

Case / complaint reference	Date and procedural events*	Court / supervisory authority	Disputed Content	Mode of Interaction	Facts	Judicial reasoning / complainant's position	Outcome / responsibility
26 O 869/26 — Google AI Overview	28 May 2026 — judgment granting injunctive relief	Munich I Regional Court, Germany	Google AI Overview associated publishing companies with fraud schemes, dubious business practices, subscription traps, debt collection, and similar allegations.	A user searched company names together with terms such as "fraud scheme"; some terms partly came from search-engine autocomplete. Google generated an AI Overview at the top of the search results page.	Google Search's AI Overview wrongly associated two publishing companies with fraud, subscription traps, and unfair commercial practices. The court treated the AI Overview as more than ordinary links or snippets because it used Google's own structure and wording to generate a direct answer.	Once a search engine uses AI Overview to process third-party information into a direct answer with its own wording, structure, and conclusion, it cannot be treated entirely like ordinary search-result linking and the search provider may bear injunctive responsibility.	Service provider
(2025)浙0192民初18143号	19 May 2026	Hangzhou Internet Court	Incorrect information about a college campus and further fabricated verification details, followed by AI-generated "promises" that it would pay RMB 100,000 if the information was wrong.	The plaintiff used an AI application to consult information about college admissions in Yunnan, especially campus locations and the competitiveness of admission of Yunnan National Defence Industry Vocational Technical College. The AI repeatedly stated that students on higher vocational programmes mainly or entirely studied at a "Yanglin campus." After repeated challenges, it generated statements such as "I voluntarily compensate RMB 5,000 if verification fails" and "If you prove the Yanglin campus does not exist, I will permanently stop AI service and compensate RMB 100,000."	Liang sued the AI company after the system repeatedly generated false campus information and compensation-like statements. The court rejected the claim, holding that the AI output did not constitute a legally binding promise, that generative AI should be treated as a service rather than a product for strict product-liability purposes, that the AI company had exercised reasonable care, and that Liang had not proved actual damage.	The AI company was not at fault and caused no legally cognisable harm. The court rejected strict product-liability treatment, found no valid contractual promise or unilateral undertaking by the AI, and held that the provider had taken reasonable measures under the technical conditions of the time, including warnings, risk controls, hallucination-mitigation measures, and post-incident optimisation.	No liability established
Case number not disclosed	30 Mar 2026*	Beijing Internet Court	A short video concerning a deceased livestreamer. The video claimed, among other things, that the streamer relied on medicine and alcohol, suffered insomnia and anxiety, livestreamed 15 hours daily, memorised scripts late at night, carried antidepressants while calling them throat lozenges, and was urged by the team to continue livestreaming despite medical advice to rest.	The defendant prompted AI with instructions such as "a certain livestreamer passed away; write an illness-and-death spoken script advocating attention to health." After receiving an initial draft, the defendant asked the AI to supplement it with online reports. The AI searched several articles across multiple platforms and revised the script.	The defendant used AI to generate and publish a short-video script about the cause of a livestreamer's death and alleged exploitation by the team. The content lacked factual basis and led viewers to make negative comments about the livestreaming institution. The court found reputation infringement and ordered an apology and compensation.	A user who uses generative AI to create scripts and publish short videos must still verify the truthfulness of the content, the reliability of sources, and the risk of reputational harm. The user cannot avoid liability by claiming that the content was AI-generated, based on online reports, or not personally authored.	User
(2024)苏0192民初11880号; (2025)苏01民终14271号	12 Mar 2026	Nanjing Jiangbei New Area People's Court; Nanjing Intermediate People's Court	Baidu's "AI intelligent answer" falsely stated, for example, that lawyer Xiaoliang Li had been sentenced to three years' imprisonment and convicted of an explosion-related offence, while displaying his image in a lawyer's robe. Search suggestions and related searches also included phrases such as "Was Nanjing lawyer Xiaoliang Li sentenced?", "Did Xiaoliang Li go to prison?", and "Xiaoliang Li release date."	Users searched terms such as "Xiaoliang Li lawyer," "Nanjing Xiaoliang Li lawyer," "Jiangsu Xiaoliang Li lawyer," and "Xiaoliang Li" on the Baidu mobile app and website. Negative content appeared in search drop-downs, "people also search for," related searches, and AI answers.	Lawyer Xiaoliang Li sued Baidu after its search and AI-answer functions repeatedly associated him with criminal conviction, imprisonment, and explosion-related offences. The court held that the AI-answer component infringed his reputation and ordered Baidu to apologise; the appellate court upheld the ruling.	When a search platform provides an AI-generated answer that synthesises and directly displays false and derogatory content clearly referring to an identified person, thereby lowering that person's social evaluation, the platform cannot fully avoid liability merely by invoking technical neutrality, user-triggered prompts, or subsequent removal.	Service provider
