Joseph Landau:

Thank you so much, Judge Ally. It's wonderful to be part of this conference. This is really just a fantastic, fantastically amazing set of questions and participants, and so it's wonderful to be here. So as Judge Ally said, the title of this panel is AI and the Legal Profession: The Implications for Law Schools and Access to Justice. AI, as we know, is going to change the practice of law as we know it, with significant implications for legal education, and for access to justice as well. This session brings together leaders from legal education, from legal technology innovation, and from civil legal services, to discuss generative AI, and its impact. Panelists today will explore how the teaching of AI can be incorporated into the law school curriculum, to ensure that future lawyers have the knowledge and the skills to use technology to help bridge the access to justice gap. I want to remind everyone that this session is being recorded, and that the recordings will be posted online, sometime shortly after the conclusion of the conference. I also want to read the CLE code for this session, for this panel, and I'm going to read it twice. The code, and you can see it on the screen here, is AIA24, AIA24.

Thank you. Now I'm going to introduce the panelists briefly, all too briefly. They have incredibly distinguished careers, and you can see more about their full biographies in the online agenda. So briefly, today we have with us James Sandman, who is a distinguished lecturer, and a senior consultant to the Future of the Profession initiative, at the University of Pennsylvania Carey Law School. Next, we have Pablo Arredondo, who is the vice president and co-counsel at Thomson Reuters. Third, we have Mark J. Williams, who is co-director of the Vanderbilt AI Law Lab. And last but not least, Dana Bowen Matthew, who is the dean and Harold H. Green Professor of Law at the George Washington University School of Law. We will go until 2:05 P.M., at which point we will take a 10-minute break, and after the panelists make their relatively brief, six-to-eight-minute remarks, I will ask questions, and I will also be taking questions that you post in the Q&A. The Q&A is open, please make use of the Q&A, and we will do our best to go through the questions, and pose those questions to our panelists. And with that, I'm going to turn things over to our first speaker, Jim Sandman.

James Sandman:

Thank you, Joe. Good afternoon. I'd like to do two things in my brief remarks. First, I'd like to say something about the need to educate law students today about the use of generative artificial intelligence. And second, I'd like to say something about the need to educate the legal services community, about how they can use generative AI today to improve access to justice. First, about legal education. I believe that as of this fall, the fall of 2024, every law school should be including instruction in the use of generative AI in their mandatory first year course, on legal research, legal skills, legal practice, whatever the particular law school calls it. Generative AI is already a basic research and practice

tool. I think as a matter of professional competence, required by comment 8 to rule 1.1 of the rules of professional conduct, it is essential that lawyers know about the use of generative Artificial Intelligence.

Comment 8 says, "A lawyer shall keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology." Knowledge of generative AI is already a matter of basic competence. Rich clients already get this, and are requiring their lawyers to use generative AI. For example, last June, the general counsel of Ford Motor Company, Steve Crowley, wrote to all of Ford's outside law firms, asking them to tell him how, not whether, how they are using generative AI in their practices to improve their efficiency and reduce Ford's costs. Steve is not some corporate elitist, he's a former academic, he was a chaired professor of law and associate dean for academic affairs at the University of Michigan Law School, and has written compellingly about access to justice. This is the state of the art, and I think legal education needs to adapt to it.

More importantly, I think it's critical to educate the access to justice community, particularly legal services providers, about how they can be using generative AI today to improve access to justice. I was formerly president of the Legal Services Corporation. I believe that generative AI offers the potential, the potential to improve access to justice in ways that will be enormously significant. And I think that there are two ways already today in which the use of generative AI can improve access to justice.

But what I detect when I talk to legal services lawyers about generative AI is discomfort and unfamiliarity with it. In that, they're no different from other lawyers. My sense, talking to lawyers across the practice spectrum is that what most lawyers know today about AI is that if you use it to write a brief, it will make up cases, and get you sanctioned. End of story.

Well, there's a lot more to it than that, and let me give you a couple of examples. I think it's important to develop what are called use cases — practical examples of the use of generative AI to make it concrete for legal aid lawyers — so that they can understand how AI can help them in their practice. The first is improving efficiency in the practice of law and a legal services environment, to allow legal aid providers to serve more clients, significantly more clients, and to serve them with higher levels of service.

