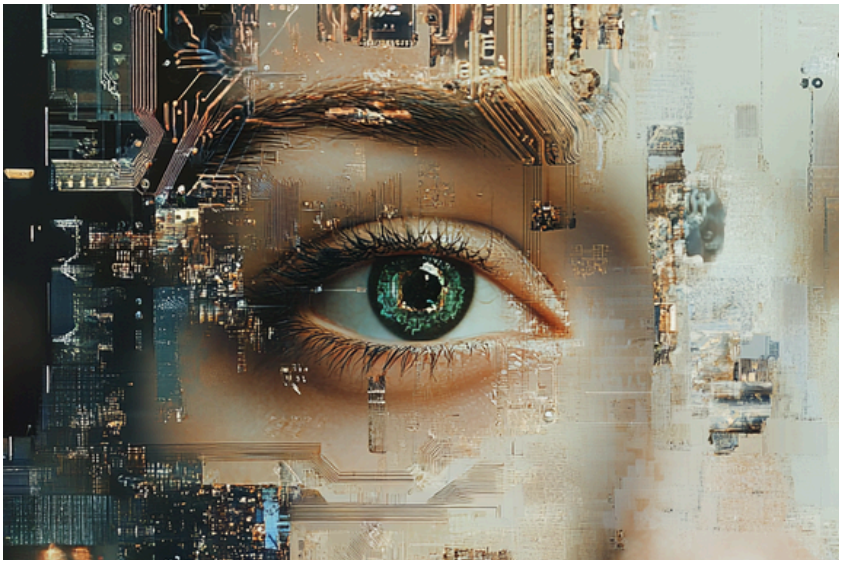
Chatbots, the Press and the Limits of EU Copyright.

The Opinion in Like Company v Google

In July 2023, a small Hungarian news portal published a story about a singer who wanted to bring freshwater dolphins to Lake Balaton. A user later asked Google's Gemini chatbot to summarise it. Gemini added a purpose-built pool on the Tihany peninsula, measuring 25 meters by 15 meters. None of that was in the article. Google calls such inventions "hallucinations".



UROŠ RAJIĆ
SENIOR ASSOCIATE



The portal's owner, Like Company, sued Google in Budapest, and the Hungarian court referred four questions to the Court of Justice of the EU. On November 3, 2026, Advocate General Maciej Szpunar delivered his opinion in Case C-250/25, the first generative AI copyright case to reach Luxembourg. It does not bind the Court, but the judges follow their Advocates General in most cases.

The dolphins are a good story, and the case is small. The lines the opinion draws are not. They matter to publishers, to companies that use AI and to those that build it, and they expose a problem the opinion only touches: where, exactly, does an act of artificial intelligence take place?

Training and retrieval

A chatbot can use text in two very different ways, and the case turns on keeping them apart.

Training is the studying phase. Before launch, a model reads an enormous library of text and learns how language works, the way a student who has read a thousand newspapers learns how journalism sounds. Copies are made along the way, but the finished model is meant to carry patterns, not articles.

Retrieval is the open-book exam. Known as retrieval-augmented generation (RAG), or "grounding", it means that the chatbot, once asked, searches the live web, fetches the relevant pages, and writes its answer from them. One specific article is copied, at the moment of the question, to answer that question.

The law has a tool for each. The text and data mining (TDM) exception allows copying for automated analysis if access was lawful and the rights holder has not opted out in machine-readable form. The press publishers' right (Article 15 of the 2019 Digital Single Market Directive) is not copyright: it protects publishers' investment and covers the "online use" of press publications by online services. Like Company relied on that right alone.

What the Advocate General said

Training is not what this case is about. The article appeared in July 2023, and Gemini 1.0 was released in December 2023. Training takes months, so it is implausible that the article was in it. Gemini most likely fetched the page live, as the user had named the website. And training is not an "online use": it happens in the background, for no one in particular, so the press publishers' right does not apply to it. The AG therefore advises the Court to leave the training questions unanswered. He answers them anyway: training involves copying, but as a rule the TDM exception covers it, provided access was lawful and the rightsholder has not opted out.

