Donna Lee:

Good afternoon and welcome to the AI Goes to Law School concurrent panel, which will run from 3:15 till 4:30. My name is Donna Lee and I'm a member of the Permanent Commission on Access to Justice. I teach at CUNY Law School. My co-organizer, Elyse Diamond, who directs the Public Interest Law Center and teaches at the Elizabeth Hobbs School of Law at Pace University, and I are delighted that you could join us for a panel and interactive discussion about several innovative programs and courses at law schools around the country that are exposing students to AI-powered legal resources and tools and the ways AI applications can impact access to justice. We're excited to share ideas and early lessons from these varied approaches.

To save time, we decided to let our five presenters introduce themselves. But first, a bit of housekeeping. We're going to start by giving each of them time to present their programs and share some insights. Please feel free to put your questions in the chat. We're reserving time after their prepared remarks for Q&A and Elyse and I will monitor the chat and try to assure those questions are answered during the Q&A.

Also, New York and New Jersey CLE credits are available for those attending this session thanks to Buffalo Law School. I'll announce the CLE codes at the 35 and 65-minute marks. This session is being recorded and recordings will be posted online after the conference.

So, without further ado, Dan, can you please get us started?

Dan Jackson:

Sure thing. Hi, everybody. My name is Dan Jackson. I direct the NuLawLab at Northeastern University School of Law. We're an interdisciplinary legal innovation and legal design laboratory. I've been around for about 11 years as a staff member and probably about 13 years before I was brought on board as the first executive director back in the day. And we work at the intersection of creative arts and law to come up with new ways of giving people access to their legal rights.

I'm also a proud graduate, a 1997 graduate of Northeastern Law School, and I want to start with a quick anecdote if I may. First couple weeks of class, first year, you all remember that for those of you who are lawyers. My property law professor was a fellow by the name of Don Berman, and Don was actually quite a luminary in the area of artificial intelligence at the time. None of us knew what the heck that was about, but Don was expounding on technology in those first sessions.

At some point in time, he took the position and stated that, "By the time we are ready to retire," so he was talking to a group of 22- to 28-year olds, "by the time all of you are ready to retire from the practice of law, computers will be doing the vast majority of the work of lawyers," and that was in 1994. And it was scandalizing. We were all absolutely

outraged all over. We spent weeks talking about that, many people were. A crisis of confidence, "Why the heck am I coming to law school?" about this.

And, certainly, as we look at the trajectory, the recent trajectory with generative artificial intelligence, that doesn't necessarily seem too far off. I saw an interesting article as well, I think, in yesterday or today's New York Times about how a lot of the drudgery work of Wall Street is now being basically done by artificial intelligence. And that may be one of the upsides, perhaps one of the downsides as well.

But as I look at that anecdote, and I often tell students that anecdote as well, I wonder about what it's going to look like 30 years from now because that was 30 years ago for me. And so, I'm 55, 65 is hopefully my outer limit for retirement, 10 years from now. If we're talking about the pace that we're going at and Moore's Law still hasn't been broken, who knows where we're going to be at? It should be a very interesting space and certainly poses, I think personally, myself, a tremendous amount of opportunity for access to justice, a tremendous amount of opportunity for achieving justice.

We, at my shop, don't really use the term "access to justice." We prefer to go with "achieving justice" because we want to go right to the heart of the matter. And I'm going to talk to you a little bit about what we've been doing and how we've been introducing law students to artificial intelligence.

So, as I mentioned, we're an innovation laboratory. We're a relatively small shop, but we work in three interlocking areas. The first is that we do redesign of existing imperfect systems of the things like creating games, creating new court forms, things like that, a lot of what a lot of folks will recognize very clearly as access to justice initiatives. And we love doing that work. And we've partnered with court systems. We partnered with a whole bunch of different entities primarily through the Legal Services Corporation Technology Initiative Grant Program.

