HAUSFELD
FOR THE CHALLENGE

AI FOR LEGAL PROFESSIONALS: HALLUCINATIONS

One of the most significant challenges currently facing legal professionals is the phenomenon of AI hallucinations which occur when an AI system generates information that sounds plausible but is entirely false, such as fictitious case law. Unlike human error, these inaccuracies are not due to carelessness. The second instalment in our series on AI for legal professionals, this article explores how hallucinations stem from the way AI models predict text based on patterns in data, without a true understanding of factual accuracy or legal validity.

What are hallucinations?

A hallucination in the context of AI refers to an output generated by a system (usually a large language model ("LLM"), such as ChatGPT) that is factually incorrect or entirely fabricated, despite in some instances sounding plausible. For instance, in the case of Harber v Commissioners for His Majesty's Revenue and Customs [2023] UKFTT 1007 (TC), a litigant in person cited nine supportive authorities in her appeal submission, none of which were genuine. They were in fact fabricated cases which had been hallucinated by a generative AI programme.

More recently, in R (on the application of Ayinde) v The London Borough of Haringey and Al-Haroun v Qatar National Bank [2025] EWHC 1383 (Admin), two cases were referred to the Divisional Court after lawyers were found to have submitted court documents citing fake legal authorities generated, or likely to have been generated, by AI tools. The President of the King's Bench Division of the High Court of Justice issued a stark warning to lawyers, explaining while generative AI tools can produce apparently coherent and plausible responses to prompts, "those coherent and plausible responses may turn out to be entirely incorrect". These fake authorities would have been the product of hallucination.

Hallucinations typically arise due to the probabilistic nature of how LLMs generate text. Rather than consulting a verified database in real time, LLMs produce responses based on patterns in the data they were trained on. If a user requests a case citation, for example, the AI:

1. processes the prompt as a string of "tokens" (units of text);
2. calculates, based on its training, what the most likely next tokens should be in response (rather than consulting a legal database) based on probability calculations; and
3. may not verify whether the statement it has produced is true.

The AI therefore returns what it calculates is a likely response, even if the case cited is fictitious. In this sense, one can think of an LLM like a very articulate person with no memory of specific books, just an idea of how stories are usually told. If asked to write a summary of a book they have not read, the person will produce something stylistically convincing but possibly full of inaccuracies. That is what it means for an AI to hallucinate.

Legal AI tools

To address the limitations of general purpose AI tools, particularly the risk of hallucinations and concerns over the accuracy of their output, many law firms are now developing or adopting purpose-built legal AI solutions tailored specifically for internal use. Unlike ChatGPT, for example, which is trained on internet-scale data, these in-house systems are trained or fine-tuned on curated, authoritative legal sources, such as the firm's own precedent documents and legal research databases. By grounding the AI's outputs in verified legal content, firms can significantly reduce the risk of the model generating inaccurate or fabricated information.

Comment

The implications of AI hallucinations in legal practice are profound. Accuracy in law is essential; a misstatement of fact or a fabricated authority can compromise a case, damage a lawyer's credibility and result in sanctions. These risks underscore the importance of human oversight, verification and professional judgment when using AI in legal practice.

Furthermore, internal legal AI tools, with safeguards that general tools lack, can be used to bridge the gap between technological innovation and the professional standards central to legal practice.

In the next edition in our series on AI for legal professionals, we will explore the current role of AI in the disclosure process and what the future might hold for lawyers willing to embrace new technology.

About Hausfeld's Digital Markets expertise
The Hausfeld Digital Markets Hub

Your point of contact

If you would like to discuss how we can assist you, please contact one of our following experts:

LUCY PERT
Partner, London
lpert@hausfeld.com

EDWARD NYMAN
Senior Associate, London
enyman@hausfeld.com

"
A serious advocate firm for claimants, the team regularly pursues actions across Europe, particularly in the financial, pharma, tech, and automotive sectors."
Chambers, 2025