

Ron Lazebnik:

Welcome. This is The Beginner's Guide to AI: What You Need to Know from the Get-Go. I am Ron Lazebnik. I teach the Law and Technology Clinic at Fordham Law School. And I'm joined by Natalie, Sam, and Sonja. I'll give them each an opportunity to introduce themselves. Before I do that, though, let me just share the CLE information. On the screen right now, you should see the CLE code, which I will also say out loud. It is ABG24. That's ABG24. Please take a moment to write it down so that you can receive CLE credit for this session. Once again, it is ABG24. Okay. Welcome. And now to introduce my great panelists who I'm very happy to speak with about this topic. Why don't we start with Natalie? Please introduce yourself.

Natalie Knowlton:

Sure, everyone, thanks so much for being here today. And I appreciate the opportunity to talk with you all and certainly to be here with my panelists. I am Natalie Knowlton. I spent a very long time in academia at the Institute for the Advancement of the American Legal System at the University of Denver where I looked at access to justice issues in the civil and family space and also worked on regulatory reform and regulatory innovations. I left several years ago to found my own company, Access to Justice Ventures, where we work with individuals, whether they're entrepreneurs or lawyers or both, who are working on scalable justice solutions and ways to leverage technology in an effort to bridge the justice gap, which we'll talk about today being very broad and very big. And I think maybe with that, I could just turn it to Sonja because she is one of those such entrepreneurs.

Sonja Ebron:

Hi, everybody. I'm Sonja Ebron. I'm a co-founder, Chief Executive at Courtroom5. We've built an online case management platform for self-represented litigants, handling, more serious claims, home foreclosure, medical debt, wrongful termination, that sort of thing. I am probably amongst the few non-lawyers in the room. My background is in engineering. I'm a PhD electrical engineer. I left academia to start Courtroom5 because I was one of our customers. I was a self-represented litigant who didn't find the tools to be able to represent myself and couldn't afford to hire a lawyer. Happy to be here and share what we can about what we're doing with AI at Courtroom5.

Ron Lazebnik:

Thanks, Sonja. And Sam, last but not least.

Sam Harden:

Hi, everyone. I'm Sam Harden. I'm a program manager at Pro Bono Net. I'm really thrilled to be here with y'all and to talk about artificial intelligence. Before Pro Bono Net, I was a practicing attorney for a number of years doing large liability cases on both the defendant's side, mostly in medical malpractice, as well as a plaintiff's side. I've also practiced in criminal and state and federal court but have been in the pro bono sector and the tech sector for a good number of years now. It just always gets always more years when I look back on it. Thanks for having me.

Ron Lazebnik:

Thank you. With that, I figured that, and maybe this is the educator in me, the best way to start off our panel is to get a sense of you, the audience, and what your level of interaction with AI in the access to justice space is. I've created a quick poll that I will start right now. Just a one question poll so that we can see and get a sense for everybody in the room. The question is where are you in the process of creating an AI tool to help your services? And folks are answering it right now. Very quickly, I might add. Thank you so much, everybody, for participating. We have 111 in the room and 80 responses so far, so very soon I'll close the poll. I'll give folks just a bit more time to answer. Okay. I think we've hit the critical mass of people who will answer.

To share the results with everybody, it looks like 26% of you have an idea but are not sure where to start, 6% of you have started work on it but have questions. And I'm going to divert all those questions to Sonja. We'll see. And then the majority of you, 68% need to know more about what AI can do before you can really start down this road. Hopefully our conversation today will help each of these groups at least a little bit. And please feel free to put questions in the Q and A to help us as we go through a variety of topics.

And just to give you an overview of the topics that we'll discuss today, we're going to talk about identifying the potential applications that AI might be able to help you with, thinking through the human component associated with that, thinking about the design process and then testing and validation, talking a little bit about the issues and challenges, and then future opportunities associated with using AI in this space. With that, why don't we jump into our first topic, which is identifying potential applications. And I'll start with you, Sam, about what's a good way for people to think through the various AI tools that are available to them? Or how would you identify them?

Sam Harden:

Yeah. No, thank you. In the first session, Jim Sandman really hit on the idea of use cases, centering thinking around use cases, not, hey, this is a really cool thing. What in the

world can I do with it? But think of it as I have a use case. What's out there that can actually meet my use case? Right now, the big attention is around ChatGPT and the models like Claude and Gemini and other things like that that we see in the media. ChatGPT passed the bar exam, or whatever. It's all text generation. That's what we typically think about. But there are many, many other AI tools out there that do a number of different things, things like classification, tagging phrases or paragraphs or documents with a set of taxonomy terms.

Things like translation, which Terry Ross hit on, that's been an AI application for a number of years now. We've just seen generative AI have the capability to do that. Sentiment analysis has been big for a number of years, especially in customer service or in social media. But now we're seeing real applications of generative and predictive models not just in text, but things like image generation, vision. This is really interesting in the access to justice space, the idea that you can pull data out of an image. If somebody takes a picture of a document or a court filing, it's able to pull out information about that case like case number, things like that. And we're also seeing a trend toward what they call multimodal AI applications where one model can do many different things. Things like extract text data, generate images, analyze images, classify text, and then do your text generation as well. To answer your question in a really long way, I'd think about it as those are the general buckets, but if you're looking at those, think use case first before you go and see what can I mash around to make AI do interesting things?

Ron Lazebnik:

That is certainly helpful. And I guess part of the use case. We saw a lot of use cases in the lightning round right before this, and I do see the requests for the CLE code, so what I'll do is between switching topics, I'll share the CLE code one more time. But to stay on this topic for a little bit, one of the things, and feel free to push back, or one of the other panelists can push back. Is I wonder, separate from the use case of oh, it's on the internet, it's a portal to do X, and I think a lot of the generative AI is pushing people to think about it that way is what space should we give the programs we're already using? And thinking about how to leverage AI with our case management system, our document management system. What do we want to say about that? And Natalie's nodding, so I'm going to let Natalie go ahead and answer first.