We also do a lot of work in legal participatory action research. We're one of the first shops to do legal participatory action research. That is working directly in underserved communities to co-design the research questions. Very different from how most academic research is done. It takes a lot more time. You co-design the research, you go through these iterative cycles of research, value-building, realignment, research, value-building, realignment, and eventually you get to some empirical data that says a few things about what you've been trying to achieve.

And the last thing that we do is we engage in radical imagination, and that's what I'm going to talk about mostly here. Radical imagination for us is introducing law students to tools and ideas like speculative fiction, often referred to as perhaps science-fiction might be the most common example, but we often use a lot of authors like Afrofuturists, as well to give students the opportunity to think about what our legal system might look like 50, 100, 150, 200 years out in the future. And then we design it backwards so that we can get back to today.

Our laboratory seminar on applied and critical legal design is the primary place where we've been engaging with artificial intelligence. This is of course what we've been teaching now for 11 years. We did that right off the bat. We decided to start teaching, because why not? And we have been teaching that now for, like I said, 11 years. It has gone through several different iterations. Right now, it is a four-credit, 15-week seminar. It has a different topic every semester. And, so, when art AI hit and ChatGPT broke, we decided our next semester we would actually engage with that.

So, we started teaching law students how to use ChatGPT and other AI tools within about, I think, four to six months out of the gate from November 1, 2022. And our topic, instead of having a really directed design question, our topic was open and engaging. We let students design their own, around the topic of artificial intelligence and justice.

And we use creative arts methods to give them an opportunity to imagine, that was the whole point, to imagine what artificial intelligence could do to achieve justice. Each of them got to identify what justice meant to each of them. One student came up with a tool, an idea for a tool that was intended to help people get restraining orders and harassment prevention orders. Another student came up with a tool that was intended to help translate basic legal documents like leases, employment agreements and the like, both from languages, but also from legalese, into regular, plain language. Another student came up with a really great tool, which you may have heard about if you heard Jalon Fowler talking in the rapid-fire presentation earlier for the AI-EP tool, which then another entity that I was involved in, The Burnes Center over at Northeastern, came up with the idea as well on a separate track. And those two tracks sort of aligned really nicely and it was a tremendous amount of fun.

We told them they had to do a ChatGPT $20, GPT-4, once GPT-4 came out. And I would note that for those of you who don't have a lot of resources within legal education, that $20 a month is a lot less than a case book. So, it doesn't take a lot of resources to do this sort of thing. And what we did is, over and over again, gave them the opportunity to engage with the tools, both as a creative co-creator partner, but also as just a trial-and-error thing to start doing it, start playing with it, start working with it.

The student outcomes were extraordinary. The basic outcome is just a paper prototype, not an actual working technology. There's nothing underneath the hood. We are not a technology-first shop. We are a legal design shop. So, when we build tech tools, I hire techies because I can barely use my smartphone at this point, which is a really horrifying thing to be saying for a guy who runs an innovation lab.

And we really focused, again, as I'll say this, on giving law students the chance to dream, to dream about what this can look like and what it can do to achieve justice in the space in a way that they understand justice.

And I'll close with this. Last semester, I co-taught a course with The Burnes Center, and that was the part that Jalon Fowler just presented on, AI for impact, and Jalon was the

only law student. There was one other humanities student, and the rest were coders, were technologists and computer science students. And tremendous talent, incredible vision, incredible energy. Almost all undergraduates, so lots of energy, most definitely.

But I was really struck by two things. First, that number one, very few of them had any conception of what legal ethics are or some of the core concepts around confidentiality or why that would even be relevant. A lot of the tools that we were ideating around, building around, had a legal component, and that troubled me. The second thing is our computer science school at Northeastern University does not have a required ethics course for their undergraduates or their graduate students. And the research I recently did to prepare to give a CLE in Pennsylvania indicated that very few computer science schools in the United States have required mandatory ethics courses. And I think that's troubling as artificial intelligence gains in its adhesion into our society and into commercial activity and into access to justice.

