Tim Baran:

Hi, and welcome to the Rapid Fire Tech session on the innovative uses of Gen AI to expand access to justice. I'm Tim Baran, Program Officer for Civil Justice Initiatives at Pro Bono Net. Today our six presenters will be sharing their innovation using a 20 slide, 20 second per slide format. The slides will auto-forward, so this will be a bit of an adventure for our very game presenters. For those seeking CLE, be on the lookout for a CLE code. It will only be announced once. Registrants will receive the CLE forms and materials via email on the conference webpage.

This session is being recorded and the recordings will be posted online after the conference. Now, this is a pretty fast-paced fifty-minute presentation. However, we should have some time at the end for questions. So please enter yours in the Q&A window and we'll try to get to as many of them as we can, and I'll introduce each presenter very briefly. You can find more information about them on the conference webpage, and so I would advise you to take a quick look just to find out more about our presenters. All right, let's get to our first presenter, Teri Ross. Teri is the Executive Director for Illinois Legal Aid Online. Take it away, Teri.

Teri Ross:

Thank you, Tim. All right, good afternoon, everyone. For the next six minutes, I will share some of what my organization has learned about using generative AI in the context of translation, both in complex to plain language and in English to Spanish. My name is Teresa Ross and I'm with Illinois Legal Aid Online. I go by Teri and my organization goes by ILAO. So why me and why ILAO? Here is a bit about the organization I have worked for for 5 years. Also included here is my avatar, care of ChatGPT-4.0's Photo Multiverse. How did it do? Not so good. You'll be disappointed when you meet me in person. So while this is intended to give some credibility, it's important to point out that I am not an expert. This image looks more like me, by the way. While I have dabbled in many technology systems and platforms, I am by no means an expert and I don't play one on TV.

I'm here to share with you what we've tried, learned, and are likely to pursue further as an organization. My focus today is on large language models, or LLMs, which are a component of the broader category of generative AI. How can LLMs be used to streamline and improve Spanish language translation and plain language readability? Can they be leveraged to increase efficiency and reduce costs? So, to understand its capabilities and limitations, it's important to remember that LLMs are built on data, ginormous amounts of data. LLMs operate on predictive analytics. Our society is generating more and more data at an increasingly exponential rate, and the more useful data we generate, the better LLMs will become. For instance, ChatGPT-4 is markedly

better than the 3.5 version. Here's how our current translation and plain language practices work.

For Spanish, we use a professional translation service, which is then reviewed by a fluent speaker on our staff. Our annual cost is about $40,000, so using an LLM could save us a lot of money. Our explorations of LLMs have focused mostly on how they can increase efficiency of routine but critical business activities. If you haven't used it before, here's what ChatGPT looks like. I find the suggestions helpful, but they can be somewhat daunting. It's like facing a blank white page when you're typing a brief. I'll start with plain language and note that our explorations have been primarily with ChatGPT-3.5 and 4.0. Here's my prompt on how it can be used to improve complex text. There are not yet many examples of scholarly work on the use of generative AI for plain language. This is one. This is a proof of concept in the medical profession on how it could be used to lower the reading level of patient information in two medical journals.

Note that the JAMA text is at an 11th grade reading level, even though it's patient information. So we're not the only profession that has this issue. Generative AI has given rise to a new phrase, which sounds otherworldly, but it's crucial to ensuring that you're getting what you want from LLMs. The important takeaway from this is to experiment, to test it out, see what works well and what doesn't. And I'll give you some examples. In the left column, I took texts from the Illinois State Bar Association webpage, which is meant to educate the public about lawyers. Its original reading level is an 11.7-grade level. You can see the prompt I wrote and the results it gave me to the right. I was pretty happy with this result. I wanted to see if I could do better. Next prompt, I clarified that it should be not for a seventh grader, but for somebody who reads at a seventh-grade-reading level, I didn't like this one as well.

And then on the right, is the result when I prompted it to apply best practices for digital usability, and gave it a set of rules to apply, including usability. It's not perfect, but it's much better than what we started with. So a few tips for doing your own evaluation of large language models for plain language. Take the time to experiment with different prompts and pay attention to the results to identify what works best. If you have style guides, include them with your prompt and ask the LLM to apply them to its response. That is really important. So plain language is the easy part. Now we'll shift quickly to actual translation, which holds more controversy. To start with, let me say that the legal community is really of two minds on machine translation. On one side are people who say something is better than nothing.

