AI GOVERNANCE

A company answers for what its AI publishes.

Regional Court of Munich I, 2026

AI Overviews are Google's own statements, not search results.

#AIassisted using HAIA Ecosystem

A Munich Court Rejected the AI Disclaimer Defense.

A Frontier AI Company Answers for What It Publishes.

A Munich court held that Google's AI Overviews are Google's own statements, not merely search results, so Google answers when those statements are false. The disclaimer did not transfer the duty, which is the same accountability principle New York's Part 161 reaches from the other end.

Basil C. Puglisi, MPA

A Human-AI Collaboration

The frontier AI companies all attach a version of the same disclaimer to their products. Google's reads, "Gemini is AI and can make mistakes." They ship the product, add that line, and claim that anything it publishes is the user's responsibility to verify. That was their shield. A regional court in Munich, Germany just rejected that position.

On May 28, 2026, the Regional Court of Munich I issued a judgment in preliminary-injunction proceedings against Google in case 26 O 869/26. The court held that Google's AI Overviews, the summaries the system generates above ordinary search links, are not merely search results. They are Google's own statements. When those statements are false, Google answers for them.

That holding is narrower than the slogans around it, which is why it matters more. The court did not rule that anyone who uses an AI tool owns everything it produces. It ruled something more precise. When a company deploys a feature that generates independent, self-contained factual claims in its own voice, and presents them to the public as answers, those claims are attributable to the company that deployed them.

What the court actually held

The facts are unflattering, and they begin with two Munich-based publishers who found that when users searched their names, Google's AI Overviews described them in connection with scams, subscription traps, and dubious business practices. None of it was true, and none of the cited sources said it. The system blended information about unrelated bad actors into a confident summary about the publishers, including a statement effectively saying that the company was known for dubious business practices and often perceived as a scam. The publishers sent a cease-and-desist letter, and when Google did not adequately fix the output, they went to court and obtained an injunction.

The reasoning the court used, more than the result, is what matters for everyone else. German precedent has long treated search engines as intermediaries, liable only indirectly because they make third-party content findable, on the theory that forcing them to police every result would break search itself. The Munich court found that this shield does not reach AI Overviews. A traditional result points to an outside page with a title, a snippet, and a link. An AI Overview evaluates and combines content from several sources, then rewrites it, in the court's words, "in its own words and according to its own structure," into a new substantive statement. To an average user, that reads as information from Google, not a forwarding of someone else's content. Because Google built, trained, and controls the system, and because only Google can compare the generated summary against the sources it drew from, the court treated Google as a direct party to the false statement rather than a neutral conduit.

Munich was not the first German court to take up the question of AI Overviews. In September 2025, the Regional Court of Frankfurt heard a competition-law challenge from a plastic surgeon whose procedure an AI Overview had described, and it dismissed the request for an injunction, finding that the specific summary, read in full context, was not actually false. The principle the court recognized was the part that mattered. The court accepted that an objectively incorrect AI Overview could amount to an unreasonable obstruction under competition law, that AI Overviews can materially affect competition by pulling clicks away from the sites they

summarize, and that liability for such a feature is not excluded in principle. Frankfurt opened the door on the law and closed it on the facts. Munich walked through that door on facts where the statements were false.

Two limits belong in plain view before the ruling is read too broadly. The decision is a preliminary injunction, not a final or appellate ruling, and Google has said it is reviewing a decision it considers not yet final and intends to appeal. It is also one court in one jurisdiction, and Germany is a civil-law system where a single ruling does not bind other courts the way precedent does in the common-law world. What the Munich court did was still concrete. It prohibited specific false statements, refused to limit the injunction to Germany under European civil-jurisdiction rules, required Google to carry roughly 80 percent of the costs, and placed standard German enforcement behind the order.

The tool did not speak, the company did

Once the specifics of German procedure are set aside, the governance lesson is direct. The court located accountability for the AI output in a named, responsible party, and refused to let the tool absorb the blame. The AI generated the summary, but the AI does not answer for it. Google does, because Google chose to deploy the system and present its output as an answer.

