

# Short Circuit 313: Memo From a Robot

📅 Thu, Mar 07, 2024 9:27PM ⌚ 53:34

## SPEAKERS

Paul Sherman, Anthony Sanders, Ed Walters

---

### A Anthony Sanders 00:24

Hello, and welcome to Short Circuit, your podcast on the federal courts of appeals. I'm your host, Anthony Sanders, director of the Center for Judicial Engagement at the Institute for Justice. We're recording this on Friday, March 1, 2024. And this is a special Short Circuit. So longtime listeners may remember that about three years ago, we had a special Short Circuit on robot law, on all the ways that law interacts with what we call robots (broadly defined). So we talked about tort law, contract law, constitutional law as it interacts with our robot friends. And today, we have our expert, who we had on then, back because there's been a lot more things happening in robot law, specifically artificial intelligence and the law. And so we wanted to talk again to our good friend, Ed Walters, about what's the latest in AI, legal research, the law, you name it. So I'm going to be introducing Ed in just a moment. But also, I am very happy to say IJ has our own little expert on AI and the law, not to call him little, but to say that he knows a lot more than the rest of us do at IJ about these things. He's always encouraging us to use AI products and beef up our legal writing, and Luddites like me are often a little skeptical, but I think he's moving us along. And he has a voice well-known to Short Circuit listeners, and that's Paul Sherman. So Paul recently wrote a letter to the 5th Circuit, which is considering a rule about AI and legal research and disclosure ... that attorneys may have to give a disclosure in the future when they file briefs. And so Paul is going to talk about that letter that he wrote and the issues more broadly. So Paul, welcome back to Short Circuit.

### P Paul Sherman 02:30

It's a pleasure to be here, as always.

### A Anthony Sanders 02:33

Now, on to Ed. So Ed Walters is a man of so many things. Back in the day, he was a speechwriter in the George H. W. Bush administration. After that, he graduated from University of Chicago Law. He worked in big law for a while, he was a clerk for Judge Garza on the 5th Circuit, and he's done many other things. He's on all kinds of nonprofit boards. But his big deal

that he's done recently, and which makes him quite an expert on this subject, is he was the CEO and co-founder of Fastcase, which is a legal research platform that many of you have heard of. And recently, Fastcase has merged with an outfit called vLex, and so he is the chief strategy officer at vLex. Every day, he is working in this area of using AI to try to assist us in our legal research and AI more broadly. He also, for years, has taught robot law. He teaches it at Georgetown Law once here, and he also now teaches it at Cornell Tech in New York City. And so we are so excited to have back Ed Walters. Ed, how are you?

E

Ed Walters 03:59

I'm great, Anthony. Thank you for such a kind introduction.

A

Anthony Sanders 04:02

Well, so hear me out, and I wasn't too kind. And tell us what's up with AI and the law, AI and legal research. We hear all these scare stories about, you know, lawyers who don't know any better getting ghost cases and citing them to courts and getting reprimanded. And then, of course, the future that outfits like ChatGPT or now Google Gemini or whatever have for us. But let's separate the hype out while we're talking with you. What is the reality of what does AI bring us when it comes to trying to figure out what the law is and in telling courts what the law is, and what are maybe some of the dangers that are involved?

E

Ed Walters 04:50

Awesome. So that's about a semester's worth of questions. I'm gonna break this into 13, hour-long lectures.

P

Paul Sherman 04:58

Great.

A

Anthony Sanders 04:59

I hope everyone's ready.

E

Ed Walters 05:02

So we'll do a semester. No, let me break it basically into two separate histories that have run in parallel. The first is about generative AI. And these are kind of general foundation models that you're familiar with like open AIs, GPT-4, or Gemini from Google. A huge breakthrough, right? These are like now about a year and a half in public, starting with ChatGPT in November 2022. And they're amazing. I mean, they're shocking in what they're able to do. The versatility of these tools, the number of things you can use them for that they're not programmed for is incredible.