On the other side are people who say translation quality has too many inaccuracies and the risk of harm is too great. There is a happy medium, which is to use machine translation as a first pass, and then human review by a native speaker for context, meaning, and localization. I found one resource on training a chatbot with Spanish language legal data, specifically in Spain. And this was interesting, because the problem

arises that it's one of not just linguistic variety, but also jurisdictional difference, and in the end, there just wasn't enough data for the Republic of Spain legal information for them to create a useful chatbot.

So, a few key elements to consider when translating generally that also apply to better translations for LLMs. The extent to which you can build a glossary of legal terms, and a style guide of how certain words or phrases should be translated, the less cleanup will be required. Better translations results from texts that had already been translated into plain language. And also, some tips for doing your own evaluation of large language model. Some I already mentioned are experimentation with translating Spanish for someone in the Chicago area captured the localization context quite well. So, we're really happy with that and we're going to pursue it.

One final point, one of the most useful things about software models like LLMs, is that they are consistent, unlike humans. ChatGPT does not have a bad day. It's not thinking about what it's going to have for lunch, lamenting the small number of hours of sleep it got last night, or how the weather is affecting its sinuses. So happy exploration, gen AI is happening, and we are all in this together. Let's go forth and explore, and I encourage you to reach out to me if you want to learn more.

Tim Baran:

Awesome, thank you, Teri. And could we go back a second to the CLE slide? I think a few folks missed that, so if we could show that just for a second. Okay, so the CLE code for this session is R, Robert, Frank, T as in Tom, 24. Awesome, thanks. Let's move on to Conor Malloy. Conor is an intervention director at the Law Center for Better Housing. And Conor, I think the slides will be up in a second.

Conor Malloy:

Thank you, Tim.

Tim Baran:

Yes.

Conor Malloy:

Nearly five years ago, Rentervention embarked on its mission with a single chat. Conceived as an expert system, it was designed to scale and democratize access to justice through an innovative chatbot, transforming how legal support is delivered. To appreciate Rentervention's evolution, we must first confront America's justice gap. As efforts unfold to bridge this divide, Rentervention emerges as a pivotal piece, blending technology with advocacy to forge new paths to legal access. In the U.S., the justice gap sharply impacts housing, with millions lacking the resources or representation needed to

defend their rights. This disparity turns housing into a battleground for fairness, demanding innovative solutions to bridge this gap and ensure equitable access to justice for all.

Traditional attempts to close the justice gap by simply increasing staff have consistently fallen short. Despite the best efforts and intentions, scaling human resources alone hasn't been sufficient to meet the vast growing need for comprehensive legal aid and housing. Rentervention is our response, blending people, process, and technology. This trinity aims not just to bridge, but to transform the justice gap in housing, making fairness accessible through an innovative ecosystem. And so, Rentervention thrives on three pillars, our process, our tech, and most importantly our people. Our dedicated staff, they're the heartbeat of our mission, pushing justice forward. In our service vertical, technology and human insight combine to tailor assistance levels. It's a seamless integration, ensuring every user receives the right support at the right time.

Our conversational AI is more than just technology. It's the bridge to justice for many. Intelligent, responsive, and always learning, it ensures no call for help goes unanswered. Beyond our bot's reach, we bring humans into the loop, blending tech with the irreplaceable touch of human assistance. It's not just about building technology; it's about ensuring we serve others with empathy and precision. At Rentervention, AI is pivotal. From understanding complex inquiries through natural language processing, to crafting responses with generative AI, and enhancing accuracy with retrieval augmented generation, each technology layers our service with unparalleled sophistication and empathy.

Natural language processing, or NLP, is the brain behind understanding human language in our system. It allows Rentervention to grasp the nuances of each query, ensuring every individual's concern is accurately interpreted. Generative AI empowers Rentervention to not just understand, but respond creatively. It crafts tailored information, drawing from vast legal knowledge. The data doesn't just answer, it adapts, providing personalized assistance that evolves with each interaction. Retrieval augmented generation enhances RAI by blending stored legal knowledge with new tailored information. It ensures responses are both precise and current, marrying expertise with innovation. Our journey unfolds through key areas, blending, learning with action, building for growth, iterating based on feedback and fostering community. These steps outline our path to making justice accessible for everyone.