Natalie Knowlton:

Sure. Well, Sam's going to have a much better answer, I'm sure. But what I was going to add, that I think increasingly now and certainly six months and a year from now, we're going to be using generative AI and not even know it. That's already the case. It's being built into the tools that we use every day. Word. If you use Zoom, the AI companion is right there. I don't know if you've tried it. If you haven't and have a Zoom account, try it

sometime. Check it out. I think there are definitely use cases to talk about because we have an audience who are interested in those. But a lot of it is just going to be streamlining and supercharging the work that we're doing in the existing tools that we're working in.

Ron Lazebnik:

Sam, do you want to, you were nodding to that if you want to-

Sam Harden:

Yeah, just to riff on that a little bit, like Natalie said, we're already seeing it in tools like Word, tools like Zoom, Outlook. And one of the panelists before us was talking about using Copilot to do different things, which is built into the Office suite now if you have a certain subscription level. We're going to see it not as a separate tool but as something that's integrated as a helper, almost, and I hate to bring this up, but if anybody's as old as me, Clippy 2.0 in some respects, but a Clippy that's useful, hopefully. And even in Google Drive, we're seeing things like that; really, really useful things that you can do. Google has text analysis, text generation tools that they're working on in Drive. I don't have access to them, but I read the news and I, we'll see more and more of that. And I think those aren't legal-specific applications, they're very generic. And they're still going to be useful in legal applications.

Ron Lazebnik:

And Sonja.

Sonja Ebron:

I would just add to Natalie's point, I agree that we're going to be using AI without knowing it. And I can confess that we're doing that at Courtroom5. We've been using AI. And we haven't announced it for a number of reasons. We'll talk a little bit about UPL a little later. There are courts around the country banning the use of AI particularly for self-represented litigants. And it's not a big deal. Our goal is to make work easier for self-represented litigants, which is a wonderful use case for generative AI. But I do think we'll begin to see in legal and outside of legal a lot more use of AI without it necessarily being a big announcement.

Ron Lazebnik:

Yeah. And I think part of it, and this goes to seeing it versus not seeing it, is thinking about whether the AI that you're integrating and thinking about using, is it something that is really client-facing, really going to face the people you're assisting? Or is it something more internal to your organization? And we saw a variety of those in the use

cases. And I think that is important to think about. Not every solution that will help you improve access to justice is one that is so publicly visible. And Sonja's giving a great example of that, that internally they're doing it, but it's not that visible. And the truth is, if it's not as visible, it might actually be less of a challenge or a liability for you because it's really, internally, there's the people to check on it to know how they're using it. You can think of a recent example of that Canadian airline that had a chatbot that promised a reduced fee for airfare, and the court held that is something that they're responsible for. Query, whether or not you agree with the opinion, but that is one of the confidential consequences with a client-facing AI.

I was told very explicitly that I cannot repeat the CLE code. I'm assuming that is part of the CLE requirements. If you missed the CLE code at the beginning of the session, please email lawcle@dordham.edu to help resolve that, and they will help you resolve that issue.

Switching over to., and this ties to that idea of thinking about the front end versus the back-end use of AI in providing legal services. There is always that human component as well about, in these use cases, what is the role of the actual person that's interacting with the AI? And Natalie, would you like to speak more about that as people are planning the use of AI in their own services?

Natalie Knowlton:

Sure. I've said this a couple times to the people on this webinar, but the more that I talk about generative AI and its impact on legal services delivery, the more I find myself talking about humans and human mindset. I just want to cue a couple things that have really come out in conversations I've had with people and what I've heard other people talking about in this field as well. The first, I want to just talk about expectations. We often have unrealistic expectations for these technologies. And it drives me nuts seeing the clickbait across legal publications. ChatGPT is incorrect in case law and makes things up. That's not what it was necessarily designed for. It was not designed for a case law research tool. What are the expectations that we can realistically have? And part of that is the expectation of how fast these tools are advancing.

I remember when GPT-3 came out, and it was, "Oh, this is really cool. Maybe not so much for now." Then GPT-3.5 came out, everybody was blown away. 4 came out, "Oh my gosh, sky is falling down." 4 Turbo is out; we're talking about imminent launch of GPT-5. We're going to get to a point where it's going to exceed our expectations, and so we just need to be a little more fluid in what we expect of these tools, whether it's we're expecting them to be more sophisticated than they are or we're expecting them to be less sophisticated than they are.

Talking about accuracy, going back to the issue of hallucinations and people criticizing some of these open-source tools for not doing well, what they're not supposed to do,

what does accuracy mean? And I would encourage all of us to think about accuracy along a spectrum, particularly in the legal services industry. Is accuracy the accuracy of a law firm partner? Is it a first-year associate or a law student? Is it a high school student? Is it the accuracy of Google, which is obviously terrible for a lot of self-represented litigants. Thinking about accuracy along those lines I think is really helpful for us. And it's safe to assume that accuracy issues are going to become less and less of an issue as these tools get more sophisticated and as more legal domain-specific tools are introduced into the market.

I've already talked a little bit about tool conflation. Let's be very careful not to put the realities of what case techs we heard earlier, what their tool can do on what ChatGPT can do. Similarly, to the different models, some are more sophisticated than the other.

Sam actually has talked about this quite a bit, this reality of what consumers are dealing with in the market right now versus what they're getting on ChatGPT. As more and more courts are talking about how they're trying to bar consumers and legal consumers and self-represented litigants from using these technologies when they come to court, the reality is these technologies are sitting on top of a broad internet that litigants are already using that is imperfect, irrelevant, not helpful for them. And so, when we're assessing the accuracy and the efficacy and whether or not people should be able to use these tools, we should be doing it from the baseline of what they already have, which is really bad information, not assuming that using ChatGPT is going to be the difference between prevailing and not prevailing in their case. The self-represented status is often going to be the difference between that. I would encourage us all to do that.

