CaseStrainer: Help with Hallucinations

By Barbara Engstrom, Executive Director

For the last several years we've been regularly reading about lawyers getting into hot water by citing AI-generated hallucinated cases to courts. In a recent federal court case in the Northern District of Mississippi, lawyers from both sides were removed from the *Withers v. City of Aberdeen* case and fined for citing AI-generated fake cases in court filings. While *Withers* seems to be the high-water mark for egregious AI-generated submissions so far, lawyers with egg on their face from playing fast and loose with AI has become all too common. In this column, I talk to **Jonathan Franklin** who, in his role as Digital Innovation Librarian at the UW Law School, developed **CaseStrainer**, a tool to help discover hallucinated cases in legal documents.

As more people adopt generative AI for legal work, it seems like hallucinations are really starting to multiply. Do you have a sense of the prevalence of AI generated hallucinations in court cases/ filings?

JF: We have data on when hallucinations (for the purposes of this discussion—citing cases that do not exist) that are caught due to a database (https://www.damiencharlotin.com/hallucinations/), but we need to do more research on raw briefs, both pre-AI (2021 and before) as well as post-ChatGPT (late-2022 and later) to compare if the rates are different and if so, by how much. The news of hallucinations is amplified by the legal media because of sanctions, juicy stories, and ultimately ethical failure. A future scholar could take a random sample of 10,000 state and federal court briefs and determine a level, but then there are motions and other documents that cite cases to consider – there's lots of research to be done.

While larger firms tend to have access to Lexis or Westlaw's vetted AI platforms, solo and small firm practitioners may find themselves getting into trouble by incorporating general AI into their workflow. What are the main issues that these attorneys should be aware of?

JF: I am still in the 'trust but verify' stage of using AI. Even Lexis and Westlaw will include citations to items that might not exist, they tend not to have a hyperlink to the underlying document the way they do for verified cases, but even then, it is not necessarily meeting an ethical standard to say, 'if there is a hyperlink, that means the case exists.' This is before we get to whether the quotation is actually in that case and/or whether that case stands for the proposition for which it is being cited. This is to say that we are not yet at the point where we can blindly say 'trust the automated system'.

You need to either know your area of law and the key cases down cold or do the legal research to check the underlying cases themselves. This is not to say it will always be this way, but for now it is. Just as you cannot delegate legal research to someone else if you are signing off on it, you cannot delegate legal research to AI. Those who say you have to go back and read every case cited in every brief might be a bit extreme, because there are cases we cite for certain propositions time and time again. For those cases, once you are confident that it says what you think it says, you can rely on that going forward, so long as you ensure it is still good law.

Beyond the court sanctioning you, or opposing counsel finding hallucinations in your filings, there is professional reputation. Due to the legal media coverage, you don't want to see your name associated with copy-pasting from AI output without using your professional judgment.

You've been very active in creating open-source, access to justice-oriented tools in your role as the Digital Innovation Librarian at the UW Law School. In that role you developed CaseStrainer, a free, open-source tool to help uncover hallucinated case citations. Could you explain in more detail what CaseStrainer does and what prompted you to develop it?

JF: When reading a brief someone else wrote, it is hard to know if a case name, citation, and year are correct just by looking at them. Citations have similar forms and case names and years often look believable. There is no way to determine if a case exists just by looking at it in a brief. This is unlike reading a headline that says the Patriots won the 2026 Superbowl and knowing it is wrong.

CaseStrainer is designed to find and separate the cases that exist from those that might not. No one can definitively say that a case does not exist, but you can get up to 99% chance of it not existing very quickly. CaseStrainer is a free tool to scan the cases cited in your document and confirm that they exist or flag them as something to confirm. It does it by breaking the citation into the three component parts, citation, year, and name. Picture this as a three-legged stool. It then tries to confirm that at least two of these three exist in the same case in a known and reliable database. It relies primarily on CourtListener from FreeLawProject but also uses six other sources. Case names can vary, but often unusual words in a case name carry over, so State v. Smith is very hard to verify based on name alone, but Franklin & Engstrom v. Taco Bell is pretty easy to verify. So, we know the name is very useful, but not sufficient. It needs to be paired with a case in the same year or with the citation. The problem with relying solely on the citation is that it might be correct, but point to a different case name and year, or it might be proprietary, like a WL or Lexis citation, where we cannot freely confirm all of them.

For the challenge of State v. Smith (2001) without a matching citation, we have another trick up our sleeve, the citation indicates the jurisdiction, so we can then try to find cases with that name and date within that jurisdiction, even if they don't have a common citation. If you are citing an Arkansas case when you mean to cite a Hawaii case with the same name and date, it will get flagged

In short, we need at least two legs of the three-legged stool to feel confident a cited case actually exists. Of course, all three is ideal.

Could you give a basic explanation of how to use CaseStrainer on the front-end and describe (for the tech un-sophisticates among us) what is going on behind the CaseStrainer curtain?

JF: When you go to the site, you can paste text, upload a document, or enter a URL. It will then do all the things I mentioned above for every case you cite and generate a report. The report will tell you which cases are verified with two or three of the data points and which need further work to confirm. Sometimes these are hallucinations. Sometimes they are cases not available in any free database. The need to verify means that instead of having to verify 100 cases cited, you only need to check three. Clearly, you need to make sure all the cases cited stand for what they are being cited for and that quotations are accurate, but that is beyond this tool and might be commonly available in a year or two. My goal was to do the best I could with the tools I had without relying on AI to do the final

check, as that seemed to defeat the purpose by relying on AI to check AI. Who knows though, AI checking AI might be standard practice in a few years.

CaseStrainer was developed with pro se litigants in mind, but how could solo and small firm attorneys who don't have access to pricey AI platforms on Westlaw and Lexis make use of it? Any little hacks or best practices that yield better results?

JF: Here are few ways I have used it.

1. Upload opposing counsel's briefs when you receive them, just to make sure the cases they cite are legit.

2. If you want to read a case, just paste the citation into the text box and use the link to the case to read it - A great way to quickly find cases on the web.

3. Courts are becoming more skeptical about cited authorities. You can use CaseStrainer to create an appendix with links to all the cases cited, so the court and opposing counsel can have confidence in it. At some point, I hope that submission of the CaseStrainer report with the brief will give everyone in the process confidence that the cases you cited exist by telling them where they can find and read them.

4. If you use paid services, you can use CaseStrainer first and then search for just the unverified cases, saving yourself time.

From your answer above, I could see a tool like CaseStrainer incorporated into court systems to help weed out hallucinations before anything is filed. How would you envision that working?

JF: Much of legal work is about workflows. If the court system has a portal where you upload documents, I could imagine that one of the steps after you upload your document is that each document would be run through a tool like CaseStrainer and the report generated could be sent to all parties. To me, this would have two benefits, alerting everyone to cases that need to be confirmed and also, ideally, having people use the tool before submitting the document to avoid embarrassment and the need to explain yourself.

Take CaseStrainer for a Spin

Thanks Jonathan, for telling us about this most interesting new tool for our MRPC 1.1 competence ("including the benefits and risks associated with relevant technology") toolbox.

CaseStrainer is free to use. Feel free to put it through its paces at https://wolf.law.uw.edu/casestrainer/.

Come Visit Us at KCLL

The law library has several databases (including Westlaw) that can assist you with researching the veracity of case citations found in your CaseStrainer validity report. To find out more about our collection and services, visit us at www.kcll.org.