Our foundation is in marrying pedagogy with praxis, teaching and applying legal knowledge, not just in theory but in action. This educational approach empowers users, fostering a deeper understanding and practical navigation of legal challenges. Rentervention is designed with scalability at its heart. Leveraging advanced AI, we're poised to meet the growing demand for justice, ensuring no issue is too big or small for our system to handle. Our journey is one of perpetual growth. Continuous improvement

through AI learning and user feedback shapes Rentervention, making it smarter and more responsive with each interaction. And feedback is our compass. It guides our development, ensuring Rentervention remains attuned to the real-world needs of its users and the evolving justice landscape. At our core is the belief in the power of community. Rentervention thrives on support and collaboration, uniting its users, legal pros, and advocates in a shared mission to democratize access to justice. As we look forward, Rentervention is more than a tool, it's a movement. Join us in reshaping access to justice, where technology, humanity and the law converge, to create a fairer future for all. And thank you.

Tim Baran:

More than a tool, but a movement. Thanks for that, Conor. All right, next up is Jonathan Pyle. Jonathan is the contract performance officer at Philadelphia Legal Assistance. Take it away, Jonathan.

Jonathan Pyle:

Good afternoon. So I'm an attorney and coder working at Philadelphia Legal Assistance in a management administration role. And unlike Conor, I haven't built any production-ready AI tools, but I have dabbled in seeing what AI can do, and that's what I would encourage other legal services organizations to do. Even if you're concerned about the confidentiality and accuracy issues, there are things you can do to experiment. For example, I tried giving Microsoft Copilot some of my aggregate case data, asking it, "I'm writing the annual report for 2023, and I'd like you to draft a few paragraphs for the annual report that highlight what's changed since 2018." Then I gave it a few pages of data in CSV format, which is easy to generate in Excel. And Copilot noticed that our bankruptcy cases were down, public utility cases up, and even invented some pretty good explanations for the changes.

Anyway, I think we should encourage staff to look for low-risk ways like this to experiment with AI, because it's the best way to learn about it. If you're willing to share confidential information with cloud-based AI providers like ChatGPT, the case notes and a legal aid case management system are a really interesting source of information for the chatbot to use. I built an integration into my case management system to test this out. I created a prompt for ChatGPT that said, "I'm going to give you the case notes from a case management system of legal services organization." Then I explained what our acronyms were and asked, "What are the next steps that the advocates need to take in this case?" And I passed the data using this JSON format, which stands for JavaScript Object Notation, which is just a format that ChatGPT knows how to digest, and it came back with a 21-point list of items, all of which were perfectly good.

Now, I wouldn't want my advocates to rely on this, but the chatbot could still be a useful helper who creates a checklist that the advocate wouldn't have time to make for themselves, and it wouldn't hurt to go through this checklist just to check things off, and the advocate would know if one of them was incorrect. But even more daring, how about asking ChatGPT to draft an advice letter to the client based on those case notes? Now you might have heard of ChatGPT's tendency to hallucinate to cite non-existing cases. But in my testing where I was just asking it to work with a specific text, I didn't [find] that it made stuff up. The main problem was that advice letters like this were just a little bit too general. But anyway, as legal aid organizations experiment, their first uses of AI will be internal for drafting things like this.

But I think we need to follow Conor's lead and work on external facing tools. In Philadelphia, the need for assistance with child custody far outweighs the supply of family law legal services, and most of our clients just receive assistance handling their case pro se. But what if this assistance could be provided with AI? I did a test focusing on the best interest of the child's standard. One of the elements in Pennsylvania is the need for stability and continuity in the child's education, family life and community life. So I made this app that asked the user to write a narrative about why they think they should have custody of the child. Here's a fake prompt where I say, "I should have custody because I'm nicer, my house has more luxuries," but I also talk about how my daughter needs to get to school on time and I say that it would be better for my child to live with me, because my husband, he lives in the boondocks.

So I sent this narrative to ChatGPT prompting it to say with a phrase, "Does the explanation address the need for stability or continuity in the child's education, family life or community life? If so, then explain that factor at a college educated reading level, and if not say not enough information."