A robots.txt file, or Google's AI-specific variant Google-Extended, counts as an opt-out.

Chatbot answers are the new part. A chatbot that retrieves an article and reuses it can infringe the publisher's rights in two ways:

1. Unlawful access. It circumvented a paywall (a lock that lets only subscribers in) or a technical block meant to prevent bots.
2. Recognizable substitution. The answer reproduces the article's form of expression so closely that the reader can skip the original.

The second rests on a distinction every journalist knows: facts are free, form is protected. The right covers how an article is arranged, structured and phrased, not what it reports. Paraphrase is no safe harbor, but a summary of the same facts, or an answer built from several sources, is not a reproduction, even if it costs the publisher a click. The more distinctive the piece, the easier it is to recognize, so analysis is better protected than a bare news item.

Responsibility lies primarily with the user who asked. It shifts to the provider where the infringing use is inherent to its business model, or where the provider has not taken the measures expected of a reasonably diligent operator, such as respecting access restrictions and filtering outputs. Nor does an open website imply consent: publishers skip paywalls to earn advertising revenue, not to give permission.

And the dolphins? Subject to the Budapest court's findings, Like Company probably loses. It had no paywall or bot block; the answer shares the facts but hardly the wording; the article is a plain, informative report; and Google appears to have reasonable safeguards.

Munich against Luxembourg: the GEMA judgments

The opinion did not arrive in a vacuum. GEMA, the German collecting society for composers and lyricists, has won two first-instance cases against AI providers before the same chamber of the Munich Regional Court.

In GEMA v OpenAI (November 11, 2025), ChatGPT returned the lyrics of nine German songs almost word-for-word in response to simple prompts. The court held that the lyrics were "memorized" in the model, that memorization is itself a reproduction outside the TDM exception, and that OpenAI, not the users, was responsible. In GEMA v Suno (July 31, 2026), the court ruled the same way on six compositions that Suno had taken from YouTube by "stream-ripping", circumventing YouTube's protection against downloading. The works were found reproducibly stored in models hosted on servers in Germany. Neither judgment is final, and the OpenAI case is on appeal.

The AG sees memorization differently. "A trained LLM is not a database," he writes: a model holds statistical relationships, not texts, and its ability to reproduce a work when prompted does not show that a copy is stored inside it. He calls that inference far-fetched. Munich sees a cupboard with the cakes in it. The AG sees a chef who has tasted thousands of cakes and can bake in their style.

Where does an act happen?

Copyright is territorial. There is no single global copyright, only a bundle of national ones, each governing acts within its borders. That worked when a copy was made in one place by one person. It strains when a model is trained in one country, deployed from a second, and used in a third.

Gemini was trained mainly in the United States. Google argued that only the country hosting the main computing infrastructure can regulate training. Most governments that intervened said that EU law applies wherever training happens, once the model is offered in the EU. Germany proposed a narrower rule: EU law for scraping from servers located in the Union.

The AG is cautious. Offering a model on the EU market does not, in itself, justify applying EU copyright law to training done abroad, because the harm in the EU is an indirect and diffuse consequence.

He does not exclude applying it to scraping from EU servers and prefers that the legislator settle the question. Chatbot answers are different: a chatbot offered to EU users produces immediate, substantial effects there, wherever its servers are.

Munich went further. In Suno, the court took jurisdiction over the training in the United States, assessed US fair use itself and found it inapplicable. The AI Act pulls the same way: Recital 106 says that providers placing general-purpose models on the Union market must respect EU copyright rules, including rights reservations, regardless of where the training took place. So a national court reached across the Atlantic, while the Union's own advocate general counseled restraint and asked for legislation. Each position is defensible. Neither is a framework.

When software crosses borders by itself

The opinion concerns a chatbot that fetches a page when asked. The next step is agents: systems that work through a task in several steps, browsing, collecting, comparing and drafting, sometimes passing parts of it to other agents and tools. The AG does not address them. What follows extends his logic.

Take an employee in Belgrade who asks an agent to follow German and Austrian media and send a daily briefing to colleagues in Belgrade and Vienna. The agent runs in a US cloud, relies on a model trained in a third country, and reads a publisher's website hosted in Frankfurt. A temporary copy is made, a summary is written, a message is sent. Which act took place where? And who acted: the employee, the agent's provider, the model's developer, the company that deployed it?