Generative AI can relieve lawyers and legal aid staff of burdensome administrative tasks. It can, for example, accurately summarize intake information. It can expedite forms completion, it can help to identify patterns and data sets, it can summarize depositions and documents. We need to develop training programs for legal services providers that showcase these uses. And in that, I think there is a role for law schools to play, in educating legal services providers, but I'd suggest to the Permanent Commission that they not depend only on law schools to do this. I think they should enlist help from the big law firms and the corporate legal departments that are today using generative AI to

improve their service to their clients. We can learn from them. Their uses are directly relevant to, and can be applied to the practices of legal services organizations as well.

The second thing I think we need to do is to demonstrate direct-to-consumer uses of generative AI. I don't think that generative AI is yet at the place where we can use it to provide person-specific, case-specific, legal advice to people, but it can be used very effectively to provide helpful legal information to people. Let me give you an example. I live in Washington DC, I'm vice chair of the DC Access to Justice Commission. We have 40 legal service providers in the District of Columbia, a city of about 710,000 people. Every one of those 40 legal services providers has its own self-help materials for people to use directly to get information about their legal needs. We have a project underway, using generative AI to consolidate all of the information that has already been prepared by knowledgeable legal aid lawyers about every subject that comes up routinely in a legal services practice. And to take the best of it, and put it together in a readable, understandable, helpful format, for people to use directly.

And what this does, using generative AI this way, controlling the inputs, the materials that you're drawing from, to create your work product, and limiting the inputs to curated materials that have been prepared by legal services lawyers themselves, reduces risk enormously. You're not drawing from the internet; you're drawing from materials you already have confidence in. And we're optimistic that what we'll be able to do, using generative AI, is to produce self-help tools in a uniform format across the District of Columbia, across all 40 legal services providers, that are simplified and clearer than anything that the lawyers have been able to do themselves. But all of this starts with education. It's not self-evident to people how generative AI can improve access to justice. We need to be preparing these use cases, to make it concrete and specific for the legal aid providers, how generative AI can improve access to justice, now. Thank you.

Joseph Landau:

Thank you so much, Jim, for those remarks. I'm now delighted to call on our next panelist, Pablo Arredondo, who is, as I mentioned previously, vice president and co-counsel at Thomson Reuters. Pablo?

Pablo Arredondo:

All right, thank you. And what a privilege to be here with such esteemed co-panelists. I know conferences are more interesting if people disagree, but I can do nothing but applaud everything Professor Sandman just said, and the efforts that he described underneath it. Let me perhaps frame things like this. Every year, at the end of the year, the Chief Justice of the Supreme Court writes a year-end report and, on two occasions, has Chief Justice Roberts used that opportunity to discuss technology and its role in the law. The first was 10 years ago, roughly I think nine or 10 years ago, where he almost, I

don't want to say bragged about how slow the law is with technology, but sort of took the approach that, it's a good thing that the law is very slow to adopt technology, given the stakes, and that that's always how they've been very conservative.

And he actually pointed to the fact that the Aesop's Fable of the turtle, the rabbit, and the hare ... sorry, the turtle and the hare, was actually etched onto the Supreme Court building, and he pointed to them and saying, look, the turtle is the way to go. And so the second time he wrote about tech was last year, on writing about these new developments, although they're a little bit older than the public attention to them, but relatively new developments. And I was quite surprised and delighted to see him take a different stance, and to say something that I've said many times, that I think is really worth meditating on, which is, that the first rule of civil procedure exhorts the just, speedy, and inexpensive resolution of matters. And Chief Justice Roberts wrote, "Surely this technology, if used well, could help us realize that goal," that goal that's sort of encoded in the very first thing we say in our rules about how we would govern civil procedure.

And I completely agree, and I've had the privilege of getting very early access to GPT-4, six months before it was publicly released. My colleagues and I at Casetext were able to start using it, testing it, and with certain data clients, exploring how it might be applied. And the first and foremost, and then as Professor Sandman alluded to, it's critical to understand that applying these technologies with the right architecture, is very different than doing it in the sort of reckless way. And specifically, if you don't anchor it to something real, and put these guardrails so that it's basing what it says on real materials, ideally curated, absolutely it's prone to hallucinate. And we've all seen those sort of horror stories of the attorneys citing completely fabricated cases. But I submit to you that that's sort of a distraction, or will increasingly be a distraction, to what is the real oversight that these models need, which is, they're not hallucinating because we put these architectures on them, but they're also not infallible, right?