There's a problem, okay? The problem is that generative AI, these kinds of generative pretrained transformers (that's what GPT stands for), these GPT tools, they're not truth engines. They're sort of completion engines. They're designed to tell you what is most statistically likely to come next. And so it adds like the next words or phrases to a likely answer. Now, this is great. If you are trying to figure out what an itinerary for a vacation in Helsinki might look like, you can ask it, a GPT tool, which isn't, you know, programmed in some way to create itineraries, and it will generate like a pretty good itinerary. But if you ask a factual question, you know, what's the capital city of something? Or, you know, how many times does something happen? How many cases have been filed in federal court dealing with X or Y or Z? GPT tools aren't designed to answer the question. They're not designed to find the fact. On the other hand, they will give a complete, very confident, very real sounding answer. This is what they're designed to do. And this is where lawyers kind of got in trouble, because they would run like a legal research question into these general tools. The tool isn't designed to find the answer. It's designed to create an answer that looks very convincingly like a right answer. So Steve Schwartz, a lawyer in New York, did this. He got torched in the New York Times for it. He asked a question about tolling the statute of limitations when you're suing a defendant who has filed for bankruptcy. ChatGPT, not a fact engine, wrote an amazingly convincing answer to the question. It cited cases. It even Bluebooked the citations to the cases, so the citations look perfect. The problem is the cases didn't exist, as we all know now. They were statistically likely answers; they weren't factual answers, okay? And so this hallucination problem is like kind of a feature of these foundation models in generative AI generally. Now, let me tell you a parallel story, a whole different stack of technology. This is called retrieval augmented generation. It's basically a method of using those same foundation models, but the way it works is effectively two steps. Step one, you go find some factual basis for answering a question. And then step two, you pass both the question and the documents that you would use to answer it into a prompt window, into a question window. And then one of those tools (GPT-4, Gemini, Claude 2 for Anthropic, Llama if you're really nerdy, Kelvin from 273 Ventures), it doesn't really matter what the kind of engine is there. But you pass the question and all the supporting materials, and you basically say answer the question using the following and only the following to substantiate it. And so, this retrieval augmented generation has created a whole new generation of tools. My company, vLex, has one called Vincent. It's not the only one. It's just the best one. There's tools like this from Casetext and Thomson Reuters, from Lexis AI and others. But this is really promising because these tools don't hallucinate. They are built to answer the question using only the supporting materials that are provided. And so it's really cool. You can put together legal research behind the scenes and then pass supporting documents into these engines that will create all kinds of stuff: answers to questions, briefs, memos, 50 state surveys. This is truly revolutionary. So I guess I would just say here that I want to distinguish the kind of general foundation models, the open AI, Google style tools, these general GPT tools, which are really interesting. Lawyers can use these, by the way, all the time, if you want to create marketing copy, or something for your website, or an email draft of a letter, or something really good. But for fact engine questions, they're a disaster. The general foundation models aren't ready for prime time. They should never be used for legal research. On the other hand, we now have these really interesting retrieval augmented generation tools. These are specialized for law. And they use primary sources that you can see and audit and verify. And so this class of tools is creating these really interesting products in law right now. I'm talking about the legal research parts of them. There's tools that do this for electronic discovery, that do it for contracts. It's a really amazing time.

A

Anthony Sanders 11:15

So it sounds like, Ed, to my novice years, that maybe what's happened is a lot of people who got excited about ChatGPT but didn't know very much about it, they thought that those programs were the latter kind that you were just talking about, that you put a bunch of facts in, because they were not advertised, but the buzz about them was that, you know, all known information on the Internet was fed into them. So, of course, they're going to know everything.

E

Ed Walters 11:43

Yeah, I think that's right. And, you know, I have to say, I have some sympathy for people who use the general foundation models to try and do this sort of research. Every piece of software that we've ever used ... In the past, if you ask questions and it doesn't know how to answer it, you get back an error message. You get back a 404 like this page can't be found. If you run a search in the traditional legal research tools like Fastcase or Westlaw or Lexis and there's nothing there, you get back a message saying there's no results found, right? We know what happens when software reaches the end of its rope, when I can't answer the question. Except these GPT tools really, for the first time, don't ever have an end to the rope. So if you ask them a question that they don't know the answer to, they don't get back a no results found message. And not only do they not, they will respond with a complete answer. And a complete answer that sounds extremely confident.

A

Anthony Sanders 12:52

And when you ask them is that actually true, they will say oh, yes, of course it is.

E

Ed Walters 12:56

Yeah, yeah, of course. Now, look, I will say that, you know, the people who got in trouble for this, for the most part, I have sympathy for that use and for what they do when they receive a very competent response to these answers for the first time, up to a point, right? I mean, the reason Steve Schwartz got in trouble was that he didn't actually then go read the cases that he was citing to the court. Always read the cases, no matter what.

A

Anthony Sanders 13:22

Or even look them up.