And then, the last piece I think, and a lot of people have talked about it, is just having this learning mindset. When it comes to generative AI and the technology more broadly, the learning mindset is using mindset. You have to use these tools. Ask ChatGPT to plan a vacation for you. You don't need to start with these tools in your legal practice, or when you're assisting self-represented litigants, or in your legal aid organization. Use it personally. And use different ones so that you can appreciate what the pros and cons are. We have to be learning.

And I was thinking earlier actually, maybe generative AI and AI will revamp our dying and really highly criticized CLE industry. Maybe this is what that mandatory CLE has been waiting for, more information about these technologies. We have to be learning all the time in this world. And I would love to hear what my co-panelists have to say on that.

Ron Lazebnik:

Yeah. Sonja, please jump in.

Sonja Ebron:

All I have to add is that I'm glad this is being recorded because I'm going to share exactly what Natalie said with the next regulator that I have to talk to. Thank you for that, Natalie. It's perfectly said.

Ron Lazebnik:

Within that, Natalie, I think something that I heard and I want to want to make more explicit, but feel free to push back or comment on, is a little bit of how the users that you're going to be working with, even if it's an internal AI view, the surety of the AI, we do have a tendency of thinking about, for instance, if we think back to IBM's Deep Blue, and when it finally beat the chess masters, you're like, "Oh, this is a really impressive AI". There's a potential fallacy of thinking that, oh, this AI is really smart because it did that. And I think part of the human element is recognizing that there is still a limitation, that no, it was really smart in doing that thing and not necessarily doing everything.

And to me, one of the things that to highlight as people are thinking about adopting these AIs is that skepticism, and especially with ChatGPT and whatnot. And certainly, people have heard the horror stories of the lawyer who didn't understand the hallucinated cases at this point. But I think it goes more to that, that even once you have an AI tool at your placement, day one, it's not going to be 100%, and even two years in, it's not necessarily going to be 100%, and you need a bit of skepticism there.

Natalie Knowlton:

Yeah, and I agree, it's not going to replace human expertise. The human is still required in the loop, and it's our professional responsibility to provide that. I think on the good front, that humans are inherently skeptical of AI. We are skeptical of things that are not like us, and AI is top of the list there. I think we are all built with this natural skepticism, and I think that's healthy.

Ron Lazebnik:

Transitioning to our next idea, which is the design process, and I'm keeping an eye on the questions to see if the questions are relevant to the topic we're on. I see some questions that are going to be relevant to a little bit more, but we do have one question that ties to design process. I'm going to start with Sonja. As somebody who is looking at this from the technical perspective in the first instance, rather than the legal perspective, what guidance would you give the people in the audience about when they're thinking about adopting AI? What does the programmer's design process look like to some degree?

Sonja Ebron:

Yeah, I think you want to start with, first of all, the ideal outcome. What is the function? What is your goal? What are you adopting AI or building AI for? How do you anticipate it being used? How is it going to improve your ultimate user's experience? Might be helpful to just share what we do at Courtroom5 from a user experience perspective. Again, we serve self-represented litigants handling really tough claims. These are folks that, by and large, have an experienced litigator representing their opponents, and they're handling serious claims. It's not a Judge Judy's type of situation where you can just fill out a form and go and tell the judge what happened, these folks have to litigate. And our job is to prepare them for that.

Our users have two primary challenges in the flow of their litigation. At each step, they need to decide what to file and serve, what procedure to engage. And then once they've made a decision on that, they need to do that effectively. If they're responding to a complaint, they've just been served a complaint, there are a range of options that lawyers know about; many laypeople just have never heard of. And so, our job at each step is to understand the recent procedural history and then to point the user to a range of reasonable options to consider at that step. And we've used AI for a good while on that task just to explain the various options to educate the consumer. We can't tell them what to do for a number of reasons, but we can at least educate them and help them make a good decision on that. That's step one. And then step two, once the user decides what procedure to engage, we have a curated filing experience that we walk them down: Coming up with a good legal position in that procedure, researching the law to find support for that position, and then drafting an effective document that a judge can understand their position.

And so all of that user experience, as you might understand, is very difficult for laypeople. Lawyers went to school to do this. And so, our customers don't have the leisure of going to law school or just absorbing all the necessary information. We've always focused on providing just-in-time education, if you will, making sure the user has just the information they need to get to the next step capably. And, so, the work has been very difficult for laypeople. We're just really excited to be able to implement generative AI on some of those tasks.

And so, with that said, our goal is to reduce the friction all along that path for our customers. And that whole cycle of deciding what to do and then doing that effectively, there's a lot of work involved. There are lots of places where people need help. And we have our limits in helping them because we're not lawyers and just can't go over the line there between information and legal advice. And so, we have found places in that journey for AI to be able to reduce the friction to make the work a good bit easier. For example, when we're helping a user find a legal position for whatever procedure they're engaged in, a motion to dismiss, for instance, we walk them down a path of connecting

their unique facts, the facts in their case, to the legal elements of the claims against them. If you've ever talked to a pro se litigant, they don't necessarily know what's relevant, so we have to educate them on what the law in their jurisdiction requires be proven against them to make them liable for a claim. But they think in narrative, and not all of it is relevant. And so, we found an opportunity for generative AI to parse their narrative, whatever it is, into relevant, very simple statements that then are easily associated with the elements of the claims against them. That's one example.