In its response, ChatGPT was able to find the facts relevant to that specific legal concept, filter out extraneous info, and rephrase the argument in a professional manner. And that's pretty similar to the type of assistance our paralegals provide. I especially like how it replaced the reference to the boondocks with, "my husband's residency in a more remote area considerably limits her social interactions." So another possibility is a client-facing app that uses data sources to give plain language explanations of where a client is in a legal process, kind of like GPS for self-represented litigants. Now clients can go online to read their dockets in an eviction case for example, but these sites are designed for lawyers and there's a lot of confusing legalese.

So I built a prototype app that asked the user for their name and address, and then sends that information to an API I developed that provides the docket entries and the text of a complaint and an eviction case. And then I fed all that data to ChatGPT with the prompt, "Tell the user, who's the defendant in the case, what's happening, and if there's an upcoming hearing, mention the importance of attending the hearing."

And one of the good things about being able to structure the prompt this way, is that I can precisely control the information that ChatGPT sees, so it doesn't get extracted by extraneous docket entries or its background knowledge. And ChatGPT came back with a pretty good explanation of the history of the case. It explained what the complaint alleged and noted that the court initially found a problem with the complaint. It did go a little too far in speculating as to why the plaintiff filed an amended complaint, but it did a really good job explaining the judgment by agreement that the tenant entered into with the landlord's attorney.

And I've done some experiments since I made these slides and found that if I teach ChatGPT more about the Philadelphia eviction process, it does an even better job explaining things and giving advice. I also tried a second example, where a tenant had an upcoming hearing, and ChatGPT all on its own did a great job explaining when and where the upcoming hearing would take place, and stressing the importance of attending that hearing.

I think the consensus of the legal aid tech community is that before we give clients AI-based tools that replicate what a human helper does, we need to measure the performance of humans as baseline, and then measure the AI's performance and see if the relative error rate is acceptable. We need to test for incorrectness, hallucination, bias, and equity, and I'm sure that we will encounter all those problems, but the question is whether we can iteratively improve performance through better prompting. And I'm continuing to experiment with this stuff. Last weekend I was at a hackathon where we worked on that eviction app some more, and we integrated it with other APIs so that it could summarize the code violations on the tenant's property and schedule the tenant for a housing counseling appointment with a human down the line using a Google calendar integration. And I used to be really skeptical about ChatGPT, but now I think it has enormous potential for helping us provide effective assistance to a greater number of people.

Tim Baran:

Awesome. Thank you, Jonathan. And Jonathan is very humble. He's actually the creator of Docassemble, the widely popular open-source system for guided interviews. So, thanks for that Jonathan. All right, up next is Jalon Fowler, a JD candidate at Northeastern University School of Law, and CEO Action for Racial Equity fellow at John Hancock. Take it away, Jalon.

Jalon Fowler:

Thanks, Tim. Welcome to AI-EP, a tool that uses AI to ingest, translate and summarize a family's individualized education program, or IEP. AI-EP improves families engagement and participation in their children's education by facilitating better communication and

access to educational resources. Comprehending and navigating complex and lengthy IEP documents can be challenging. We partnered with Innovate Public Schools to support their mission to help parents simplify and summarize IEPs. All improvements directly benefit students on IEPs. By understanding parents' experiences and perspectives, AI-EP was designed to simplify their access to educational justice. AI-EP's tailored AI solutions address problems that were identified by the families. The upcoming slides will delve into the identified problems.

Look at these couple of sentences from an actual IEP. For those of us with a strong command of the English language, deciphering these terms can be quite challenging. Now imagine how overwhelming they must feel for those with lower literacy levels, or for non-native English speakers. In addition to navigating complex medical, educational, and legal jargon, IEPs often span dozens of pages and new pages may be appended throughout the IEP year further complicating an already dense document. Slower delayed translation, along with inadequate or flawed translations, can lead to significant learning disruptions and hinder informed decision-making. These challenges highlight the importance of addressing language barriers promptly to ensure effective communication and equitable access to information.

These National Center for Education Statistics underscore the significant number of families with IEPs. With 28% of students identified as English language learners, it suggests that their parents may encounter comparable language barriers, which complicate communications for families. We learn that parents want clear translations and simplified explanations, which enable them to fully participate in the IEP process. This empowers active collaboration with teachers and specialists, facilitating well-informed decisions regarding their children's education. AI-EP leverages the ingestion translation and summarization powers of AI, which are checked by a human review process.

