive of the diverse proposals advanced by leading experts, commentators, and policy advocates regarding mechanisms to ensure fair remuneration for content creators, whether through collective management schemes or other innovative statutory remuneration models. This study therefore engages critically with the doctrinal questions before the CJEU while maintaining the conviction that copyright law must evolve to protect creator interests in an AI-driven information ecosystem.

2 | TECHNICAL ANALYSIS

There are various working definitions of the term 'artificial intelligence'.¹³ In the context of this study, the focus is on large-scale machine-learning systems, especially so-called foundation models such as LLMs and large multimodal models, that are trained on vast corpora of digital material to produce general-purpose capabilities.¹⁴ This section explains how such models, and specifically Google's Gemini system deployed in *Like Company*, function, in order to ground the subsequent legal discussion on reproduction, communication to the public, TDM and platform liability in that concrete technical setting.

2.1 | The Gemini model

Google's Gemini system is a GenAI chatbot built on an LLM architecture. In the *Like Company* scenario, the relevant system is Gemini v1.0. Rather than storing static, human-readable copies of training data, Gemini processes input texts by breaking them down into minimum units known as 'tokens'.¹⁵ During training the model learns statistical patterns that govern how tokens co-occur. The resulting system does not retrieve pre-existing documents upon request; it generates new text, code, or summaries by predicting the next token in a sequence on the basis of probability distributions learned from the training corpus.¹⁶

Operationally, Gemini thus behaves less like a database and more like a probabilistic language generator.¹⁷ Because it relies on predictions, it is capable of producing both accurate summaries and 'hallucinations', where it generates inaccurate or fabricated content. In *Like Company*, the Hungarian-language summary of the Balaton Környéke article may therefore be analysed as a new linguistic realisation guided by patterns in data and (under the primary hypothesis explored here) by retrieval, rather than as the system 'playing back' a stored copy of the press publication. From a technical perspective, Gemini's core model is not an accessible archive of works but a set of learned parameters.¹⁸

2.2 | Training stages

At a high level, an LLM's life-cycle can be divided into (i) data gathering and pre-processing, (ii) training, and (iii) post-training and deployment. These stages are directly relevant to the *Like Company* dispute because the claimant alleges reproduction of its article during the input phase.

In the input phase, large data sets are assembled from multiple sources.¹⁹ Raw data are then pre-processed: duplicates, corrupted files and irrelevant material are filtered out; data are cleaned and normalised;²⁰ and texts are tokenised into words or sub-word units, while other modalities (audio, video, images) are segmented into workable chunks.²¹ This stage may also involve annotation or labelling.²² Pre-processing serves to improve data quality, reduce bias and noise, and make the corpus computationally manageable so that the model learns meaningful patterns rather than artefacts of messy input.²³

Training begins once this processed data set is assembled into a training corpus.²⁴ At this point, temporary copies of the data necessarily exist within the technical pipeline, during ingestion, caching and tokenisation, to permit learning. These copies are functional and short-lived but, strictly speaking, the data consists of a transient copy. During training, the system converts input into numerical representations and adjusts its internal parameters (weights) to minimise prediction error. Gemini v1.0, like other LLMs, is therefore an architecture plus a very large array of numbers that encode statistical regularities such as how often words or visual elements co-occur.²⁵

Two observations are here crucial for the subsequent legal analysis of the questions referred to the CJEU in *Like Company*. First, the model does not contain accessible copies of the works in the training data.²⁶ Models can memorise some material, but they are not static repositories from which users can browse or retrieve particular works at will. In normal operation a trained model such as Gemini v1.0 cannot be regarded as making specific training works available to the public. Second, each individual work contributes only a negligible fraction to overall model capability: what matters is the aggregate information extracted from billions or trillions of tokens.²⁷ The marginal impact of a single news article, or even an entire regional news site such as Balaton Környéke, on the model's capabilities is vanishingly small.

In addition, there is a temporal constraint that is decisive in *Like Company*. Gemini v1.0 has a knowledge cut-off of February 2023, while the disputed article was published in July 2023.²⁸ The article therefore could not have been present in the training corpus of the specific model version that generated the impugned output. As a matter of technical fact, no reproduction of that press publication occurred during Gemini v1.0's training.

2.3 | Memorisation

Nonetheless, LLMs like Gemini can under some circumstances memorise and regurgitate specific items from their training data. Memorisation describes the situation where a model retains concrete snippets rather than only abstract patterns, so that it can reproduce them substantially when suitably prompted.²⁹ This phenomenon is associated with overfitting, where the model has been trained too closely on particular examples, often in highly duplicated or repetitive data.³⁰

Empirical studies show that memorisation is relatively rare and tends to affect highly repeated or salient sequences (famous poems, lyrics, well-known images), rather than ordinary texts such as a single regional news story.³¹ It can be substantially reduced by deduplication of training data,³² regularisation to discourage over-reliance on specific features,³³ and other techniques such as differential privacy, early stopping, and gradient clipping.³⁴ Memorisation therefore exists but is statistically bounded and technically mitigable; it does not convert Gemini or similar models into a database of their inputs.

2.4 | Post-training and retrieval-augmented generation

After pre-training, models like Gemini are often further optimised through post-training techniques. This phase does not fundamentally change the model's nature as a parametric function, but it shapes how it responds to prompts. One widely used method is Reinforcement Learning from Human Feedback, in which humans rate or rank outputs and their preferences are used to fine-tune the model so that it produces more helpful, safe and 'aligned' answers.³⁵ Post-training techniques include also processes based on so-called 'inference',³⁶ where a trained model applies its learned parameters to new data to produce an output, as in the case of retrieval-augmented generation (RAG).³⁷

A RAG-based operation offers a primary hypothesis for understanding the *Like Company* facts. When users interact with Gemini, their prompts and histories of the conversations are packaged into a 'context window', that is, the temporary working memory for that interaction.³⁸ This is not part of the model, but an instance of interactions between the model and the user. Within the 'context window', RAG supplements what is in the model's weights, at the knowledge cut-off point, with information retrieved in real time from external sources (search indices, document repositories), or the web, at inference time.³⁹ The RAG component selects the most relevant passages, which are then injected into the 'context window' and supplied to the LLM. The model is instructed to generate an answer conditioned both on its learned parameters and on this fresh retrieved material. The effect is to bridge traditional information retrieval and GenAI: the system behaves partly like a search engine (finding relevant documents) and partly like a generator (synthesising an answer in natural language). This design allows models to stay up to date despite a fixed training cut-off and reduces hallucinations by grounding responses in external sources. It also enables citation or attribution of those sources and modular updates to the retrieval layer without retraining the entire model.

In Google's publicly described 2025 Gemini architectures, this retrieval-based workflow is complemented by an explicit 'grounding' phase.⁴⁰ The model issues one or more search queries to Google Search, processes the returned results into candidate answers, and selects a response that is explicitly anchored in those search results.⁴¹ The search phase itself is a classic information-retrieval operation on an index; no additional 'making available' of content occurs inside the model beyond what is already possible via Google Search.⁴² The generative component then reformulates and condenses that content for the specific user, within the confines of the context window.⁴³

2.5 | Technical characteristics of the *Like Company* case

Against this background, the *Like Company* dispute turns on two technical assertions. As mentioned, Like Company alleges, first, that Gemini was trained on its content during the input phase, thus allegedly infringing its right of reproduction; and second, that the chatbot's detailed summaries amount to making it available to the public.⁴⁴ The claimant's core intuition is that such summaries would only be possible if Gemini had first reproduced the article in its entirety during training.

But as already noted, the publicly reported timing of Gemini v1.0's training and the publication date of the *Balaton Környéke* article point in the opposite direction: the article appears to post-date the relevant model's

knowledge cut-off.⁴⁵ On that basis, the page could not have been included in the model's training corpus as a matter of timing. Accordingly, on the best reading of the public record, there was no reproduction of that article at any stage of training for the model used in these proceedings. Nor could Gemini v1.0, as a set of parameters, make that specific article available to the public, because it neither contained the article nor allowed users to retrieve it as such.

Moreover, even apart from timing, it is highly unlikely that Gemini would reproduce an ordinary regional news article verbatim or in substantial part. Google reports that verbatim memorisation rates for the Gemini family are extremely low, and independent studies of similar large models find that regurgitation tends to be limited to highly duplicated or famous material.⁴⁶ Combined with the fact that the *Like Company's* article was not in the training data at all, this makes substantial memorisation in this case technically implausible.

Instead, one plausible interpretation of the available technical evidence suggests that the contested summary was generated through RAG and grounding. By explicitly supplying the URL and asking for a summary of 'the online press publication that appeared on *balatonkornyeye.hu*',⁴⁷ the claimant's prompt triggered Gemini's retrieval mechanisms. A retrieval-based system would typically issue one or more search queries to Google Search, retrieve the relevant page, process it into embeddings, and use those retrieved passages within the context window to generate a condensed narrative in Hungarian.⁴⁸ The underlying model provided linguistic capability and summarisation skills, but the factual content was drawn in real time from the indexed web page, not from training-time copies.

Technically, then, under this framework, Gemini v1.0 operates in *Like Company* as an LLM equipped with RAG and grounding. It does not contain or reproduce the news article in its weights, nor does it, by virtue of its architecture, make that article available to members of the public beyond those who already have access via the Balaton Környéke website and Google Search. The output at issue is best characterised as a probabilistic summary generated on-the-fly from lawfully accessible online content, rather than as a communication of stored copies from within the model itself.

3 | THE PRESS PUBLISHERS' RIGHT

To answer Question 1, that is whether chatbot responses that display text partially matching press publisher's web pages constitute communication to the public, particularly when the reproduced text is substantial enough to fall within the scope of the PPR, it is necessary to start with Article 15 CDSM's internal threshold. The notion of a 'very short extract' matters here because it operates as a built-in limitation on the PPR's scope. Indeed, the referring court explicitly frames the dispute around chatbot answers that 'partially' match press web pages and exceed what can be regarded a 'very short extract' within the meaning of Article 15 CDSM.

In this context, the Hungarian referring court signals that liability for communication to the public arises only once the output falls outside this statutory carve-out. We will return to the general reproduction and communication to the public questions in Sections 4 and 6, respectively. For present purposes, the preliminary task is to clarify the scope of Article 15 CDSM, that is, including whether it extends to AI-generated content, and its structural relationship with Articles 2 and 3 InfoSoc.

3.1 | Justification for the introduction of the PPR

The PPR's aim is to recognise and further encourage publishers' organisational and financial contributions to the production of press publications, supporting the sustainability of the publishing industry and the availability of reliable information.⁴⁹ It is believed that a harmonised legal approach at the EU level would assist publishers of press publications facing difficulties in licensing the online use of their content to providers of new online services,

such as news aggregators or media monitoring services, allowing them to recover their investments; and that the lack of such a right would make the licensing and enforcement of rights in press publications related to online uses by information society service providers complex and inefficient.⁵⁰ Importantly for this study, GenAI chatbots like Gemini offer information society services and are thus subject to Article 15 CDSM.

3.2 | The relationship between Art. 15 CDSM and the InfoSoc regime

Crucially, the relationship between this specific regime and the general copyright framework is central to the *Like Company* referral. The referring court explicitly asks whether the application of the general communication to the public right under Article 3(2) InfoSoc 'may exclude' the application of Article 15 CDSM.⁵¹ The referral's 'may exclude' formulation brings an architectural question to the surface: if conduct that is, in substance, the online use of a press publication could be re-labelled and assessed solely under the broader making available right in Article 3(2) InfoSoc, Article 15 CDSM's internal limits, above all 'individual words or very short extracts', risk becoming easy to plead around. This suggests a structural hierarchy: Article 15 was designed as a calibrated interface for press content online, containing specific balances, most notably the exclusion of 'very short extracts', that were heavily debated by the EU legislature. If claimants could bypass these internal limits simply by invoking the general InfoSoc rights for the same conduct, the legislative compromise would be effectively neutralised. Consequently, the analysis of press content usage by information society service providers must arguably be governed primarily by the specific provisions of Article 15, preventing the general law from swallowing the specific regulation.

3.3 | Does the PPR extend to AI-generated content and to content resembling the one covered by the PPR?

Despite the fact that the PPR is not expressly linked to the existence of copyright in the protected subject matter, the philosophy behind it is not to extend protection where protection does not apply under copyright, or the content has fallen into the public domain. On the contrary, one of the publishers' arguments for its introduction was that this additional related right was needed because it is not always clear from the chain of transfers or the contracts/assignments with creators whether a publisher owns or possesses the rights in a particular piece of content, to block the use or press charges for unauthorised use based on a copyright licence.⁵²

If the PPR were interpreted to cover content that is not protected by copyright, it could also be invoked in relation to AI-generated content that inherently lacks copyright protection, because, under EU copyright's human-authorship requirement, output generated autonomously by an AI system will generally fail the 'author's own intellectual creation' standard and therefore fall outside copyright in the first place.⁵³ In other words, even where AI-generated content falls outside copyright and would normally be free for anyone to use, it could be pulled back under exclusive control through the PPR. That would be conceptually incoherent: it would treat the same content as 'unprotectable' for copyright purposes, yet still protecting it via a related right defeating its original purpose for non-protection.⁵⁴ For that reason, AI-generated content that is not protected by copyright should not be brought within the scope of the PPR.

This is also in line with Article 14 CDSM, which provides that when the term of protection of a work of the visual arts has expired, any material resulting from an act of reproduction of that work is not subject to copyright or related rights. Thus, a related right cannot be invoked to prohibit the use of works or other subject matter for which protection has expired and has therefore fallen into the public domain.⁵⁵

In light of the above, the PPR does not cover 'very short extracts' and should not cover AI-generated content or content that is not or is no longer protected by copyright. In addition, the protection afforded by this related right should apply to content used verbatim or extracted verbatim from the publication, that is word for word, as is the

case, for example, for other types of related rights, such as those for performers and for the producers of sound and visual recordings. These related rights pertain, respectively, to the publication, performance, or recording of a particular work as such, and not to any content resembling the original. It is the entrepreneurial act (or, in the case of performers, the interpretative act) as such that is protected and not variations of what this act would look like if the content published, interpreted, or recorded were similar. Any other interpretation would unjustifiably block other uses or creations that could express the same facts in a different manner. If the wording produced by chatbots like Gemini are based on the facts of the publication but are technically the outcome of predicted language and hallucination by the chatbot (e.g., as a person in real life would narrate the same things), then no infringement could take place as this narration would be unlikely to infringe the content of the publication; on the contrary the new narration (if done by a human being) may also invoke a separate copyright protection or even infringe the creator's original copyright in the work, and if published it may also attract a separate PPR.

In conclusion, chatbots like Gemini, which generate language that is the outcome of predicted language and hallucination, without being a word-for-word reproduction of content protected by the PPR, cannot give rise to any infringement. Even more if the extract appearing in the outcome qualifies as a 'very short extract'.

4 | INPUT AND REPRODUCTION

This section addresses Question 2: that is, whether the training of an LLM-based chatbot such as Gemini, built on the observation and matching of linguistic patterns, constitutes an act of reproduction within Article 2 InfoSoc particularly regarding press publications, whether such technical acts are pre-empted or governed by the specific regime of Article 15(1) CDSM.

This study suggests that, on the better reading of current EU law, model training, properly characterised, cannot, as such, be treated as reproduction of protected expression. That conclusion preserves the idea/expression dichotomy,⁵⁶ and avoids creating a de facto control right over 'learning'. Learning statistical regularities is better understood as distinct from the making of copies of a work, and stretching reproduction to cover statistical parameterisation risks being both doctrinally unsound and economically counter-productive in the EU acquis.⁵⁷ That position is compatible with an important qualification. As noted in Section 2 above and developed further below, models can, in some rare circumstances, regurgitate passages from training material that they have effectively memorised,⁵⁸ often under conditions of duplication, rarity, overfitting, and, in practice, adversarial or extraction-style prompting. However, it is crucial to emphasise that the incidental memorisation of a single work does not imply that the model systematically retains copies of its entire training corpus. Nor is there a reliable, ex ante way to infer from the trained weights whether any particular work has been memorised, or to extract that information directly from the trained model. The proof lies solely in the output.⁵⁹

4.1 | A preliminary note: Keeping training and RAG legally distinct

As explained earlier, the factual record indicates that no reproduction occurred at the training stage in this case, because the website at issue was published after the knowledge cut-off date for Gemini v1.0, that is, the first generation of the Gemini model family, announced in December 2023 as Google's new flagship multimodal foundation model. Any allegedly infringing conduct must therefore be ascribed, if at all, to RAG at inference time (on which see the analysis in Section 6). Put bluntly, the factual matrix behind the reference is about what Gemini displayed to users when answering a prompt, that is, a RAG-based response that allegedly incorporated fragments from a press article, rather than about the offline training phase in which the model's parameters were learned.

Keeping these two stages distinct helps avoid conflating 'learning from' a text' with 'showing' that text to users. RAG and model-training are technologically and legally distinct: the latter raises communication to the public and

output-layer reproduction/attribution issues⁶⁰; the former raises the question addressed here, that is, whether the *learning process itself* fixes protected expression. Given the wording of Question 2, it is appropriate to analyse training *as such*; however, conflating inference-time retrieval with training risks inviting the CJEU to opine on matters not necessary to resolve the national dispute. The Court presumes the referring court's questions are relevant, yet it may reformulate or decline aspects that are hypothetical or bear no relation to the facts essential to judgment.⁶¹ Our discussion of 'training ≠ reproduction' therefore proceeds on a clean technical footing, without relying on inference-time behaviours, and is intended to offer a useful interpretation should the CJEU consider the training issue necessary to the disposition of the case.