Another example is on legal research, which is really difficult for even lawyers to do. You can imagine the challenges for pro se litigants. We have a large learning center, as we call it; video-based lessons on a variety of topics. And the large body of that education is on how to do legal research. This isn't something that comes naturally to anybody, but pro se litigants have to spend way more time than most lawyers do issue spotting and reading appellate opinions and that whole bit. Generative AI helps us take the unique facts in the user's case, it helps us pair that with the procedure they're working on, and it helps us pair that with the position that the user is taking to provide them with long-tail keywords that help them search case law effectively. And, so, rather than searching on "motion to dismiss" or "adverse possession" or whatever it is the user might initially want to search on, we can give them, with generative AI, long-tail keywords that narrow the amount of opinions they have to review. And then when they find some opinions that are helpful, rather than spending a week or two trying to understand what this judge or this panel of judges is talking about, we can summarize that case for them very succinctly in layman's terms, to help them really understand whether the opinion supports their position or not.

There are a number of ways that we have thought granularly about the user flow in our application and ways to simplify that work, and in some cases complete the work for them so they can just move on to the next step. Our north star, we don't know that we're going to release it yet, for a number of reasons, it's in testing, is actually creating a first draft of the document for the user. Gets us as close to the BL as you can imagine. We've got a compliance team that's working closely with our developers to see what's going to be permissible there.

But frankly, to Natalie's point earlier, folks are going to ChatGPT to write legal documents right now. We think of what we're doing is building AI with guardrails for prosaic litigants so that they can use this technology effectively. But that's the way that we think about designing AI for our users. And, so, the short answer to your question, and for anybody interested in building AI, start with the user experience, start with the user experience without AI, and then find ways to use the technology to minimize the work, to simplify the work so that it's profitable for everybody. I'll stop there. I can go on for, as you can imagine, go on for a good while with this.

Ron Lazebnik:

All right. As part of that, and Sam, as somebody who's worked on it in Pro Bono Net to help build an AI tool, and I think this is something that a lot of the audience is curious about is how do you learn about the potential AIs that are out there that might solve this. Because, obviously, one option is you hire a programmer and build from scratch, but most people don't have that kind of funding in the world that we're talking about. What is a good way to identify the resources available for people who are thinking about using AI or just learning about it in the first instance?

Sam Harden:

Yeah, I'm a big advocate of hands-on learning. You can go and you can read about what the different tools can do, but to me there's no substitute for going and actually trying things out yourself and seeing what happens. Now, obviously, I'm not saying, "Take your client's files and upload them to ChatGPT, the free version, and see what happens." You should go to ChatGPT and make up a scenario and ask it what it would say. You should take some sample documents that are public record or something like that and use those. But test it out and see what it can do.

There are a lot of tools out there. It's like this weird horse race right now of GPT-4 versus Claude versus Gemini, things like that. And then you have the specialized tools like CoCounsel, like in the big law space, the contract analysis tools, the eDiscovery tools, stuff like that. You can really get bogged down in thinking I need a specialized tool for my use case.

But what some of the research has shown, I think there's a paper coming out about this, is what people find useful is really generic stuff. People said, when surveyed, "I would go back and use the paid version of ChatGPT versus really highly specialized tools because it's like a Swiss Army knife. I don't always need a super specific specialized tool like an Apple corer, I may just need a Swiss Army knife to cut an apple. And that's good enough, and it works, and I don't need to worry about doing hyper-specific stuff." Go and use it; see what it does. It's really good at certain things; it's terrible at others. Weirdly, it's very bad at guessing random numbers. It defaults to 42. And you can probably guess why. But it's really good at analyzing the sentiment from an email, things like that. Use it; see what it does. Do hands-on learning. I don't think there's a substitute for that. And I don't think you should be worried about going and just trying it out yourself.

Ron Lazebnik:

And just because it's related to a question that somebody posted and you highlighted, don't upload confidential documents to the free ChatGPT. What are confidential friendly tools that people can play around with? I know that Lexis and Westlaw are pushing their

Als, and I assume people who have access to them, that those are subject to the same confidentiality that people generally have with Lexis and Westlaw. But what are some of the other tools that people can feel more comfortable with if they want to play around with that you or anyone else on the panel know about?

Sam Harden:

Yeah, the first thing I would do is read the terms of use of some of these tools. They may say, the paid version of ChatGPT, I believe, says, "We don't use your inputs for training data, but we do have the ability to go in and QA, have a person look at what you put in." I think you, as the attorney or the legal services provider, have to decide is that a make or break [the] deal for what I'm thinking about this use case? That's for the paid version of ChatGPT. I think Anthropic's Claude is similar. But then there are versions of these tools that can be self-contained now. Some law firms are using a version of ChatGPT that sits on their private cloud service, cloud computing service, and isn't exposed to anything else other than what they want it to be exposed to. It doesn't communicate outside. That is doable. That's a heavyweight solution. But these law firms are very concerned about client confidentiality and waiving that, obviously.

Google offers similar types of architectures. But these things require development talent, frankly. But I think we'll get to the point, hopefully soon, where they don't, where you're able to say, "I wanted ChatGPT, but here's my use case. Here's what I want." The big question for me is, like Microsoft's Copilot, what are the privacy concerns around that? And seeing where that takes us when law firms really start using that.

Ron Lazebnik:

Yeah. And that's something that I'll say even within the law school before Copilot existed. There's Grammarly, which is essentially also an AI, a simpler AI. But students started using it in the clinic, and I had to have a conversation about, "Hey, have you thought about what it means to use this AI to help correct the letter you just drafted? Do you know where that data is actually going?" Definitely thinking about that. And I will say, for those interested in reading a bit more background on AI resources and whatnot, the CLE materials for this session have a very nice list of resources you can read online at a variety of levels related to this topic.

But going back to the design and involving Natalie in the conversation now, Sonja had mentioned, think about the user flow and let that be the driving principle. And that makes complete sense. How much should the members of the audience already have a conception of that flow before they talk to the tech person? Do they need to come in and be able to whiteboard for the person, "This is what we're envisioning"? Or should they rather anticipate that that's a conversation to have with the tech person? What level of vision should they have for the tool that they thought might be helpful?