Additionally, AI-EP is enhanced by collaborative features in a social layer, fostering greater engagement and knowledge sharing among users. AI-EP utilizes AWS translate, an AI translation service to translate from English to grade school level Spanish. Given the complexity of the terminology, it was crucial to ensure translations were at a level accessible to parents. Once translated, Innovate collaborates with parents to review the translated IEPs. The chatbot feature allows parents to ask questions about their IEP, or the IEP process, or to receive simplified explanations for complex terms. The chatbot references a knowledge base containing information from authoritative sources that have undergone human review, then it uses the power of AI to return a simplified explanation. Using AWS translate, the IEP Digest presents translated text informed by feedback from parents as part of the human-centered design approach of the project.

AI-EP's knowledge base features an evolving knowledge base that continuously expands its collection of IEP terms as it trains on more IEPs. The terms are sourced from

authoritative references and reviewed by Innovate Public Schools. Parents can access the knowledge base to learn, to explore and to gain valuable insights. The Q&A feature uses a combination of the knowledge base and generative AI to provide simplified answers to questions asked by parents. Innovate Public Schools also reviews questions and answers to ensure the quality and accuracy. This is how AI-EP harnesses the power of AI to help families overcome language barriers. AI-EP promotes independent and direct access for parents to matters relating to the children's education, empowering them to navigate the educational landscape with confidence.

The text-to-speech feature will allow parents to listen to the content of their IEP, enhancing accessibility and accommodating different learning preferences. This feature will ensure that individuals with visual impairments or reading difficulties can independently access and comprehend important educational information. We plan to expand our language support to cater to a wide range of communities. Our goal is to ensure that all parents, regardless of their native languages, can access essential resources and information in a manner that is clear, accessible, and inclusive.

AI-EP was developed with many partners. The fully functional tool is being tested by Innovate Public Schools' parent community in California. Innovate Public Schools is a nonprofit organization dedicated to helping families to organize, advocate, and demand high quality schools for their children. AI-EP was developed with many partners. The fully functional tool is being tested, as well as being guided by the generous knowledge of the people on the screen. The original team consisted of engineering students and law students. The ongoing effort is being managed by AI for Impact co-op students, who use the power of AI to make social change. Thank you. Back to you, Tim.

Tim Baran:

All right. Thank you, Jalon, and a shout-out for mentors. I love that last slide. All right, next up our fifth presenter, five of six is Quinten Steenhuis. Quinten is the Co-Director of the Legal Innovation and Technology Lab at Suffolk University Law School. Take it away, Quinten.

Quinten Steenhuis:

All right, thanks, Tim. So today I want to talk to you a little bit about how older AI technology expert systems, those old flowchart-based tools, can have a happy marriage with these newfangled generative AI tools. And I think there's a lot of promise in the combination of the two. First of all, it may be helpful to talk about why we even want to get better than those ChatGPT interfaces that just let us talk to the system in a natural way. In essence, they're doing something like math on words. So, you can see here where, if we subtract from the word queen the word woman, and add the word man, we can get king, but they can do some negative things too. Here you see that, if we take

from the word nurse the word woman, and add man, we get doctor, and that's revealing a built-in bias in these systems, where they're associating gender concepts with words that don't have an inherent gender.

So if we want to get past that and do math in a safer way, where we're not allowing it to be infected with that same level of bias, we have to think about things we can do with generative AI that don't require things to be true or false, and that don't have that room for bias to creep in, as well as the other problems you heard talked about such as hallucinations. Here are a few ways that in the Suffolk Law School Legal Innovation Technology Lab, we've been using generative AI to help improve document automation, which is very safe and predictable.

The first is that we're using it to ask follow-up questions. Traditional document automation is really rigid, asking follow-up questions just where we need them with generative AI as a way to make it a better user experience. We can also use generative AI to get information out of the user's response, and that can let us provide a better active listening model than traditional document automation can handle. You heard a little bit of, in the earlier sessions both from Jonathan and Teri, how you can use generative AI to translate the tone or actually the language of someone's input. And finally, I'm going to talk to you about an exciting way that we've been experimenting with using generative AI, GPT-3 and 4, to help build draft document automations. So it's speeding up that work that can be very time-consuming.