E

Ed Walters 13:23

Yeah. And then, when the court said to him, hey, we've tried to find these cases and we couldn't, what does he do? I mean, does he go back to run traditional legal research? Does he go into Fastcase or Westlaw or Lexis to try and find the cases? He does not. He goes back to ChatGPT and says, hey, these are real cases, right? And what does ChatGPT say? It says of course they are real cases. You can find them in reputable legal research services. Go check for yourself. And he says, that's good enough for me.

A

Anthony Sanders 13:58

So going back to the products that you're developing and that, as you noted, others are developing. If you would ask one of them for something and they just didn't have a case, they would own up to it, essentially?

E

Ed Walters 14:12

The best tools own up to it. So like I can't speak for anybody else, but I'll tell you that one of the biggest parts of the training of the vLex Vincent AI tool was to do exactly that. To train the tool to say, I don't know. You know, or like here's what I do know, and here's the extent of it. And I can't give you anything more than that. But it's really surprising, you know, that the dataset is so large that it's trained on. The number of answers you can find is really pretty impressive. It's possible to stump the expert, but you have to try pretty hard. You know, if you ask for something that was nonsense, you'll get back something that says like we can't find anything on this. But for actual legal questions, pretty incredible.

A

Anthony Sanders 15:02

Maybe you could give us a quick example. People like me may be wondering how this works, so a lot of, you know, all the lawyers listening, maybe even some non-lawyers know about if you go on a product like Lexis or Westlaw and you type what's often called a natural language search ... You know, what is the standard of liability for negligence in Ohio, say. You're in the Ohio courts database. You'll get a bunch of cases. And so you can click on the first three cases and chances are, you probably will find the case that's closest to trying to give you that kind of answer. What if I type that into the models that you're talking about? What kind of response will I get?

E

Ed Walters 15:53

Yeah, and thank you for such a generous explanation of traditional legal research. There's so little contempt in your voice. There's no like anger or scar tissue. I think if you asked a lot of lawyers about their traditional legal research experience, you would get like, you know, an eye roll or something else because I have to say, I've been building these tools for like 25 years; they're incredibly stupid. I had to explain to my kid the other day like, you know, sort of how legal research has worked. And the process is, I'm not making this up for people listening to this in the future, the way it has worked for the last 30 years, you imagine language that a court would use in discussing your issue or a legislature would use to discuss your issue. What would a sentence look like if someone was describing it? And then you formulate a key word search, this exact phrase, and then this word within four words of this word or this word or this word, but not this word.

P

Paul Sherman 16:59

Yeah. I mean, if I can interject, just from like a daily practitioner standpoint.

E

Ed Walters 17:05

Please.

P

**Paul Sherman 17:06**

I have long found the natural language search functions on, you know, major legal databases to be essentially useless. And I have since law school, and I graduated in 2006. I've used essentially exclusively Boolean searches of, you know, like standard within three of liability, you know, within paragraph, you know, tort, you know. And, as you said, you have to imagine if I were a judge, I guess, how would I phrase this? And just kind of try and cobble together like what proximity would those words be in? If you enjoy solving puzzles, it's kind of a treasure hunt. But it can be extremely tedious, particularly if you've got a very, very discrete question that has not arisen many times.

E

**Ed Walters 18:04**

Yeah, and so then, the stupidity continues. You get back a list of documents that meet those criteria. If it's within five words, but not four words, not on the list. So you have the closest overinclusive and underinclusive, includes a whole bunch of documents that have nothing to do with your case but meet the keyword criteria you asked for. There's a bunch of things that are important that aren't in the list because the court didn't use the precise language that you figured. And then the traditional experiences, you get hundreds of documents. How many of those do you read? Nobody knows. Like I can't tell you. I'm an expert. Westlaw can't tell you. Lexis can't tell you how many to read. When I was in law school, the guidance was, you know, try to keep narrowing the list. So there's 50 of them, and then read all 50. When I was in practice, the guidance was don't ever think, just print out hundreds of cases as quickly as you can because all the time you spend thinking online, you're being charged for. And then read them all in paper, in analog, and don't ever miss anything. So this is, I mean, this was the state of the art of legal research in 2022. The best systems in the world, you know, I built some of them, like this is how you did it. And then you read them, hope you didn't miss anything, and try and synthesize an answer for yourself based on what happened in those documents. So this is the difference. So in legal research now, using tools like Vincent, I won't speak for others, but I'll tell you how we do it. There's two steps, as I said before, in a retrieval augmented generation. The first is to do all that research behind the scenes, not like old natural language where it's based on synonyms and things like that. But on this new technology called vector search, it indexes the whole legal research database based on the concepts in the documents, and then takes your question and converts it into concepts, and then finds all the documents that deal with that concept, but in a very targeted way. So instead of getting, you know, 3,000 results or 300 results, you get back like 16 of the most highly relevant documents that matter for you. Think like if you boil down the entire ocean of documents in your result set to find the ones that really mattered, it would look like the 16 results that we come up with with this vector search. And then we pass those 16 documents along with a question into generative AI to say, what would a brief look like based on making this point of law? Or write me a memo that says what the state of the law is in this place. Your negligence standard in Ohio, you get like a really nice memo saying what the answer is. And that whole process I'm talking about takes about seven minutes. You know, research tasks that were like one week are now like a point one on the billing sheet. I'm not even trying to sell it, right? I'm not trying to hawk software or something. I'm telling you though the process is so much better. I think we have killed Boolean search dead. Sorry, George Boole.