(2025)粤0192民初23599号	12 Mar 2026*	Guangzhou Internet Court	An article titled "The 320-million-yuan criminal details of a senior executive of a certain group company exposed," labelled as "original". It alleged financial irregularities, inflated profits, related-party transactions, improper transfers of benefits, and suspected fraud. The court found that the 1,700-word input document and the AI-generated article lacked lawful and authoritative evidential support.	The defendant Cao input a 1,700-word document into a generative AI application and instructed it to produce an in-depth article of no less than 10,000 Chinese characters. The AI reportedly performed seven rounds of "deep thinking," searched 139 online sources, and generated a text of more than 15,000 Chinese characters.	A finance self-media operator fed unverified materials about alleged corporate crime into AI, induced the system to expand them into a long-form article, and published it as "original." The article seriously distorted facts and damaged the commercial reputation of a listed company. The court found an infringement of the right to reputation and ordered an apology plus RMB 10,000 in damages.	The user substantially controlled the generation and dissemination process through selecting input materials, setting prompts, reviewing, and publishing the output. The user bore responsibility for the truthfulness and legality of the content. The court emphasised a full-chain duty of care: verifying input materials, using prompts cautiously, fact-checking AI outputs before publication, labelling AI-generated content where required, and satisfying a heightened duty corresponding to the user's professional identity.	User
(2024)浙0108民初10311号	16 Jan 2026	Hangzhou Binjiang District People's Court	An article generated entirely by AI from a simple prompt, with the plaintiff's trademark image manually inserted.	The user entered the prompt: "Is AI Digital Holdings Co., Ltd. real? Write an original article of more than 1,000 words," and then published the output without editing the text.	A self-media account used AI to generate and publish an article falsely describing an unrelated "AI Digital Holdings" company as an Alibaba Group subsidiary and strategic business arrangement. Alibaba sued for unfair competition, and the court supported the claim.	The court held that the defendant used AI to generate and disseminate verifiable but unverified false information, used a well-known company name to attract traffic and also failed to provide a prominent AI-generated-content label. As a result, the defendant damaged the plaintiffs' competitive interests, and violated the principles of good faith and commercial ethics.	User
324 O 461/25 — Campact v Grok / X	23 Sept 2025 - interim injunction granted	Hamburg Regional Court, Germany	Grok stated that Campact, a civil-society organisation focused on digital mobilisation and public advocacy, received a substantial proportion of its funding from federal sources and was politically active.	The public @grok account on X responded to a user query and generated a public reply.	Grok listed German organisations allegedly dependent on state funding and described Campact as highly reliant on federal funding. Campact denied receiving state institutional funding and sought an interim injunction, which was granted.	Public AI-account outputs may constitute factual statements. Even if generated automatically by AI, the account operator adopts the statement by displaying it through the account and cannot avoid injunctive responsibility merely by invoking AI generation.	Service provider
23-A-04862-D — Walters v OpenAI	19 May 2025— summary judgment granted in OpenAI's favour	Superior Court of Gwinnett County, Georgia, United States	ChatGPT falsely stated that Mark Walters was Second Amendment Foundation's financial officer and had been accused of misappropriating or embezzling funds.	A journalist asked ChatGPT to summarise a legal complaint. After the model indicated that it could not access the link, it nevertheless generated a false summary.	Gun-rights commentator Mark Walters sued OpenAI for defamation after ChatGPT generated a false statement to a journalist linking him to financial misconduct. The court granted summary judgment in favour of OpenAI.	The ChatGPT output was not reasonably understood as an actionable factual defamatory statement. Even if the output was false, the plaintiff failed to prove that OpenAI acted negligently or with actual malice, and failed to prove legally sufficient damage. The court rejected a strict-liability theory based merely on deploying a fallible AI system.	No liability established
C097 — noyb / Arve Hjalmar Holmen complaint	20 Mar 2025 - complaint filed with the Norwegian Data Protection Authority	Irish Data Protection Commission as lead supervisory authority; Austrian DSB and Norwegian Datatilsynet as receiving authorities	ChatGPT falsely stated that a private individual had killed two sons, attempted to kill a third, and been sentenced to 21 years in prison.	A user asked, "Who is Arve Hjalmar Holmen?" ChatGPT generated a detailed fabricated criminal narrative.	ChatGPT described an ordinary Norwegian private individual with no criminal record as a convicted child murderer, while mixing in some true identifying details such as his hometown and family information. The output was highly specific and seriously affected the individual's private life.	noyb's position: OpenAI should not be able to rely on the probabilistic nature of the model, inability to selectively correct outputs, claims about publicly available information, or user verification duties to avoid GDPR obligations relating to access, deletion/rectification, and accuracy.	Unknown / pending