Natalie Knowlton:

I'm going to give one answer, but then I would obviously love to hear Sonja's thoughts on this. I would encourage anyone, whether you're in law practice, or court, or legal services organization, to map out the process first, to see if there are any redundant, confusing, irrelevant, unnecessary tasks. Because one of the things we need to be aware of, and this is with technology and automation as well, is if we're just taking these incredibly complicated core processes or internal processes, and then putting AI over them, that's not going to necessarily reduce the complexity on the end of the user. I think it's helpful internally for whoever's thinking about building, to sit down and look at the internal process flow, what a client might experience, what a self-represented litigant might experience, figure out where you can clean that up first and then take that to a developer.

Ron Lazebnik:

Sonja?

Sonja Ebron:

And I think, yeah, that's absolutely right. You have to start with the user and you have to really 10 X your understanding of the work that they're putting into using your product or your solution. And for us, it's something that we do. As a tech startup, we have to really walk hand in hand with our user without their participation; we just watch what they do. And whenever we see friction, long before AI came along, whenever we see friction, we have teams, product teams that go and see if we can make that work simpler. That's just our job. That's how we make money, frankly, and keep our customers. And, so, with AI, this has given us an opportunity to more sharply address those friction points. That is where your design process has to start and end.

I would just add one other thing to Sam's point. I know we're going to talk a bit more about testing, but that is really where the rubber meets the road. You can put a product out there, you can build AI and break down your concept to user flows and all of that, but you've got to watch the product and the while, and adjust it, be prepared to adjust and adapt it based on the experience of the users. There's just no way to adequately test internally. You have to watch what is happening with your users and be prepared to rapidly walk with them in adapting your product.

What we do at Courtroom5 is all of those little bots that I described earlier; we actually have two versions of them. We've built them both on OpenAI or ChatGPT's framework and on the Anthropic Claude API because there are very granular differences in the way they respond. I think OpenAI hallucinates a lot more, for instance, than Claude does, at least at this version. But they're both evolving so rapidly that we've had to build the infrastructure, if you will, to be able to rapidly swap between the two of them. And

there's lots of other tools available that we could have chosen, but it's going to be the experience in the wild that really determines [in] the long-term what your solution looks like.

Ron Lazebnik:

I think this is actually a good moment to switch to that discussion about testing and validation. Sam, springboarding off of what Sonja said, what would you say [are] the best practices for people to think about in terms of, after they have the AI, they think they like, what next? What does that look like?

Sam Harden:

Yeah. And I think the most recent example to me in the news, and just to be clear, AI and AI chat agents have had a long and storied history of embarrassing failure going back, I don't know if anybody remembers the Tay bot that Microsoft released on Twitter that ended up spewing extremely hateful rhetoric because people figured out they could make it do that and things before that and things after that.

And the most recent one in my mind is the NYC chatbot where people figured out that it would say things like, "No, you don't have to pay overtime," or you don't have to have [a] written agreement for something; things that were clearly legally wrong even though it was marketed as something that would answer questions about business law for people. Testing is important, obviously.

And I think the fundamental question for me is if you are building one of these tools, are using AI to help build something that's user-facing, people that aren't you or your attorneys or your staff are going to interact with in the wild, what harms are you worried about? What are the potential harms that things could cause? And what are you not conceiving of? There are known knowns, known unknowns, and unknown unknowns. What are the unknown unknowns that you could conceivably figure out that people will use this for? And how can they abuse it?

There's a lot of methods that you can use for testing. I think they're in the session materials. I put together a quality assurance cheat sheet for this. But test it, test it, test it, test it. As much as you can test it, test it, and don't test it from the perspective of a user that's doing the expected. If you have a chatbot that is supposed to answer legal questions based on articles about legal information, don't get on there and just test, "How do I avoid foreclosure? How do I avoid eviction? How do I file for divorce?" Ask it weird stuff and see what it does. Ask it, "How do I explode the moon? How do I make a birthday cake out of sea coral" or something completely random because people are going to do stuff like that. People are going to abuse it if you put it out in the wild. And at least for me, I would much rather know how it's going to behave before releasing it into the wild. And then you can adjust it.

There's also concerns around privacy, exposure of confidential data to the end user. There are jailbreaking methods that people have come up with that you should think about, things like prompt injection, which I won't get into the details of, and then just your general harmful outputs like trolls out there trying to make a tool do bad stuff so they can screenshot it and say, "Haha, this thing is really bad." You have to think of those scenarios as well, that kind of malicious user intent.

But I think the bottom line to me is test, test, and test again. Don't just test it yourself, have people in the office test it, have your friends and family test it, and then test it again. And redo it and test it again. And you have to decide what level of concern do you have around the thing going awry or slightly awry? Is it make or break? Or is it something that's mission critical? Or is it something that's like it gets the opening hour wrong of your office? Something like that.

Ron Lazebnik:

Yeah, I think that's right. And one of the things to highlight in all that, you used the term jailbreak, which I think the panel knows, but I'm not 100% sure everybody else knows. We're basically talking about instances where you get the software to do something it wasn't intended to do, and what kind of additional liabilities can that create? Beyond the airline scenario I talked about before, a standpoint of that, is there a way to reveal data that wasn't supposed to be revealed by figuring out a set of inputs into the system?

And I guess maybe Sonja, you're best positioned to answer this because you're doing it right now, is how do you mitigate during this testing phase, that you're not giving bad information or that the AI is working like you intend? What are the safety valves that you've put in place that people should think about emulating?