The territorial rule assumes that an act can be pinned to a single place and a single actor. Agents scatter acts across multiple locations within milliseconds, with no human involved at each step. The AG concedes that a server's location is an imperfect connecting factor in the age of the cloud. His test of responsibility is more promising, because it turns on design and diligence: did the provider build the system to respect access restrictions and avoid recognizable reproduction? That shifts the law from where an act occurred to who designed the system and for which market. It is the right instinct, but not yet a rule on applicable law.

These are the classic questions of private international law: which court has jurisdiction, which law applies, and whether an order can travel.1 The current answers were built for a world in which an infringement had a location. Meanwhile, courts improvise with server locations, market effects, and special forum rules, introducing unpredictability, forum shopping, and conflicting remedies. The ingredients of a framework are visible:

- connecting factors based on the market a service targets, rather than on where a server sits;
- a clear way to identify the responsible party in a chain of agents, models and tools;
- a shared reading of machine-readable rights reservations, so that an opt-out means the same in every jurisdiction;
- mutual recognition of remedies.

No national court and no single EU judgment can supply this. The international legal community, through forums such as the Hague Conference on Private International Law or WIPO, should start now. At the same time, the case law is still young enough to be shaped.

What it means for Serbia

Serbia is not an EU member, so the Court's judgment will not be binding on its courts. EU law can still reach a Serbian business in three ways. Through the market: a company that offers a chatbot or news-digest product to users in the EU produces effects there, which, on the AG's reasoning, is enough for the output. Through the source: a company whose bots collect articles from servers located in the EU cannot assume that training done at home keeps those activities out of the EU's reach. Through the supply chain: providers of general-purpose models placed on the EU market must respect EU copyright rules and rights reservations, and those obligations can be expected to be passed on by contract to the businesses that build on their models.

Home law adds a twist. Serbia has not yet transposed the 2019 DSM Directive, the source of both the TDM exception and the press publishers' right. Its current copyright law has no TDM exception, and the draft Copyright and Related Rights Act put to consultation in March 2026 does not introduce one either. Copying protected works for data mining generally requires the rightsholder's permission and has no domestic counterpart. A Serbian developer may face a stricter environment at home than an EU competitor does in Munich or Luxembourg. Montenegro, by contrast, amended its law in October 2024 to introduce both the TDM exceptions and the press publishers' right, so the region is already diverging.

In practical terms, the opinion sorts the players as follows:

- Publishers. A paywall or a working bot block protects against retrieval attempts that circumvent it, and a machine-readable opt-out matters for training. An open page is protected only against answers that recognizably substitute for it. How a publisher guards its front door is now a legal decision as well as a commercial one, and blocking the bots may cost visibility.
- Companies that use AI tools. The user who directs a chatbot is treated as the primary actor, and private, non-commercial users fall outside the right. Using a general chatbot to understand the news looks unproblematic. A product whose value lies in paraphrasing individual articles does not, and the AG names that kind of service as conduct inherent to a provider's business model.
- Companies that build AI services. Lawful access, respect for access restrictions and filters that keep answers sufficiently removed from any specific source are what a reasonably diligent operator does. Where they are in place, occasional slips should not create liability. Where they are absent, it should.

What to watch

The Court's judgment will be delivered on a date not yet set. The judges may follow the AG, trim him, or go their own way, above all on training, which he asked them not to decide, and on memorization. Around it, the picture keeps filling in: the OpenAI judgment is on appeal in Munich, and further generative AI references are already pending in Luxembourg.2

The AG ends modestly: copyright, he writes, will not solve every problem that chatbots raise. One of those problems is not about copyright at all, but about whose law governs software that has no idea where it is. The dolphins never reached Lake Balaton, yet the story about them was retold across several countries in a single afternoon. The sooner the international community works on an answer, the less of it will have to be improvised in courtrooms.

GECIĆ | LAW

BELGRADE | BRUSSELS

WWW.GECICLAW.COM