Underneath the procedure sits the question that separates real governance from its appearance. Ethical AI asks whether something should be done. Responsible AI asks for internal controls and lets the machine validate the machine, with no named party answering personally when it fails. AI Governance asks harder questions: who decides, by what authority, and who answers. It requires a named party whose accountability survives audit. A disclaimer lives in the ethics layer at best. It expresses a value, that outputs may be wrong, without attaching anyone to the consequence. The Munich court treated that as insufficient and supplied the missing accountability layer: a party who answers, enforced through civil liability.

Why the disclaimer cannot cure the harm

The deeper point in the ruling is about how these answers are received. An AI Overview is written in a confident, structured, authoritative voice. It presents a self-contained answer with no visible signal that it might be wrong or incomplete, creating an illusion of completeness. The harm happens at that confident output, before a reader ever consults a source, the way a misleading headline can mislead even when the article beneath it is accurate.

The behavior data on how people read these answers makes the disclaimer thinner still. A July 2025 analysis by the Pew Research Center found that users clicked a link inside an AI summary in roughly 1 percent of visits to pages that carried one. The instruction to go verify assumes a step very few users take. Google argued that users know AI output should not be blindly trusted and can check the linked sources, and the court answered that the ability to falsify a statement through further research does not, on its own, release the party that made it.

The court added a point that reaches beyond search and into every AI answer feature. An AI Overview is, in the court's words, by no means absolutely necessary, since traditional search already organizes the web for people. The summary is an optional layer Google chooses to

inject, and because the deployment is a choice rather than a necessity, Google carries the full risk of what the added layer states. The same logic can apply to any chatbot or answer engine that a company elects to put in front of the public.

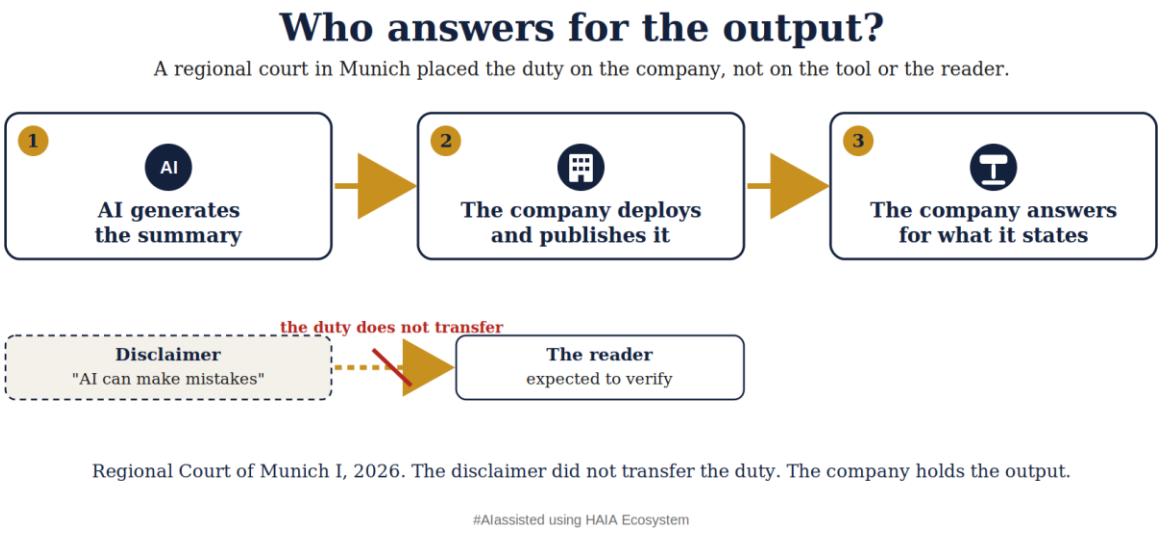

The other defenses the court rejected

Google's defenses did not stop at the disclaimer, and neither did the court. Google pointed to the small label that marks the feature as created with AI, the kind of notice that is supposed to put readers on guard. The court held that the note does not change the attribution. An untrue factual claim is governed by the ordinary law of false statements, and a company cannot hide behind its own automated process by tagging the output as machine-made. The label describes how the statement was produced. It does not change who answers for it.