Sonja Ebron:

Yeah, our primary focus obviously is on reliability, on making sure the bot does what you want it to do reliably and accurately. But the negative testing that Sam's referring to is just as important. And basic prompting is giving the bot a personality, if you will, a skillset. And that's really important to narrow the domain, but just as important as telling it what it can't do, telling it the skills that it does not have. And, so, if you try to talk to one of our bots about your robbery charge, a criminal case, or an eviction case, the bot knows it doesn't know anything about that. We're focused on certain types of litigation in US courts. And, so, if you're in Canada or Mexico or somewhere else, the bot just says, "Sorry, I don't know anything about that." And that's because we've told it, "You don't have any knowledge there." And it's really important to be able to train your bots that way. That's probably the biggest bang for the buck in negative testing.

But we also actively, internally, try to break the bot, try to jailbreak it, and try to make it do things that it's not supposed to do. That's just another testing protocol that we have.

But again, the rubber meets the road in the wild, and so this isn't a one-stop thing. We can't guarantee that because our bot is performing as we asked it to right now, six months down the road that it's going to be doing that reliably as well, so it's an ongoing process and protocol.

Ron Lazebnik:

On that training and testing element, one thing that we haven't touched on, but I think is helpful for the audience to know is what's a volume of data that they should expect in training the bot that they should make available? Giving it two pages as the sample of here's what a letter looks like, and here's another letter is obviously too small. But what's a realistic volume of, if part of this is to help an organization in parsing a document or helping to generate a document, what kind of volume of data should people be thinking about as part of the training and testing for that? And so-

Sonja Ebron:

It depends on the application. It really depends on the application. One of the reasons that we've gotten so granular with our bots and train the bot to do a very specific thing, as opposed to having one bot that does a whole bunch of stuff, is because we wanted to use the lowest level of the technology, which is context windows, just giving it a maybe 500-page Word document as training. Two pages would probably not do a whole lot of good. But we've got bunches of data, just historically, and we've used transcripts of some of the videos we've produced, for instance, that sort of thing, to be able to train these things. But that's still very low tech in the world of AI. And, so, on our road map, we're going to get into RAG and embeddings and being able to use a lot more data and therefore make much more capable bots.

But right now, on this first iteration, we're trying to keep it simple. But you do need a good bit of data. It doesn't have to be actual user data. It's possible. And I think there are some efforts to be able to train bots on US case law, for instance. That's a whole lot of data. Wouldn't probably be able to use Claude's API or OpenAI, API on a context window. But nonetheless, those are some of the opportunities going forward to use a whole lot of the legal corpus to be able to train some of these bots.

Ron Lazebnik:

Let's switch over to the last big topic that we're going to cover before we really open it up to questions, which is the issues and challenges. And, obviously, if it's a client-facing AI tool that you're using, and Sonja alluded to this earlier, is you have to think about the unauthorized practice of law versus the authorized practice of law. And there are cases out there, like related to the legal Zoom and whatnot, to some degree, that answer where this lands. But the truth is this is still a very gray area. Because Sonja's company is

not a law firm, it would very easily fall under the unauthorized practice of law, and so Sonja does a very good job of making sure she's staying on the right side of that. But what if it's a legal service provider that has this bot? Do they also need to worry about the unauthorized practice of law? And I'll turn to Natalie if you have thoughts on where that line for the folks in the audience should be.

Natalie Knowlton:

Well, where the line should be is radically different from where it is. And I will tap on you, Sam, just as I'm sure you figured I would, to talk about the research that you've done on the advice and information coming out of some of these large language models. UPL is a really big concern. I've read a disgusting amount of UPL case law across the states. And in the federal courts, the number one issue is often optionality. When you're developing tools or the bots or whatever this is, it limiting the optionality of the people to do something else? And if it is or if it's suggesting a course, UPL flags come up.

I'm not trying to talk people out of it. There are great ways of leveraging legal information. Sonja talked about the corpus of data that you can train something on in order to provide answers and provide information, so there's a lot of potential there. But just be sure that however the output is translated, if it's a direct-to-consumer tool, that it is not suggesting one route or the other, but rather a number of different routes that someone can take. That should be different, but that's where it is. Sam, I think your insights here would be great.

Sam Harden:

Yeah, thank you. The distinction between legal advice and legal information is very much of the court saying, "We know it when we see it." I think that was even in the New York report, the Task Force report on AI where they said, "It comes down to we know it when we see it." There's no clear definition. And the closer you get, the hazier it gets. That's really stifling for these companies trying to innovate around what they can and can't do with AI tools especially.

One of the interesting things that I found; I did a research project a while ago where I asked legal professionals to rate AI responses to anonymized legal questions from an anonymized data set of people asking legal questions. And one of the questions I asked the professionals to answer was, is this legal advice or legal information? And on a lot of the responses, there was a lot of uncertainty. But, interestingly enough, if the response that the AI created contained an adequate disclaimer, what the survey taker considered to be an adequate disclaimer, it was overwhelmingly in favor of it being legal information and not advice, so the disclaimer to the people taking the survey was key, at least from the data that I gathered. You'll see a lot of these tools now include like a

stock disclaimer. If you ask it to role play as an attorney, it will say, "Pretending to be an attorney, I would do this, but you should talk to a real flesh and blood attorney." We'll see if that goes anywhere.

But yeah, the UPL question is a big, big hazy thing. Some jurisdictions have even suggested that chatbots could create an attorney-client relationship without the attorney getting involved. You could conceivably have scenarios where a legal aid organization puts a chatbot on their site, a person goes and interacts with that chatbot and never talks to anyone else and says they have an attorney-client relationship and the court upholds that. Now, I think that's unlikely, but it is possible according to what some of the state bars are saying.

Ron Lazebnik:

Over this past weekend, I was discussing unauthorized practice of law in the context of clinical students in law school. And somebody told me, I haven't verified it, so don't hold me to say whether this is true, but he said that originally, before some of these statutes were passed, the way that law students 1900s, mid-1900s were able to practice law without the licenses because technically it wasn't considered practicing law because they weren't charging. And if there wasn't a fee, it clearly wasn't legal advice. But sadly, that's not how we measure it anymore.