4.2 | The EU concept of reproduction and its limits

Article 2 InfoSoc confers on rightholders the exclusive right to authorise or prohibit direct or indirect, temporary or permanent reproduction by any means and in any form, in whole or in part, of their works. The CJEU has given this provision a technology-neutral but expression-oriented reading. In *Infopaq*, it held that the fixation of 11-word extracts may be a reproduction 'in part' if the extract contains elements which are the author's own intellectual creation; crucially, the finding turned on the existence of copies embodying protected expression, even if brief and transient in the scanning and storage stages of a news-clipping workflow.⁶² In later cases the CJEU insisted that protection attaches to expression, not abstract ideas, styles or methods (the 'idea/expression' dichotomy remains a structural limit), and that recognisability of the original protects the boundary between lawful influence and infringement.⁶³ In *Pelham*, for example, the Court stated that a sound sample may escape infringement where it is unrecognisable to the ear in the new context.⁶⁴

Two immediate consequences follow for the *Like Company* reference. First, analysis *as such* should not be treated as reproduction. An analytical process, even one that ingests protected works, does not itself constitute reproduction unless it fixes expression. *Infopaq* condemns the making of copies, not the mere act of processing or 'reading' inputs. The decisive question is whether training fixes protected expression. Secondly, EU law has consistently declined to extend rights to functional or abstract features, that is, program functionality, programming languages and formats are outside copyright.⁶⁵ That logic, rooted in the idea/expression dichotomy, matters when assessing what is stored inside a modern model. Models internalise patterns rather than store human-readable sequences.

4.3 | What training does (and does not) store: From text to statistics

Modern LLMs ingest large text corpora and use gradient-descent optimisation to adjust parameters (weights) in a neural network so that, given a context, the model minimises prediction error for the next token.⁶⁶ The learned artefact consists of a high-dimensional parameter matrix and associated embeddings.⁶⁷ It is not a database of strings but a statistical representation of co-occurrence patterns captured via latent-space interpolation.⁶⁸ In information-theoretic terms, the process is lossy: it discards vast amounts of surface form, that is, verbatim character strings, and retains compressed correlations that improve prediction.⁶⁹ Retrieval of specific protected passages from the parameter matrix is neither the training objective nor a generally available capability absent special circumstances, for example, overfitting to rare, duplicated sequences combined with attack-style prompts—a point to which we return later in this Section.⁷⁰

Doctrinally, this matters because parameters are not expression. They are numeric encodings of statistical relationships⁷¹; they do not, in ordinary operation, fix or embody the protected form of any given work.⁷² On the CJEU's own terms, recognisability is pivotal⁷³: what the model retains is not recognisable expression, but abstraction.⁷⁴ Treating the parameter matrix as a copy of every input would collapse the idea/expression

boundary⁷⁵ and by parity of reasoning imply that human learning, which also internalises regularities,⁷⁶ infringes whenever one studies protected works: that is, an outcome EU law has never endorsed. Here the model 'memorises' patterns through lossy compression and latent-space interpolation: it discards verbatim detail and retains only statistical patterns, not the original.⁷⁷ That looks much closer to human reading and learning, that is, internalising ideas and styles, than to storing expressive copies.⁷⁸ In this regard, the point could be made that an AI model's internalisation of data, like human learning, does not produce expressive copies but instead encodes abstract patterns or statistical relationships. A recent US decision provides a useful external confirmation of the same basic intuition: in *Bartz and others v Anthropic PBC*, the US court characterised LLM training on lawfully acquired books as 'quintessentially transformative' and expressly likened it to a reader learning in order to write, training 'not to [...] replicate or supplant [the works]', but to 'create something different'.⁷⁹ While that fair use holding is not transferable to EU law, the description of what training does is: it frames training as learning from works rather than warehousing them, which is exactly why equating parameterisation with 'reproduction' is doctrinally overbroad.

4.4 | Why comparing training to reproduction is doctrinally overbroad

An automatic equivalence between training and reproduction may sit uneasily with three well-settled strands of CJEU jurisprudence. First, the CJEU has repeatedly anchored protection in authorial expression, not in ideas, methods, styles or functional elements. Read together, cases such as *Infopaq*, *Painer*, *Football Dataco* and *SAS Institute* confirm that functionality, programming languages and data formats fall outside copyright, while protection attaches only to the author's original expressive choices.⁸⁰ Against that backdrop, LLM training does not store expressive sequences but abstracts style and statistical regularity from a corpus, precisely the kind of non-expressive regularities that EU law has refused to monopolise.

Secondly, the Court has now made recognisability of creative elements the practical limiting principle for reproduction. As mentioned, in *Pelham*, the phonogram right was not infringed where a sample had been altered so as to be 'unrecognisable to the ear': copyright protects expression, not every trace of influence.⁸¹ The recent CJEU's decision in *Mio/Konektra* generalises that logic to works of applied art and, by implication, to works under Articles 2–4 InfoSoc: infringement requires that the original creative elements of the protected work are reproduced in a recognisable manner in the later subject matter, whereas the 'overall impression' test belongs to design law.⁸² Stylistic similarity, that is, 'this looks like Eames/sounds like Dylan/reads like Woolf', may matter for design, but it is not the copyright test.⁸³ Nor do fragment size or degree of originality modulate scope: once a work qualifies, any recognisable reproduction of its creative elements is covered.⁸⁴ Transposed to LLMs, this points away from treating a trained parameter matrix as a 'copy': the learned weights are not recognisable as the text of any particular article; they encode correlations, not the protected form. The infringement risk instead lies at the output layer, where memorisation and regurgitation may produce recognisable passages from training works, that is, an issue to be addressed with targeted safeguards, not by reclassifying the entire learning process as reproduction.

Thirdly, the *Infopaq* line draws a clear distinction between analysis and fixation. The finding of reproduction there depended on the creation of actual string copies (snippets scanned, stored and printed) however brief their duration.⁸⁵ Where a process does not generate or retain human-readable fragments, but yields only statistical traces in a parameter space, it is better characterised as analysis than as reproduction.⁸⁶ To extend Article 2 to such non-expressive traces would be to regulate learning itself, rather than copying, and to collapse the idea/expression dichotomy that structures EU copyright.⁸⁷

A contrary rule would carry far-reaching and technology-distorting consequences.⁸⁸ It would equate any statistical model trained on protected materials, whether in language, music or vision, with a copy of all its sources, thereby creating a de facto control right over learning.⁸⁹ That result is alien to the EU tradition and difficult to administer in practice.

4.5 | Transient copies as technical ephemera in the pipeline, not training itself

Thus, on the best reading of the *acquis* and the CJEU's expression-centred jurisprudence, the process of training an LLM, understood as the analytical adjustment of parameters to capture statistical regularities in lawfully accessible text, may not constitute an act of reproduction under Article 2 InfoSoc. The trained artefact is a parametric model, not a store of human-readable sequences; it does not fix protected expression and so falls on the 'analysis' side of the idea/expression dichotomy that the Court has consistently maintained.

This conclusion does not ignore engineering realities. Training pipelines may produce fleeting, functional artefacts, that is, ingestion buffers, caches, tokenisation shards,⁹⁰ that, taken strictly, can amount to acts of reproduction.⁹¹ But their existence may not make *training* itself a 'reproduction' for the purposes of Question 2.⁹² These artefacts are ancillary, transient and incidental to the technological process; they are not the learning event.⁹³ EU law resolves them with the instruments built for such technical ephemera: where inputs are lawfully accessible and no reservation is breached, Article 5(1) InfoSoc (which, it should be noted, is mentioned in the *Like Company* referral but not in the questions posed to the CJEU) covers the temporary, integral copies; where the analysis runs beyond that narrow shelter, Article 4 CDSM (as we will see more in depth in the next section) supplies the dedicated, technology-neutral TDM mechanism, including a machine-readable opt-out. This two-step design provides legal certainty and avoids the high transaction costs of litigating every cache or batch, or worse, chilling research through inaction.⁹⁴ At the same time, Article 4's opt-out gives rightholders a practical route to internalise value where licensing is economically sound, without converting learning into a property-rule control, as mentioned already above. The default is permission to mine; a rightholder may switch a given use onto a licensing track by making a machine-readable reservation, enabling remuneration in markets that can bear it. That approach aligns with copyright fundamentals, namely the idea/expression dichotomy: mining/training extracts patterns, trends and correlations rather than expressive form, and so it is analysis, not copying.

Framed this way, the answer to Question 2 is negative in the sense that matters for authorisation: *training* as *such* cannot be a rights-relevant reproduction. At the same time, legal certainty is preserved because any ancillary, short-lived pipeline copies arising in a compliant workflow are either exempt under Article 5(1) InfoSoc; or, where applicable, authorised under Article 4 CDSM. That preserves the settlement under the latter Directive, respects rightholders' choices through reservation and licensing where they wish to exercise control, and avoids creating a *de facto* monopoly over the act of learning. On that footing, Section 5 of this article will turn to the exceptions-and-limitations framework: if a court nevertheless treats some pipeline elements as reproductions, the correct analysis lies in applying the legislature's chosen tools, that is, temporary copying (Article 5(1) InfoSoc) and TDM exception (Article 4 CDSM), rather than expanding Article 2 InfoSoc.

4.6 | Memorisation, regurgitation and output-layer responsibility

As an additional clarification to avoid talking past the technology, it is useful to distinguish rare 'regurgitation' phenomena from the nature of training itself.⁹⁵ As anticipated in Section 2, computer-science research shows that memorisation, that is, the verbatim emission of training snippets, does occur, but it is statistically bounded and materially mitigable through data set deduplication, regularisation, sampling controls and similar engineering choices⁹⁶; leading providers also deploy safety filters and anti-extraction measures that further reduce the likelihood that training results in expressive regurgitation.⁹⁷

Even where string-level memorisation can be induced by adversarial prompting,⁹⁸ that behaviour concerns the output layer and service design rather than the training process as 'reproduction'. Put differently: the occasional appearance of long passages is not proof that the learned parameters are copies; it is an operational risk to be addressed by provider diligence (e.g., filters, guardrails, notice-and-action) and, in exceptional cases, targeted injunctions against specific extraction behaviours. Those questions of diligence and attribution are properly

analysed under Question 4 (see Section 7 below), not as a reason to collapse the doctrinal boundary between analytical training and rights-relevant reproduction.

Crucially, it is vital to distinguish targeted regurgitation driven by adversarial prompting from continuous, systematic market substitution for ordinary readers. A recent study illustrates this precise distinction.⁹⁹ While the researchers demonstrated that text extraction from a book is technically possible, their results were achieved only through a highly specific, non-standard 'sliding window' attack that involved thousands of successive prompts designed to bypass security filters,¹⁰⁰ a process far more resource-intensive than simply purchasing the book.¹⁰¹ Furthermore, the attack's success depended heavily on the specific target, *Harry Potter and the Sorcerer's Stone*, a work repeated thousands of times across the public web, making it a statistical outlier particularly susceptible to memorisation.¹⁰² Consequently, while media coverage has understandably amplified the headline risk,¹⁰³ the paper itself treats the result primarily as a safety/security vulnerability addressable by technical means,¹⁰⁴ not as a freestanding copyright conclusion that LLMs are 'free libraries'. Rather, these findings confirm that while determined actors using sophisticated tools can elicit memorised sequences from highly duplicated data, standard guardrails such as data deduplication and output filtering remain effective in preventing the model from serving as a market substitute in ordinary operation.

Recent national litigation, however, illustrates how easily the debate can slide from 'rare regurgitation' to a general theory that treats learning itself as copying. The recent decision by the Munich Regional Court in *GEMA* is widely being presented as a landmark for 'training-as-reproduction', but much of the public takeaway rests on conflating proof of memorisation (via prompted lyric outputs) with a general claim that works are stored 'inside' a model as such.¹⁰⁵ That framing is technically contestable and doctrinally risky, because it substitutes an output-driven inference (recognisable passages can be elicited) for the Article 2 InfoSoc question whether the trained artefact *as such* fixes protected expression.¹⁰⁶

In sum, legally and economically, the implications of such memorisation (regurgitated via adversarial prompting) are narrow. Verbatim emissions are rare, prompt-dependent and concentrate on duplicated training items; they fall sharply with deduplication and alignment. Occasional output that matches training text does not prove a latent internal copy of each work; it shows memorisation at the margin, which can be mitigated. Treating these tail-events as proof that the entire learning process equals reproduction would turn Article 2 InfoSoc into a property rule over learning capacity. The fact that a model *can* be forced, through sustained adversarial interaction, to emit recognisable passages is evidence that output-layer guardrails matter and that liability analysis should focus on the concrete instance of recognisable reproduction in the output, rather than treating training or ordinary conversational summarisation as market-substituting 'publication' by default.

4.7 | A cautionary note: Harmonisation with the PPR's scope

The second referred question explicitly links Article 15(1) CDSM and Article 2 InfoSoc, asking whether training constitutes reproduction under either provision. It is crucial to insist that the definition of 'reproduction' remain consistent across both.¹⁰⁷ While related rights typically protect the specific 'fixation' of a publication rather than its 'originality', interpreting 'reproduction' under Article 15 to cover non-expressive technical acts (such as statistical parameterisation) would lead to a paradoxical result: the related right would become functionally broader than the copyright it is meant to complement.¹⁰⁸

Specifically, if Article 15 were read to capture any technical processing of the fixation 'as such', even where no original expression is preserved, it would grant publishers a monopoly over the informational content of news, a protection explicitly denied to authors under copyright.¹⁰⁹ This would contravene the legislative intent that the rights granted to press publishers should have the same 'scope' as the reproduction right provided for in InfoSoc.¹¹⁰ Consequently, to avoid inverting the hierarchy between the author's primary right and the publisher's related right, 'reproduction' under Article 15 must be interpreted as requiring the same focus on recognisable expression as

Article 2 InfoSoc. Therefore, Article 15 must be interpreted restrictively to exclude acts of machine learning that do not preserve the protectable substance of the publication.

5 | TDM EXCEPTION

The related and crucially important Question 3 concerns the applicability of the TDM exception to the *Like Company* factual scenario. This section examines not only the scope of the CDSM's TDM exception in that dispute, but also the explicit legislative intent of EU institutions, particularly as codified in the AI Act and Recital 105 CDSM, to subject AI training to the TDM exception framework.

As mentioned, the CJEU is asked whether, in case an LLM-based chatbot like Gemini constitutes an instance of reproduction under Article 15(1) CDSM and Article 2 InfoSoc, this reproduction of lawfully accessible works falls within the TDM exception under Article 4 CDSM. As we have exposed before,¹¹¹ we consider that in scenarios such as in *Like Company*, reproductions in the sense of Article 2 InfoSoc may not materialise; however, should the transient copies and any related step in the training process qualify as 'reproductions' by the Court, we are of the opinion that the TDM exception applies in any case in this context, legitimising such use in the absence of a right reservation by *Like Company*.

Before analysing the TDM exception, we preliminarily note that a hierarchy of defenses applies here: the RAG of the article in question likely constitutes a transient, incidental act under Article 5(1) InfoSoc. The issue is of interest, even though as mentioned such exception falls outside the express scope of the questions referred to the CJEU in *Like Company*, which suggests that the referring court may have treated this defense as settled or secondary to the copyright infringement analysis. That the RAG in question is such a copy under Article 5(1) InfoSoc is confirmed by the technical analysis within Section 2 of this article, which shows that the grounding mechanism through which Gemini has retrieved the Balaton article did operate in real time, creating temporary buffers and cache copies that persisted only for the duration of the inference process. These ephemeral copies were integral to the technical functioning of the search and summarisation process itself, much like the transient reproductions occurring in RAM or cache memory during any ordinary web retrieval.

Said that, how the training of an LLM like Gemini can be evaluated from a copyright perspective, in particular when the algorithm is 'fed' with copyright protected works,¹¹² is here the crucial question. Different from other jurisdictions such as the US, where the legality of machine learning is discussed in the framework of the fair use exception¹¹³ or Japan, where GenAI training is legalised under a specific 'non enjoyment' exception,¹¹⁴ the EU legislator took the decision in the AI Act to subject the training of GenAI models to the CDSM's TDM exception. Indeed, Article 53, 1-c) AI Act imposes obligations on providers of general-purpose AI models like Gemini to 'put in place a policy to comply with Union copyright law, and in particular to identify and comply with, including through state-of-the-art technologies, a reservation of rights expressed pursuant to Article 4(3) of Directive (EU) 2019/790 [CDSM]' (emphasis added).¹¹⁵ This provision needs to be read together with Recital 105 CDSM, which clarifies the EU legislator's will by making explicit reference to GenAI models and their training on copyright protected works. According to this recital:

general-purpose AI models, in particular large generative AI models, capable of generating text, images, and other content, present unique innovation opportunities but also challenges to artists, authors, and other creators and the way their creative content is created, distributed, used and consumed. The development and training of such models require access to vast amounts of text, images, videos and other data. Text and data mining techniques may be used extensively in this context for the retrieval and analysis of such content, which may be protected by copyright and related rights. Any use of copyright protected content requires the authorisation of the rightsholder concerned unless relevant copyright exceptions and limitations apply. Directive (EU) 2019/790

introduced exceptions and limitations allowing reproductions and extractions of works or other subject matter, for the purpose of text and data mining, under certain conditions. Under these rules, rightsholders may choose to reserve their rights over their works or other subject matter to prevent text and data mining, unless this is done for the purposes of scientific research.