The court narrowed a second exit: the claim that AI output deserves the same weight as human expression. Across recent cases, companies have begun to argue that what their systems generate is expression entitled to strong protection. The Munich court weighed it differently. German and European law balance free expression against personality rights, and the court gave Google's interest reduced weight, treating the false claims as algorithm-generated and as primarily an expression of Google's commercial activity, then setting that against the publishers' interest in not being falsely described. In that balance, the commercial interest gave way. An output a company ships as part of its product remains tied to the company, and in this ruling it received less weight than expression made for its own sake.

Google also asked to be treated as a hosting provider under the Digital Services Act, shielded as a passive host for third-party content, or alternatively as a search engine subject only to notice and takedown. The court refused both. It read Article 6(4) of the Digital Services Act to preserve a court's power to order a violation stopped through a civil injunction, and it found that a system generating its own statements is not the passive host the shield was written for. Taken together, the court closed the exits Google offered: the disclaimer, the label, the expression claim, and the hosting shield. What remained was the company holding the output.

The same principle New York wrote from the other end

Six weeks earlier and an ocean away, New York's court system reached the same destination from the opposite direction. Part 161, the New York rule that took effect on June 1, 2026, lets lawyers use AI to prepare court papers without disclosing that a tool was involved, while leaving accountability on the signature. The lawyer who signs owns every fact the filing contains, including any a model invented.

Set the two rulings side by side, and the practical symmetry is hard to miss. New York fixed accountability on the human who submits the work, and Munich fixed it on the company that publishes the work. Neither let the tool carry the responsibility, and neither accepted a disclaimer or a notice as a substitute for a party who answers. This is the core of what Checkpoint-Based Governance describes and what the Named Human Test asks of any AI-assisted process: whether a named, accountable party can approve, modify, or reject the output, and whether that accountability survives audit. New York supplied the answer at the signature. Munich supplied it at the point of publication. The instrument differs, professional duty in one case and civil liability in the other, but the governance principle is the same. Accountability attaches to a party, never to the software, and it cannot be disclaimed away.

The cost of ungoverned output

Beyond the principle it establishes, the ruling also puts a number on the gap. Under standard German enforcement, repeating the prohibited statements can draw an administrative fine of up to 250,000 euros per violation, with coercive detention against responsible management available for noncompliance, which is a compliance measure under German procedure rather than a criminal sentence. Google must also bear 80 percent of the litigation costs. The price of an output the company did not adequately govern is now visible on a court order.

That visibility is the mechanism, because insurers and sophisticated buyers operate on measurable risk. They extend trust to organizations that can document verification, and they price against organizations whose AI use remains opaque. The pattern underneath is familiar. Deployment incentives push capability into the market faster than control is built, and the disclaimer is what stands in for the control that was skipped. Courts and insurers are now making skipped controls expensive, which is one of the few forces that reliably closes the gap between what a company can ship and what it is willing to stand behind.

