

WHY DOES MY LAWYER BRAIN DEFER TO A I?

A recent post on an attorney listserv included the following message – having never encountered an expert in a particular field before, the lawyer explained that they had “used A.I. for a suggested cross-examination since I had never cross-examined [that type of expert] before...”

The danger to a client may be readily apparent, along with ethical concerns. And I wrote an article about them for the regional legal community newspaper. But when I sent that piece to Brain Lessons co-editor Grant Rost to ask if it belonged in our column, he wrote back the following.

Could it be a brain lesson? Well, in its current state I think it's more “lesson” than “brain” if you catch my drift. Perhaps with some prefatory linking you can tie the two together more clearly. It currently looks more like a “perils and/or limitations of A.I.” lesson than a study about the brain. The “linkage” that immediately popped into my head is from Kahneman himself. The brain prefers effortlessness—which is clearly manifesting in brain of the lawyer in the post

Grant, unsurprisingly, nailed it. What I didn't know, and needed to research, was why a lawyer would trust AI. It turns out that there is ample research.

What I found first was the article **“The Secret Reason People Trust Artificial Intelligence More Than They Fear It”** by the [Editors of ScienceNewsToday](#) from [August 28, 2025](#). <https://www.sciencenewstoday.org/the-secret-reason-people-trust-artificial-intelligence-more-than-they-fear-it> (last visited April 30, 2026). That article described, and led me to, research by Phillip Brauner, a Ph.D. who studied computer science with specializations in psychology and human computer interaction.

Brauner and colleague surveyed 1,100 German adults about AI. The primary purpose was to assess attitudes – for each of numerous AI applications, the subjects were asked about whether they had potential utility but also whether they were viewed as risky and detrimental. As described in SECRET REASON, the findings included the following:

If participants saw an AI scenario as tangibly helpful—such as supporting healthcare, easing everyday tasks, or providing companionship for the elderly—they tended to rate it positively, even when some risks were acknowledged...

Interestingly, perceived risks were consistently rated higher than perceived benefits across the board...***But here's the twist: when it came to shaping overall evaluations, benefits carried far more weight than risks. A potentially useful application could win support despite acknowledged dangers...***

That last point, that a useful application “could win support despite acknowledged dangers” especially when it “eas[ed] everyday tasks” drove me to contact Dr. Brauner.ⁱ His response, prompt and exceptionally detailed, provided more articles and added insights. One involved “automation bias” – the tendency to “favor suggestions from automated decision-making systems and to ignore contradictory information made without automation, even if it is correct.” Hoffman, **Automation Bias: What It Is And How To Overcome It, Forbes, March 10, 2024.**
<https://www.forbes.com/sites/brycehoffman/2024/03/10/automation-bias-what-it-is-and-how-to-overcome-it/>

Dr. Brauner described another study of his.

We also did a study on that using a business simulation game and found that people tend to over-rely on automated systems (an effect or bias called automation complacency). <https://www.tandfonline.com/doi/full/10.1080/0144929X.2019.1581258> . The error of the Simulated Support System , clearly noticeable by the participant, did not really change their behavior and they still relied on the system, despite its evident malfunctioning.

Automation bias is informed by “miscalibrated trust.” Romeo, G., Conti, D. Exploring automation bias in human–AI collaboration: a review and implications for explainable AI. *AI & Soc* **41**, 259–278 (2026). <https://doi.org/10.1007/s00146-025-02422-7> Romeo and Conti explain that “[i]n situations with high workload, users may reallocate attention to manual tasks, neglecting the automation. Such behavior increases the risk of delayed detection of automation errors. A direct consequence of this behavioral pattern is the cognitive error of [automation bias].”

So – we can now explain how the brain led this lawyer down the path to AI-reliance. Hopefully, that satisfied the Brain Lesson component of this article. What follows is the original law-focused article, the story of the lawyer’s choice, my rejoinder, and what AI [Claude] had to say about it.

The case in which the lawyer sought AI help was a child sexual assault prosecution; and the type of expert was one on the behaviors that might be exhibited by children who are victims

of sexual abuse, including delayed reporting. The simple response to the use of A.I. is – bad idea.