This has always been the position of the European Commission also prior to the AI Act: when asked to clarify its position, Commissioner Thierry Breton on behalf of the Commission considered that the existing EU legislative framework fit for the purpose of GenAI as the 'creation of art works by AI does not deserve a specific legislative intervention'; and this was the case because the TDM exceptions give content creators opt-out chances, 'providing balance between the protection of rightholders including artists and the facilitation of TDM, including by AI developers'.¹¹⁶ It has been equally the position of a majority of scholars that the training of GenAI models like Gemini is subject to the CDSM's TDM rules.¹¹⁷ It can nevertheless be regretted that this important legislative decision in 2024 did not follow an informed debate in the European Parliament. The relevant provision of the AI Act was introduced last minute in the very last round of discussions over the TDM exceptions. Those exceptions, even if fairly recent, also predate the GenAI boom.¹¹⁸ Therefore, as determined by the EU legislator, the technical steps and reproductions necessary for training GenAI models fall within the scope of the CDSM TDM exception. This is because AI systems learn from data sets through TDM.¹¹⁹

Crucially, the framework described above seems to encompass the specific technical process at issue in the *Like Company* referral: that is the 'grounding' or RAG process. When the model retrieves the Balaton article to generate a summary, it is performing an 'automated analytical technique' under Article 2(2) CDSM to extract semantic information, thereby triggering the exception. This real-time retrieval and analysis of lawfully accessible content, that is the inference-stage processing, constitutes TDM regardless of whether the training phase itself involves reproductions, ensuring that derivative works and summaries generated from publicly available sources remain lawfully exempt. More specifically, Article 2(2) CDSM defines TDM as 'any automated analytical technique applied in order to analyse text and data in digital form in order to generate information'. The grounding process precisely matches this definition: it is automated, analytical, applied to digital text, and generates new information, that is, the summary. This functional alignment grounds the TDM exception not in theoretical disputes about parameterisation, but in the concrete observable act of content analysis.

As is known, Article 3 CDSM introduces an exception for TDM for scientific research, which solely benefits research organisations and cultural heritage institutions. By contrast, Article 4, which was added later in the legislative process that led to the CDSM, is not limited to specific types of institutions and therefore applies more broadly.¹²⁰ The latter is the provision effectively relied upon by the AI Act, as GenAI systems like Gemini are usually developed and operated by private commercial companies that do not fall within the scope of Article 3. According to Article 4(1) CDSM, 'Member States shall provide for an exception or limitation [...] for reproductions and extractions of *lawfully accessible works* and other subject matter for the purposes of text and data mining' (emphasis added). Interestingly, under Article 4(2) CDSM, 'reproductions and extractions made pursuant to paragraph 1 may be retained for as long as is necessary for the purposes of text and data mining', which is relevant to solve any possible question of storage or memorisation of protected works by GenAI models in the learning process.¹²¹

In the *Like Company* case, the Balaton article was freely available on the open web without paywalls, authentication walls, or technical restrictions, that is, a critical factual foundation for establishing lawful access under Article 4(1) CDSM. The publisher made the content publicly indexable and retrievable by standard web protocols and search mechanisms, thereby consenting to its accessibility in the information ecosystem. This contrasts with scenarios involving hacked databases, circumvented paywalls, or deliberately obscured repositories. The absence of technical or legal barriers to access means that the retrieval of the article for purposes of analysis and summary generation, whether during training or inference, occurs in a context where the content's availability is lawfully assured. This satisfaction of the Article 4(1) prerequisite is fundamental to triggering the TDM exception's protections.

A radically different approach has been taken recently by some scholars, arguing that the TDM exception does not cover training activities¹²² or at least not all copyright-relevant uses of GenAI training activities¹²³; or that from a policy perspective, it should not cover them because of structural imbalances this would introduce.¹²⁴ In this sense, the decision of the Munich court in *Gema* caught a lot of attention, as it decided that only parts of the training activities are covered by the TDM exception, while the memorisation/storage of song texts via training involves reproductions that are not covered and therefore infringing.¹²⁵ This ruling, however, is very debatable on the TDM side (input), as we have seen that the AI Act clearly intends to submit GenAI training to the TDM exception under Article 4 CDSM and the opt-out mechanism. Splitting the training process into different technical parts seems to be a *contra legem* interpretation which will be detrimental to a unified approach in the EU member states which the regulation of the AI Act meant to achieve.¹²⁶ For this reason, we believe the Munich court should have referred the matter to the CJEU for clarification as it clearly involves the interpretation of EU law.

The outcome reached by the Munich Court in *Gema*, which divided the training process into separate technical phases and denied TDM protection to the storage/memorisation component, underscores the critical importance of the *Like Company* referral to the CJEU. Unlike the German court's *contra legem* approach, which seems to conflict with the AI Act's explicit subjection of GenAI training to Article 4 CDSM, the Hungarian referring court has framed its questions to encompass both the training process and the inference-stage (RAG/grounding) operations. This broader scope affords the CJEU an opportunity to establish authoritative, unified guidance that TDM protection covers the entire training pipeline, including memorisation and parameter storage, as an integrated analytical technique under Article 2(2) CDSM. In other words, by anchoring its analysis in the concrete facts of *Like Company*, where the alleged infringement pivots on real-time content retrieval and lawful access, the Court can establish binding EU-wide clarity: that is, that Article 4 CDSM applies holistically to GenAI development absent valid opt-outs, and that no artificial subdivision of technical stages can circumvent the exception's protections. Such a ruling would pre-empt fragmentation and preserve the digital single market's coherence.

In sum, also considering what the EU legislator meant by passing the AI Act, the reproductions aimed at training GenAI chatbots like Gemini (assuming reproductions do materialise) are subject to the TDM exception under Article 4 CDSM. As we have underlined, this solution is surely not ideal and might call for a revision *de lege ferenda* in order to implement a model more favourable to GenAI innovation in the EU, while safeguarding the core European values such as a fair remuneration for creators. But according to the current legislative framework, there is little doubt that the use for training should be exempted unless the rightholders have opted out. The *Like Company* referral presents the CJEU with an ideal vehicle to crystallise this legislative intent into binding judicial doctrine, thereby foreclosing national-court deviation and anchoring the TDM exception as the governing framework for lawful GenAI development across the EU.

Finally, while the *Like Company* referral does not explicitly turn on whether the publisher validly exercised an opt-out, the legal uncertainties surrounding Article 4(3) CDSM's opt-out mechanism remain central to the CJEU's broader interpretative task. The opt-out provision has been extensively critiqued for generating legal insecurity¹²⁷: when precisely is content 'lawfully available' online for TDM purposes? What constitutes 'an appropriate reservation by machine-readable means', particularly in the absence of harmonised technical standards? May derivative rightholders (such as press publishers) exercise the opt-out independently of authors, or must consent align vertically within the rights chain? The absence of both a uniform technical standard and a centralised EU-level registry where opt-outs can be consulted compounds these uncertainties. The European Commission is actively addressing these gaps: the November 2025 launch of the Copyright Knowledge Centre at the EUIPO, tasked to 'explore solutions for copyright to keep powering creativity and innovation in Europe', signals institutional recognition of the problem.¹²⁸ Given the Centre's existing management of the Orphan Works Database and the Out-of-Commerce Works Portal, and the Commission's commissioning in January 2025 of a feasibility study on a centralised TDM opt-out registry, a unified, interoperable opt-out infrastructure may eventually emerge as a practical complement to the CJEU's doctrinal clarification in *Like Company*.¹²⁹

A lot has been written on the lack of ambition of the CDSM's TDM provisions, and this is not here the place where to discuss them at length.¹³⁰ These exceptions faced criticism for insufficiently taking into account the importance of TDM in many sectors, in particular with regard to the development of GenAI activities in the EU.¹³¹ The *Like Company* dispute, while not explicitly focusing on whether an opt-out was exercised, nonetheless seems to expose the structural inadequacy of Article 4(3) CDSM: absent a centralised registry and uniform technical standard, GenAI developers cannot reliably identify which works have been reserved, creating legal uncertainty detrimental to both innovation and enforcement. It seems obvious that the opt-out will be used as a new bargaining power to license the use of existing works for training purposes,¹³² which can potentially inhibit the development of this technology and thus make Europe unattractive for GenAI developers, as in particular start-ups or individual innovators will have great difficulties to manoeuvre in a fully licensed environment, not having the financial means to deal with the high transaction costs that such a situation might generate. At the other end, the individual author whose works are used in the training process will not necessarily benefit directly from this situation as it will likely be the rightsholders that will license the uses, with the authors having to (re)negotiate successfully with their producers to get additional remunerations.¹³³ The *Like Company* referral thus arrives at a critical juncture where proposals for a 'permitted-but-paid' future legislative reform have been put forward by scholars in the debate, proposals which deserve close attention when reflecting on the possible revision of the AI Act and the CDSM.¹³⁴

6 | OUTPUT AND COMMUNICATION TO THE PUBLIC

Question 1 concerns whether the display, by a GenAI chatbot such as Gemini, of text that partially reproduces press publisher web pages constitutes a communication to the public within the meaning of Article 15(1) CDSM and Article 3(2) InfoSoc. In essence, the Hungarian court asks whether AI-generated summaries or responses that incorporate verbatim segments or close paraphrases of online press articles amount to a new act of making those publications available to an indeterminate public, thereby falling within the PPR's scope. This section examines that question in light of the CJEU's existing communication to the public case law, the specific architecture of RAG systems such as Gemini, and the internal balancing mechanism built into Article 15(1) CDSM through the exclusion of 'individual words or very short extracts'.

6.1 | Summaries, paraphrase, and when an output becomes 'reproduction in part'

Having established that the *training* phase, that is, the offline internalisation of statistical patterns into model parameters, cannot be characterised as reproduction under Article 2 InfoSoc, it is crucial to distinguish this from the separate issue of *inference-time* generation. The legal questions raised by chatbot summaries, such as that in *Like Company*, concern an output-layer (often RAG-mediated) phenomenon: summaries and paraphrases arise at inference time, when the system responds to a user prompt and may retrieve and display (or restate) material from external sources. While this section addresses whether such displays constitute a 'communication to the public', we must first resolve the antecedent question of reproduction at the output stage: does a summary or paraphrase generated by GenAI at the moment of inference constitute a 'reproduction in part' of the original article?

The answer generally leans towards the negative. A summary expressed in fresh language should not be treated as a *reproduction* of the press article. As mentioned, EU copyright (and the PPR under Article 15 CDSM) protect *expression*, that is, the author's original choices in wording, structure, and presentation, not the underlying facts, themes, or 'information' conveyed by a report. A chatbot output that restates an article's informational content in different words will therefore ordinarily fall on the 'ideas/facts' side of the boundary, even if it is accurate and even if it is prompted by a user who has access to the original publication.¹³⁵ The difficult cases arise where the output

crosses from *paraphrase* into *recognisable copying*, for example, by reproducing distinctive turns of phrase, sentences, or other concrete expressive selections that embody the author's own intellectual creation.¹³⁶

That 'hinge' is increasingly well-articulated in the CJEU's jurisprudence. As mentioned above, in *Infopaq* the Court made clear that reproduction 'in part' turns on whether the material taken contains elements that are the expression of the author's intellectual creation¹³⁷; and most recently, in *Mio/konektra*, the CJEU sharpened the infringement inquiry for works of applied art: infringement requires that the original *creative elements* of the protected work are reproduced 'in a recognisable manner' in the later subject matter, and the design law criterion of 'overall impression' must not be imported into copyright analysis.¹³⁸ Although those cases arise outside journalism, the limiting principle travels well: what matters is not that an output 'resembles' the article in topic or tone, but whether protected expressive choices survive *recognisably* in what is shown. The same logic underpins *Pelham*: where a fragment is altered so that it is unrecognisable (there, 'to the ear'), it falls outside reproduction of the protected subject matter.¹³⁹ In the summary setting, this helps explain why a probabilistic paraphrase will ordinarily not count as reproduction: it yields a new, distinct work and, in ordinary cases, does not undermine the publisher's ability to obtain a satisfactory return on its investment.¹⁴⁰ The doctrinal point is modest but important: transformation is not a free-standing defence, yet recognisability functions as a practical boundary-marker between protected expression and non-protected abstraction.¹⁴¹

Several features sometimes emphasised in the present factual matrix may be relevant. The referring court in *Like Company* describes scenarios where different responses may display different parts of a press publication (including summaries), and notes the risk of inaccurate or 'hallucinated' content.¹⁴² It is also said that (i) the chatbot output points to, or is assessable only by consulting, the original site; (ii) the full article remains available at the source; and that (iii) the output is generated only after a user prompt that already signals awareness of the publication.¹⁴³ Those considerations may bear on *market substitution* and on the practical function of the output (e.g., whether it operates as a navigational aid rather than a substitute): also, they help characterise the output as a probabilistic *answer* that often pushes users back to sources: hence why disputes of this kind typically turn on Article 15 CDSM's calibrated interface for press content online (and, where applicable, quotation), rather than on treating summarisation as if it were the showing of the work itself.¹⁴⁴ As will be detailed shortly, Article 15 CDSM applies to extracts that are reproduced verbatim and do not qualify as 'very short'.

This distinction between abstract paraphrase and recognisable reproduction is where the 'may exclude' question from the referral becomes decisive. It is analytically cleaner to separate the inquiry into two stages. The first is the threshold question of reproduction under general copyright principles: has the output reproduced protected expression in a recognisable way? If the answer is negative, as with most generative summaries that maintain sufficient distance from the source text, no rights are infringed. However, if the answer is positive, that is, where the system outputs verbatim or near-verbatim blocks of text, or strings together distinctive phrases so that protected expression is reproduced recognisably, the second step must be to apply the specific filter of Article 15 CDSM.

To treat such displays solely under Article 3(2) InfoSoc would ignore the legislature's intent to create a dedicated regime for press aggregators. Therefore, if the displayed content qualifies as a 'very short extract' under Article 15(1), this specific limitation should effectively displace or 'exclude' liability under the general communication right, ensuring the fair balance between rights holders and technology providers intended by the CDSM is preserved. The precise boundary of this liability, specifically regarding the 'very short extracts' exception, is the subject of the next subsection, while the broader question of whether such displays constitute an act of dissemination is addressed later.

6.2 | Exception for short extracts and fair balance

The specific balancing mechanism built into Article 15(1) CDSM, that is, the exclusion of 'individual words or very short extracts', sets the general thresholds for copyright protection within the concept of 'reproduction in part'

under Article 2 InfoSoc. More specifically, the referring court seeks clarification on whether AI-generated summaries containing multiple verbatim phrases constitute 'very short extracts'; and how to assess extract length when content appears across multiple chatbot interactions. As mentioned earlier, in the specific case before the Hungarian court, Gemini's summary incorporated elements from Like Company's article about Kozsó's dolphin project, but there is no evidence of direct quotations; there were also references to related articles on the singer's criminal history and health issues.¹⁴⁵ The national court noted this might exceed typical snippet lengths allowed under national implementations of Article 15(1) CDSM.

It goes without saying that the PPR does not extend to mere facts or data reported in press publications since this is a basic premise of copyright law. Also, it does not apply to individual words or very short extracts of a press publication, as this has also been set out in the CJEU's case law with regard to copyright works.¹⁴⁶ In particular, in *Infopaq*¹⁴⁷ it was famously held that words as such are not protected since they do not constitute elements covered by protection, nor are they the intellectual creation of the author who employs them. It is only through the choice, sequence, and combination of these words that the author can express their creativity in an original manner and achieve a result - an intellectual creation. Isolated sentences, or even certain parts of sentences in a text, are protected if they are suitable for conveying to the reader the originality of a publication, such as a newspaper article, by communicating to that reader an element which is, in itself, the expression of the intellectual creation of the author of that article. It is only such sentences or parts of sentences that may therefore be protected.¹⁴⁸ It would be inconsistent for the free use of those extracts to be allowed under copyright law and barred via the PPR. That would render the related right broader than copyright and would unjustifiably block the permitted use of unprotected works.¹⁴⁹ This train of thought can also be derived from the fact that there is an express provision in the CDSM clarifying that journalists are to receive an appropriate share of the revenues for the use of their works: this is Article 15(5) CDSM. It would make no sense for journalists to be remunerated for uses for which they are not entitled, that is, if the extract used was so minimal that it could not actually qualify for copyright protection.¹⁵⁰ The *Like Company* referral implicitly raises important policy questions about how to reconcile the PPR with the EU's objectives for GenAI development within the digital single market. The CJEU will need to consider whether existing exceptions sufficiently accommodate machine learning workflows, and how to apply the CDSMs 'fair balance' principle between rights holders and technology providers. Clarifying the threshold for 'short extracts' will directly affect the permissible level of detail in output tokens generated by transformer-based systems.

Based on our analysis, Gemini-style summaries are structurally aligned with Article 15(1) CDSM's 'very short extracts' threshold because, even when they incorporate brief quotations, they primarily re-express factual content in independent language and compressed form rather than re-presenting the press publication's protected expression. They function as dynamic, query-specific snippets, comparable in length and purpose to search result previews that Recital 58 CDSM and national implementations treat as permissible, so long as they do not substitute for reading the original article.¹⁵¹ By providing high-level overviews, such extracts preserve incentives for users to click through to the publisher's site while facilitating information discovery and AI-assisted research, thus maintaining the CDSM's intended fair balance between press publishers' monetisation interests and broad public access to news.