I think that the reason why we are so invested in giving all students the potential to dream and the tools to dream in this space is the importance of getting lawyers involved and being creators because the future of the profession really does hinge on that creative imagination, on that connection, on getting genuinely invested in the direction of AI.

I have a couple other points that I'll save for the takeaways. And with that, I think I turn it over to Keith, right?

Keith Porcaro:

Great. Thanks, Dan. Hi, everyone. My name is Keith Porcaro. I am a senior lecturing fellow at Duke Law School where I'm affiliated with the Duke Center on Law & Tech. Before I joined the faculty, I spent a fair amount of my career helping organizations who depend on data and technology learn how to govern it, learn how to design it responsibly, everyone from legal aid to national governments, to small-boat fishermen, to children's hospitals and labor unions.

What I want to talk about, and I want to offer a provocation, which is that you don't necessarily need a lot of fancy tech tools in order to teach students how to be good critical thinkers about technology. Simple games and simulations might do the trick. And the class that I'm going to talk about is one that we debuted this semester, called Algorithms and the Law, and it's a 50-student course that's one credit and eight sessions, that meet for two hours. Our goal here was to maximize the number of students who took it, so we'll be teaching it both semesters next year. We have 100 students in a 750-person law school. That's a big deal.

The hypothesis behind the class is that as future lawyers, our students are going to need to know how to critically analyze and dissect algorithms that they or their clients are going to see out in the wild, and especially to think about ways that those algorithms

might break or cause harm. It's designed for people with no tech background, it's not a programming class. One way to think about it is we're trying to help lawyers of all stripes apply their legal skills to algorithms.

And so, just from a high-level structure of the course, and then I'll talk about one specific lesson as an example, over the eight sessions, we moved from a really simple notion of what an algorithm is to more complex notions. So, we start with Boolean expressions, if-then, and we end with large language models. And those six classes in the middle are bookended, one with data and one with the future of advice, which is what my research focuses on.

Each class is a self-contained activity. It's a game or it's a simulation. The goal is for it to be simple and memorable and evocative and to invite students to apply their skills in unexpected ways. I think the dirty secret for most of us at law schools is we know that students are going to forget 90% of what they learn at school. My hope is that they'll remember how they feel when they're doing this.

Many of these are super low-tech. We use printouts, we use paper cards. And I'm going to tell you about one of them, which is an exercise that I call "Sandwich Court."

So, in our second class of the year, we focus on logical expressions. So, if-then-else, the simple rules that make up the building blocks of really, really, basic algorithms. Their homework before the class is to take a couple of statutes and to translate them into machine-readable logic, so pseudocode. We do one where it's the Michigan statute to do automatic expungement, and then we have students look at a bunch of requirements for getting government benefits and identifying ways that they overlap.

And the learning goal of this class, when we're in class, is to help students understand how the application of a rule changes how we code the rule. So, elsewhere in law school, it's an intuitive thing that we ask students to do. The blackletter rule that's in a statute is affected by all the various cases that are layered on top of it, so the rule that is written in a statute is not always the rule as it's applied in the real world. And, so, what we need students to take away from this is that algorithmic rules, how we code rules into algorithms, need to account for those applications, those messy collisions with the real world. And those applications change how we actually need to code the rule in the first place. And it changes how we evaluate whether or not an algorithm is working, or not working, or causing harm, or doing its job.

Sandwich Court taps law students' natural perfectionism and competitiveness. The prompt is very simple. We ask students to use the logical expressions that they learned prior to class to come up with a thorough, objective algorithm that describes what a sandwich is. They do this in teams of three and then we have students from another group come up and grade their algorithm using a series of sandwich candidates that I present. So, everything from a hot dog and a taco and a club sandwich, to three slices of bread, to two slices of bread with a rock in the middle. And, as students are grading the

algorithms, the rest of the students are actually using just simple yes-no cards to vote on is this a sandwich or not? And the algorithm gets a point if it matches the overall student vote.