They're not able to ... they need oversight. And so anyone working with this technology, I would urge you to adopt a trust but verify, to borrow from the Cold War, the trust but verify approach. And companies like mine, Thompson Reuters, are ensuring that in everything we build, we're streamlining the ability to provide exactly that oversight, because they're not infallible.

Once you are using them correctly, and now we're getting into sort of speed. And with speed and justice, are of course related, because justice delayed is justice denied, the gains can be spectacular. Things that would take much longer to do, without these tools, can now be done in a fraction of the time. And the truth is, it's a pretty general technology, and so in just our early months using it, we've seen it do things like help the Innocence Project triage applications to decide which cases to take.

And then we've seen multinational law firms using it to do due diligence on their contracts. We've seen Fortune 50 companies using it to better triage materials that might trigger the need to alert government agencies. And so it's one of these general-purpose technologies, where there's very few things where we're not going to be able to apply it once we do it well. And I think that, more than anything, we should be judged on how well we bring these LLMs to bear on access to justice, on raising the ability for everyone to get access to these things. And so I'm just very reassured to see communities like this coming together to focus exactly on that issue, because I think it will be one of the most important ways that it's applied.

Now, inexpensive is the last part of that. And here the news I have for you guys is not all great. And let me be specific on this. Well, sorry, let me start with another colorful analogy if I may. In the early days of electricity, only Wall Street and Madison Avenue had streetlamps. And that was because, at the time, in those early days, it was quite expensive to make it. And if we fast forward to, I think somebody alluded to that moment when Westlaw and Lexis were first coming on, Paul Lomio, who was the director of Stanford's Law Library, once told me that there was a time when Stanford law students weren't allowed to use Westlaw between certain hours, because that's when New York was heavily using it, and you just couldn't have them both on at the same time. And in a strange way, these large language models, as they're called, have brought us back to this world, because they're very computationally expensive, and there's a scarcity of the underlying chips, and systems that you need to actually run them.

And so one example of how this plays out, and one of the ways that it harms, is that Casetext, we've always made our tools completely freely available both to the judiciary, and to law schools. Just .edu or .gov, you were in, you didn't have to even let us know. We can't do that with this technology, we can't afford to do it, we couldn't do it, it would break the servers if we were doing that. And so, one of the things to appreciate right now is that these models are currently expensive. The dangers of course, are that there are other models that are less expensive, but may not be as good. And so you have this quality, but there's balance between just, meaning you're finding the right things, and you're actually improving the outcomes, versus expense. And so you do have to be wary, when somebody says they're using this AI, that can now mean a lot of different things.

And Casetext, we were exposed to earlier models like GPT-3, and though we found them neat, we emphatically did not think that they were ready to be used in professional context. GPT-3.5 made gains, so I guess the point is to look at specifically, at which models are being used, and how they're being used, but the overall picture right now is one of these things being quite expensive to run. Now the good news is that capitalism is putting its shoulder towards trying to reduce the cost on this. I was just at a session where Sam Eltman, the CEO of OpenAI, was noting that, look, this is one of the areas

where they expect to see big progress. And so just like now, electricity is not something that's only limited to certain streets in New York, but everyone has it. Hopefully it'll be a much faster journey with this technology, to get it so that it can be deployed at scale, in all of the various ways that it should.

And then I want to add the other aspect. I think that's important to think, as you're framing how we're going to apply this and what use cases we have, is that although it's called generative AI, and the focus has been sort of look how it can write text, I actually think that that's misleading and a misnomer, especially in the near term.

It's not that this technology can generate text that really excites me. It's that it can read text and analyze it, and annotate it, and classify it, because so much of what a lawyer needs to do is to interrogate large corpora, and find the things that are relevant. And if you think about other ways that you could apply this at a high level. Ten years ago, my landlord company sent me this letter, saying they were going to raise my rent, and pointing to a California statute, citing it with a statute symbol and everything, "Dear Pablo, nobody likes a rent increase, but the law's the law, and we can do this."

And they chose the wrong Pablo to send that to, because I had a legal research company, and I went to law school and litigated patents at Kirkland and Ellis. And I looked up the statute, and it did not apply at all. When I called them, they didn't even defend it. They just said, "Oh, I'm sorry, that never should have gone out." And you wonder how many other people got that letter, and how many of them saw that statute symbol, and immediately thought, well, I can't, it's the law. And so the ability of models like this to not just adjudicate disputes as they arise, but to monitor, and look for shenanigans like that, I think is another really powerful application of it, as well. So I guess I would urge you guys, as you're thinking of the different ways to apply this, note that a lot of what it can do is to read, and that can actually be a huge amount of wind at the sails if it happens.