Let's dive right into some of those examples. What you're seeing on the screen is a simulation of what might be a really tricky question. We could do a lot better than the language that we have on this screen. Let's pretend this was a real question that we asked a tenant to explain why they wanted to get some repairs from their landlord. They might do not such a great job at first answering this. And we're going to provide both the user's answer and that rubric to ChatGPT.

Here you can see that it's read that rubric and it's read the user's answer, and ChatGPT, the API, is generating this follow-up question. And we can keep doing that until the user gives us a complete response to this open-ended tricky question. We can make sure it's good enough to persuade the judge. And if we trust ChatGPT, we can have it take the user's words and just rephrase them a little bit, so that series of questions and answers makes it complete response. Or we can go as far as you see here, as having it generate a full complaint. In this case, a demand letter to the landlord.

This is an example of us doing that extraction. So here we only care about three bits of information in this user's answer. Their name, their address, and their favorite movie, but they can tell us a lot more, and that might let us skip some of the questions that come later in the interview. We should be good active listeners and do that. Here you can see the person has given only an answer to two parts of that and we still need to get the person's name. So we're not using a question that's written by ChatGPT. This is one of

our questions that we've prewritten and tested, and we know is good, and we're using that when the user's answer didn't include this bit of information. Finally, we give them a chance to review it, and make sure everything is correct. They get to decide if our interpretation of what they said was right, but we can still skip that annoying set of questions that makes them repeat themselves.

One of the things that I'm really excited about using GPT-3 and 4, is to help improve the user's ability to communicate. And Teri talked about some of the ways that you can use translation. You can do it inside a guided interview as well. Imagine that a landlord only speaks English and the tenant only speaks Spanish. We want to have a guided interview in Spanish, but we can give the person's answer in English in the letter. Maybe we want to do it side by side at first at least.

We can take all of this work one step further. Imagine what Xzibit would think if he thought we were really big fans of document automation. How much better would it be if we could automate our document automation, and we could actually build some of that automation with an automation tool? Well, we've done that at the lab. We built a tool that takes all the different pieces of a document automation platform and it makes them themselves automatable. What you see on screen right here is an example of a sample letter that comes from the Attorney General in Massachusetts. And underneath this is a draft of just a section of that document that's been generated completely by generative AI. That lets us skip a really time-consuming step at the beginning of automation. You can take it a little bit further and we can take that template that we've created with the help of generative AI for the cost of about 20 cents, and we can actually build a draft automation.

That automation is going to be editable by a professional before it gets shown to a real user. You see here a mix of screens that were generated by generative AI and ones that kind of come from our library of questions that were part of the more traditional document automation approach to document automation. And again, the important thing here is we're building a draft that can be reviewed by a person. I can talk about document automation all day, but I'll stop there and you can hear if you want to learn more about what we're doing, here's my contact information.

Tim Baran:

Awesome, thank you so much, Quinten. And our final presenter of the day is Adrian Palma. He is the pro bono manager at Microsoft. Take it away, Adrian.

Adrian Palma:

Thank you, Tim. And thank you to all the other presenters. My name is Adrian Palma. I am Microsoft's Pro Bono Program Manager. My role at Microsoft focuses in the intersection of pro bono digital transformation and AI, and I support Microsoft's 2000

legal and business professionals in our legal department, including our 800 volunteers. And so, today you're going to be learning about an AI use case that we have developed for veteran discharge upgrade file reviews. So, through this presentation you're going to get an overview of NVLSP, the work that they do, the file review process, the opportunity for digital transformation, and some of the tools that we delved into to develop this opportunity, including Microsoft's Purview eDiscovery, Clearbrief, and Microsoft Word CoPilot. So NVLSP, or the National Veterans Legal Services Program is a DC-based org. They help veterans and they have an opportunity for discharge upgrade file reviews, where program of volunteers review a veteran's case, and they determine if they have opportunity to change the characterization of their discharge, and then be able to access benefits that they are entitled to.

And so, the file review process is extensive. The volunteer reviews the veteran's files, which could be up to 20,000 pages long. They produce an Excel sheet that maps about 120 keywords to each responsive page. And then they produce a narrative summary with timelines of the key statement of facts and a narrative summary of their analysis. And so, this is what that final work product looks like. On the left-hand side, you see an Excel sheet, every single page of the veteran's file is included as a row in the Excel sheet. And then in column D, you have the keyword tagged. And then on the right-hand side, you have that file review summary that has a key timeline and the key statements of facts for the veteran's case. And there are many challenges and pain points to this process. It could take program volunteers upwards of 50 hours to complete a file review, just given that some of these client file sizes are very large.