We do three rounds of this and one of the things that we do here also is the student that's grading another team's algorithm, we send them back with the team that they just graded. So, we're mixing up the student teams and we're helping students figure out how writing an algorithm and writing rules collide with how people actually interpret those rules in practice. And, so, we do three rounds and then we talk about it.

And it does a couple of interesting things. One is that you start to see students anticipate the weird sandwich candidates that I'm going to come up with, but you also see them start to respond to how they think their classmates are going to vote. So, the algorithm quickly goes from something that is objective and thorough, to something that is really, really socially constructed. So, I'm betting that from one semester to the next, how my students build these algorithms is going to depend on how my students vote.

And then, to bring it all home, we tie it back to the law. So, we talk about how classification problems are really at the heart of what a lot of legal work is, from Supreme Court cases from 100 years ago that decided whether a tomato was a vegetable, to rules on the books now that decide what a gun is, for example, and deciding that a gun is actually just the center receiver with the serial number on it, and how gun-destruction companies get around that or navigate through loopholes by destroying just the center frame and selling the rest as gun kits. We link it back to the Harmonized Trade System, which is a system that uses a set of rules and categories to classify every good that comes into the US, and how it depends on interpretation to help people understand that, say, rum mixed with grape juice is not actually wine.

And this becomes a jumping-off point for future lessons. This Sandwich Court is supervised machine learning done really, really simply. But to riff off of what Dan said, the goal of this class is to help lawyers think about structured imagination and build structured imagination because they're eventually going to need to be able to take a disorganized client story and to add structure and detail, and it's going to be their job as lawyers to imagine possible futures for their clients, whether they're good or bad or complicated. And you don't need a lot of fancy infrastructure to teach this effectively. In fact, if you've listened to my talk, you can probably replicate this without needing really any more materials. All you need are some pictures of things that may or may not be sandwiches.

Thanks very much. Thanks. Alice and Drew?

Donna Lee:

Great. Alice and Drew? Yes.

Drew Amerson:

Thanks, everybody. We're going to split this up into four minutes each. I'm Drew Amerson. Alice Armitage is here with me. We collectively run LexLab, which is a center at UC Law San Francisco.

First of all, I want to say, Keith, thank you for that. I'm absolutely stealing the Sandwich Court idea next year. Sounds easy and I think it'll really get the students engaged. So that's a brilliant idea.

Let me tell you a bit about LexLab. Lexlab was founded about five years ago. Our chancellor at the law school at the time knew that legal tech was really starting to boom, didn't know what that meant, but he knew that we were in the middle of San Francisco, and we should try to do something there, so he tapped Alice to get this program started and she brought me over. We both have backgrounds where we were entrepreneurs and started companies after we'd practiced. So, a natural fit for us to get this up and running. And he gave us free rein. It was fantastic.

And what we landed on is LexLab really has three primary pillars. The first pillar is educational. We're a law school. It all comes back to the curriculum for the students. The second pillar is events. We want to bring more innovative events into our law school community. And the final one is an accelerator for legal tech companies that I run.

And, unlike the universities that we just heard from, our school is a standalone law school, so we have a unique challenge. I don't have a computer science department I can go down to and tap. We don't have a business school that I can tap. So, we created an accelerator to bring those technologists and those entrepreneurs into our community. We've run the accelerator for five years now. The first two years, it was fantastic. We had a community right on the campus, everybody was there, and it was buzzing. And then COVID happened, and it changed the dynamic, not just of our program, but of the entire world. And we've gone hybrid and we're trying to bring it back to in person.

But another shift that we've made is this past fall, we made our accelerator explicitly a justice technology accelerator. We partner with the Justice Technology Association, which is a relatively new trade organization for startups, and more established companies who are trying to provide greater access to justice tools, and we have also collaborated with Village Capital, a social impact fund.

