

Prevalence of AI hallucinations in the legal context

By **Oksana Romanov**

Law360 Canada (October 7, 2025, 1:55 PM EDT) -- This is the second article in a series addressing AI hallucinations in the legal context.

When it comes to legal research and the use of generative AI, the amount of false information being generated is alarming. However, the data varies depending on the study, the AI tools analyzed and how AI hallucinations manifest. For example, the first Canadian case to use a citation hallucinated by ChatGPT was *Zhang v. Chen*, 2024 BCSC 285, at para. 38. It mentions the January 2024 study by Matthew Dahl et al., which looked at the prevalence of AI hallucinations in a systematic manner.

In the words of this decision:

The risks of using ChatGPT and other similar tools for legal purposes was recently quantified in a January 2024 study: Matthew Dahl et al., "Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models" (2024) arXIV:2401.01301. The study found that legal hallucinations are alarmingly prevalent, occurring between 69% of the time with ChatGPT 3.5 and 88% with Llama 2. It further found that large language models ("LLMs") often fail to correct a user's incorrect legal assumptions in a contrafactual question setup, and that LLMs cannot always predict, or do not always know, when they are producing legal hallucinations. The study states that "[t]aken together, these findings caution against the rapid and unsupervised integration of popular LLMs into legal tasks."



Oksana Romanov

A recent article published in *Mashable* points to an upward trend, citing a global collection of over 100 AI-fabricated cases available on display in a newly created database by Damien Charlotin. For example, regarding the country of origin of AI hallucination cases in the last two years,



Aparna Sinha: ISTOCKPHOTO.COM

Charlotin's database attributes nine errant cases to the U.K., 13 hallucinated cases to Canada, 17 fake cases to Australia and 123 non-existent legal authorities to the U.S. alone. At least, these were

the stats the last time I checked.

According to the recent study of large language models (LLMs) conducted by the Stanford Institute for Human-Centered AI (HAI), LLMs hallucinated at least 75 per cent of the time. In another HAI study, the researchers looked at general-purpose chatbots and found they hallucinated between 58 per cent and 82 per cent when running legal queries. When it comes to the specialized legal AI tools, researchers found that AI-driven legal research tools also hallucinate. As suggested by the HAI study, they produced incorrect information either more than 17 per cent of the time or more than 34 per cent of the time, depending on the AI-assisted research tool used.

For instance, a simple keyword search for "AI hallucinations" of the Canadian legal content on CanLII returned four relevant cases and four commentaries for the period of 2024-2025. When searching for "ChatGPT," the query returned 27 cases and 39 commentaries. The keyword search containing "generative AI" returned seven cases and 41 commentaries. These are publicly available reported cases and commentaries on the topic. There is no data currently available on AI hallucinations in unreported decisions or oral submissions.

Interested in learning more and experimenting? If you have access to one of the subscription-based legal research databases, such as Lexis+, try searching for similar content. Take note of your results.

In the next article, we will examine the seriousness and scope of errors found in a criminal law case. To reiterate a key point from my first article: lawyers must curate legal outputs rather than delegate everything to AI.

Oksana Romanov, BA (Hons), MA (Comm), JD with Distinction, is a sole practitioner practising criminal law as Law Office of Oksana Romanov.

The opinions expressed are those of the author(s) and do not necessarily reflect the views of the author's firm, its clients, Law360 Canada, LexisNexis Canada or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.

Interested in writing for us? To learn more about how you can add your voice to Law360 Canada, contact Analysis Editor Peter Carter at peter.carter@lexisnexis.ca or call 647-776-6740.