So I sort of want to limit my remarks there, and get back to questions. And unlike our panelists here, I can't point to a career focusing on these issues. So I am actually of limited use to this conversation, but I can sort of report to you guys, what we've seen when applied correctly, is improved precision and recall, which are just fancy ways of saying, it's finding the right stuff, and it's not missing stuff, over even the most well deep-pocketed law firms in the world, working on nine figure litigation. So, we are seeing absolutely that it can raise the standards, and raise the outcomes, at the most expensive areas of law, and I have absolutely no doubts that, as we get the server costs down, and as we start figuring these things out, there'll be vast, oceans more value created, applying it in much broader context, and especially to folks who suffer from this justice gap. So with that, I think I'll turn it back to you guys, and thank you so much again.

Joseph Landau:

Thank you so much, Pablo, for those remarks. Next, we have Mark J. Williams, who is co-director of the Vanderbilt AI Law Lab. Take it away, Mark.

Mark Williams:

Thank you, Joe. And I also, I know it's probably more interesting to have a level of disagreement, but similar to Pablo, I completely agree with everything that our previous panelists have said. So I think what I can do is talk a little bit about my approach, our approach at Vanderbilt Law School, and some of my experiences with this technology. I've been teaching legal technology courses for 10 years now, and, in that time, it had always had a section on AI in those courses. And even over the last four to five years, we would bring in earlier versions of these large language models. So Google had one called BERT, GPT1, GPT2. And we would bring them in and we'd have the students play around with them, and sometimes if it answered a legal question somewhat coherently, it was really fun and interesting. But there was this sort of incremental sort of iteration every year where, maybe it was a little bit better, it was a little more interesting every year, but certainly nothing more than an interesting party trick. One that I really liked, but it wasn't useful in any meaningful sort of way.

And then in January 2023, I'm teaching this course, and ChatGPT, GPT 3.5 had been out for a month, and teaching that same class, our students had probably finished up their finals in early December. They'd been away for a month, they come back into my class that January, and they had just clearly been, during that month, experimenting with it. And now it's a self-selecting group, they chose to take a course in legal practice technology, but it was just something in the water was different, and they had thoughts about it. And whether I liked it or not, that class went from being a one unit, one week AI class to a, this was now an AI class, a generative AI class. And that sort of drove this home to me too, it was just a lot of times in law with technology, we can be very slow to adopt, and we have mechanisms at our disposal to sort of, where we can decide how much to bring in and how much not to, we can ignore things if we want to. But what struck me in that moment is, whether we liked it or not, the external forces and pressure that this technology was creating, based on what it could do, was unavoidable, and in an exciting way, in a scary way too, but in an exciting way.

And this was just chat GPT3.5. When GPT4 came out in March, for those of us didn't get early access, and then Casetext CoCounsel came out, it was profound, and our dean at our law school, Chris Guthrie, saw this as well. And immediately we began to plan in earnest, what can we do as a law school to prepare our students to be leaders, both as users from a practical sense, but also in terms of being sophisticated about the policy choices that we can make about this technology, and how it's deployed, so that they can

be leaders in the field. So there's a couple of things we need to think about with the technology, and how we approach teaching it to our law students.

The first is we do implement it in our first-year legal research instruction, in a low stakes way, just to give our students a baseline understanding of what generative AI is, versus other types of AI, and give them a low stakes environment to play around with things, what is a hallucination? why might a hallucination occur? Basic low-level education like that, all the way up to more. I teach a whole semester long class now, I've taught it twice, called AI and Legal Practice, where we do simple things like basic prompting techniques, and just giving them a structured way to understand how the models work, and what they can do and what they can't do, to the more larger implications of confidentiality, privacy, long-term implications about unauthorized practice of law, consent, in terms of how much do you tell a client about what you've done that constitutes generative AI, but not in the terms of that we have the answers.

Because I think this technology and the policy choices we're making around it, change so rapidly, if we were trying to come up with answers, the questions would change. So I think we often talk a lot about, we are in the business of just formulating what the right questions to ask are at this time, because the technology is changing so rapidly and I have the sense that in the coming months, we're going to see an even larger exponential increase in the capabilities of some of these frontier models, that will, once again, if we were approaching this as if we currently had the answers to these problems, the questions would change in a pretty profound way.