One of the big reasons I've done that is in the five years that I've run this accelerator, I've seen the legal tech boom. There [are] tons of money pouring into this space, but so much of it goes to B2B, contract lifecycle management companies, e-discovery companies. The startups that are coming through that are trying to solve these really unique access to justice problems weren't getting funded, so we want to give them a boost up. So, we're trying to create that ecosystem with the accelerator.

And I see I'm running almost my halfway point, so I'm going to pass it over to Alice who really runs the curricular side of LexLab.

Alice Armitage:

Hi, everybody, and thanks for having us. Yeah, so Drew and I split things up. He runs the accelerator and the events, and I've concentrated more on creating classes and creating a concentration in technology and innovation.

So, I've been teaching about the intersection of law and technology since 2018, but I've tried to do it in the context of complex problem-solving so that the students have a project and it's been building a legal tech startup that they have to create and learn how to pitch and defend in front of an audience of legal tech experts. I've taught about homelessness and the debt collection system, and I've asked the students to come up with some sort of tech-based solution. The sad thing about a 13-week semester is that there's hardly time to even come up with a solution, never mind actually build it, and that has been frustrating.

But for the students, they've had the chance to interact with other people in other disciplines, which again, we don't have on our campus, because I've had hackathons. I ran a hackathon for three years. And again, we have these demo days. So, the students end up having to actually talk, to learn to also talk and present their ideas, and to have worked with people in other disciplines in order to come up with their solution.

But lately, I have been focusing on generative AI. So instead of just teaching technology and law, from talking about blockchain and automation and all the things that we had before November '22, I've changed dramatically in my focus on generative AI. And this year, I have really, this past semester, I have decided that it is very important for the students to understand how the technology of LLMs work. So, not just so they are comfortable using it, but because it's going to impact not only the way they practice, but every substantive doctrinal issue there is in the law. I mean, we can already see it in copyright law and in regulation, and I think that that's just going to go down through the line. And I wanted them to really understand how the working of it impacts what you can say. And what's so interesting for me was that right before I started teaching, the New York Times sued OpenAI, and in their argument, did get into the technology of how it works to prove their point.

And I found that the students in the beginning of the course were not comfortable using LLMs because they'd been largely banned at UC Law. I mean, it was up to every professor to decide what use they would allow in their course. And most people said, "None." But I said, "Every week, you're going to be using it." And I gave them assignments from the beginning. It was just like, "Oh, ask it any question." But we went on to write a cover letter, write a persuasive argument about why UC Law SF should only be an online school, and then I had a midterm in which they had to come up with an

answer on a copyright issue using LLMs first to draft it and then editing it themselves to create a business email that they would be proud to send to a partner. That was an interesting... That latter part turned out to be harder than using LLMs.

But in general, I found the students are very excited about this because they know that this is coming and maybe because we're in San Francisco, but maybe not, they know that their careers are going to be impacted in a major, major way by this, and that also the practice of law, the business models of law, all of that is up for change, as well as perhaps the regulation of lawyers that has kept others out of it. And that's how I think that generative AI is going to have the biggest impact in access to law is that we are going to be able to provide so many, not only educational, but basic form-filling tasks that will allow people who can't afford a lawyer to begin the process and perhaps even finish most of it before having to pay the fees that they need for a lawyer.

So, I'm excited for the future for our students and I'm hoping that law schools are going to really jump on board with this.

Basem Aly:

Hi, everyone. My name is Basem Aly. I'm at Columbia Law School and I am the Assistant Director of Instructional Technology. I've also taught for about 10 years at the grad and undergrad level, mostly things having to do with technology, programming, VR, things like that. So, I have a lot of, I guess, experience demystifying technology for students and for faculty.