A

**Anthony Sanders 21:24**

So I think a lot of lawyers listening ... either their jaws just dropped and they're thinking of seven minutes for a memo, or they're shaking their heads and like this can't be, whatever you get is going to be garbage. How can that be true? So I think this is a good time to go to Paul. And because of, you know, all kinds of reasons maybe, including fear, courts are looking at rules or bar boards or whoever they are, are looking at rules about this newfangled stuff. And Paul saw some dangers in one of those rules. So tell us a little bit, Paul, about what is going on in the 5th Circuit and elsewhere, too. And why you submitted some comments about this proposed rule.

P

Paul Sherman 22:21

Yes, I'd be happy to, Anthony. And yeah, so I'm quite bullish on AI and the effect that it's going to have on legal research, but not just legal research, also legal writing. You know, lawyers have been criticized for generations for the quality of their writing, and there are a number of tools that are coming out that I think can really assist lawyers in writing clearer, more persuasive prose. And that's not just good for their clients, it's good for courts, it's good for clerks, and ultimately, I think it's good for the law because we would rather have legal issues decided on a clear understanding of the merits and not just based on, you know, well, one guy is a naturally gifted writer and the other guy, or maybe the other pro se litigant, is not a naturally gifted writer. And, you know, we would like both arguments ideally ... I mean, I want all of my opponents to be terrible writers, but at least, you know, outside of the context of my individual cases, I want cases decided on the basis of the merits. So, as you mentioned, the 5th Circuit, in response to some of these phantom, hallucinated case citations, proposed a new certificate of compliance that people would have to file with their briefs. And it took the form of two checkboxes, one of which was no generative AI was used in the creation of this document. And the other checkbox was generative AI was used in the creation of this document, but a human has verified the accuracy of the citations and the legal arguments. You might think, well, what's wrong with that? Well, we were concerned about it because it is poorly calibrated to the actual problem. The actual problem is the hallucinated citations. But dealing with the hallucinated citations doesn't require you to make people affirmatively out themselves as using generative AI. And I think having people do that is going to have a tendency to prompt judges ... You know, judges tend to be older and less tech savvy, and I think it's going to create a bias against the use of generative AI in the production of legal documents. And I think that's going to discourage people from using generative AI in ways that could be really good for them and their clients. So you can imagine sort of two case situations. One is the one that we read about in the New York Times where this fellow thought that ChatGPT, he described it as a super search engine, had access to cases that, you know, he couldn't find on Lexis and Westlaw. And so he just sort of outsourced all of the writing and thinking to ChatGPT. So that's one potential use of generative AI. And I think we would all agree it's an irresponsible one. But another potential use is you go out, you do research, you write a brief, and then you put the brief into something like ChatGPT. And you say, how can I make this clearer or more persuasive? How can I make this more concise, if you're working under a word limit, without changing the meaning? I think these are all very valuable uses of AI. And, in fact, there are tools out there, not just ChatGPT, there's a great tool called BriefCatch, which I use at IJ, that has begun to incorporate AI. And it helps people improve the quality of their prose, which I think is completely laudable. Another issue with the 5th Circuit's proposed rule, I think, is a lot of this technology is going to be essentially invisible to people. They're not going to realize that, you know, it's one thing if you say to ChatGPT like, hey, write a motion to dismiss for me. Then obviously, you know, you're getting a brief that's written by ChatGPT. But if you run something