Here are concerns that I raised in my listserv response:

I write to urge extreme caution in using A.I. in this fashion. There are a number of reasons:

- 1. There are better sources - NACDL and other well-credentialed organizations have materials*
- 2. The suggested cross might end up opening the door to damaging testimony. For example, in Pennsylvania an expert may never opine on a particular case/child; but an A.I.-suggested cross, by emphasizing that, could arguably open the door.*
- 3. A.I. relies on caselaw that is not necessarily current.*
- 4. A.I. relies on all sorts of sources, some of which may be really bad.*
- 5. Depending on the A.I. system used, what you upload may become public information.*

But that's not all. I decided to ask an expert – CLAUDE -about the virtues and pitfalls of using A.I.

Here is my first prompt:

If I were to ask you to design a defense cross-examination of an expert witness in a child sexual assault prosecution where the expert would say that "the fact that the child had no particular behavioral problems does not mean there was no abuse," could you do it and what risks of error would there be in the results? And would my input into Claude be public or private?

Claude responded quickly and thoroughly. It had some admittedly cogent ideas about such a cross, but volunteered these concerns:

Risks of Error in My Output

This is important to be direct about:

- ***Outdated case law.*** *My knowledge has a cutoff, and evidentiary standards, Daubert/Frye rulings, and state-specific rules about expert testimony in child abuse cases evolve constantly. A line of cross I suggest might be inadmissible in your jurisdiction.*

- **Jurisdiction-specific rules.** Some states have rape shield analogs or specific statutes limiting what can be asked of certain experts in child abuse cases. I may not know the specific constraints in your court.
- **I don't know your record.** Good cross-examination is built on what's already in the record — the expert's report, prior testimony, deposition, and the specific facts of the case. Without those, my suggestions are generic, and generic cross is often ineffective or counterproductive.
- **I'm not a lawyer.** I can help you think through the structure and the science, but strategic decisions — what to ask, what to leave alone, when to stop — require a licensed attorney's judgment.
- **Backfire risk.** Cross-examining a sympathetic expert in a child abuse case carries real risk of alienating a jury. I can't assess that risk for you without knowing far more about your specific case.

I decided I needed to know more. Claude was surprisingly candid in replying to my next prompt, which was “One other question, please. How do I know the quality of the resources you would use for this type of inquiry?”

Here is the A.I. confession:

This is one of the most important questions you could ask, and I want to give you an honest answer rather than a reassuring one.

The Short Answer

You largely can't verify my sources directly — and that's a real limitation you should factor into how you use my output.

What I Actually Do (and Don't Do)

I don't retrieve sources in real time for this kind of task. When I draft a cross-examination or analyze expert testimony methodology, I'm drawing on patterns learned during training — from legal textbooks, law review articles, case law, forensic psychology literature, trial advocacy materials, and similar sources. I don't pull up a specific article and cite it to you unless I use a web search tool. This means:

- *I can't hand you a footnote you can verify*
- *I may synthesize accurately from good sources, or I may reflect outdated, minority, or even incorrect views without signaling which it is*

- *I can sound confident when I shouldn't be*

I can hallucinate citations. *This is a known and serious problem with AI systems. If you asked me to name the leading peer-reviewed studies on behavioral symptom presentation in child sexual abuse victims, I might give you author names, journal names, and publication years that feel authoritative but are partially or entirely fabricated. This has caused real harm when lawyers submitted AI-generated briefs with fake case citations.*

To be clear, the lawyer who turned to A.I. didn't just regurgitate what was offered but adapted it. And Dr. Brauner wrote me again to add these comments:

[M]ore complex systems have evolved tremendously over the past few years, especially with agentic AI. While the mentioned points, such as the cut off date of the training data or the lack of information on the jurisdiction is certainly true for the base LLM models, smarter tools are able to automatically send queries to the right databases, summarize or filter semantically (instead of an exact word match as a decade ago), and present the results in a meaningful manner. I do see many benefits of this across a broad range of disciplines and jobs (who even can read every case on a specific law anyway), but the main problem remains: what happens if errors occur and who takes the responsibility for it.

For this writer there remains the risk of 'garbage in, [adapted] garbage out.' I am with Claude. If there is a listserv, ask for help from humans. Find out what the reputable sources are. And check your jurisdiction's law. Otherwise, what has superficial appeal may lead to disastrous consequences.

ⁱ To learn more about Dr. Brauner, go to <https://www.researchgate.net/profile/Philipp-Brauner> To learn about how public perception of AI differs from that of AI experts, see his 2026 article **Charting the AI perception gap: divergent views on risk, benefit, and value between experts and the public challenge the societal acceptance of AI.** <https://link.springer.com/article/10.1007/s00146-026-03023-8>