6.3 | Act of communication in the context of GenAI systems and the notion of 'public'

The CJEU has consistently understood 'communication' to cover any transmission or making available that enables the public to access a work from a place and at a time individually chosen by them, provided there is a deliberate intervention by the person who organises that access. The focus falls on an intentional act which makes protected subject matter perceptible to others and which goes beyond mere provision of physical facilities. This is equally true for Article 3(2) InfoSoc and the PPR in Article 15 CDSM, which builds on that provision.¹⁵²

Two aspects are central: Whether protected content is actually communicated, that is., whether the service makes the original press publication, or a substantial part of it, perceptible to users; and whether the communication

is directed at a 'public', an indeterminate but fairly large number of potential recipients, rather than an isolated individual or closed, predetermined circle. In the *Like Company* scenario, the chatbot's core contribution is to generate a probabilistic summary in response to a user's prompt. Technically and legally, that summary is a new text. It draws on information that may be retrieved at inference time (via RAG or grounding mechanisms) but is not a reproduction of the original article as such. As explained above in this article, the model does not store press publications in a human-readable format; it operates on learned statistical regularities and, in the RAG pipeline, on transiently retrieved passages that are processed into the context window. In *GEMA*,¹⁵³ by contrast, the Munich court proceeded on the factual finding that ChatGPT had generated output containing substantial segments of specific song lyrics; this made it significantly easier for the court to characterise the output as a direct communication of the phonogram or musical work, as the user was confronted with text that effectively was the original song, rather than a mere synopsis of its content. That factual matrix diverges sharply from the summarisation of short news pieces that is at issue in *Like Company*.

Where the chatbot produces a condensed narrative in its own words, we are dealing with an act of communication of a newly generated work, not a communication of the original press publication. The mere fact that the summary concerns the same events as the article does not transform that summary into a communication of the original protected expression. Facts and data remain unprotected.¹⁵⁴ What matters is whether the chatbot reproduces the author's original selection and arrangement of words in a manner that remains recognisable.

Even if one were to identify isolated verbatim passages in the output, these would typically be limited to single phrases or sentences, embedded in an otherwise reformulated narrative. Based on the CJEU's own logic in *Infopaq* and subsequent case law, only those parts of the output that convey the author's intellectual creation, through a recognisable choice and arrangement of words, could be treated as reproducing the work in part. In the typical chatbot summarisation context, that threshold is rarely met and in any event it affects only fragments, not the summary as a whole. In this context, *GEMA* serves as a cautionary example without displacing the general rule: it confirms that where a GenAI system does cross the line into substantial verbatim reproduction of protected text, courts will have little difficulty applying the communication and reproduction rights. But it also implicitly acknowledges that such cases are at the extreme end of the spectrum and do not define the norm for AI-assisted summarisation.

As regards the notion of the 'public', chatbot interactions such as those triggered by Gemini are ordinarily individual and on demand: a single user submits a prompt and receives a response generated specifically for that session. However, under the jurisprudence of the CJEU, the assessment of 'communication to the public' is not confined to the one-to-one technical structure of the transmission. The Court has consistently defined the 'public' as an indeterminate number of potential recipients and implies a fairly large number of persons, assessed cumulatively rather than individually.¹⁵⁵ It has therefore accepted that successive, individual acts of making available may be considered together in determining whether there is a communication to the public. On that approach, the fact that access occurs through discrete, user-initiated requests does not preclude a finding that the same protected subject matter is communicated to an indeterminate public where the service enables access by a potentially unlimited number of users.

Yet, the scenario highlighted in this article, that is, a GenAI platform which provides an interactive tool rather than broadcasting pre-selected content to a mass audience, is factually different. In other words, such scenario is conceptually distinct from streaming, broadcasting or the permanent provision of content to an undifferentiated public, exactly because the GenAI platform generates responses dynamically and transiently for each user, without making predefined works available or enabling simultaneous access to identical content by multiple users, thereby maintaining the inherently individual nature of each interaction. Accordingly, at the level of primary qualification, there are good reasons to conclude that the GenAI summarisation in question does not amount to an act of communication to the public of the underlying press publication. Rather, it is the communication of a newly generated text, which only incidentally overlaps with the original expression.

6.4 | The notion of 'new public'

Even if one assumed *arguendo* that some chatbot interactions could amount to an 'act of communication' within the meaning of Article 3 InfoSoc, the separate requirement that such communication be directed at a 'new public' would not be satisfied in the *Like Company* scenario. The 'new public' criterion, developed in *Svensson* and refined in *GS Media* and *VG Bild-Kunst*, requires that the communication reach a public that was not taken into account by the rightholder when the initial communication to the public was authorised.¹⁵⁶ Where a work is freely accessible on the internet without technical or contractual restrictions (e.g., paywalls or login barriers), the rightholder is presumed to have taken into account all internet users as the relevant public. Hyperlinks, including embedded links, that merely facilitate the finding of such content do not generally address a new public; they act as a 'pointer' to an existing communication.¹⁵⁷

In the Hungarian reference, the *Like Company* article about Kozsó was made available free of charge on a public website, without technological restrictions. The relevant public thus consists of all internet users who might access the site. Gemini's RAG-based operation, when prompted with the URL, essentially performs a targeted search and summarisation of content already accessible to that same public. It does not circumvent technological protection measures, nor does it unlock content that was previously limited to a restricted audience. Moreover, the behaviour of the user in this case is telling: by explicitly including the URL of the article in the prompt ('the online press publication that appeared on balatonkornyeke.hu'), the user demonstrates awareness of the original publication's existence and accessibility. The chatbot thus presupposes, instead of creating, the user's connection to the underlying content. This could be seen as the opposite of reaching a new public; the service facilitates interaction with content that the user has *already* identified.

To treat such an interaction as addressing a new public would be to conflate any mediation (e.g., search, indexing, recommendation and summarisation) with communication to a new public whenever the intermediary provides added value. That would extend the notion of 'new public' far beyond its justificatory basis and would risk capturing commonplace acts of online navigation and information retrieval, which the CJEU has so far carefully excluded from the scope of Article 3 InfoSoc, especially in cases where no circumvention of access restrictions is involved.¹⁵⁸ In *Getty Images v Stability AI* ('Getty'), the UK High Court did not accept that the mere provision of a general-purpose generative model constituted infringement per se, instead directing attention to whether infringing reproductions could be shown to subsist within the model or its deployment in the UK.¹⁵⁹ The focus remained on whether, in concrete instances, the output reached users as recognisable reproductions of the claimant's works. By contrast, the Munich Court's decision in *GEMA*¹⁶⁰ illustrates a situation where the 'public' element was easily satisfied because the allegedly infringing outputs, in particular song lyrics, were made available to ordinary users in a form that was indistinguishable from the original; yet even there, the court's reasoning turned on the specific expressive character and commercial importance of lyrics, not on any general hypothesis of AI-mediated access as such. That logic militates against any blanket characterisation of chatbot summarisation as communication to a 'new public'.

6.5 | Hyperlinking case law revisited for GenAI outputs

The hyperlinking jurisprudence of the CJEU offers an indispensable reference point even if chatbots do more than simply reproduce URLs. In *Svensson*, *BestWater*, *GS Media* and *VG Bild-Kunst*, the Court sought to balance authors' rights with the architecture of the web by drawing a line between (i) links that enable users to access lawfully available content without reaching a new public and (ii) links that circumvent restrictions or target clearly unlawful sources.¹⁶¹

Chatbots that like Gemini operate via RAG can be understood as combining two functions: Information retrieval, which is functionally equivalent to a search engine querying an index and returning snippets and links; and

generative summarisation, which synthesises retrieved information into a new text. From the perspective of communication to the public, the first component aligns closely with the hyperlink cases. Where a chatbot surfaces a link (or other in-text reference) to the original press publication, the situation is indistinguishable from that of a search engine results page displaying the article title, a short snippet and a clickable URL. In both instances, the user must still visit the original site to access the full content; the public reached remains the same. The second component, that is, summarisation, adds a layer of transformation. Yet, this added functionality does not, by itself, create a new public or a new communication of the underlying work. On the contrary, the summary is a distinct work generated by the system, typically of a transformative nature, which does not fully or substantially reproduce the original expression. As explained earlier in Section 3 of this article, the PPR cannot be interpreted in a manner that would render it broader in substance than the underlying copyright, particularly where the extracts used fall beneath the threshold of originality, are 'very short', or amount to unprotected facts and ideas.

Seen in this light, the hyperlink cases remain conceptually fit for purpose. The mere fact that the 'pointer' is embedded in a richer, AI-generated narrative does not change the basic structure: the chatbot, like a search engine, facilitates discovery and understanding of content already lawfully accessible on the web, and does so without expanding the circle of persons who can access the original article. The CJEU can therefore build on its existing case law, adapting its language to account for generative features but preserving the core insights about lawful access and lack of a new public. The approach taken by the UK High Court in *Getty* supports the analogy: rather than treating GenAI services as infringing *per se*, the court focused on whether the specific acts of reproduction and use could be established on the evidence in the relevant jurisdiction. And the *GEMA* judgment by the Munich Court shows what happens when the factual pattern changes: once the GenAI output ceases to be a pointer or paraphrase and instead delivers large quantities of recognisable expressive material (song lyrics) on demand, the mediation begins to look much closer to a classic act of making available. Crucially, however, nothing in *GEMA* suggests that summarising non-creative factual reporting, without extensive verbatim reproduction, should be treated in the same way.

6.6 | Parallels between chatbots and search engine results pages

Like Company presents, in fact, a paradigm instance of how chatbots and search engines converge. By explicitly asking Gemini to 'provide a summary in Hungarian of the online press publication that appeared on balatonkornyeke.hu', the user is functionally performing a highly specific search query. In a retrieval-based scenario, the system's RAG or grounding module would then retrieve the article (or relevant cached/indexed versions), analyse it, and provide a condensed account, frequently accompanied by an in-text reference or link to the source. From the user's perspective, this behaviour is close to what a modern search engine already does when answering 'rich' queries: it may show a snippet, a 'knowledge panel', or a short extract together with the URL. What is new is not the basic function, that is, mediating access to web content, but the sophistication of the natural-language interface and the degree of paraphrasing.

From the standpoint of EU copyright law, chatbots employing RAG therefore operate much like enhanced search engines. Search engines routinely display verbatim or paraphrased snippets as part of their indexing functions, and the CJEU has never held that such displays amount to acts of communication to the public, provided no substantial protected expression is reproduced. If one were to classify such summarisation as an independent communication to the public of the press publication, it would be difficult to avoid the conclusion that search engine snippets and many forms of enriched search results likewise constitute communication to the public. That reading would stand in tension with the careful calibration reflected in Article 15 CDSM, which as mentioned addresses search engines and aggregators by granting press publishers a related right but expressly carves out individual words and very short extracts, hyperlinks, and private/non-commercial uses. It would also risk undermining the CJEU's insistence that technological intermediation which does not expand the circle of recipients or circumvent access restrictions should not be treated as an autonomous communication.¹⁶²

Gemini's output in *Like Company*, insofar as it arises from real-time retrieval, is thus conceptually adjacent to a search engine snippet. Treating it as a communication to the public would produce an unprecedented expansion of Article 3 InfoSoc and would effectively override the CJEU's balancing in the search-engine context. The *GEMA* decision¹⁶³ confirms the limits of this analogy: the Munich Court treated long verbatim reproductions of song lyrics as communication to the public, but only because the outputs contained substantial expressive copying. The reasoning turned not on a rejection of the search engine analogy, but on the presence of direct expressive appropriation. As the output in *Like Company* is merely a summary of a 579-word article generated via probabilistic prediction, the *GEMA* logic does not apply. Read together, *GEMA* and *Getty* support a calibrated standard: while *GEMA* shows that GenAI services may incur liability where outputs closely replicate protected works, *Getty* illustrates that such services are not treated as infringing *per se* in the absence of evidence that the model or its outputs embody reproductions of the claimant's works. The *Like Company* facts fall well short of that threshold.

Accordingly, a principled handling of *Like Company* requires the Court to recognise the functional parallel between chatbot summarisation and conventional search engine results. Both operate as discovery and navigation tools; both rely on content that is freely accessible; both, at most, reproduce brief excerpts to assist users in deciding whether to consult the original source. The presence of an additional generative layer does not alter this basic communicative structure.

7 | PLATFORM'S LIABILITY FOR AI-GENERATED OUTPUTS PROMPTED BY USERS

This section addresses Question 4, that is, whether, when a user instructs an LLM-based chatbot like Gemini to reproduce or refer to the text of a press publication and the system responds by displaying part or all of that text, the service provider thereby commits an act of reproduction under Article 2 InfoSoc and Article 15(1) CDSM. On the best interpretation of the *acquis*, the answer should be no. Any rights-relevant reproduction in this configuration is attributable to the user who issues the prompt and causes the specific output; any potentially infringing passage appears only because the user deliberately requests it. The provider operates a general-purpose, content-agnostic, analytical tool whose functioning is automatic and neutral with respect to any given work. That allocation accords with (i) the CJEU's focus on who performs the protected act and on the difference between neutral technical intermediation and deliberate, content-specific intervention; (ii) the design of the PPR, which is aimed at uses by information society service providers while excluding private or non-commercial uses by individual users, and (iii) the intermediary-liability framework (i.e., e-Commerce Directive and the DSA), which predicates provider responsibility on knowledge or control over specific items and rejects general monitoring obligations.¹⁶⁴

7.1 | Who performs the 'reproduction'? attribution, volition and the nature of the tool

The scenario is user-initiated: a prompt instructs the system to reproduce, or closely paraphrase, a specific press passage. The model then generates a response probabilistically from that input, and in some deployments from documents retrieved at inference time. The attribution question is therefore who makes the copy. At the output layer, the proximate actor is the user: she selects the work, specifies the extent, and triggers the generation event, that is, the prompt. The provider, by contrast, supplies a general-purpose tool that tokenises input and predicts continuations; it does not exercise content-specific choice over the article or passage at issue. In short, volition and selection lie with the user; the provider's role is instrumental.

Multiple arguments support this conclusion. First, EU copyright, and the CJEU case law, do look at who performs the act in functional terms and whether the alleged infringer plays a deliberate role in bringing about the protected act. That logic tracks the CJEU's insistence, across rights, on a deliberate intervention by the person said

to perform the protected act, and its distinction between neutral, automatic intermediation and content-specific control. In communication to the public cases, liability turns on intentional intervention directed at a public and, in some lines, on knowledge of illegality.¹⁶⁵ In hosting and platform cases, the Court distinguishes a passive, neutral technical service, benefiting from safe harbour, from an operator that plays an active role giving it knowledge or control over specific unlawful content. In *Google France* which focused on AdWords and trademarks, the Court emphasised the provider's 'mere technical, automatic and passive' role; and in *L'Oréal/eBay* and in *YouTube/Cyando* it framed liability around knowledge and control over specific content and the presence, or absence, of an active role.¹⁶⁶

Second, in this context, the role of user agency and system unpredictability matters. LLM outputs are user-initiated and stochastic. The probabilistic nature of LLM outputs means providers cannot reasonably foresee every verbatim reproduction *ex ante*, especially with complex prompts. Even diligent providers cannot foresee, *ex ante*, every possible chain of tokens that a creative or adversarial prompt may elicit, especially where the prompt itself embeds protected text, a common pattern in extraction experiments. The computer-science evidence dovetails with this allocation. Empirical work demonstrates that verbatim training-data emissions, or regurgitation, are uncommon, tend to arise under attack-style prompting and high-duplication training examples, and are substantially reduced through data set deduplication, privacy-oriented training, output filters, and retrieval policies that favour linking/summarisation over long verbatim spans. Carlini et al. show verbatim extraction is possible but conditional¹⁶⁷; Lee et al. document an order-of-magnitude reduction in verbatim emissions after deduplication¹⁶⁸; later studies quantify how model scale, duplicate frequency and prompt length interact with memorisation.¹⁶⁹ These findings support a targeted, notice-and-mitigate governance model over strict provider-side primary liability for user-prompted reproductions. These are precisely the guardrails a negligence-based standard should reward.

Third, that the provider's servers instantiate the bits does not decide attribution: the legally relevant act is user-directed and user-specific. In this regard, two familiar objections fall away: that the provider is the reproducer merely because its infrastructure instantiates the output; and that its design and tuning of the model make it the reproducer of any protected text the system may emit. First, a 'server-location' theory, that is, whoever's hardware hosts the output is the reproducer, has no footing in the case law; attribution turns on deliberate, content-specific intervention, not on where the bits reside.¹⁷⁰ Second, model design and tuning do not make the designer the reproducer of any particular text: in *Infopaq*, the reproducer was the party that organised and ran the copying process for its own ends; in *PRCA v NLA*, transient technical copies integral to a process were held exempt to preserve technological development.¹⁷¹ Where the user specifies the subject and extent of any reproduction and the provider offers a general-purpose tool with reasonable guardrails, recasting the provider as the primary 'reproducer' asks Article 2 InfoSoc to do work it was not designed to do. The reproduction right protects fixation of expression; it is not a general control over analytical capacity. Stretching it would collapse the *acquis'* neutral-tool/content-actor distinction and, in practice, invite *ex ante* screening of prompts and outputs, that is, general monitoring that EU law rejects on proportionality grounds.¹⁷² The result would be to convert copyright into near-strict liability for learning and generation system design, a function the EU legal order assigns instead to horizontal due-diligence regimes such as the DSA and the AI Act and to targeted, output-layer remedies.