So, what do we think about in terms of training our lawyers, our future lawyers now, and law students, and how they can use this technology? And it's about curiosity, and learning how to learn about it, about an adaptable mindset. Because even before large language models came in vogue, in my technology courses, I would often tell my students, half of the substantive material we're going to cover here in this course, might be outdated by the time, five years from now, four years from now, by the time you graduate. But what we can learn is an adaptable framework to continue to learn about how this technology can be applied, how it evolves, how you can continue to learn about it.

And what we also do a lot, with all of our courses here, and just as a law school, particularly with my co-director, Kat Moon, and myself, is the technology is often just a hook. And what is really behind it, is a sort of a human-centered focus on what it means to be a legal professional, and also what that means for access to justice implications, in ways that we might use technology to close the justice gap. Even before GPT4 came out, a lot of our programming and projects were oriented towards those kind of things. What's interesting for us now is that framework of thinking about the nature of skill acquisition and what it means to develop the next generation of lawyers, and that sort of apprenticeship model that we have, where many of the tasks and low level sort of

automatable things that AI we know is already pretty good at and will continue to get better at, how do we train the next generation of lawyers to be experts in the field, if the tasks that we train them on are increasingly going away?

So for me, that's actually a very exciting opportunity, for us to intentionally think about what it means to learn, and what's the best way to learn, and what's the best way to do skill acquisition? And that pipeline has never been more important, from that transfer from when you're a law student, to when you're a first, or second, third year associate, or even a summer associate. And the intentionality of how we develop those skills, because to me it's never ... it's completely clear that the best and highest use of any of these tools is when you're already operating at an expert level. If you're not operating at an expert level, it's not a replacement for your own judgment, for your own learning.

I'm a pragmatist about these tools, and there is a real danger that we start to use them as a replacement for our own judgment. And I would hope, through forums like these, through intentional education like we're doing at Vanderbilt, and other schools are doing this as well, through intentional basic education about what these tools can do, when the appropriate time to use them are, we can create the next generation of lawyers who will make a lot of the policy choices about how we deploy this technology, and how we implement it into our practice. And that is where it has profound implications on access to justice and closing the justice gap, and the potential for it to take that legal aid model out of the one-to-one model that we currently have, and potentially make it a one-to-many model, which the statistics tend to bear out that we kind of desperately need, whether we are comfortable with that or not.

Professor Sandman's here, the Legal Services Corporation statistics are 92% of Americans who have some kind of civil legal problem can't receive adequate legal help. And even to me that rings true. I grew up in a very middle class, but blue-collar family, and even today, I worked at Cornell and now I work at Vanderbilt. I've worked at these well-off law schools, but even in the back of my mind today, that blue-collar part of me is like, "You can't afford an attorney, attorneys are for fancy people, people with money." There's still this part of me in the back of my head that’s like that, and I know there are other people out there that feel this way. It feels unattainable, even if you're someone who's reasonably well off. A lot of us have that mindset, and this technology has a way to make those things more approachable, even if our minds always go to the legal research or hiring use cases.

But when I talk to legal services, legal aid clinics, and then also the technology firms that will often be working with them to sort of prototype use cases, it's usually the vanilla stuff, just the business of intake, of triage, of summarization, of knowledge management, that we don't even have to touch the practice of law side of it. And these tools can have an immense impact on scaling up our ability to help people in a variety of legal tasks. So our other mission, in conjunction with the technology itself, and my

co-director is long been doing this work, it's just keeping it front and center of our students that this gap exists. Because as law students, attorneys, we're often busy, where we've got our own problems to solve, we're incentivized to look in certain directions. That doesn't make us good or bad, it's just we all have our direction that we're going, and we don't always have time or the inclination to be reminded that this issue is out there, and to be intentional about focusing on it through our coursework, through our application of the technology, and through, again, I think it's going to be not very long now where it's not a question of whether the technology can do it, it is whether we as a profession are both open to, educated, and inspired enough to find ways to make that happen. So with that, I will turn over my time to our next speaker.

Joseph Landau:

Thanks so much, Mark. And now I'm going to turn things over to Dayna Bowen Matthew, who is the Dean and Harold H. Green Professor of Law, at the George Washington University School of Law. Take it away, Dana.