But another challenge with the use of AI, and somebody alluded to it in the Q&A questions, was the concept of bias within the use of AI. And bias can come in a variety of ways. I do think it's important to highlight that. The term that's often used is bad data in, bad data out. If you have biased data going in, you're going to have biased data going out. But biased data going in can look different than you might think. There's obviously data that is prejudiced explicitly, and that of course creates prejudiced outputs as well. But it could also simply be sampling error. That when you created the training set for your system, you weren't thinking about the real universe of people that might be interacting with it or the real universe of matters. And that's a different kind of bias that's being inserted. It's a less malicious one, but one that still can occur.

And then there's the one that you can't really account for and have to keep in mind, which is the systemic bias, that if the system itself that you're dealing with, if you want to think about the criminal justice system, and to the extent that there is bias in there just from a historical perspective, if you base your data on that system, it's going to repeat those errors of the past. You do want to think about the sample set that you're creating, and is it reflective of the true universe that you're trying to represent? And what's in there? And I'll pause there to see if any of my other panelists want to comment on bias and how to mitigate it once you're aware of these kinds of potential biases. It looks like Sonja might have something to say.

Sonja Ebron:

I think it's one of those things that you're aware that it's going to happen, happen, but you can't anticipate it or how it's going to appear. And, so, it's one of those things that comes in and watching it in the wild and addressing it and trying to mitigate it as you see it happening. That's at least the way we're trying to address it, that we know that it's going to happen. And, obviously, civil justice is different from criminal, and so dealing with a different set of biases and some biases against prosaic litigants and that sort of thing. We know that it's there, we know that it's going to show up even in artificial intelligence responses, but there's no way for us to practice or mitigate it beforehand, so preparing to watch it in the wild.

Ron Lazebnik:

Yeah, I mean this is where this, oh, go ahead, Natalie. Please.

Natalie Knowlton:

Well, I was just going to suggest that for those who are interested in the issue of bias, I highly recommend you follow what's happening in the medical field because there you see real bias in these diagnostic tools based on the data and the studies that have been done in the past. And I think there's a lot of real interesting stuff going on there. That would be my recommendation.

Ron Lazebnik:

Yeah. And I will say that the National Institute of Science and Technology has a very nice paper about AI and bias. It was written before generative AI. But what they talk about I think is just as true for this round of AI tools as it was for the AI tools that they were discussing a few years ago.

Another potential challenge that we thought of for the organizations here that is worth highlighting is to the extent that you guys are becoming more productive and more efficient and increasing your bandwidth by using these tools, so these are more internal facing tools than outward facing tools, this is potentially affecting your equivalent of the billable hour. Now, obviously, legal service providers don't do the billable hour the same way that law firms do, but it is important to highlight that. What does it mean to have a more efficient system, that you're helping more people?

And I raise that because often there is the request of grant funds associated with legal service providers. And what are you reporting as the metrics for success for your grants? And how will being more efficient in this way affect that? Do you want to have a conversation with the people who are providing the grant or asking for the metrics? That some amount of change has to occur.

And I'll pause there to see if anybody wants to talk more about that. The billable hour in general is something that all lawyers love to lament on generally, at least those who were in private practice. But it's not going to disappear yet, but it is worth thinking about, what does that mean when AI is helping you become more efficient in what you're doing? There's obviously also the challenge of hallucinations that we've talked about. And there's no complete, easy way to do that other than good training and putting good limiters on what it is that the AI can say or the disclaimers that it's going to talk about.

And I always like to highlight that separate from all of this, you do want to think about the environmental impact that training these AIs and running these AIs have. It's a lot of energy and power. And you want to think about the global impact all of us running these AIs might have. And just to the extent that that is something that is occurring right now, how are we contributing to it? And how are we mitigating our contribution to it, is worth thinking about a little as part of that discussion. Any other issues or challenges that I have not covered that my fellow panelists want to touch on really quickly? There's always more, but we'll leave it at that for now.

Before we open it up to full questions from the audience, we've touched on a lot of different aspects, and I think either Sonja or Natalie mentioned that the next ChatGPT version is coming out, so what are the potential opportunities in this space that we haven't seen through the use cases of the lightning rounds earlier today that you're predicting if people are still thinking through tools? What are some things that we might see in the next, we'll say, five to 10 years in terms of the tools that might be able to help legal services? What do you think? Sonja, go ahead.

Sonja Ebron:

I'm not sure I want to address that question, so I'm an expert at changing the question. But I want to talk about, as a person who doesn't practice law, never stepped foot in a law school. But one of the opportunities that we have found at Courtroom5 is to actually help our customers make good use of limited-scope services. And I know there are a lot of lawyers these days who are struggling to keep the lights on in their practices because they can't find enough paying clients. And, so, we actually have paying clients, but they need different terms. Our customers aren't looking for full representation; can't afford generally full representation.

And the ideal solution when I look at the access to justice crisis is a market for limited-scope services. There are just a real ton of reasons why more lawyers don't do limited-scope services. I think the risk of giving advice to people who are handling their own matters, particularly when you're a small practice, thinking about what's this person going to go off and do with the advice I gave them or the document I ghostwrote for them, or whatever it is. The risk of malpractice or bar complaints, it tends to outweigh

the economic opportunity of offering limited-scope services. We've been lucky to find a couple dozen intrepid lawyers around the country who are willing to work with some of our clients who can demonstrate the capacity to do well with their advice.

But I think the way we look at generative AI and the potential there is to prepare more pro se litigants to be able to demonstrate a capacity to profitably use limited-scope services and potentially reduce or mitigate the risk for a lawyer of providing those services. In other words, we're looking at making a market for limited-scope services and having a bigger impact on access to justice that way. That's just one of the opportunities that I think is staring us in the face when we look at the power of technologies like gen AI.