through BriefCatch, you may not know that its text recommendations are based on generative AI. Or another possibility is you have a junior associate write a memo, and the junior associate uses ChatGPT or, you know, one of these tools like Vincent, and the memo ends up with text from a generative AI. Are you going to have firms require some kind of internal monitoring of like how each memo was created so that if you end up quoting it in a brief, you know, you know it came from ChatGPT, as opposed to something else? So we just thought the rule was poorly calibrated. And at the same time, courts have tons of tools at their disposal for dealing with this kind of inappropriate and unethical lawyering already. You know, under the Federal Rules of Civil Procedure, we've got Rule 11. Courts have inherent power to discipline these kinds of things, and we see this from the way that courts are actually responding to this. None of these courts is throwing up their hands and saying, oh, my God, I can't do anything about these briefs that are being submitted with phantom citations. They're relying on their traditional powers to discipline these lawyers. So, you know, we commented we were critical of the 5th Circuit's rule, proposed rule (they have not taken any action on it yet). Overwhelmingly, the comments that they received were this is just not necessary; you already have tools at your disposal for dealing with this. So that's one approach, which is sort of rooted in the skepticism of AI. I think another more optimistic approach comes from the Florida Bar, which in January, released an ethics opinion about the use of generative AI in legal practice. And I thought it was very clearheaded and focused on the real problems. And it basically said generative AI is a technology like email or cloud computing, and just as we have with those kinds of technologies, you can adopt it, but you have to satisfy your ethical obligations. And so, in the same way that it would be unethical for you to, you know, if a junior associate writes a memo, and then you put the content of that into a brief and you sign it with your name and submit it to a court without having looked at any of those cases, you're the one who's ethically on the hook for that because you didn't do your diligence. And it's no different if you use a product like vLex or Thomson Reuters' competing product. You got to do the work and confirm the accuracy of what you're submitting to the court. And then they, you know, talked about other standard ethical obligations of lawyers. So Ed mentioned that, you know, this research can be much faster with these tools. I had an experience with this recently where I was writing a response to a motion to dismiss in New York, and I wanted to know what's the standard for determining whether it should be with prejudice or without prejudice. Using the typical Boolean search, it would have been like standard within sentence, dismiss within sentence, prejudice. And then you'd get all kinds of stuff where it's like, oh, if there'll be prejudiced by the dismissal, which is not quite what you're looking for. I put it into an AI legal research tool, and it immediately said like, oh, you know, there are two tests. One of them is called the Zagato factors, which has these five things, and it provided a list of highly relevant cases. I mean, it took what would have been a couple hours of research probably, and I had the answers I needed in 10 minutes. So it was really remarkable. So one of the things the Florida Bar pointed out is your billing now has to reflect the fact that that took 15 minutes and not two hours. You can't bill based on what it would have taken you. You have to bill based on what it actually took you. And similarly, you have to do things to maintain client confidences. So if you're working with some kind of generative AI that is trained on input and things like that, you have to take reasonable steps to make sure that you're not disclosing attorney-client privileged information. But all of this stuff is just new applications of existing ethical rules. It's not saying that we need to upend things and adopt a bunch of new rules to deal with this new technology. I think that's the much more sound approach. And are there going to be growing pains? I mean, sure, because you're going to have people who are unfamiliar with this tool. As Ed said, we can have sympathy for this guy in New York, who, you know, just overestimated what the technology was, but our sympathy also has to have a limit. He should have read those cases, and it's not like these stories are not getting highly publicized. I mean, like these are horror stories among lawyers. We hear these, and we're just like, oh, you know, there but for the grace of God go I. So I think the growing