Dean Dayna Bowen Matthew:

Thank you so much. I am honored to be here as a part of this really important discussion, and although I am not going to necessarily create the dissonance that great conferences have when people flat out disagree, I am going to use my time to be a little bit of a provocateur and ask whether we should be as enthusiastic about the promise of AI, especially with respect to access to justice. So let me start with the big picture. As a legal educator, I feel as though our responsibility, our social contract if you will, is to provide a well-educated citizenry, that are equipped to make society better. That's my big picture view of what the promise of higher education generally is. With respect to AI, the way I'm going to focus my attention with that calling, if you will, in mind, is to talk about our obligation with respect to educating lawyers in two ways, in terms of their professionalism, and in terms of their professional excellence.

And by professionalism, I mean their integrity is our responsibility. The legal education experience should be focused on producing lawyers that are not only capable lawyers, but also lawyers with the highest values of our noble profession. And secondly, those lawyers should be capable of engaging, and deploying, all the tools of our profession to the benefit of their future clients, and to the benefit of society. And one of those tools is AI. So first, with respect to professionalism, my comments are briefest on this. We have to protect against the misuse of AI in our educational experience, and we do that right now, by limiting or prohibiting the use of AI in areas in which it can be misused. In law school, really that means substituting one's work for the work of another, making sure that AI or generative AI, does not replace the work that someone in our student body submits as their own work and their own thinking.

So, that's a pretty easy way to make sure we protect against the misuse, and to explain why that's true, that's not complicated, professionalism has been our concern for the 157 years of our law school's existence, and I dare say of every other law school as well. More complicated, or nuanced, is the question of the professional excellence that we are teaching and training, with respect to the proper use of AI in discharging the work of a lawyer. And so here I want to step back for a moment, and describe a thought experiment that defines or outlines how we think about what we are called upon to do in training future lawyers in the proper use of AI. So, I just called up ChatGPT, you can do it with any generative AI platform, and ask a legal question. And I said, in my thought experiment in preparation for these comments, help me write a memo that describes the state of the law on book bans in public schools.

So, I gave it the topic book bans. Let's say my partner called me at four o'clock on a Friday, and I made the mistake of picking up the phone, and he wanted to know what's the state of the law on book bans. And ChatGPT wrote me a memo. It was a very short memo, but it was a memo that did three things. One, it told me what the first amendment said. Secondly, it cited three primary cases, and if you do this area, you know Tinker vs. Des Moines Independent Community Schools, the 1969 case, is very important. Board of Education Island Trees Union vs. Pico, very important, and Hazelwood School District vs. Kuhlmeier, is very important. So those are the three big cases. And then the memo ended by saying, essentially this, challenges to book bans may result in legal proceedings on cutting to the chase, where courts assess the constitutionality of such actions. Consult a lawyer, right? You need a lawyer in order to solve that.

And I said, well, all right, I have the three most important cases probably to outline the landscape, but I still don't know the state of the law. What I mistakenly did was not tell my ChatGPT, my inquiry didn't include a request for detailed information that would help a client. So, I re-entered my query, and I said, "I represent the state that is considering legislation to ban books. I'd like to know what my state should know about the law concerning book bans," right? So, here I'm now giving more information about my posture, the fact that I'm interested in making a decision between two competing views, and I'm asking for strategy that will help inform my choices. And this is what I think my partner is going to want me to cover in the memo.

And I re-entered the query. What came back was essentially the same memo, same three cases, same reference to the first amendment. Interestingly, there was another feature of the answer that came back, and that was an opinion that basically said, book bans are bad, they shouldn't be used, and some reasons why they limited free thinking and free choice. Okay, here's my point, ChatGPT and AI generally do a very good job of making research more efficient at the start. It is very clear from the fact that not anywhere mentioned, and I actually continued this experiment with three other

iterations of the question, I never got to, for example, the fact that between 2022 and 2023, about 3,362 versions of book bans have been introduced in state legislatures, that have resulted in the banning of over 1,500 books. It is up to me, as the lawyer that is writing this memo, to talk about the types of books that were banned, the differences among the states, and how each of these authorities will play out in each of the states, or in the state, for example, that is my client.

There was no mention of any of the bills that are pending, and how they compare with one another. No background on Senate hearings, no mention of the U.S. Department of Education strategies with respect to how I sort out the nuances of which authorities will have an impact in which jurisdictions. I could go on, but you're beginning to see my point, I'm quite sure. And that is, in teaching our law students how to properly use AI, in giving them the tools for professional excellence, what they must understand, is that AI is no substitute for the critical thinking of a lawyer. It may make one more efficient, it may make one more comprehensive more quickly, in their initial research, but it certainly does not give the critical information that would allow one to represent a client with excellence. And so, what we're doing at GW, and I now agree with the prior speaker's thinking, about what we're doing in our legal research and writing program, which is called Fundamentals of Lawyering, I think it would satisfy each of the speaker's criteria, with respect to a responsible legal education.