Ron Lazebnik:

Sam?

Sam Harden:

Yeah, I don't know if I can prognosticate the next five years with AI. It depends on the day for me. Some days I think we're in the peak of the hype cycle, and now we're going to figure out what it's actually good for, and then more hype comes out. I think what we're seeing now, especially with companies pushing out what they're calling autonomous agents, I don't know if anybody's heard of this, but Google just released this agent builder, which is a no-code way to create a chat agent that has access to your company data. It's a customer service focus tool that can look up orders if you're an e-commerce store, that sort of thing. But the fact that they can do that with you just telling it, "Here's what I want you to do," is remarkable. And I think there are applications for legal aid offices, for law firms, things like that, to take advantage of some of those capabilities. Now, to do it responsibly, obviously, but to take advantage of them.

I also think we'll see a world where AI is not its own thing, that it's part of other things. I've always thought lawyers, especially, find it hard to go and use new tools, not because lawyers are Luddites, but because lawyers already have a set of tools that they use day to day. And it's hard to break that workflow. I think the successful AI methods and products are going to be part of what lawyers already do and help them streamline those processes versus having to go and use a second research suite or use a second word processor or something like that. I don't know, we could all, in five years, not have conferences, it could be AIs talking to each other. There's always that possibility as well.

Ron Lazebnik:

Well, I didn't want to reveal it until now, but I'm actually an AI. No. Natalie, any additional-

Natalie Knowlton:

Yeah. I think we're going to be able to see and develop a lot more tools that are capable of delivering legal advice direct to consumer. And before that scares people, consider that helping a self-represented litigant determine what 15 business days means, what that actual date is, that's considered legal advice. We have tools that are ready, can do legal advice, I just believe we're going to have a lot more and they're going to become a lot more sophisticated. And I don't think this is a threat to lawyers because I think it will actually allow lawyers and law firms to move down market to begin offering more one-to-many solutions that are tech enabled, as opposed to that one-to-one traditional law service delivery model. That's where I think we're going and also where I hope we're going.

Ron Lazebnik:

Great. Any other questions from our audience? I think we've answered all the ones that were typed up. And we have about five minutes if anybody wants to type in an additional question into the chat.

Sam Harden:

There was one from Kim Susser that I think I would love to hear Sonja and Natalie's take on this. She asked, "Can you address, big picture, how firms can use AI to increase their pro bono?" Big picture thinking, that kind of thing.

Ron Lazebnik:

Natalie, you're already unmuted, so go ahead.

Natalie Knowlton:

Oh, sure. Well, I think companies like Paladin are certainly going to be on top of using generative AI to increase the connections between firms who are interested in providing pro bono and clients who need pro bono, so I think that will only get more sophisticated. But in terms of what individual lawyers within a firm can do to increase their efficiency when providing pro bono, it's just all about back to those tools. What are the case-type-specific tools that can be leveraged in order to take the time that an attorney would've, taken an hour of the attorney's time and condense that into five minutes? That means more pro bono. I would just focus on those internal tools in that instance.

Ron Lazebnik:

Yeah. And Sonja?

Sonja Ebron:

Just briefly, I would go back to limited-scope services. If there's interest in the firm in doing limited scope and offering that pro bono and preparing the client with some justice tech ahead of that actual delivery of service would be a great application if they were interested.

Ron Lazebnik:

Yeah. There is a question about the role of the courts in all this. And if courts are willing, there are many courts that have FAQs to help pro se; some have pro se offices. But if a court is willing to experiment a little about helping the pro se litigants understand the pathway, the more procedural elements, and maybe pointing them towards where the substantive questions can be answered, that would be great because an AI chatbot at the court, as Natalie points out that you can ask, "What does it mean for me to serve this within 15 days?" and the chatbot can say what the court believes is the answer to that, that would be amazing.

Separate from that, within the court system itself, you could envision AI that's internal to the court to help summarize materials that have been submitted so that the clerks and the judge are positioned to process certain motions faster than they're docked currently now. Some of those kind of lower-hanging fruits that might come in high volume but still take up a lot, don't require a lot of thought, but still require a lot of time for the court to do that now. I've been asked-

Natalie Knowlton:

Can I just-

Ron Lazebnik:

Oh, go ahead, Natalie.

Natalie Knowlton:

Oh, sorry, Ron. No. Are we out of time?

Ron Lazebnik:

I was just going to say, people have asked once again for the CLE code. I'm not allowed to repeat it, but if you did not manage to catch it, you can email CLE Fordham. I'm sorry, it is, I'm going to look it up right now. The email is lawcle@fordham.edu. That's lawcle@fordham.edu. Go ahead, Natalie. Last word.

Natalie Knowlton:

Oh, just to the judge's question. A couple points that I think are important. When courts are thinking about this, there's a couple of threshold steps that I would encourage anyone, especially court leadership, to consider. The first is that the self-help information that's available could very well be inaccessible for any number of reasons. It could be too complex, it could contain legalese, it could not contain certain wayfinders, so simplifying the self-help information is a first threshold step before you add any generative AI on top of it.

And then also with respect to litigants finding the information, there's a sea of junk out there. Courts don't even come up number one when someone is Googling, I have a divorce in Philadelphia. So, make sure if you develop tools that are helpful to people to get the word out, SEO for your website, send the information to self-represented litigants, make sure people know that it's there or else it won't get used. That's all I got.

Ron Lazebnik:

Absolutely. And with that, I want to thank my panelists for what I think at least is a great session. And there is an opportunity to give feedback. Barbara has shared the link to the forum if you would like to share your feedback on the session. And now we ask that you go back and use your original Zoom link to rejoin the plenary that will restart soon. Thank you, everyone, for joining us.