So, it's a very time-consuming process, and it is also very much prone to human error if a program of volunteers is reviewing this manually. And so, this opportunity is ripe for digital transformation, just given that we at Microsoft, we're not able to scale our work with NVLSP, just given how manual tedious this was. And so, we delved into the exploration of Microsoft Purview eDiscovery, Clearbrief, and Word Copilot to digitally transform this opportunity. And so, Microsoft Purview eDiscovery is like any other eDiscovery system that you've probably already used. You can ingest data into it and then you can automate the production of data, the automation of keyword searches and tagging. And this has been a really helpful technological tool for large client file reviews. And so, in terms of the program volunteer workflow, we use scripts to communicate with your computer to split that 20,000-page PDF into individual pages.

We then ingest all of those pages into Purview eDiscovery, we automate the keywords and the tagging, and then we export that Excel sheet as the final work product for NVLSP. And again, this is what that final work product looks like, with the keywords contained in column D. So we were able to automate that first work product, but then we still had the need to think about a digital transformation opportunity for the other work product. And that's where we turned to Clearbrief and Microsoft Word Copilot.

And so, Clearbrief is actually an AI legal tech company; they live as an add-in inside Microsoft Word. And Clearbrief can help you generate hyperlink timelines from the client files' statement of facts or narrative summaries. And every single AI work product has a citation hyperlink, so that you can verify and validate the data that it's producing directly through the Clearbrief user interface.

And here you see an example of the AI timeline. And so, I fed the 20,000 pages into Clearbrief. It then produced a 170-page timeline with a date, a description, and a citation to those key statement of facts. And so, if you click on that citation hyperlink again, it then gives you the document, so that you can validate the data. And so we then fed that timeline into Clearbrief. It then produced an 11-page narrative summary for us, a statement of fact. Again, this includes that citation hyperlink, so that the pro bono attorney always has the opportunity to validate the work product that's being produced by AI. And so, these are still large work products. The pro bono attorney still has to review them. And given that Clearbrief lives inside Microsoft Word, we thought, "Hey, why don't we think about incorporating Microsoft Word Copilot into this workflow and see what we can produce from this?"

And so, one of the uses that we had for Microsoft Word Copilot is to use prompt engineering, take that Clearbrief AI statement of facts, and reword it and reorganize it in a way that's just easier for the pro bono attorney to consume and to understand. So, you see here in the prompt, we asked for it to organize the statement of fact into categories of discharge recent and type, discharge data and branch, service record, etc. We also used Word Copilot to be able to validate the data. So with Word Copilot, we could ask for example, for the technology to summarize the entire document, to summarize the entire 170-page timeline. It's just another way that we can be validating the facts that are produced by the AI work products.

You can also use Microsoft Word Copilot as you see here on the right-hand side to ask questions. So, in this instance I asked, "What is the client's date of birth?" And Microsoft Word Copilot then went through the timeline, went through the narrative summary, and produced a response, and then referenced that, so that the pro bono attorney can then validate that data. So that's how we sort of stack the technology together. And the final work product again, is that file review summary. Now you have the automated or the AI produced timeline effects on the narrative summary, and then we take that and through Clearbrief produce a web-based version for NVLSP, so that they can also access the documents directly through their browser as well, and validate the facts.

And this has been a beautiful pro bono opportunity for us. We've been able to engage diverse pro bono volunteers, e-discovery experts, individuals that we had never engaged with before. And so, our pro bono volunteers are not only helping scale access to justice for veterans, but they are also up-skilling and learning about AI too.

And just some closing thoughts. Experimentation is key. Every single AI product has its strengths and its limitations. So identify those through experimentation and if possible, stack them so that you can build an end-to-end workflow that is holistic and is responsible as well. And as always with AI, you need a human agent to validate every single thing that's produced by AI. Thank you.