Just last night, I did a workshop with faculty members on prompt engineering, but I also assist in a class called Lawyering in the Digital Age, at Columbia, which is taught by Professor Conrad Johnson, and it's tied to Columbia's legal clinic. And over the years, the class has evolved to a type of project-based learning, and often, they will have a client for the semester. So, in previous years, it was the Buffalo school system or transportation system. This year, we're working with the Legal Aid Society on housing law, and it's been a very fruitful area to explore the use of AI in that because it's a very specific area of law. It has to do with New York City. It's a place where everyone has a guaranteed or a right to counsel before eviction, so there's tons and tons of demand. The Legal Aid Society has about 200 lawyers, but they are always inundated with calls and inquiries.

So, there are two projects that we have been working with, and the students have been doing a very good job. In many cases, it's their first experience in a client-producer relationship. So, they have to meet with clients, manage their expectations, have deliverables, et cetera, anticipate their questions. So, one of them builds up on a previous set of projects, which is called HCOR, which is the Housing Case Online Resource. It's basically giant Wiki housed in, I think it's in a Microsoft-type of environment, SitePoint or something like that. And it was meant to aid the intake

people, the people answering the calls, and then they are on a chat with a lawyer. So that's one area, a first line of response is how to produce and how to give very helpful, very practical advice to the advisor or to hand off to the client. And another project is a chatbot for the intake lawyers to help generate pleading documents, so complaints or answers to eviction notices, et cetera.

So, in both cases, it involves taking in lot of information, synthesizing that information, matching it up with existing precedent or law or document templates, et cetera. And there's always a massive demand.

So last year, one of the students did a Python program that generated pleadings and it was well received, et cetera, but it's very difficult for lawyers to modify or change or adapt. So, this year, we changed it into a custom ChatGPT chatbot, put in all of the knowledge information in a couple of PDFs, and then worked on tuning the right set of instructions for the chatbot. And it's also been received fairly well, and it helps generate their pleadings, it helps busy lawyers construct these documents to find all the objections to a particular motion, et cetera.

And I think, well, first, the most important thing I think is we do a tremendous disservice when we call younger people "digital natives." I think it puts a lot of pressure on them to be instantly and ineffably adept at all sorts of technology. And I'm as guilty of it as anyone else, but I think they grew up in a sense in the golden age of consumer electronics or the golden age of graphical user interfaces. And lo and behold, we have a new technology that is basically just a text box. Some people say that the entire graphical user interface is going to be obsolete. But the disservice occurs because very rarely do these students understand what's going on under the hood, and in some cases, they might have mystical or semi-mystical ideas of what is actually happening.

I think similar to Keith's point about an algorithm, I've had very, very well-educated people, provosts, certain professors who profess to not know what an algorithm is and are proud of that. So, I think one aspect about this is it's an opportunity for students to demystify and to really kick the tires on these types of systems. It's one thing to have a simple-to-use app that essentially has one function, and be good at it, and then it's an entirely other thing when you're dealing with something that is language-based, something where the technology can produce all sorts of symbols and outputs that might seem interesting or that might seem coherent, but is more or less just empty rhetoric.

So, I think their experience trying to get a chatbot to do what it's supposed to do, to be able to test it thoroughly, and to then advise a client on its proper uses and train the user on it is a much bigger and much more... I think it's a moment where they might feel that "Oh, maybe a computer isn't going to take away my job in the near future."

And by the way, we're always told that in six months, in three years, in 30 years, everything's different, and technologists do that as much as anybody else. But I think

one interesting thing that we find is no matter how interesting or how magical a type of technology is, when we're confronted with the actual product, there's always the temptation to want a little bit more. There's always the temptation to try and completely eliminate a step or two of labor or scrutiny in order to have that magical output. And I think many times, people feel that it's just a matter of pressing a button and getting exactly what you need.

But as we found working with The Legal Aid Society and with the students, there's a lot more testing that has to happen and a lot more involvement with the client. But then in many ways, you might have to hold the client back or manage their expectation because the guardrails you anticipate that you might need is something that they might, for one reason or another, they might want to put the chatbot in front of the general public, for example.