Fourth, assigning the act to the user preserves the functioning of exceptions and limitations where appropriate. Article 5(3)(d) InfoSoc enables quotation for purposes such as criticism or review, subject to fair practice and acknowledgment of source; whether a particular reproduction qualifies depends on context and purpose, which are facts known to the user who issued the prompt.¹⁷³ By contrast, a regime that imposed strict primary liability on providers would encourage blunt *ex ante* filtering that cannot reliably distinguish quotation from verbatim copying for other purposes, threatening the fair balance the CJEU polices across exceptions.¹⁷⁴

Finally, from a law-and-economics perspective, placing primary liability on the prompting user minimises both error costs (over-blocking lawful uses such as quotation) and administrative costs (monitoring and assessing billions of prompts). In that environment, the least-cost avoider¹⁷⁵ for one-off, prompt-engineered reproductions is typically the user who chooses the target sequence and drafts the instruction, not the provider, absent evidence of

design choices that encourage regurgitation. The user knows whether her prompt targets a protected text and whether an exception, for example, quotation, criticism/review or parody, is implicated. Providers, as repeat players, bear a duty of care to reduce foreseeable, systematic regurgitation but are not taxed for idiosyncratic, one-off prompts that only a general monitoring obligation could catch. Assigning strict primary liability to providers would effectively require generalised anticipatory filtering, which the *acquis* disfavors.¹⁷⁶ It would also amplify anti-commons effects.¹⁷⁷ While both users and providers face anti-commons risks, a provider-primary liability rule creates a 'Schrödinger's-cat' uncertainty: before any prompt, at the model's design and tuning stage, each prospective output sits, *from the provider's vantage point*, in a superposition of lawful and unlawful, collapsing only when the user's purpose materialises at the point of use. This is because the provider is not itself using the works at that stage; it builds a neutral, general-purpose instrument that merely enables innumerable potential uses of the corpus of human knowledge (some lawful, some unlawful) by users, so the legal character of any given output turns on context the provider cannot know *ex ante*. Yet the provider, a neutral toolmaker rather than a user of the works, would be forced to behave as if collapse had already occurred, pre-clearing or blocking *ex ante*. Lacking the user's context, it cannot rely on purpose-dependent exceptions (e.g., quotation), so any arguable overlap must be treated as potentially infringing, multiplying veto points and inflating transaction costs.

7.2 | Intermediary liability under the E-Commerce directive and DSA: Knowledge, neutrality and no general monitoring

Under the EU intermediary liability regime, chatbot service providers' position is analogous to that of a hosting service protected by the E-Commerce Directive and the DSA. As noted earlier, the E-Commerce Directive establishes hosting safe harbours for information provided by recipients and prohibits general monitoring.¹⁷⁸ The DSA preserves the no-general-monitoring principle and adds structured *ex post* accountability, i.e. notice-and-action, statements of reasons, and for Very Large Online Platforms (VLOPS) and Very Large Online Search Engines (VLOSEs) systemic risk assessments, and audits.¹⁷⁹

The CJEU has consistently rejected general and permanent filtering obligations as incompatible with fundamental-rights balancing¹⁸⁰; however, it has accepted specific and targeted injunctions (including 'equivalent content' in limited conditions),¹⁸¹ and has framed platform responsibility around knowledge, control and active role, not the mere availability of a general-purpose service.¹⁸² The Court's platform jurisprudence across rights, from *Google France* through *L'Oréal/eBay* to *YouTube/Cyando*, sets out a fault-sensitive allocation of responsibility. Hosting providers are not directly liable merely because illegal content passes through their systems; they risk primary liability only when they actively select or promote unlawful content, or when, having actual or constructive knowledge of specific illegality, they fail to act expeditiously.¹⁸³ Nothing in that case law suggests that a provider becomes the maker of the copy simply because its infrastructure executes the user's instruction.

For outputs generated by chatbots like Gemini, that translates to: (i) no strict *ex ante* policing of all possible verbatim emissions, that is, general monitoring is barred; (ii) expeditious action on notice concerning concrete passages; and (iii) proportionate, system-level safeguards under the DSA without collapsing into strict primary liability for user-prompted reproductions. This reading aligns with the CJEU's reasoning in *YouTube/Cyando*, which shields diligent platforms whose business model is not predicated on copyright infringement.¹⁸⁴ The Court draws a sharp line between a neutral host and an operator that, through knowledge/control or a model that incentivises infringement, plays an active role.¹⁸⁵ Where a provider implements proportionate guardrails and preventive tools and operates a reliable notice-and-action mechanism, the hosting exemption remains available and the provider is not treated as performing a protected act on users' behalf. Transposed to LLM chatbots, a service that deploys reasonable output safeguards and acts expeditiously on specific notices meets the E-Commerce Directive and DSA standard; any rights-relevant reproduction triggered by a prompt is attributable to the user, not the toolmaker. However, while the general-purpose model remains a neutral tool, liability could re-emerge if a provider specifically

fine-tunes or curates the model to 'regurgitate' protected works, thereby playing an essential role in the communication of that content.¹⁸⁶

7.3 | Interplay with the PPR and exceptions

The PPR under Article 15 CDSM is specifically designed to target online uses by information society service providers, while expressly exempting 'acts by individual users' and the use of 'individual words or very short extracts'.¹⁸⁷ Reading these statutory definitions alongside the CJEU's platform jurisprudence just mentioned implies a clear division of labour: service-provider liability attaches only where the provider *itself* performs the act of communication, for example, by actively curating or pushing protected content to the public, and not where an individual user elicits a specific extract ad hoc from a neutral, general-purpose tool.¹⁸⁸

To hold otherwise would be to nullify the 'individual user' exemption by re-characterising user-initiated acts as provider acts. This connects directly to the referring court's query regarding whether Article 3(2) InfoSoc 'may exclude' the application of Article 15 CDSM. This framing is decisive: it highlights the risk that a broad, unanchored application of the general making-available right could bypass the specific legislative compromises of Article 15. If the dispute is not channeled through the *lex specialis* of Article 15, with its tailored thresholds for 'very short extracts' and user exemptions, claimants could effectively ignore the limits the EU legislature imposed on the related right.

A contrary reading would invert the logical hierarchy of the *acquis*. It would risk treating the related right (Article 15 CDSM) as a strict-liability regime that is functionally broader and more burdensome than the author's core rights (Articles 2 and 3 InfoSoc) upon which it is built. By design, Article 15 is framed as a complementary related right, not a mechanism to exceed the protection afforded to the underlying author's right.¹⁸⁹ If Article 15 CDSM were interpreted to impose strict primary liability on a provider for every user-elicited output, it would create a regime functionally stricter than the author's core rights under Articles 2 and 3 InfoSoc. As detailed earlier, under established case law, a neutral platform is not liable for user-posted content absent knowledge or an active role. If Article 15 were to bypass this neutrality defense and hold providers automatically liable for user prompts, it would grant the derivative rightholder (the publisher) a 'super-right' more powerful than that of the original creator (the author), thereby inverting the system's intended balance.

Therefore, the question of whether a chatbot provider 'reproduces' press content must be answered in a way that preserves the integrity of Article 15 CDSM. A fault-conditioned, neutrality-preserving approach to provider responsibility helps maintain that architecture. It preserves meaningful remedies for press publishers where a provider plays an active role or fails to act after concrete notice, without letting the mere provision of a general-purpose conversational tool turn Article 15 into a route to *de facto* strict primary liability.

7.4 | Liability allocation, remedies and workable standards

On a principled view, four allocations of liability are conceivable. First, a provider-primary liability model would treat any regurgitation as the provider's reproduction; it sits uneasily with hosting safe harbours and the ban on general monitoring, and invites heavy *ex ante* filtering with collateral damage to quotation, criticism and research. Second, a user-primary model with provider secondary liability, or loss of safe harbour, upon fault aligns with the *acquis*: the user who elicits the extract performs the reproduction, while the provider is engaged only where it knows of a concrete infringement, for example, via notice or evidence of systematic regurgitation, and fails to act expeditiously, or plays an active role; this calibrates incentives (users internalise the cost of unlawful prompts, providers invest in guardrails) without drifting into strict liability. Third, joint primary liability risks a back-door strict-liability regime: if 'providing the model' is enough for primary liability whenever a user elicits a protected passage, compliance will

tend toward generalised, prior restraint. That is precisely what *Scarlet Extended* and *SABAM/Netlog* caution against.¹⁹⁰ Fourth, no primary liability; injunctions only would under-protect rightholders where the user's intent to reproduce is clear. Of these, the second model best fits *YouTube/Cyando*, the E-Commerce/DSA settlement, and the EU Charter's proportionality.

Some may worry that placing primary liability on the user under-protects rightholders. But the horizontal framework already supplies effective remedies without general monitoring. This is not a plea for laissez-faire but for calibration: provider duties should track the DSA/E-Commerce model, that is, notice-and-action, narrowly proportionate *ex-ante* checks for long verbatim passages, rate-limits and guardrails on prompt patterns associated with regurgitation, transparent redress, and where necessary targeted injunctions against specific extraction behaviours.¹⁹¹

Nothing in this allocation deprives right-holders of remedies. They retain: (i) primary claims against users who elicit or disseminate infringing excerpts; (ii) DSA notice-and-action tools to trigger expeditious removal and algorithmic adjustments, for VLOPs and VLOSEs¹⁹²; (iii) injunctive relief against intermediaries whose services are used by a third party to infringe, crafted with *UPC Telekabel* proportionality to avoid general monitoring¹⁹³; and (iv) where a provider's design choices systematically produce regurgitation and it fails to adopt obvious mitigations, loss of safe harbour and damages aligned with *YouTube/Cyando*.¹⁹⁴

These levers remove unlawful items and prevent recurrences without turning providers into strict-liability guarantors for every private prompt. That is precisely the compliance pattern endorsed in *YouTube/Cyando* (diligent platforms are shielded; actors whose business models hinge on infringement are not) and it preserves space for innovation and for lawful uses such as quotation and very short extracts.¹⁹⁵

8 | CONCLUSION

This study has examined the *Like Company* reference to the CJEU to elucidate how Articles 2 and 3 InfoSoc, read alongside Article 15 CDSM, apply to GenAI systems that summarise press content. Beyond the specific facts of the Hungarian dispute, the analysis tries to clarify fundamental doctrinal questions with EU-wide implications and draw conclusions: training cannot, as such, constitute reproduction under Article 2 InfoSoc as it fixes, in ordinary operation, parameters (weights) not protected expressions; anyway, if it is accepted that technical copies arise in the training pipeline, they are covered by the Article 4 CDSM TDM exception (absent a valid opt-out reservation); RAG-generated summaries that remain within 'very short extracts' fall outside Article 15(1) CDSM; and outputs that paraphrase rather than reproduce do not communicate the original publication to a new public. On platform liability, the study concludes that rights-relevant reproduction should be attributed to the user formulating the prompt, rather than the neutral service provider, with liability conditioned on knowledge of specific infringements. Applied to *Like Company*, this means Gemini's training process, operating under the CDSM's TDM framework, does not require publishers' authorisation, and its summarisation outputs do not infringe the PPR.

Critically, the interpretation proposed here preserves the structural relationship between Articles 2 and 3 InfoSoc and Article 15 CDSM: the PPR remains a targeted, complementary related right that does not extend protection beyond the underlying authors rights, but applies those rights in a defined online context, including *vis-à-vis* GenAI services. As the referring court captures in its language of whether Article 3(2) InfoSoc 'may exclude' the application of Article 15 CDSM,¹⁹⁶ the two regimes must be understood hierarchically in this context. Article 15 operates as the specific, governing interface for the use of press publications by information society service providers. While it borrows the doctrinal definitions of Articles 2 and 3 InfoSoc, its specific limitations, particularly regarding 'very short extracts', must take functional precedence. This ensures that the two regimes do not simply cumulate to trap lawful uses, but rather articulate a coherent system where the specific choices of the EU legislature regarding press content are respected.

However, this framework also reveals a critical gap: while the TDM exception enables lawful training and the InfoSoc-CDSM interaction prevents overreach, neither clearly and fully ensures equitable remuneration for individual creators whose works fuel GenAI systems. Article 15 guards publishers' investments, yet authors may receive no direct benefit from licensing revenues that accrue to corporate rightholders. This study, therefore, as mentioned in the Introduction, acknowledges that sustainable copyright governance must couple rigorous doctrinal analysis with a binding remuneration obligation that channels value to individual authors and creators whose intellectual contributions power GenAI systems, ensuring the creative workforce remains adequately compensated and incentivised across the EU.

ACKNOWLEDGMENTS

The authors are grateful to Alexander Cleveewood NG for editing assistance. This article results from a research project on Generative AI and copyright law supported by Google LLC. The project was conducted in full academic independence, and all views and recommendations herein reflect the authors' personal academic judgment, for which they remain solely responsible. This study was prepared exclusively on the basis of publicly available materials. The authors had no access to non-public evidence (including internal logs, prompt transcripts, model weights, training data sets, indexing/caching configurations, or confidential technical documentation). Any description of the system's internal technical pathway is therefore necessarily inferential and is framed as such.

DATA AVAILABILITY STATEMENT

Not applicable.

ORCID

Andres Guadamuz  <https://orcid.org/0000-0001-6092-9143>

ENDNOTES

- ¹ *Like Company v Google Ireland Limited*, C-250/25, request for a preliminary ruling from Budapest Környéki Törvényszék (Hungary), 3 April 2025, <https://curia.europa.eu/juris/liste.jsf?num=C-250/25&language=en> (hereinafter "*Like Company* referral"). For another analysis of the case, see European Copyright Society, Opinion on Copyright and Generative AI, February 2025 <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6333539>
- ² Directive (EU) 2019/790 of the European Parliament and of the Council of 17 April 2019 on copyright and related rights in the Digital Single Market and amending Directives 96/9/EC and 2001/29/EC, OJ 2019 L 130.
- ³ Directive 2001/29/EC of the European Parliament and of the Council of 22 May 2001 on the harmonisation of certain aspects of copyright and related rights in the information society, OJ 2001 L 167.
- ⁴ See Guadamuz, A., 'First Case on AI and Copyright Referred to the CJEU', (Technollama Blog, 27 May 2025) <<https://www.technollama.co.uk/first-case-on-ai-and-copyright-referred-to-the-cjeu>>
- ⁵ *Like Company* referral, para 7.
- ⁶ *Ibid*, para 8.
- ⁷ *Ibid*. The publicly available record does not disclose the precise technical pathway by which Gemini generated the contested output. In particular, it cannot be inferred from the referral alone whether the response was produced by (i) inference-time retrieval (RAG/grounding) of the claimant's page (or cached/indexed equivalents), (ii) user-supplied text (e.g., copy-pasting all or part of the article into the prompt and requesting a summary), or (iii) other combinations (including reliance on third-party sources that reproduced the article). Where our legal analysis depends on one of these pathways, we state that dependency explicitly; otherwise, the analysis proceeds on propositions that hold across scenarios.
- ⁸ *Like Company* referral, para 9.
- ⁹ *Ibid*, para 9.
- ¹⁰ *Ibid*, paras 19–20 and 23.
- ¹¹ *Ibid*, para 21.

- ¹² Ibid, para 26.
- ¹³ See, for example, Russell, S.J. and Norvig, P., *Artificial Intelligence: A Modern Approach* (4th edn., Pearson 2020), 45.
- ¹⁴ See, for example, for definitions used in a policy perspective, Bas, G. et al., 'The EU AI Act: A Pioneering Effort to Regulate Frontier AI?', 27 IA 55, 2024.
- ¹⁵ Comanici, G. and others, 'Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities', (arXiv 7 July 2025) <<https://doi.org/10.48550/arXiv.2507.06261>> The paper discusses that while memorisation of training data is possible, models do not act as static databases where any input has been stored.
- ¹⁶ Georgiev, P. and others, 'Gemini 1.5: Unlocking Multimodal Understanding across Millions of Tokens of Context', (arXiv 8 March 2024) <<https://doi.org/10.48550/arXiv.2403.05530>>
- ¹⁷ See *Like Company* referral, para 5. See also Rohan, J. and others, 'Gemini: A Family of Highly Capable Multimodal Models' (arXiv 13 December 2023) <<https://doi.org/10.48550/arXiv.2312.11805>>
- ¹⁸ *Like Company* referral, para 5.
- ¹⁹ See, for example, Frankle, J., Schwab D.J., and Morcos A.S., 'The Early Phase of Neural Network Training' (arXiv, 24 February 2020) <<http://arxiv.org/abs/2002.10365>>
- ²⁰ Ibid.
- ²¹ See, for example, Choo, S. and Kim, W., 'A Study on the Evaluation of Tokenizer Performance in Natural Language Processing', 37 AAI (2023).
- ²² See Frankle, Schwab and Morcos (n 19).
- ²³ See, for example, Hendrycks, D., Lee, K., and Mazeika, M., 'Using Pre-Training Can Improve Model Robustness and Uncertainty' (arXiv, 20 October 2019) <<http://arxiv.org/abs/1901.09960>>
- ²⁴ See, on the process of training, Goodfellow, I., Bengio, Y., and Courville, A., *Deep Learning* (MIT Press 2016), Chapter 6.
- ²⁵ See, on the process of extracting meaning, Vaswani, A. and others, 'Attention is All You Need', in Guyon, I. and others (eds), *Advances in Neural Information Processing Systems*, vol. 30 (2017) 5998–6008.
- ²⁶ See Comanici (n 15). See also Kudo, T., and Richardson, T., 'SentencePiece: A Simple and Language Independent Subword Tokenizer and Detokenizer for Neural Text Processing' (arXiv, 19 August 2018) <<https://doi.org/10.48550/arXiv.1808.06226>> (noting that LLMs do not store recognisable words; and that they are broken up into tokens, not documents from the training data).
- ²⁷ See, for example, Kaplan, J. and others, 'Scaling Laws for Neural Language Models' (arXiv, 23 January 2020), ss 1.1–1.3 <<https://arxiv.org/abs/2001.08361>>
- ²⁸ For a list of different model knowledge cut-off dates, see: <<https://github.com/HaoooWang/llm-knowledge-cutoff-dates?tab=readme-ov-file>>
- ²⁹ See, for example, Carlini, N. and others, 'Extracting Training Data from Large Language Models', in *Proceedings of the 30th USENIX Security Symposium*, USENIX Association (2021), pp. 2633–2650.
- ³⁰ See, for example, Ying, X., 'An Overview of Overfitting and Its Solutions', (2019) 1168 JP: Conference Series 022022.
- ³¹ See, for example, Carlini, N. and others, 'Quantifying Memorization across Neural Language Models', in *Proceedings of the 11th International Conference on Learning Representations (ICLR)* (2023), pp. 7–9 (on effects of model scale, duplication and prompt/context length; and deduplication as mitigation); Morris, J.X. and others, 'How Much Do Language Models Memorize?' (arXiv, 18 June 2025) <<http://arxiv.org/abs/2505.24832>>; See also Sag, M., 'Copyright Safety for Generative AI', (2023) 61 Houston LR 295, 310, 326–336, 338 (discussing de-duplication from a copyright safety perspective and noting that de-duplication makes "memorisation of the training data more likely"); Cooper, A.F. and others, 'Extracting Memorized Pieces of (Copyrighted) Books from Open-Weight Language Models' (arXiv, 30 November 2025), <<http://arxiv.org/abs/2505.12546>>
- ³² See, for example, Xia, W. and others, 'A Comprehensive Study of the Past, Present, and Future of Data Deduplication', (2016) 104 *Proceedings of the IEEE* 1681; Lee, K., and others, 'Deduplicating Training Data Makes Language Models Better', in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics (2022), 8429–8431 (noting that data set-level deduplication reduces verbatim regurgitation by up to an order of magnitude while providing quality gains).
- ³³ See, for example, Baptista, R. and others, 'Memorization and Regularization in Generative Diffusion Models' (arXiv, 18 March 2025), <<http://arxiv.org/abs/2501.15785>>