We're looking at ways to do guided exercises on how AI may be used to do ministerial research tasks more quickly. How we could categorize a set of complicated documents, that contain client facts. How to provide lists of authority. By my third query, I was able to get a longer list of authorities. However, reconciling and distinguishing them is not within ChatGPT's authority and ability, quite yet. Cautionary tale to come about that. Our legal reading, research and writing has specific training, and guided exercises on how to use generative AI, to begin to prepare an outline for the written work product, that one would submit, but not as a substitute for that written work product. So, in other words, our curriculum does help, and must help, our lawyers in training become excellent practitioners at how to use AI as a tool to support critical thinking, strategic thinking, to support the task of being a lawyer, that can't be replaced by AI.

Now the notion is that this is a democratization tool. I'll spend the rest of my comments challenging that just a bit, or at least if not challenging that notion, raising a few issues that must be considered if we are to see it as that. Increasing efficiency is one thing. Limiting the time to gather data is a very appropriate use of any tool that would be helpful in delivering excellent legal services. Doing so uncritically is a mistake. Now here I'm going to just actually access or refer to my personal area of research and expertise. It's health law, and public health equity. We worry in our field about the use of generative AI that replicates existing biases in the data and algorithms, that will inform

the output that are going to be depended upon by clinicians. I think there's an analogy, if you will, to the use of AI in the legal services.

So, what we worry about, for example, is whether or not biases in a limited patient population that sources the data for treatment of, let's say, high blood pressure, heart disease, stroke. If we have an excluded population, or a population that has less access to particular treatment modalities, then the data sets that are going to be used to predict clinical interventions, as a result of using generative AI, rather than individual discretion or individual decision-making, are going to just not only replicate, but indeed exacerbate those biases.

So, we worry about algorithms that are based on bias treatment in order to reduce answers through generative AI that will exacerbate those biases. I think those worries are important to consider in the use of generative AI as an amplifier for access to justice. I think we have to keep the focus on what law schools are very good at, in training people to think and to use their knowledge, not only to produce what is informative about the current state of play in the world, what the list of laws are today, what the exact statements of the rules are in each state today, but people hire lawyers to talk about what the world ought to look like.

There's no limit in how important critical thinking will be to answer the question of whether we should or should not adopt a book ban. I couldn't get a memo, I couldn't get a query that would help me in this generative AI to do that. Now, pause here for a moment, and recognize that the more I use it, the machine learning will learn to do better, and it will give me a memo that does more of what I am asking it to do. The danger here though is that it does so in a way that replicates any biases that are reflected in the input. I have to remind myself and my students that the art of persuasion, the art of being a strategist, the art of developing interpersonal skills, building trust, and most importantly, discharging the duty I spoke of at the very beginning of my comments, that is to provide excellence in lawyers who can make their client situations better, and our society better, is not the work of generative AI. It is the work of individuals, and that's how we are perceiving our job to educate lawyers in the use of AI, at George Washington University.

So, I'll stop there.

Joseph Landau:

Thank you so much, Dean Matthew, for those remarks. And thanks to all of our panelists, who I now invite to turn back on your cameras for a few minutes of Q&A. I have to say that I think these remarks set a great tone for the rest of the day. There are way too many questions that I'm sure we all have to get to, and we'll do our best. But if nothing else, I think this really helps situate us, across a broad landscape for additional conversation and discussion throughout the day. So, I think I want to try to find a way, in

the few minutes we have. And first of all, thanks to all of you for putting your questions in the Q&A section of this webinar. We are going to share the questions with all of our panelists, and we'll try to get as many of them answered as we can.

And again, thank you for posing them. So, it seems to me, in light of what Dean Matthew just said, and in light of starting with where Jim was, clearly, we have to teach our students something. Clearly, we have an obligation to keep up with the times. Clearly, we have an obligation to think about making sure that we're training lawyers, not just so that they have the right tools, but also that they have the professionalism and leadership, and also to what Dayna just spoke about, the values, the understanding, the recognition, the breadth of perspective on what kind of this moment means, not in a narrow sense, but in a broad and rich sense. And so, I guess I would come back, really to all of you and say, the question is, what exactly should we be doing in our classrooms around this? Jim, there's a sense in which the legal writing idea, which sounds great, might not go deep enough, in light of some of the comments that Mark and Pablo are telling us, involving the way that this technology is going to change, the impact that it's going to have on lawyers using AI as a tool, sounds like an okay first step. But I'm not sure it provides the ability for lawyers to be able to stretch and adapt the way that their clients will need them to, as the technology changes, and the profession changes around that technology.