So, we thought that that was not a great usage for it because the way it's designed, and I think the way LLMs work, is if you ask it in a certain way, it answers you in a certain way. So, if you ask it as a layman, it's going to basically just say, "Go to a lawyer," or if you ask it, "my foot hurts," it's going to say, "go to a doctor." But if you have the details and you prompt it in a way that a lawyer would prompt it with the language, with exemplars, with examples, with strictures, et cetera, and with a type of incentives, the right type of incentives, then it can come back with a useful answer.

So, New York City just recently had this problem where they prematurely rolled out a legal chatbot and it was advising small businesses with all sorts of wacky stuff, and it was just a big embarrassment. But I think that's an example of user's expectations and the client's expectations getting ahead of where the verifiable testable capabilities are. And I think a lot of times, in terms of technology and in terms of technology adoption, the proper use of a certain technology doesn't emerge until later.

The interesting thing about humans, about us is we are creative, as Dan mentioned, and our creativity often extends towards breaking or pushing the boundaries of what technology can do. So, oftentimes, that is misuse of the technology. So, the inventor of the phonograph may not have ever anticipated the rise of the DJ. And I think, just, we're in that early period now. I think when our students begin practicing in 20 and 30 years, they will redefine what it means to be a lawyer.

So, a lot of the things that we may find as crucial to the practice of law is in fact something that can be automated, and what we're asking students is, to raise their game and to exercise higher levels of creativity or organization or rhetorical facility and just let the chatbot do that grunt work. And that, I think, was something that was reserved for new graduates or for new associates in the past, and I think that was touted as a learning experience. But I think there's a lot to learn and there's a lot of ways of learning, and I think learning in a hands-on environment is probably a lot better in terms of their continuing evolution.

Donna Lee:

Thanks so much, Basem. And I'm sorry to interrupt.

Basem Aly:

No, it's totally fine.

Donna Lee:

And we definitely will have more time for talking and sharing. I do want to do my duty as a moderator. The CLE code is AJC1-041124. And I was told to say it twice and I like to do what I'm told, so the code is AJC1-041124. Okay.

There was one question in the chat. Sateesh, do you want to unmute and ask your question?

Sateesh Nori:

Yeah, yeah.

Donna Lee:

And then they say the only bad question is the question not asked, so I would encourage others, if you want to throw questions in the chat or just unmute yourself or raise your hand, however, we do have until 4:20 or so, to do the Q&A.

Sateesh Nori:

Okay. I want to be a little provocative, and thank you all for your presentations, but my question is are we emphasizing too much how this stuff works and not emphasizing enough how to use it practically to help real people? So, for example, I don't know how Google works. I don't even know how the internet works. I can't explain electricity or plumbing, but I take for granted that I use those things every day. So how much do law students really need to learn how this stuff works? How relevant is that and how meaningful will that be with the rate of advancement and the adoption of this stuff in every other sphere of our lives?

Dan Jackson:

I mean, I think you absolutely have to couple... Anytime you're going out into the world and you're creating something for a community, for an organization, I think you absolutely have to have a really good user... Well, Sateesh, you're at JustFix, right? Is that right? Yeah. So, you know all about that, right? So, you have to have very good design research, very good user research, very good testing, and all that. And you absolutely have to couple that.

I mean, that's a whole area of work that we did, and I didn't get into. Whenever we move outside of the classroom, this imagination stuff that we're doing is just one piece of the puzzle. But I think you're absolutely right that we absolutely do need to make sure that we're designing things, creating things that are going to be absolutely useful, usable.

My biggest concern in that regard is what do we, and especially in the access to justice space, what do we do about not just the digital divide, but the digital literacy divide? People who simply don't trust computer systems, they're not going to put in their private information, and especially if they're undocumented, it's just not going to happen. And so now you've got a whole other class of individuals who can't access the courtroom because they don't have a lawyer, can't access the technology because they don't have a smartphone or a laptop or a 14-year-old nephew.