- ³⁴ See, for example, Zhang, T. and others, 'Balancing Learning Model Privacy, Fairness, and Accuracy with Early Stopping Criteria', (2023) 34 IEEE TNNLS 5557-5569.
- ³⁵ See, for example, Chaudhari, S., and others, 'RLHF Deciphered: A Critical Analysis of Reinforcement Learning from Human Feedback for LLMs', (2025) 58 ACM CS. See also Church, K.W., Chen, Z. and Ma, Y., 'Emerging Trends: A Gentle Introduction to Fine-Tuning', (2021) 27 NLE 763-778 (describing fine-tuning as a way of optimising a model with specific information after training).
- ³⁶ Inference is the process by which a trained model applies its learned parameters to new data to produce an output. See Cui, P. and Athey, S., 'Stable Learning Establishes Some Common Ground between Causal Inference and Machine Learning', (2022) 4 NMI 110-115.
- ³⁷ See, for example, Gupta, S., Ranjan, R., and Singh, S.N., 'A Comprehensive Survey of Retrieval-Augmented Generation (RAG): Evolution, Current Landscape and Future Directions' (arXiv, 3 October 2024) <<http://arxiv.org/abs/2410.12837>>
- ³⁸ See, for example, Paulsen, N., 'Context Is What You Need: The Maximum Effective Context Window for Real World Limits of LLMs' (arXiv, 21 September 2025) <<http://arxiv.org/abs/2509.21361>>
- ³⁹ See, for example, Lewis, P. and others, 'Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks', in Hugo Larochelle and others (eds), *Advances in Neural Information Processing Systems* 33 (Curran Associates 2020), 9459-9474 (explaining that in RAG mode the model looks up relevant documents when answering and then writes using both its learned knowledge and those retrieved texts, so the response stays grounded in external sources).
- ⁴⁰ See Google, 'Grounding with Google Search' (2025) <<https://ai.google.dev/gemini-api/docs/google-search>> It should be noted that public documentation describes these mechanisms at a general level and across model iterations; the publicly available record in *Like Company* does not specify the exact model/version and retrieval configuration deployed at the relevant time.
- ⁴¹ Google (n 40).
- ⁴² Ibid.
- ⁴³ Ibid.
- ⁴⁴ See *Like Company* referral, paras 17-18.
- ⁴⁵ See (n 28).
- ⁴⁶ See Mesnard, T. and others, 'Gemma: Open Models Based on Gemini Research and Technology' (arXiv 13 March 2024) <<https://arxiv.org/abs/2403.08295>>; Cooper, A.F. and others, 'Extracting Memorized Pieces of (Copyrighted) Books from Open-weight Language Models' (arXiv 18 May 2025) <<https://arxiv.org/abs/2505.12546>>
- ⁴⁷ See *Like Company* referral, para 8.
- ⁴⁸ See Google (n 40).
- ⁴⁹ See Recital 55 CDSM.
- ⁵⁰ See Recital 54 CDSM.
- ⁵¹ See *Like Company* referral, para 26.
- ⁵² See indicatively the Study for the JURI Committee, 'Strengthening the Position of Press Publishers and Authors and Performers in the Copyright Directive', Legal and Parliamentary Affairs, Policy Department for Citizens' Rights and Constitutional Affairs, Directorate General for Internal Policies of the Union, PE 596.810- September 2017 <https://www.europarl.europa.eu/RegData/etudes/STUD/2017/596810/IPOL_STU%282017%29596810_EN.pdf?utm_source=chatgpt.com>
- ⁵³ See, for example, Daniel Gervais, 'The Machine as Author', (2020) 105 Iowa LR 2053-2106, esp. pp. 2053-2054 and 2065; Hugenholtz, P. B.; Quintais, J. P., 'Copyright and Artificial Creation: Does EU Copyright Law Protect AI-Assisted Output?', (2021) 52(9) IIC-IRIPCL 1190-1216, esp. pp. 1195-1196, 1211-1212; Frosio G., 'Four theories in search of an AI(l)uthor', in Ryan Abbott and David Geffen (eds), *Research Handbook on Intellectual Property and Artificial Intelligence* (Elgar 2022), 155-177.
- ⁵⁴ Gaidartzi, A. and Stamatoudi, I., 'Authorship and Ownership Issues Raised by AI-Generated Works: A Comparative Analysis' (2025) 14(4) Laws 57 <<https://www.mdpi.com/2075-471X/14/4/57>>; Stamatoudi, I., 'Copyright Protection for AI-generated Works: Is this Something We Want?' (2025) Cyprus Review.
- ⁵⁵ See Recital 53 and Article 14 CDSM. See Stamatoudi, I., and Torremans, P., 'The Digital Single Market Directive', in Stamatoudi, I., and Torremans, P., (eds), *EU Copyright Law. A Commentary* (2nd edn, Edward Elgar Publishing 2021), 718 seq.

- ⁵⁶ On the dichotomy idea/expression, see Article 2(1) of the Berne Convention for the Protection of Literary and Artistic Works (Paris Act of 24 July 1971, as amended on 28 September 1979); Article 9(2) of the Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS), annex 1C to the Marrakesh Agreement establishing the World Trade Organization, 15 April 1994; Article 1(2) of Directive (EU) No 2009/24 of the European Parliament and of the Council of 23 April 2009 on the legal protection of computer programs; Case C-406/10, *SAS Institute Inc. v World Programming Ltd*, EU:C:2012:259, paras 39–41 and 45; Masouyé, C., *Guide to the Berne Convention for the Protection of Literary and Artistic Works* (Paris Act, 1971), WIPO, Geneva (1978), pp. 11–12; Rosati, E., *Originality in EU Copyright Law: Full Harmonization through Case Law* (Oxford University Press 2013), 49–62 (explaining the idea/expression dichotomy as a structural limit in EU copyright).
- ⁵⁷ Such an interpretation, recently adopted by the Regional Court of Munich in *GEMA v OpenAI* (see LG München I, Urteil vom 11.11.2025–42 O 14139/24 <<https://aifray.com/wp-content/uploads/2025/11/42-O-14139-24-Endurteil.pdf>> hereafter 'GEMA'), seems to rely on a misinterpretation of the technical functioning of the AI models as generally described in scientific literature, even though some ambiguities seem to persist. For an in-depth critique of *GEMA*, see Frosio, G., 'Copyright in Formaldehyde: How *GEMA v OpenAI* Freezes Doctrine and Chills AI-Part 1 and 2' (*Kluwer Copyright Blog*, 10–11 December 2025) <<https://legalblogs.wolterskluwer.com/copyright-blog/copyright-in-formaldehyde-how-gema-v-openai-freezes-doctrine-and-chills-ai-part-1>>; Guadamuz, A., 'Victory for GEMA in Germany against OpenAI' (*Technollama*, 22 November 2025) <<https://www.technollama.co.uk/victory-for-gema-in-germany-against-openai>>
- ⁵⁸ See Cooper, A.F. and others, (n 31).
- ⁵⁹ See, for example, Guadamuz, A., 'The Persistence of Memorisation' (*Technollama Blog*, 28 November 2025) <<https://www.technollama.co.uk/the-persistence-of-memorization>>
- ⁶⁰ The legal questions raised by chatbot summaries concern inference-time retrieval and display (RAG/output-layer use) and are addressed below in Section 6 (including on Article 15(1) CDSM and the concept of "very short extracts").
- ⁶¹ See, classically, Case C-379/98, *PreussenElektra AG v Schlesweg AG*, EU:C:2001:160, para 39 (clarifying that the Court may refuse to answer where the question bears no relation to the facts of the main action or is hypothetical); Case C-621/18, *Wightman and Others v Secretary of State for Exiting the European Union*, EU:C:2018:999, para 26 (holding that the Court may, in order to provide a useful answer, reformulate questions referred).
- ⁶² See Case C-5/08, *Infopaq International A/S v Danske Dagblades Forening*, EU:C:2009:465, paras 37–51, 64–66.
- ⁶³ See, for example, *SAS Institute* (n 55) paras 32–46 (confirming that functionality, programming language and data formats are not protected by copyright); Case C-604/10 *Football Dataco Ltd v Yahoo! UK Ltd*, EU:C:2012:115, paras 37–39 (on originality standard).
- ⁶⁴ See Case C-476/17, *Pelham GmbH v Hütter*, EU:C:2019:624, paras 31–40 (clarifying that use "unrecognisable to the ear" falls outside reproduction of a phonogram).
- ⁶⁵ *SAS Institute* (n 56), paras 39–46.
- ⁶⁶ See, for example, Goodfellow, Bengio, and Courville, (n 24) Ch 8–10, 294–296, 388–390 (regarding standard reference on gradient descent and parameter optimisation in neural networks).
- ⁶⁷ See, for example, Bengio, Y. and others, 'A Neural Probabilistic Language Model', (2003) 3 *JMLR* 1137, 1141; Vaswani and others, (n 25) paras 3.2.2 and 3.4.
- ⁶⁸ See, for example, Bengio and others, (n 66), 1141; Kingma, D.P. and Welling, M., 'An Introduction to Variational Autoencoders' (2019) 12(4) *FTML* 58–60 (noting that variational autoencoders provide a continuous latent space that enables smooth interpolation between inputs, capturing patterns rather than fixed copies of works). See also, confirming this key point, *Like Company* referral, para 5.
- ⁶⁹ See, for example, Hinton, G.E. and Salakhutdinov, R.R., 'Reducing the Dimensionality of Data with Neural Networks', (2006) 313 *Science* 504–506.
- ⁷⁰ See, for example, Carlini and others, (n 29), 2633–2635; Lee, K. and others, (n 32) 8424–8437, at 8424–8426 (on unprompted memorization and role of duplicates), 8429–8430 & 8430–8431 (on 10× reduction in *verbatim* emissions after deduplication).
- ⁷¹ See Bengio and others, (n 67) 1141; Vaswani and others, (n 25), paras 3.2.2 and 3.4.
- ⁷² See *SAS Institute* (n 56), paras 39–41 and 45 (clarifying that programming language, functionality and formats are not protected "expression"; and that copyright attaches to the author's expressive form, not functional/abstract elements).
- ⁷³ See *Pelham* (n 64), paras 31–40.

- ⁷⁴ See, for example, LeCun, Y. and others, 'Deep learning' (2015) 521 Nature 436–437 (noting that deep models “learn representations of data with multiple levels of abstraction”; and that higher layers capture increasingly abstract features); Infopaq (n 62), paras 37, 48 and 51 (stating that reproduction “in part” requires fixation of expressive elements, not mere ideas or abstractions).
- ⁷⁵ This error lies at the heart of the GEMA decision by the Munich Court (see above footnote 58), which conflated the probabilistic regeneration of text with the retention of fixed copies, a view that ignores the “lossy” nature of training and effectively freezes the doctrine of reproduction. See, for example, Guadamuz, A. (n 59).
- ⁷⁶ See Saffran, J.R. and others, 'Statistical Learning by 8-month-old Infants' (1996) 274 Science 1926–1927 (noting that infants track transitional probabilities in syllable streams, learning distributional regularities from exposure); Frost, R. and others, 'Domain Generality of Statistical Learning' (2015) 19(3) TCS 124–134, at 124–126 (noting that humans acquire statistical structure across modalities; and that learning reflects sensitivity to co-occurrence and sequence regularities); Saffran, J.R. and Kirkham, N.Z., 'Infant Statistical Learning' (2018) 69 ARP 182–185 (on survey of evidence that human learners extract patterns and regularities from continuous input, forming abstract representations rather than memorising surface forms).
- ⁷⁷ See (n 66) and (n 67).
- ⁷⁸ See, for example, Tenenbaum, J.B. and others, 'How to Grow a Mind: Statistics, Structure, and Abstraction' (2011) 331 Science 1279–1280 (noting that human learning builds abstract structural representations from statistical evidence rather than storing surface details); Clark, A., 'Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science' (2013) 36(3) BBS 181–204, at 181–183 (noting that, on a predictive-processing picture, cognition encodes generative, abstract models of causes, not verbatim inputs).
- ⁷⁹ See *Bartz and others v Anthropic PBC*, N.D. Cal., Case No. 3:24-cv-05417-WHA, Order on Fair Use, 23 June 2025, p. 12 (holding that copying and storing pirated books to build a central, permanent library was not fair use, but describing training as “quintessentially transformative” and comparing it to “any reader aspiring to be a writer”).
- ⁸⁰ See Infopaq (n 62), paras 37, 48 and 51; Case C-145/10, *Painer v Standard VerlagsGmbH and Others*, EU:C:2011:798, paras 89–92 (finding that a work is protected where it is the result of authors' free and creative choices and reflects their personality; and that protection attaches to those expressive choices); SAS Institute (n 56), paras 39–41 and 45; Football Dataco (n 63), paras 37–39 (confirming that protection concerns expressive choices, not mere data or effort).
- ⁸¹ Pelham (n 64), paras 31–40.
- ⁸² See Joined Cases C-580/23 and C-795/23, *Mio AB and Others v Galleri Mikael & Thomas Asplund AB and USM U. Schärer Söhne AG v konektra GmbH and LN*, ECLI:EU:C:2025:941 ('*Mio/konektra*'), paras 86–87.
- ⁸³ *Id.*, para 87. See also Opinion of Advocate General Emiliou in Case C-590/23 CG, *YN v Pelham GmbH, SD, UP*, ECLI:EU:C:2025:452, para 64 (noting that style “is not, as such, protected by copyright” and belongs to the realm of ideas, not expressions; and that overall stylistic features and common motifs are not individually protected).
- ⁸⁴ *Ibid.*
- ⁸⁵ Infopaq (n 62), paras 50–66.
- ⁸⁶ See Recital 18 and Article 4(1) CDSM (clarifying that TDM consists of automated analytical techniques that derive “patterns, trends and correlations” from digital text/data, *i.e.*, a recognition that such operations are analysis, not the copying of expressive strings).
- ⁸⁷ See (n 56).
- ⁸⁸ See Hugenholtz, P.B., 'Against “Data Property”', in Ullrich, H. et al. (eds), *Kritika: Essays on Intellectual Property: Vol. 3* (Edward Elgar 2018) 48–71 (warning that new exclusion rights over information would distort markets and innovation).
- ⁸⁹ See Lemley, M. and Casey, B., 'Fair Learning' (2021) 99 Texas LR 743, 748–49 (arguing that copyright should not create a *de facto* control right over the ability of machines, or people, to learn); Margoni, T. and Kretschmer, M., 'A Deeper Look into the EU Text and Data Mining Exceptions: Harmonisation, Data Ownership, and the Future of Technology' (2022) 71 GRUR Int. 688–690, 696–702 (warning that allocating control upstream over analytic uses risks regulating AI via property rights in inputs).
- ⁹⁰ For the engineering reality that training pipelines generate *ephemeral artefacts* (buffers, caches, tokenisation shards and sharded files) as part of efficient preprocessing and input streaming, see, for example, Kudo, T. and Richardson, J., 'SentencePiece: a Simple and Language-independent Subword Tokenizer and Detokenizer for Neural Text Processing', in *Proceedings of EMNLP 2018: System Demonstrations*, Association for Computational Linguistics (2018), 66–71 (focusing on tokenisation as a distinct preprocessing stage producing model files and segmented corpora used downstream); Murray, D.G. et al., 'tf.data: A Machine Learning Data Processing Framework' (2021) 14(12) *Proceedings*

VLDB E 2945–2949, 2951, 2954–2955 (noting that standard pipeline operators include caching, prefetching and buffering to overlap I/O and compute, *i.e.*, transient, functional copies integral to training throughput).