pains' period with these tools is going to be relatively short. People are going to see other lawyers get nailed for this, and they're gonna say, you know, I'm not gonna let that happen to me. Now, if that discourages them from using these tools, I think they're going to be at a competitive disadvantage. I think that their research is going to be a lot slower. And I think that their prose is not going to be as lively as it could be. And so, you know, I think they're great tools. And I think the way that I've tried to explain it to other lawyers who haven't tried them is if you've played around with ChatGPT and looked at where ChatGPT was a year ago versus where it is now, I think we're gonna see a similar jump with these legal tools. The tools that I have tried, I think, have their limitations. I think for issues that recur hundreds or thousands of times, I think they're very good at producing quick, accurate answers to things like what's the standard for the, you know, dismissal with prejudice that's occurred in 1,000 case? So I think they're very good at answering that. If it's something that's only occurred in two cases or requires a lot of discriminating judgment about like a nuanced constitutional issue, they're not there yet. And they, you know, for example, I asked one of these tools about what is the standard that the Supreme Court uses for distinguishing speech from conduct for First Amendment purposes. It's very important to IJ's litigation involving occupational speech, people who speak for a living, because in a lot of these cases, the government will say, oh, that's not speech, that's professional conduct. And there is a test for this set forth in a case called *Holder v. Humanitarian Law Project*. And the tool completely missed that, didn't get that case. And instead, it gave me the similar, but not exactly the same, test for distinguishing expressive conduct from non-expressive conduct. So, you know, these things are still going to struggle with these more nuanced issues, at least in the short-term, but as I said, the leap that we saw in ChatGPT over the last year, I think we're going to see that with these legal tools. And I imagine 10 years from now, people coming out of law school are just ... They're going to look at, you know, people like me and Ed as dinosaurs. I mean, it's going to be kind of the way like ... I mean, so I'm in my 40s, and the idea of doing like legal research with just the case books, you know, seems hopelessly archaic. It, frankly, seems like malpractice to me.

A

Anthony Sanders 34:52

It's funny, Paul, we actually talked about that on last week's show. And we had Dan Knepper and Bobbi Taylor on, both IJ attorneys. Bobbi's just out of law school, so I'm sure to her, it's like the Stone Age. But I asked Dan, you know, back when you just got out of law school, were some people still pulling the books off the shelf to read the cases? And he said even he didn't remember someone like that. I think I remember a couple older partners doing it, but yeah, I think you're right in terms of how things are changing.

P

Paul Sherman 35:26

Yeah. I mean, when I was in law school my first semester, they made us use the books before they would allow us to use Westlaw. It was like, okay, now ...

A

Anthony Sanders 35:37

Yeah, us too.

P

Paul Sherman 35:38

Now, we'll allow you to use these electronic tools.

A

Anthony Sanders 35:41

Yeah, it was like hazing.

P

Paul Sherman 35:42

Yeah. If you were to rely on just the case books now, I think that's essentially malpractice. There's no way that you're getting the universe of cases you need. And I think 10 years from now, if you're doing just Boolean searches and guessing what phrases courts may have used to announce these legal issues, instead of using tools like this, which really parse the language and seem to understand it, I think you're going to be similarly seen as a dinosaur.

A

Anthony Sanders 36:11

Ed, so I'd love to ask you a couple of questions based on what Paul said, but do you have fears of overreaction, of regulation, in this regard? And have you seen, you know, any other instances of those dangers?

E

Ed Walters 36:30

Well, first, let me say, I totally agree with Paul that we do have rules in place that deal with these questions. In 2019, I wrote a law review article basically making that point, talking about the Model Rules of Professional Conduct and how rules like Rule 5.3, the duty of the lawyer to supervise, not just lawyers on the staff, but everyone on the staff and I would say including AI tools, you know, meet the mandate of the professional rules' duty of candor. Ultimately, we already have very good rules in place that deal with most of the issues behind legal AI tools, so I'm very much in sync with Paul. I will say that we see this trend in many, many courts right now. There's, I think, 17 or 18 different courts that have issued this brand of local rule. It's very burdensome for litigants to try and figure out the disclosure rules in each individual court. And I'll just say, maybe for IJ listeners in particular, one of the things that really gets me is I'm starting to see self-represented litigants using these general AI tools and coming up with hallucinated citations, and then the individual self-represented litigant being fined by a court. And I think that's the biggest tragedy of all, you know. It's very difficult for individuals to navigate litigation, navigate their way through the courts. This makes it much harder. I do think that there's great potential for AI tools to make that process easier for people, but this problem of hallucination in the general tools is a real problem. It's a problem for lawyers, of course, but lawyers can then go read the cases if they want to and act as sort of a check. Lawyers, for the most part, are going to have access to these tools from vLex or Thomson Reuters or someone else. But self-represented litigants will not. And so the people who need the help the most might be getting tools that hurt them the most. And that's something that I really worry about.

A

Anthony Sanders 38:52

So one issue, this is a different angle on all this, but one issue that I've read about is copyright concerns with some of these databases. So I believe, I haven't heard more of it since, but a couple months ago, there was a lawsuit, I think the New York Times, took against one of these services, because essentially, when you feed, you know, the entire Internet or the entire history of a paper, whatever, into a tool and then you ask it to give you information to sum up the argument, is to some extent, that is just copying and pasting from our copyrighted works. I don't think that's a problem with case law because case law, the law, thank God, isn't copyrighted, but it could be with secondary sources, right? And so are these worries that also are added onto this, Ed, or is that kind of just a peripheral thing and it stems from some misunderstandings?