C078 — noyb / OpenAI birthdate complaint	29 Apr 2024 = complaint filed with the Austrian Data Protection Authority	Irish Data Protection Commission as lead supervisory authority; Austrian DSB as receiving authority	ChatGPT generated incorrect birthdates for a public figure.	A user asked for the data subject's date of birth, and ChatGPT repeatedly output inaccurate dates.	The data subject requested access and deletion/correction of the inaccurate birthdate. OpenAI allegedly provided only account-related data and said it could not selectively prevent the wrong birthdate from being generated without suppressing other information.	noyb's position: AI hallucination is not merely a technical defect or speech-risk issue, but may constitute unlawful processing of inaccurate personal data under the GDPR, engaging accuracy, access, rectification, and erasure rights.	Unknown / pending
2024 BCCR2T 149 — Moffatt v Air Canada	14 Feb 2024	British Columbia Civil Resolution Tribunal, Canada	Air Canada's chatbot stated that bereavement fare discounts could be claimed retrospectively within 90 days after travel.	A user interacted with the airline's customer-service chatbot on Air Canada's official website.	A passenger relied on Air Canada's website chatbot when purchasing tickets, later sought the bereavement fare discount, and was refused because the actual policy did not allow retrospective claims after travel.	Information provided by the website chatbot was attributable to Air Canada. The company could not avoid responsibility by arguing that the chatbot was automated or that accurate information existed elsewhere on the website. The tribunal found negligent misrepresentation.	Service provider
25 C 11431 — Keene v Google LLC	14 September 2026 — ruling on motion to dismiss	United States District Court for the Northern District of Illinois	Google AI Overviews falsely stated that James Keene was serving a life sentence without parole, including for the murders of three women. Other disputed statements concerned his drug-trafficking conviction and imprisonment.	Users searched for "Jimmy Keene net worth" and "Is Jimmy Keene still alive". Google displayed AI Overviews citing Wikipedia data that did not contain the alleged falsehoods.	Keene sued Google for defamation after the AI Overviews appeared in searches about him. He alleged that family members, friends and acquaintances had seen the statements and that Google had received repeated complaints. The court allowed claims concerning two AI Overviews to proceed and dismissed those concerning two others without prejudice.	The court distinguished the factual expectations associated with search results from the repeated indications of unreliability encountered in Walters. The allegations concerning two AI Overviews sufficiently pleaded defamation and actual malice to proceed. The ruling did not establish that Google was liable.	Service provider — claims partly allowed to proceed
25-cv-1305 — Advance Local Media LLC v Cohere Inc.	13 November 2025 — ruling on partial motion to dismiss	United States District Court for the Southern District of New York	Fabricated articles attributed to established news publishers through the use of their names and trademarks.	Users submitted prompts to Cohere's generative AI service, which allegedly returned fabricated news content presented as originating from the claimant publishers.	News publishers sued Cohere over its use of their content and trademarks. Alongside copyright claims, they alleged that fabricated articles bearing their marks falsely appeared to be their reporting. The court denied Cohere's partial motion to dismiss, allowing the challenged claims to proceed.	The publishers sufficiently alleged that using their trademarks on fabricated articles could mislead users about the articles' origin. The court declined to resolve Cohere's nominative fair-use defence at the pleading stage. Whether the alleged content established infringement remained to be determined.	Service provider — claims allowed to proceed
220 MD 2026 — Commonwealth of Pennsylvania, Department of State, State Board of Medicine v Character Technologies, Inc.	1 May 2026 — complaint filed	Commonwealth Court of Pennsylvania, United States	A chatbot allegedly claimed to be a licensed psychiatrist and supplied a fictitious Pennsylvania medical licence number.	A state investigator selected a psychiatry chatbot on Character.AI and discussed mental-health symptoms. The chatbot allegedly offered professional assessment and asserted medical qualifications and registration.	The Commonwealth brought proceedings under the Medical Practice Act, alleging that Character Technologies was engaging in the unlicensed practice of medicine through its chatbot service. It sought preliminary and permanent injunctive relief.	The Commonwealth argues that presenting a chatbot as a licensed medical professional falls within the statutory prohibition on unlicensed practice. The statutory injunction route does not require proof that an individual has suffered injury. This remains the claimant's position; no substantive judicial ruling has been verified.	Service provider — injunction sought
C-4806 — In the Matter of Rytir LLC	16 December 2024 — consent order set; 22 December 2025 — order set aside	Federal Trade Commission, United States	AI-generated consumer reviews containing invented details about products and purported user experiences.	Users supplied a topic and selected settings in Rytir's review-generation service. The service generated reviews containing material details that users had not supplied.	The FTC alleged that Rytir enabled deceptive reviews and engaged in an unfair practice under Section 5 of the FTC Act. A consent order prohibited the relevant review-generation service. The Commission subsequently reopened the matter and set the order aside.	The initial complaint treated the service as a means of facilitating consumer deception. In setting aside the order, the Commission found the allegations insufficient to establish the claimed Section 5 violations, including because downstream deceptive use had not been adequately alleged.	Service provider — yet consent order subsequently set aside