- ⁹¹ See, for example, Geiger, C. and others, 'The Exception for Text and Data Mining (TDM) in the Proposed Directive on Copyright in the Digital Single Market—Legal Aspects' (European Parliament, In-Depth Analysis, 2018) (explaining that TDM processes inherently trigger restricted acts—*i.e.*, reproduction/extraction—and recommending an exception framework, precisely to avoid treating machine *learning* as *copying*, while still addressing ancillary copies); Margoni and Kretschmer, 'A Deeper Look into the EU Text and Data Mining Exceptions' (2022) 71(8) GRUR Int. 685–699 (emphasising that TDM entails copying steps but that EU law channels them through specific exceptions/opt-outs rather than recasting *analysis* as *reproduction*).
- ⁹² See, for example, Dornis, T.W. and Stober, S., *Urheberrecht und Training generativer KI-Modelle—technologische und juristische Grundlagen* (Nomos, Baden-Baden 2024), 1–2 (stating expressly that during pre-training and fine-tuning "copyright-relevant reproductions [...] materialize 'inside' the AI model"; and that this constitutes copying in the legal sense); Dornis, T.W., 'Generative AI, Reproductions Inside the Model, and the Making Available to the Public' (2025) 56 IIC—IRIPCL 916–922 (sections 2.2.1 & 2.2.1.2: arguing that works used for training are reproduced "inside" the model in a vector/latent space and that regurgitating outputs imply an internal copy); Lucchi, N., *Generative AI and Copyright: Training, Creation, Regulation*, Study for the European Parliament (Policy Department for Justice, Civil Liberties and Institutional Affairs), PE 774.095, July 2025, p. 30 (noting that "AI training entails systematic reproduction of protected works"); pp. 34–45 and 78 (maintaining that ingestion by generative systems "typically involves acts of reproduction beyond analytical use"), and p. 125 (pointing out that reproduction right is engaged even if use serves a new purpose).
- ⁹³ See (n 90).
- ⁹⁴ See, for example, European Commission, Commission Staff Working Document, *Impact Assessment on the modernisation of EU copyright rules*, SWD (2016) 301 final, Part 1/3, Brussels, 14 September 2016, s. 4.3 "Text and data mining", pp. 105–106 (on legal uncertainty and diversity of licensing practices generating transaction costs, including for compliance).
- ⁹⁵ See Emanuilov, I. and Margoni, T., 'Forget Me Not: Memorisation in Generative Sequence Models Trained on Open Source Licensed Code' (2024), SSRN <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4720990> (showing that problematic verbatim "memorisation" is the exception, underscoring why models should not be equated with their training corpora).
- ⁹⁶ See (n 31–34).
- ⁹⁷ See, for example, Microsoft, 'Protected Material Detection Filter' (Microsoft Learn, 16 September 2025) <<https://learn.microsoft.com/en-us/azure/ai-foundry/openai/concepts/content-filter-protected-material?tabs=text>> (describing output-side filters that flag and block content closely matching copyrighted text or code).
- ⁹⁸ See, for example, Kumar, P., 'Adversarial Attacks and Defenses for Large Language Models (LLMs): Methods, Frameworks & Challenges' (2024) 13 IJMIR 3.
- ⁹⁹ Ahmed, A. and others, 'Extracting Books From Production Language Models', arXiv preprint, arXiv:2601.02671 (2026) <<https://arxiv.org/abs/2601.02671>>
- ¹⁰⁰ Ibid.
- ¹⁰¹ Ibid, 19 ("for the cases in which we retrieved whole copies of near-verbatim books, it was often quite costly (> \$100) to do so").
- ¹⁰² Ibid.
- ¹⁰³ See, for example, Robison, K., 'AI Models Can Memorize Books, and That's a Problem,' *The Atlantic* (7 January 2026) <<https://www.theatlantic.com/technology/2026/01/ai-memorization-research/685552>>
- ¹⁰⁴ See Ahmed and others, (n 99). Some further joint interdisciplinary research involving tech experts and legal scholars would be useful to firmly clarify all issues, though (see in this spirit Guadamuz, A., 'Why We Need Tech Lawyers to Shine Again' (2025) 56 IIC 1797), as the current technical results give rise to apparently very different interpretations among legal scholars.
- ¹⁰⁵ See GEMA (n 57).
- ¹⁰⁶ See, for example, Guadamuz (n 59).
- ¹⁰⁷ Recital 57 CDSM.
- ¹⁰⁸ Ibid; see also Reto M. Hilty and others, 'Position Statement of the Max Planck Institute for Innovation and Competition on the Proposed Modernisation of European Copyright Rules' (2017) para 28 (warning against related rights creating super-protection for non-original content).

- ¹⁰⁹ See again Recital 57 CDSM (clarifying that the “rights granted to publishers of press publications should not extend [...] to mere facts reported in press publications”).
- ¹¹⁰ *Ibid.* (mentioning that the “rights granted to the publishers of press publications under this Directive should have the same scope as the rights of reproduction and communication to the public provided for in Directive 2001/29/EC...”).
- ¹¹¹ See above the technical analysis (Section 2) and the legal assessment of the right of reproduction issues (Section 4).
- ¹¹² See Guadamuz, A., ‘A Scanner Darkly: Copyright Liability and Exceptions in Artificial Intelligence Inputs and Outputs’ (2022) 73 GRUR Int. 11; Lucchi, N., ‘ChatGPT: A Case Study on Copyright Challenges for Generative Artificial Intelligence Systems’ (2023) European JRR 1.
- ¹¹³ Around 80 copyright-related lawsuits have been filed thus far against GenAI system operators in the US. Several scholars argue it is s likely fair use. See Sag (n 31) 295; Lemley, M., ‘How Generative AI Turns Copyright Law Upside Down’, (2024) 25 CSTR 21; Samuelson, P., ‘Fair Use Defenses in Disruptive Technology Cases’ (2024) 71 UCLA LR 1484; Lemley and Casey (n 89) 743. This position however is not undisputed: see Brauneis, R., ‘Copyright and the Training of Human Authors and Generative Machines’ (2024) 48 CJLA 1; Ginsburg, J.C., ‘AI Input Data and Fair Use: A View from the U.S.’ (2024) AM 165.
- ¹¹⁴ Japan has established one of the most permissive legal environments globally for TDM, and has therefore been sometimes called by Japanese scholars a “paradise for GenAI training”. Indeed, Article 30-4 of the Japanese Copyright Act allows the use of copyrighted materials for “non-enjoyment purposes”, which includes TDM for commercial purposes, without the need for rights holder’s permission or compensation (see Ueno, T., ‘The Flexible Copyright Exception for ‘Non-Enjoyment’ Purposes—Recent Amendment in Japan and Its Implication’ (2021) GRUR Int. 145).
- ¹¹⁵ For a comment see de Champris, B., ‘Artificial Intelligence and Copyright: The EU Should Preserve the Copyright Directive’s Delicate Balance to Safeguard and Promote Innovation’ (2024) 1 AIRE 113; Peukert, A., ‘Copyright in the Artificial Intelligence Act—A Primer’ (2024) 73(6) GRUR Int. 497.
- ¹¹⁶ Answer given on 31 March 2023 by Commissioner Breton on Behalf of the European Commission (Parliamentary written question for written answer, E-000479/2023, posed on 15 February 2023: “How Does the Commission Plan to Regulate This Use of AI, Which Harms Artists and Rights Holders?”) <https://www.europarl.europa.eu/doceo/document/E-9-2023-000479-ASW_EN.html>
- ¹¹⁷ See, for example, Senftleben, M., ‘GenAI and the Copyright Three-Step Test—Do TDM Exceptions for AI Training Conflict with a Work’s Normal Exploitation?’ (2025) 74 GRUR Int. 1–2 <<https://doi.org/10.1093/grurint/ikaf113>> (also denying a possible conflict of the TDM exception used for AI training with the three-step test); ‘Copyright and Generative AI: Opinion of the European Copyright Society’ (2025) 16(1) JIPITEC 121.
- ¹¹⁸ The argument that the TDM exception has not been tailor-made for GenAI purposes (which is correct) and therefore might not cover all reproductions relevant to the process seems therefore not relevant to determine the legality of training, as the EU legislator decided to submit the training to the exception.
- ¹¹⁹ See on the issue more broadly Manteghi, M., ‘Readjusting the EU’s Regulation of Text and Data Mining in the Age of Artificial Intelligence’, PhD defended at the University of Turku, 28 Nov. 2025, *Annales Universitatis Turkuensis*, 2025.
- ¹²⁰ On these exceptions see Geiger, C., Frosio, G., and Bulayenko, O., ‘Text and Data Mining: Articles 3 and 4 of the Directive 2019/790/EU’, in Saiz García, C. and Evangelio Llorca, R. (eds.), ‘*Propiedad intelectual y mercado único digital europeo*’ (Tirant lo Blanch 2019), p. 27. For an analysis prior to the CDSM see also Stamatoudi, I., ‘Text and Data Mining’, in Stamatoudi (ed.), “New Developments in EU and International Copyright Law” (Leiden, Netherlands, Kluwer Law International 2016), 264–265.
- ¹²¹ Question which was addressed by the Munich Regional Court in *Gema* (on this decision see (n 57) and again *infra*).
- ¹²² Dornis, T.W. and Stober, S., ‘Generative AI Training and Copyright Law’, *Transactions of the International Society for Music Information Retrieval*, V (2025) <<https://arxiv.org/pdf/2502.15858>> and from the same authors, ‘Urheberrecht und Training generativer KI-Modelle—technologische und juristische Grundlagen’ (4 September 2024); Dornis, T.W., ‘The Training of Generative AI is not Text and Data Mining’ (2025) 47 EIPR 65–78.
- ¹²³ Rosati, E., ‘Infringing AI: Liability for AI-Generated Outputs under International, EU and UK Copyright Law’ (2025) 16 EJRR 610–611 (stressing that AI training entails some copyright-relevant acts, *i.e.*, reproduction, communication/ making available to the public, are not covered by the TDM exception under Articles 3 and 4 CDSM); Dornis, T.W., and Lucchi, N., ‘Generative AI and the Scope of EU Copyright Law: A Doctrinal Analysis in Light of the Referral in *Like Company v. Google*’ (2025) IIC.
- ¹²⁴ Alonso, E., and Lucchi, N., ‘AI and Copyright ‘Hallucinations’: Does the Text and Data Mining Exception Really Support Generative AI Training?’ (2025) 47(9) EIPR 515–526.

- ¹²⁵ For an early comment on this case, see Peukert, A., 'Vervielfältigung im KI Model? Zugleich Anmerkung zu LG München I, Urteil vom 11.11.2025', ZUM 2026, p. 17. On memorization see also Cooper, A.F. and Grimmelmann, J., 'The Files are in the Computer: on Copyright, Memorization, and Generative AI' (2025) 100 Chicago-Kent LR 1 141-219. Likewise, in *Gema* the Munich court decided that the output of large portion of song texts based on prompts qualify as communication to the public, which however has not been much disputed so far, as verbatim take-over of identical parts of the protected work in the output can amount to a copyright infringement (see our analysis below in Section 5).
- ¹²⁶ See in this sense Geiger, C. and Iaia, V., 'Is There a Captain in the Ship? The EU Copyright Regulator's Quest in the Generative AI Era', *ILEO Research Paper Series* No. 25-09 (2025).
- ¹²⁷ More detailed in this sense is for example, Mezei, P., 'A Saviour or a Dead End? Reservation of Rights in the Age of Generative AI' (2024) 46(7) EIPR 461.
- ¹²⁸ The Copyright Knowledge Center at EUIPO was launched on the 21 November 2025: <<https://www.euipo.europa.eu/en/copyright-knowledge-centre>>; the Center will in particular "explore new services that enhance copyright transparency and accessibility across the Single Market".
- ¹²⁹ On the potential future role of the EUIPO, and on more generally on the need for governance on this issue at EU level, see above footnote 129; Christophe Geiger and Vincenzo Iaia, 'Towards an Independent EU Regulator for Copyright Issues of Generative AI: What Role for the AI Office (But More Importantly: What's Next)?' (2024) 2 AM 185.
- ¹³⁰ See in this regard Geiger and others, (n 91). For a critical evaluation of the Directive proposal, see Geiger, C., Frosio, G. and Bulayenko, O., 'Text and Data Mining in the Proposed Copyright Reform: Making the EU Ready for an Age of Big Data?' (2018) 49(7) IIC 814, and from the same authors: 'The EU Commission's Proposal to Reform Copyright Limitations: A Good but Far Too Timid Step in the Right Direction' (2018) 40(1) EIPR 4; European Copyright Society, 'General Opinion on the EU Copyright Reform Package', 24 January 2017, 5; Hilty, R.M. and Richter, H., in Hilty, R.M., and Moscon, V., (eds.), 'Modernisation of the EU Copyright Rules, Position Statement', MPI for *Innovation and Competition Research Paper* No. 17-12 (2017), p. 25 et sq.
- ¹³¹ See Geiger, C., 'The Missing Goal-Scorers in the Artificial Intelligence Team: Of Big Data, the Fundamental Right to Research and the failed Text and Data Mining Limitations in the CSDM Directive', in Senftleben, M. and others (eds.), 'Intellectual Property and Sports, Essays in Honour of P. Bernt Hugenholtz' (Kluwer Law International 2021), 383 sq.; Geiger and others, (n 91). For a (critical) analysis see also Hugenholtz, P.B., 'The New Copyright Directive: Text and Data Mining (Articles 3 and 4)' (*Kluwer Copyright Blog*, 24 July 2019) <<http://copyrightblog.kluweriplaw.com/2019/07/24/the-new-copyright-directive-text-and-data-mining-articles-3-and-4/>>; Kop, M., 'The Right to Process Data for Machine Learning Purposes in the EU' (2021) 34(1) Harvard JLT 7; European Copyright Society, 'Comment of the European Copyright Society Addressing Selected Aspects of the Implementation of Articles 3 to 7 of Directive (EU) 2019/790 on Copyright in the Digital Single Market', 3 May 2022, <https://europeancopyrightsocietydotorg.files.wordpress.com/2022/05/ecs_exceptions_final-3.pdf>; Senftleben, M., 'Compliance of National TDM Rules with International Copyright Law—An Overrated Nonissue?' (2022) 53(10) IIC 1477 sq.; Margoni and Kretschmer (n 89), 685; Ducato, R. and Strowel, A., 'Limitations to Text and Data Mining and Consumer Empowerment: Making the Case for a Right to 'Machine Legibility'' (2019) 50(6) IIC 649 seq; Bonadio, E. and McDonagh, L., 'Artificial Intelligence as Producer and Consumer of Copyright Works: Evaluating the Consequences of Algorithmic Creativity' (2020) 2 IPQ 112. With a similar conclusion for South America, see Jackson Berton, M., 'Text and Data Mining Exception in South America: A Way to Foster AI Development in the Region' (2021) 70(12) GRUR Int. 1145.
- ¹³² This seems confirmed by Recital 105 of the AI Act.
- ¹³³ Geiger, C. and others, 'Time to (Finally) Reinstall the Author in EU Copyright Law: From Contractual Protection to Remuneration Rights' (2025) 56 IRIPCL 10.
- ¹³⁴ Geiger, C., 'Elaborating a Human Rights friendly Copyright Framework for Generative AI' (2024) 55(7) IRIPCL 1129; Geiger, C. and Iaia, V., 'The Forgotten Creator: Towards a Statutory Remuneration Right for Machine Learning of Generative AI' (2024) 52(105925) CLSR 1; Frosio, G., 'Should We Ban Generative AI, Incentivise It or Make it a Medium for Inclusive Creativity', in Bonadio, E. and Sganga, C. (eds.), 'A Research Agenda for EU Copyright Law' (Elgar 2025), 61; Lucchi, N., 'Generative AI and Copyright', *Study for the JURI Committee*, July 2025, 122, admitting "that the long-term viability of individualized licensing across billions of works is limited" and that "a statutory licensing scheme or collective remuneration mechanism may ultimately provide a more scalable and equitable model"; Carlini, R., Menéndez, N. and Schiffrin, A., 'How to Update EU and US Copyright Regimes in the Age of AI', *IPD Working Paper* 2026, p. 17, according to whom "shifting to a default system of prior authorization coupled with a fixed payment scale seems to us to be the most likely to address the challenge of preserving journalism in the era of AI". For a similar proposal but focusing on the output, see Senftleben, M., 'Generative AI and Author Remuneration' (2023) 54 IIC 1535. In a recent economic study for the European Parliament, one of the key findings was that this regulatory solution was

the optimal to balance all interests involved: see Peukert, C., 'The Economics of Copyright and AI, Empirical Evidence and Optimal Policy', In-Depth Analysis Requested by the JURI Committee, Dec. 2025, p. 7 (noting that "[s]tatutory licensing appears optimal. Among realistic options, *statutory licensing emerges as a robustly welfare-superior framework under most calibrations*. Making the license compulsory ensures wide access to works, a royalty rate set by a regulatory authority balances interests of rightsholders, AI developers and end users, enabling royalty payments that create incentives for enough new creation to keep AI systems valuable over time") (emphasis added). In general, on the necessity to reach a balance between the rights of content creators and the need not to jeopardise GenAI innovation, see Bonadio, E. and Pollicino, O., 'From Wall to Compass: Modernising EU Copyright in the Generative AI Era' (2026) 2 EIPR.