E

Ed Walters 39:59

Yeah, I think this is the topic of my law of robots class at Cornell next week, actually.

A

Anthony Sanders 40:05

Wow.

E

Ed Walters 40:06

So there's two real approaches here that courts have taken. The first approach, I'll say, is the Napster approach. You may recall, Napster was the first online service for streaming music. It took copyrighted music and made it available online. And the Napster fig leaf was to say this is a backup for your physical media. If you own the disk, you can listen to it on Napster, but if you don't, you can't. And so if you don't have, if you don't own, the music, rights to the music, then you can't listen to our Napster. And if you do that, then you are violating copyright law, not us.

A

Anthony Sanders 40:50

I remember that. And I remember not many people thought much about that, then used it anyway.

P

Paul Sherman 40:56

I remember listening to a lot of music that I did not own on physical media.

E

Ed Walters 41:01

Right? Well, so that argument failed. But, you know, part of the argument also was this is something transformative. We're creating something beyond the individual music files. The second approach was Google Books. Similarly, Google took copyrighted material without permission, scanned it, digitized it, indexed it, and put it online. And you could read sections of

copyrighted books on Google. You still can. And the publisher sued, basically saying we didn't authorize this, we didn't authorize this under copyright, we own the copyright on this, you made a copy of it, and created something out of our copyrighted works. And the court said, this is a transformative use. This is fair use, even though they used your copyrighted materials without permission, what they created is something completely different. And so this doesn't compete with your book. No one is going to decline to read *The Tipping Point* because there's a squib of it on Google. And so we're going to allow this under copyright law. Now, the question is is generative AI more like Napster, or is it more like Google Books? I think you could make the arguments really either way. If I had to guess, you know, crystal ball, I would say courts will probably find it to be more like Google Books because this is not a service that asks people to call up *The Tipping Point* in generative AI. We're not asking to pull *Starry Night* from Van Gogh. But we might say, you know, create this family picture in the style of *Starry Night*, which uses Van Gogh and uses that style, but doesn't pull up the copyrighted image. And so this is an interesting question. It's a hard question. It applies a lot too, right? So I imagine that there's a lot of online materials that are secondary materials that were used to pull into this; copyrighted blog posts, things that are public on the web; that were used to train the foundation models like OpenAI or Claude. And you'll have the exact same copyright question with those copyrighted materials that you do with the *New York Times*, or Sarah Silverman, or others. And so we'll have to wait to see what the courts do with that. I'll say that from a publisher perspective, I would not want Fastcase's books used without permission to train these foundation models. As a software developer, we built the Vincent tools so that you can hot swap the AI model. We use GPT-4 right now. But if a court said the *New York Times* wins this case, and we're going to take GPT-4 offline so you'll have to use something else, there are a lot of legal research tools that have GPT-4 baked in, so you'd have to start from scratch. We'd have to change a configuration setting and move back online with Claude 2 from Anthropic, or Gemini from Google, or Kelvin from 273 Ventures.

P

Paul Sherman 43:31

You know, in some ways, it's an interesting philosophical question about how humans form prose, you know, because really, my perspective is like you become a better writer by reading lots of good prose. It's a major part of it. And so when something sounds good to you, it usually sounds good to you because something in your head is saying like this is the most next likely word in the sentence expressing what I'm trying to express, and in many ways, that's what these large language models do. That's not to say that they're intelligent. But maybe that's also not to say that we're intelligent.

A

Anthony Sanders 45:09

Well, they are called artificial intelligence. And to close, so I thought that we'd go in a different direction. And the, you know, for a long time, these have been called artificial intelligence, and to some extent, that means that they're much better at doing things that computers have done for a long time. And then, to some extent, of course, people think does that mean that these programs are, to some degree, becoming self-aware in a way that they were not before? And then the big question is, are they eventually going to become self-aware in a way that they're going to have their own needs (this is, of course, much bigger than AI in the law). Are they going to be friendly with us? Or are they eventually going to be like the Terminator? So what do you see as the ... You gave us kind of the outlook in the next couple of years in AI. What is the,

you know, the more medium-term looking like for where you think, not just the law, but more general, how we're going to be integrating these tools into our lives? And should we be at all worried?