Tim Baran:

Well thank you, Adrian. And hey panelists, could you all turn on your, yeah, awesome. If you could turn on your cameras and your mic. Just wanted to give a shout-out for a fantastic presentation team. And I know you were busy in the chat answering questions, but there is one question that I'd like to for you guys to answer one. And this question is: Dean Matthew alluded in the previous session that it may not be as clear cut that AI is a tool for democratizing the law. This question is for the panelists, for us here, how do you propose we avoid the formation of a two-tiered legal system, where the high-quality gen AI tools are more easily and primarily accessible and utilized by wealthy individuals and deep pocketed organizations? Anyone want to field that?

Adrian Palma:

Yeah, Tim, I can step in to provide context on Clearbrief. So we developed this AI use case with Clearbrief and with NVLSP, and through that process we were able to discuss with Clearbrief, hey, NVLSP will also need access to this technology. What options do you have to be able to give them the technology for free, or a highly reduced discount? And so, we used our leverage as Microsoft to be able to add that layer of pressure. And I'm happy to report that Clearbrief has been able to offer NVLSP the technology for free. And so, I encourage all of you, if you are developing AI use cases with other stakeholders, use that influence you have to ensure that the nonprofits and the legal services organizations have access to this technology.

Quinten Steenhuis:

I'll jump in as well. I think you can spend a lot of engineering time and talent on building solutions around generative AI, but I also think the underlying technology itself is relatively cheap. I mean, that's something to keep in mind, that compared to the cost of an additional staff person, you can run millions of queries against generative AI tools. So, once we've built the right things, and I'm confident that we have the ability to get some investment from the Legal Services Corporation and others that have been innovators in the access to justice space, we can build solutions that help poor people, just like we can build ones to help wealthy corporations.

And a lot of that special sauce and some of the commercial tools may not be necessary for the same solutions that we're looking for, because they're designed at scaling across

thousands of deposition documents, or about doing legal research that may not be necessary in every instance at the same level of use, as it is to help a litigant tell their story effectively, and to get in front of a judge. I think there's ability to scale what we have for not that much money, and I think we should just need to start investing and trying some of those solutions.

Tim Baran:

And there's a question just about how to get started, which is a question I think we all have at various stages in our education around AI. And the question is, any recommendations for how to get started using AI in legal practice? Classes, website, books? Maybe Teri, we could start, everyone could go around and give one suggestion, or however many suggestions, or if you have a link that you want to put in the chat, but maybe Teri, we can start with you?

Teri Ross:

Yeah, I mean I would say don't start with it as a professional matter. Start with it as a personal matter, right? Because that's a much easier bar, right? Ask it to plan your two-week road trip to Colorado and Utah this summer and see these five sites and see what kind of itinerary it comes up with for you. These things kind of build on each other, right? Then that generates some more thoughts and you just keep going.

Quinten Steenhuis:

I agree with Teri, just trying it out is really critical. If you want to try GPT version four, you can do that inside Microsoft Copilot, although it has some kind of funny limitations in the moderation that can throw you off. Also, the newest Gemini model is totally free, which is according to most people as high quality as GPT-4. The free version of ChatGPT is maybe not as performant. One of my favorite things is to have it help me with meal planning. Another kind of fun thing to have it do, that is going to be more like what you should maybe use it for in legal work, is to have it edit your work, have it give you feedback or suggestions for how to make your work clearer, or to improve it. And you can do that with anything that you write.

Jonathan Pyle:

I've been using ChatGPT-4 on a paid plan, and it's cost me a whopping $10 for all the work that I've done. But even the free version is excellent. I would just suggest that you just go to chat.openai.com and start talking to it but don't rely on its background knowledge because that's where it gets, like... Copy and paste a whole treatise of your own into a prompt, and then just instruct it to respond to your questions based on that treatise, and just experiment with things like that. I think you'll be surprised at how good

it is at doing things when you give it explicit instructions. Talk to it like a child who needs boundaries and direction.

Tim Baran:

All right, I think we are at time. And during a conference earlier this week at Suffolk, Judge McCormick mentioned someone, Ethan Mollick. And I happened to be listening to a podcast, the Ezra Klein Show, where it was just a practical and every day how to use AI in your everyday life, to Teri's point. And for me it was very informative. So if you want to take a look, Ethan Mollick, just look him up. But all right, well that's a wrap. Thank you, team, for a great presentation. And I think we take about a ten-minute break and then reconvene at about 3:15 for breakout sessions. So, thank you and have a good rest of the conference.