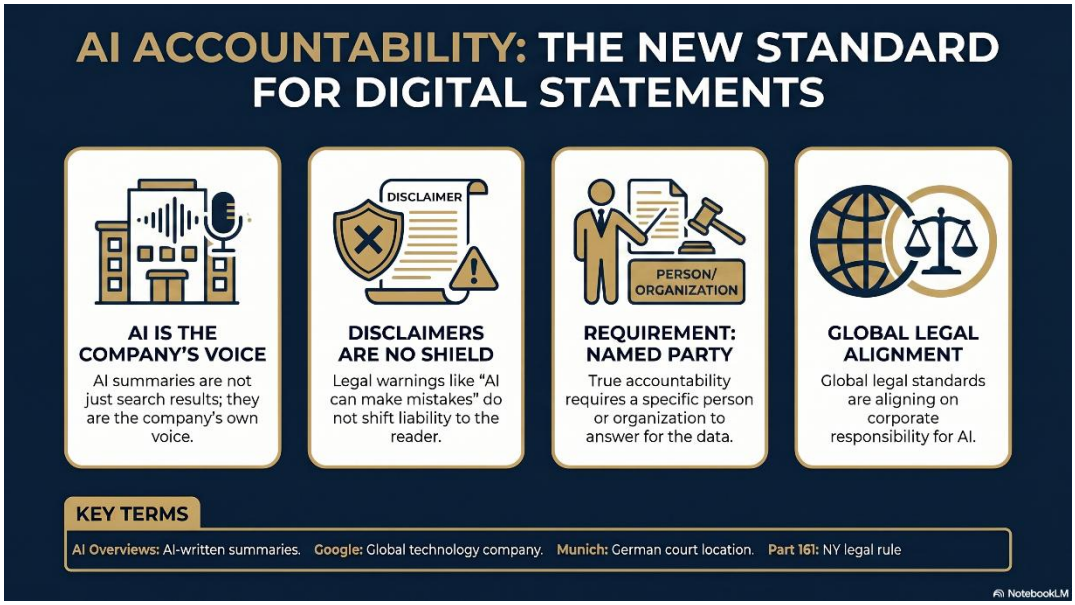

Why this outlasts the next feature

A rule that depends on labeling AI ages badly because the line between human and machine output is dissolving as these features move into ordinary software. A rule that fixes accountability on the party that deploys the output does not depend on that line. It asks one durable question that will read the same way in a decade: did a responsible party stand behind this statement before it reached the public?

None of this is settled: the Munich decision is preliminary and under appeal, the German mechanism does not transfer cleanly to other legal systems, and the United States analysis in particular is unsettled, since Section 230 shields platforms for third-party content but may not shield AI output that a court treats as the platform's own. In Europe, the EU AI Act is moving in the same direction from a different angle, attaching transparency and human-oversight obligations to AI systems as it phases in, though it sits alongside this ruling rather than driving it. What travels across all of those uncertainties is the question itself. As AI becomes a default layer rather than a separate tool, the systems that endure will be the ones that can name the party who answers for the output, and a disclaimer will not be that party. The tool can generate the answer. It cannot answer for it.

Sources

Regional Court of Munich I (Landgericht München I). (2026, May 28). Case 26 O 869/26, judgment in preliminary-injunction proceedings. Redacted text hosted by The Decoder, no official court URL published. https://the-decoder.com/landmark-german-ruling-declares-googles-ai-overviews-are-googles-own-words-and-makes-it-liable-for-false-answers/

The Decoder. (2026, June 9). Landmark German ruling declares Google's AI Overviews are Google's own words and makes it liable for false answers. https://the-decoder.com/landmark-german-ruling-declares-googles-ai-overviews-are-googles-own-words-and-makes-it-liable-for-false-answers/

heise online. (2026). LG Munich I: Google ordered to pay for false statements in AI summaries. https://www.heise.de/en/news/LG-Munich-I-Google-ordered-to-pay-for-false-statements-in-AI-summaries-11327217.html

Reuters. (2026, June 12). Google to challenge German court ruling assigning liability for AI Overviews' false claims. https://www.reuters.com/world/google-appeal-german-court-ruling-assigning-liability-ai-overviews-false-claims-2026-06-12/

Google. (2026). Gemini interface disclaimer ("Gemini is AI and can make mistakes"). https://gemini.google.com/

Pew Research Center. (2025, July 22). Google users are less likely to click on links when an AI summary appears in the results. https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/

Regional Court of Frankfurt am Main (Landgericht Frankfurt). (2025, September 10). Case 2-06 O 271/25, decision on AI Overview liability.

ppc.land. (2025). German court dismisses surgeon's AI Overview lawsuit but confirms Google can be liable for false information. https://ppc.land/german-court-dismisses-surgeons-ai-overview-lawsuit-but-confirms-google-can-be-liable-for-false-information/

technology.org. (2026, June 12). German court: Google AI Overviews liable. https://www.technology.org/2026/06/12/german-court-google-ai-overviews-liable/

New York State Unified Court System. (2026). Part 161, Use of Artificial Intelligence Technology (Rules of the Chief Administrator, added March 25, 2026, effective June 1, 2026). https://www.nycourts.gov/rules/part-161-use-artificial-intelligence-technology

Information in this article is current as of June 2026 and should be re-verified, since the Munich injunction is preliminary and Google's appeal and any later proceedings will change its practical weight.

The Other AI: Audio Briefings on Augmented Intelligence and AI Governance

Spotify; Apple Podcasts; Amazon Music; YouTube Playlist;

#AIassisted using the HAIA Ecosystem