E

Ed Walters 46:32

Well, so I do think we should be worried. I'm less worried about artificial general intelligence where the software takes on its own motives and starts to address its own needs. I don't worry about that so much. I sort of think about these things as like calculating machines. Calculators have gotten much more sophisticated since they were originally invented. But you can apply Moore's law as long as you want, the calculator will never become self aware, you know. And so the tools that we're talking about right now are autocompleting tools. They might get better at autocompleting, they might get much faster at it, but there's nothing that says that they will become sinister or become self-aware. But here's what I do worry about. I worry about us. I worry about the economic and social impacts of artificial intelligence. I worry about widening inequality being further widened by the use of AI tools. And I worry about overreliance. One of my favorite short stories is Isaac Asimov's *The Evident Conflict*. It's the last short story in the *iRobot* series, and the premise is we use AI to allocate things to run systems, and it works really well and allows us to overcome poverty and to, you know, find a great way to organize the world. Maybe for IJ listeners, it creates perfect marketplaces so that people can truck, barter, and exchange in the most efficient way possible, leading to utopia. But, at some point, people, you know, have AI systems and then AI systems to manage the AI systems, and then they start to get better, and then we achieve kind of an equilibrium. But nobody understands how they work, and then they start to go wrong. And we can't figure out why. And that might sound like weird science fiction, but like that is how the stock market operates, by the way. You know, the highest, by far the most volume on, the major exchanges is high frequency traders. And they just have algorithms that see little blips, little variations, that are unexpected and arbitrage those. They trade millions of times a second to exploit like a little market inefficiency and then sell a half a second later. Nobody knows how they work. No one knows how they interact with each other. And there's a lot of money at stake in these systems. They all use AI. So is there any reason to think that we won't do this all over the place where we have plaintiffs who file lawsuits and class actions in an automated way? And then defendants, overwhelmed with the volume, replying with AI. And then courts overwhelmed with the volume of lawsuits that are created by artificial intelligence doing first pass or maybe last pass adjudications with artificial intelligence. And then, in that world, maybe we deal with things in a very efficient way. Maybe it's right on the law the way Paul was talking about. But what does that do to us? How do we navigate a world that we don't really fundamentally understand how to use? This is why the law school instructors make you read the books in your first semester of law school. It's why we teach still arithmetic and multiplication tables to children in school, even though there are really good automated tools that resolve math. And so this is something I worry about. I worry about our overreliance on AI and our underdevelopment as a species.

P

Paul Sherman 50:29

Can I say two quick things on that? Because I think those are very valid concerns. One of the ... The last point that you made about, you know, us still teaching arithmetic to kids. I think one of the reasons why that's important is not just that people know how to do it, but it teaches you to recognize insane results, you know? Like you get a number and you're like, I don't have to do

the math to know that can't possibly be the right answer to this situation. And so I think it will always be important for us to be able to do those kinds of insanity checks. And then I just want to flag, I think a shorter term concern with AI that's not so much related to the kind of natural language stuff that we've been discussing. But I think the problem of deep fakes, particularly as it comes to political information, potentially campaign advertising or, you know, international campaign interference and things like that, I think we're gonna see a tremendous amount of that in the next few years because the tools have gotten really good at generating video and audio of things that seemed very convincing that just did not happen. And I don't claim to know what the solution to that is. But, again, evidence not that the tools are sinister or malicious but that people will turn them to malicious ends.

**A** Anthony Sanders 52:08

And I'm sure you would say, Paul, that whatever the answer is, it's not to have a department of deep fakes that would regulate how we use AI?

**P** Paul Sherman 52:18

Correct. Yeah. Yeah. I mean, you know, it's a topic that we could do a whole podcast on. And there are arguments, for example, that under tests, like the Zauderer test, the government could require disclosure, as opposed to prohibition on some of this speech. We don't have the time to get into it here. But yeah, it's gonna raise some difficult First Amendment questions.

**A** Anthony Sanders 52:46

Well, we'll leave that for another episode. But I would like to thank these guests for this episode, most of which has been above my intelligence. But I've still learned all kinds of things about the state of the law and AI today, and I hope the audience has as well. So thank you very much to Ed and Paul.

**P** Paul Sherman 53:09

Thank you.

**E** Ed Walters 53:11

Thanks for having us, Anthony.

**A** Anthony Sanders 53:13

And, in the meantime, whether you're using AI or anything else, I hope that all of you get engaged.