And then more deeply of course, how we inculcate the values concept, the responsible use concept. So, I think what I want to do is kind of give each panelist a chance to answer that if they want to. And if not, I'll move to a question that's exclusively about the profession, and access to justice. So, Jim, I'll give you the first shot, if you'd like it.

James Sandman:

I don't think it's possible, in a first year required course, covering the range of skills that the courses I'm talking about cover, to go into the depth that you're talking about. I'm contemplating an introduction to generative AI, particularly instruction on what it is, how it works, its limitations, and its potential uses with illustrative examples. A basic background in it. If you know something about how it works, the large language models, and their predictive approach to generating information, that can help go a long way toward understanding what the limitations on use are. But I think that, I perceive, in some quarters, so much concern about AI, that people are kind of burying their heads in the sand about it, and not getting law students up to speed right now on what it is, how it works, and how they might use it in practice.

Joseph Landau:

Okay, let me take that answer, and if I might try to advance the ball a little bit further then, it's a completely fair answer. But you could imagine a world in which the

excitement, and the need to not bury our heads in the sand, which is a completely fair point, is one that is easier for some to do than others. Now let me shift to the profession, and maybe some of the other panelists will want to address this. You can envision why a very wealthy law firm, or company, would be able to kind of access the technology, and hire the people to create the closed system, to invite their people to leverage this tool, in all of the ways that we are talking about leveraging it, fine. But in addition to the hallucination concept that was raised, in addition to some of the other problems, there's also the client confidences problem. So I can't just go to, in my own low bono practice, or whatever you want to call it, and put client information into a bot that interfaces with the cloud, with the internet. So I wondered if the panelists could sort of address that other kind of access to justice question, where there are firms and lawyers out there who face an impediment, based on an inability to even have access to this kind of technology, in a way that adheres to other ethical considerations around client competences and what have you.

Pablo Arredondo:

I like that. Can I briefly go back though to something Professor Sandman said on this? My big concern on education is that we lose the ability to write and specifically wrestling with the blank page. I think wrestling with the blank page is one of the most important things we can do. It's thinking, it's a journey where you actually start understanding it yourself better. And a world where we trust the AI to write the first draft and then change the adverbs and say, oh look, we did it together. I worry about that. People can spell less and less well, because their little red squiggly line, you mashing enough keys, you'll get there. You can't read maps because we have a map ... but you know what, I'm okay with those. But losing the ability to wrestle with how to organize and draft and persuade, I think is to me one of the bigger ... one of the ways I could see that's going horribly wrong for our profession, in terms of the caliber and quality we put out.

And so forgive me guys, but I have a thing I have to drop to on this. On the issue of the security though, I'll say this. You can put your stuff into the cloud if there are the certain contractual relationships. And these companies that are creating this technology, although at first they sort of wanted it like, anything you give us, we'll just train the next model with, they quickly realized that that's something they can't do. And so the fact is we have Am Law 10 firms putting very sensitive stuff, and this predates AI. They were using it in e-discovery, and other contexts as well. So while I think there are some profound imbalances right now, just from the cost, I actually don't see the security being the dimension that I would focus on there. And with that, thank you guys.

Joseph Landau:

Thank you, Pablo. And I'm told we're basically -

Pablo Arredondo:

It’s been an absolute privilege. Somebody asked about discount pricing. Boy, after this conference, it would be horrible if we don't come out with something like this, right? Come next year. And I can tell you, this might sound contrived, but literally this morning, I was having a conversation with folks, and our government team and colleagues who are looking at exactly that question, and how we can do it. It's a very, very important thing that we do, especially while we're waiting for the server costs themselves to come down. Thank you all so much.

Joseph Landau:

Thank you all. I'm unfortunately getting all kinds of messages and pings, from the organizers of this conference, that I have to call time, I'm sorry. These were incredibly informative comments from all of you, and I will respect the request that we stop. But I want to thank all of our panelists today for an amazing conversation, and I know that there's more to be had in the remaining sessions. Thank you all so much, and have a great rest of the day, and we'll take a short break prior to the next panel.