- ¹³⁵ See Football Dataco (n 63) paras 38–39 (confirming that copyright does not protect "information" as such, but the author's own intellectual creation in selection/arrangement); SAS Institute (n 56) paras 39–41 (noting that functionality, programming language and data formats are not protected expression).
- ¹³⁶ See Infopaq (n 62), paras 37, 48 and 51 (confirming that reproduction "in part" depends on fixation of elements expressing the author's own intellectual creation).
- ¹³⁷ Ibid.
- ¹³⁸ See Mio/konektra (n 82), paras 86–87 and 92–93.
- ¹³⁹ See Pelham (n 64), para 38–39.
- ¹⁴⁰ For the same reason, reliance on the quotation exception is generally misplaced. Quotation presupposes the incorporation of an identifiable extract from the protected work (in *Pelham*, a sample recognisable to the ear; for literary works, an extract i.e., in substance, taken verbatim). Where the output is a paraphrase rather than an identifiable excerpt, it is more naturally analysed as summarisation than as "quotation", much as, in ordinary writing practice, quotation marks signal a word-for-word extract, not a restatement.
- ¹⁴¹ This distinction is vital for Article 15 CDSM. Because the PPR is a neighbouring right, it applies to publications regardless of their 'originality'. This creates a specific legal risk: if the concept of 'reproduction' were interpreted loosely to cover the mere overlap of facts, themes, or information, the PPR would effectively mutate into a monopoly over information itself, something EU law has consistently refused to grant. See Infopaq (n 62), paras 45–46; Football Dataco (n 63), para 42; SAS Institute (n 56), paras 39–41; see also again Recital 57 CDSM (stating as mentioned that the "rights granted to publishers of press publications should not extend [...] to mere facts reported in press publications"). To prevent this, the test for 'reproduction' under Article 15 CDSM must remain anchored to the taking of recognisable wording or expression, not just the underlying facts. As discussed in Section 4, this alignment is essential to preserve the hierarchy of the *acquis*: if the PPR were allowed to capture informational overlap that copyright leaves free, the related right would become functionally broader and more restrictive than the author's primary right it was designed to complement.
- ¹⁴² See *Like Company* referral, paras 11 and 22.
- ¹⁴³ Ibid, paras 4–5.
- ¹⁴⁴ See, for example, Recital 58 CDSM (on the limitations of Article 15).
- ¹⁴⁵ Ibid, para 7.
- ¹⁴⁶ Stamatoudi I., 'Originality under EU Copyright Law', in Torremans, P. (ed), *Copyright Law: A Handbook of Contemporary Research* (Elgar 2017), 57; Stamatoudi I., 'More' Originality for Cypriot Copyright Law according to the CJEU's Case Law' (2018) 30(1) CR 207.
- ¹⁴⁷ See Infopaq (n 62).
- ¹⁴⁸ Ibid, para 47.
- ¹⁴⁹ Even the *sui generis* right for databases requires that substantial investment (either in terms of quality or in terms of quantity) has taken place in either the obtaining, verification or presentation of the contents, as provided in Article 7 of the 1996 Directive 96/9/EC on the Legal Protection of Databases ("Database Directive"). Also, only substantial takings are prevented and not any taking unless it is done repeatedly and systematically. In the case of the PPR there is no need for investment despite the fact that the notion of investment is referred to in the reasoning of this provision. See Recital 58 CDSM: "[...] [T]he use of individual words or very short extracts of press publications by information society service providers may not undermine the investments made by publishers of press publications in the production of content".
- ¹⁵⁰ Stamatoudi, I., and Torremans, P., 'The Digital Single Market Directive', in Stamatoudi, I. and Torremans, P. (eds), *EU Copyright Law. A Commentary* (Elgar 2021), pp. 651, and 727–728.

- 151 Recital 58 CDSM.
- 152 However, the structural question raised by the national court, i.e. whether Article 3(2) “may exclude” Article 15, remains vital here. While Article 15 builds on the concepts of Articles 2 and 3 InfoSoc, it functions as a *lex specialis* that tailors these rights for the online press sector. This implies that where the specific conditions of Article 15 are not met (e.g., because the use involves only very short extracts), a fallback claim under the general Article 3(2) right should not be available to re-capture the same conduct. To hold otherwise would render the Article 15 carve-outs nugatory.
- 153 See (n 57).
- 154 Infopaq (n 62). As acknowledged in the Evaluation Report on the Database Directive, “[t]here is a long-standing principle that copyright should not be extended to cover basic information or “raw” data”. See Commission of the European Communities, DG Internal Market and Services Working Paper. First Evaluation of Directive 96/9/EC on the Legal Protection of Databases, 12 December 2005, p. 23. The principle has been applied by the CJEU with reference to the *sui generis* right in *The British Horseracing Board Ltd v William Hill Organization Ltd*, Case C-203/02, EU:C:2004:695; in general see Karapapa, S., *Defences to Copyright Infringement: Creativity, Innovation and Freedom on the Internet* (Oxford University Press 2020), Chapter 2.
- 155 See in this regard C-135/10, *Società Consortile Fonografici (SCF) v Marco Del Corso*, EU:C:2012:140, para 84; C-89/04 *Mediakabel BV v Commissariaat voor de Media*, EU:C:2005:348, para 30; C-192/04 *Lagardère Active Broadcast, v Société pour la perception de la rémunération équitable (SPRE) et al.*, EU:C:2005:475, para 31; and C-306/05 *Sociedad General de Auditores y Editores de España (SGAE) v Rafael Hoteles SL* EU:C:2006:764, paras 37 and 38; see also in general Karapapa, S., ‘The Requirement for A ‘New Public’ in EU Copyright Law’ (2017) 1 ELR 63–81.
- 156 C-466/12, *Nils Svensson and Others v Retriever Sverige AB*, EU:C:2014:76; C-279/13, *C More Entertainment AB v Linus Sandberg*, EU:C:2015:199; C-160/15, *GS Media BV v Sanoma Media Netherlands BV, Playboy Enterprises International Inc., Britt Geertruida Dekker*, EU:C:2016:644; C-527/15, *Stichting Brein v Jack Frederik Wullems (Stichting I)*, EU:C:2017:300. Relevant cases also include: Case C-360/13, *Public Relations Consultants Association Ltd v The Newspaper Licensing Agency Ltd*, EU:C:2014:1195 (on ‘browsing’); and C-610/15, *Brein v Ziggo BV and XS4ALL Internet BV*, EU:C:2017:456; Stamatoudi, I., ‘Linking’ and ‘Browsing’ in the Light of the EU Court of Justice’s Recent Case Law’, in Pistorius, T. (ed.), *IP Perspectives on the Regulation of New Technologies* (Elgar 2018), 179.
- 157 Quintais, J.P., ‘Untangling the Hyperlinking Web: In Search of the Online Right of Communication to the Public’ (2018) 21 JWIP 385–420, at 398.
- 158 See *Svensson, C More Entertainment; GS Media BV; Wullems (Stichting I); Ziggo (Stichting II)*, (n 156).
- 159 *Getty Images vs. Stability AI* [2025] EWHC 2863 (Ch).
- 160 See GEMA (n 56).
- 161 See *Svensson, BestWater, GS Media and VG Bild-Kunst*, (n 156).
- 162 *Idem*.
- 163 See (n 57).
- 164 See Directive 2000/31/EC of the European Parliament and of the Council of 8 June 2000 on certain legal aspects of information society services, in particular electronic commerce, in the Internal Market, OJ 2000 L 178, Artt 12–15; Regulation (EU) No. 2022/2065 of 19 October 2022 on a Single Market For Digital Services (Digital Services Act) [2022] OJ L277/1, Articles 4–9.
- 165 See *Ziggo* (n 157), paras 36–45 (on deliberate intervention; knowledge of infringement; substantial infringing uses and lack of neutrality; and inducement of infringement); Case C-392/19, *VG Bild-Kunst v Stiftung Preußischer Kulturbesitz*, EU:C:2021:181, paras 33–34 (on intentional intervention and control); Case C-160/15, *GS Media BV v Sanoma Media Netherlands BV*, EU:C:2016:644, paras 34–51 (on the knowledge element).
- 166 Joined Cases C-682/18 and C-683/18, *YouTube LLC v Cyando AG*, EU:C:2021:503, paras 102–115 (on criteria indicating active role/knowledge); Case C-324/09, *L’Oréal SA v eBay International AG*, EU:C:2011:474, paras 113–124 (clarifying that active role removes safe harbour); Joined Cases C-236/08 to C-238/08, *Google France and Google v Louis Vuitton Malletier and others*, EU:C:2010:159, paras 113–120 (on hosting when neutral).
- 167 See, for example, *Carlini and others* (n 29) 2642–2646 (on rarity/duplication and attack prompts).
- 168 See, for example, *Lee, K. and others* (n 31) 8429–8431 (on order-of-magnitude reduction in verbatim emissions after deduplication).
- 169 See, for example, *Abadi and others*, ‘Deep Learning with Differential Privacy’ in *Proceedings of CCS 2016* (ACM 2016), 309–312 (on DP-SGD limits per-example memorisation); *Carlini and others* (n 31) 7–9 (on scale, duplication and prompt length).

- ¹⁷⁰ See Google France (n 166), paras 113–114 (confirming that neutral, mere technical/automatic/passive role implies no knowledge/control); Svensson, above footnote 157, paras 24–25, 29–31 (clarifying that the assessment does not turn on the site where the work resides; and that even where framing makes a work appear on the linker's site, the legal analysis is unchanged); *YouTube LLC and Cyando AG* (n 166), paras 68 and 90, 100–102 (noting that attribution hinges on an operator's deliberate intervention beyond merely making a service available); Opinion of AG Hogan in Case C-433/20, *Austro-Mechana v Strato*, EU:C:2022:217, paras 59–60, 69 (emphasising that who makes the copy is key; and that cloud providers can remain outside the "reproducer" role absent active involvement). See also Stamatoudi, I., Torremans, P.L.C. (eds), *EU Copyright Law: A Commentary* (Cheltenham, 2nd edn., Edward Elgar 2021), Art. 2 comment (noting that attribution in reproduction turns on the party that performs the act); Peukert, C., 'Copyright Levies and Cloud Storage: Ex-ante Policy and Legal Questions after *Austro-Mechana v Strato*' (2024) 53 RP article 104114 (pointing out that *Austro-Mechana* treats cloud-server copies as private copies by users; and distinguishing hosting capacity from making the reproduction); Jütte, B.J., 'Analysis: The Court of Justice Rules on the Private Copying Exception on Cloud Servers (*Austro-Mechana v Strato*, C-433/20)', *EU Law Live*, 25 March 2022 (clarifying that the user can be the reproducer even when the copy resides on a third-party server; and that the provider's role is relevant only if it steps beyond neutrality).
- ¹⁷¹ See Infopaq (n 62), paras 50–51 and 64–66 (highlighting reproduction turned on fixation of human-readable extracts); Case C-360/13, *PRCA v NLA*, EU:C:2014:1195, paras 21–27 and 52–63 (confirming that the temporary copies exception must allow the development and operation of new technologies).
- ¹⁷² See, for example, Case C-70/10, *Scarlet Extended SA v SABAM*, EU:C:2011:771, paras 40–53; Case C-360/10, *SABAM v Netlog NV*, EU:C:2012:85, paras 34–51 (focusing on "no general monitoring", proportionality and fundamental rights).
- ¹⁷³ See, for example, Case C-516/17, *Spiegel Online GmbH v Volker Beck*, EU:C:2019:625, paras 76 and 94 (on conditions for quotation; contextual assessment, fair practice and extent required by the specific purpose); Pelham (n 64), paras 67, 69 and 71–74 (on definition and purpose of quotation; intention to enter into "dialogue" with the quoted work; and proportionality).
- ¹⁷⁴ See, for example, European Commission, Communication - Guidance on Article 17 of Directive (EU) 2019/790 on Copyright in the Digital Single Market, COM(2021) 288 final, Brussels, 4.6.2021, para 171 (noting that at the present state of the art, automated tools cannot reliably recognise context-dependent lawful uses such as quotation, parody or pastiche, so ambiguous cases should not be blocked *ex ante*); Case C-401/19, *Republic of Poland v European Parliament and Council*, EU:C:2022:297, paras 80, 85, and 90–94 (noting that filtering measures must include robust safeguards to prevent blocking of lawful content and maintain a fair balance of rights; and that only manifestly unlawful content may be prevented automatically); *Scarlet Extended* (n 172), paras 36–40, 52; and *Netlog* (n 172), paras 34–38, 50 (clarifying that general, preventive filtering obligations are incompatible with fundamental rights balancing and risk over-blocking lawful uses); see also Article 8 DSA (titled "No general monitoring or active fact-finding obligations").
- ¹⁷⁵ See Calabresi, G. and Melamed, A.D., 'Property Rules, Liability Rules, and Inalienability: One View of the Cathedral' (1972) 85 *Harvard LR* 1116–1123 (on choosing instruments under high transaction costs; and least-cost avoider logic).
- ¹⁷⁶ See (n 174).
- ¹⁷⁷ See, for example, Michael Heller, 'The Tragedy of the Anticommons: Property in the Transition from Marx to Markets' (1998) 111 *Harvard LR* 621. See also empirical discussion in the European Commission, Communication - A European strategy for data, COM(2020) 66 final, Brussels, 19.2.2020, pp. 2–11 (on documents concentration of data with a few Big Tech players; fragmentation and unequal conditions harming SMEs; and B2B frictions such as imbalanced bargaining power and legal uncertainty: all classic sources of overlapping claims/transaction costs that skew markets toward well-resourced incumbents).
- ¹⁷⁸ Article 14 E-Commerce Directive.
- ¹⁷⁹ Article 8 DSA.
- ¹⁸⁰ See *Scarlet Extended* and *SABAM/Netlog* (n 172); *Republic of Poland* (n 174).
- ¹⁸¹ See *L'Oréal* (n 166), paras 131–134, 139 and 143–144; Case C-18/18, *Glawischign-Piesczek v Facebook Ireland*, EU: C: 2019: 821, paras 37, 41, 45–47 and 53 (on targeted injunctions; and "equivalent content" within limits consistent with Article 15 E-Commerce Directive).
- ¹⁸² See Article 14 E-Commerce Directive; and DSA, especially Articles 14–17, 20–24, 34–35 (on due diligence; and systemic risks).
- ¹⁸³ See, for example, *L'Oréal* (n 166), paras 116–123; *YouTube/Cyando* (n 166), paras 107–115.
- ¹⁸⁴ See *YouTube/Cyando* (n 166), paras 93–96 (confirming that YouTube had taken "credible and effective measures"; and that its business model is not based on illegality).

- ¹⁸⁵ *Idem*, paras 84–87 (highlighting that relevant factors include whether the operator deploys appropriate measures and whether its financial model encourages unlawful sharing).
- ¹⁸⁶ See, for example, *Ziggo* (n 156) (deciding that providers lose the protection afforded to neutral intermediaries if they actively intervene to facilitate infringement, for instance, by knowingly indexing and managing illicit content).
- ¹⁸⁷ Articles 15(1)–(3) CDSM.
- ¹⁸⁸ See, for example, *YouTube/Cyando* (n 166), paras 71, 102, 106 and 117–118.
- ¹⁸⁹ See Article 15(2) and Recital 57 CDSM. Article 15 is expressly “without prejudice” to authors’ and other rightholders’ rights and may not be invoked against them, signalling that the PPR is meant to add protection rather than re-allocate value. A provider-primary, strict-liability reading can still economically harm authors in two concrete ways: (i) rent-shifting: platforms will pay to clear risk through publisher licences, so money that could otherwise flow to authors via author-right licensing (or remain available for lawful dissemination that supports authors’ market value) is diverted to publishers unless it is shared; and (ii) over-blocking: because providers cannot know the user’s purpose in advance (e.g., quotation, criticism, reporting), they will block or heavily filter borderline uses, reducing lawful visibility, readership and downstream revenue opportunities for authors. The “prejudice” is therefore practical and monetary: authors may capture less value and lose audience-building uses, even though their formal rights remain intact.).
- ¹⁹⁰ See *Scarlet Extended* (n 172), paras 48–53; *SABAM/Netlog* (n 172), paras 43–51.
- ¹⁹¹ See Articles 16 and 17 DSA (on notice-and-action; and statements of reasons); Article 20 (on repeat-infringer policy); Articles 34–35 (on risk assessment/mitigation for VLOPs/VLOSEs); see also Case C-314/12, *UPC Telekabel Wien GmbH v Constantin Film*, EU:C:2014:192, paras 62–63 (on proportionate, targeted injunctions).
- ¹⁹² See Articles 16–17, 34–35 DSA.
- ¹⁹³ See Article 8(3) InfoSoc; *UPC Telekabel* (n 191), paras 56–63 (confirming that injunctions must be proportionate and avoid imposing general monitoring).
- ¹⁹⁴ *YouTube/Cyando* (n 166), paras 102–115 (on diligent platforms not being primarily liable; and criteria for knowledge/active role).
- ¹⁹⁵ *Ibid.*
- ¹⁹⁶ *Like Company* referral, para 26.

AUTHOR BIOGRAPHIES

Enrico Bonadio is Professor of Intellectual Property Law at City St. George's, University of London, UK.

Giancarlo Frosio is a Professor of Intellectual Property and Technology Law at the School of Law of Queen's University Belfast, UK.

Christophe Geiger is Professor of Law at the Luiss Guido Carli University in Rome, Italy.

Andres Guadamuz is Reader in Intellectual Property Law at University of Sussex, UK.

Stavroula Karapapa is Professor of Intellectual Property and Information Law at the University of Essex, UK.

Irini Stamatoudi is Professor of Intellectual Property Law at University of Nicosia, Cyprus.

How to cite this article: Bonadio, E., Frosio, G., Geiger, C., Guadamuz, A., Karapapa, S., & Stamatoudi, I. (2026). 'Preserving Balance in the EU Digital Single Market: How *Like Company* Could Reframe Copyright and Innovation in the Generative AI Era.' *Journal of World Intellectual Property*, 1–39.

<https://doi.org/10.1111/jwip.70032>