Drew Amerson:

Yeah. If I can follow up there, I think that's totally right, Dan. And Sateesh, you raised the example of Google. I think that's a perfect example. Google killed the search market because it was so good. Before that, you were really stumbling around in the dark. You type in a search query to Google and it would lead you to the right sources, which was revolutionary, and that's fallen apart as it's advanced in its age.

But one of the things that I've taught, so I had students pull up two, they had GPT-4 and they had Claude. And I would ask them all to type in some questions, and one of the things that I then had them do is, "Let's go try to find this on the internet." Because I think that one thing we have to teach students is how to be a critical thinker. Can we trust the answers that they're giving us? Where are these coming from? So, I think at this really early stage, it's really important for us to get to the point where we can trust these large language models and what they're spitting out, and I don't know that we're there yet, so I'm trying to have our students really be able to, and trust what they're getting.

Alice Armitage:

So, before you turn off the mute or before you mute us, I would say that what the students and I have found, and I was surprised by this as well, but by using an LLM not just to summarize cases, but to ask for their holdings, that the LLMs are very unreliable. I mean, sometimes they'll read the dissent and think that's the holding and put that up there. And students, as much as I've told them, "You cannot trust this," some of them, they just used what the LLM told them is the holding of the case.

And I think being forced to use LLMs and then try to use them in a legal task has made them realize that there are great benefits, there are great ways to use LLMs and what they can do, but they're not, and maybe they will be there two, five, 10 years from now, they're not there yet. But I do think that they're valuable enough in what they can do

that they are incredible time-savers and they're also, they're great at brainstorming. They throw ideas at you that you would never have thought of and give you a platform to jump off of.

So, I do think it's important for them to use it, just as it was important to start using the internet. People were so afraid of the internet and email, and you had to nudge people into it, and then of course, it became all they could do. So, there's that problem, that generative AI, we may not be having these conferences, the robots will, but in general, I do think it's really important for people to dive into this and try it.

Keith Porcaro:

To maybe, I guess, be the resident grump and to answer Sateesh's provocation with another, lawyers need to learn about electricity when their client gets electrocuted. People need to learn about Google when it turns out that people are using rarely searched Google search terms to spread misinformation. People need to know that large language models are probabilistic next-word generators, so they don't try and use them to generate citations.

I think for at least us in the class that I teach, it really is focused on what are the... Explaining how something works is important so that students can start to apply other skills that they already have, both to, I'd say, two things. One is to identify client issues and prevent harm to clients, which is a core of what they are supposed to be doing as lawyers, but then also to develop error-correction skills. And I think that's one thing that we, all of us, have not quite mastered when it comes to integrating large language models. And a lot of the studies that have talked about how to integrate large language models into other workflows have noticed it's just hard to spot the types of errors that LLMs generate. They're just different from the types of errors that a paralegal makes or a first-year associate makes.

And so, when we're teaching, I think we're doing that at a delicate time where students don't necessarily have some of the foundational reasoning and knowledge that we hope for them to have. I don't think that means don't use large language models. So, when I teach it, for instance, I have the students start with red-teaming it. So, let's come up with a bunch of questions that somebody might ask of a large language model that are the complicated corner cases, and then use that as a way to test how well, say, Claude or ChatGPT does answering questions about guidance about how to vote in North Carolina, for example.

Basem Aly:

You stole my thunder. I wanted to be the resident crank. I'll just say that at this point in time, LLMs are mansplaining engines. They simulate confidence and they produce fluid language, and that fools some people. But I think their true facility, I think, is yet to be

discovered. And a lot of these technologies, I think, it's really a matter of discovering its use rather than designing its use.

So, I think over time, its true métier will be revealed. I think a lot of times, people may think of an LLM as a replacement for Google, but it's not. It's a little bit different. But over, just in the last six months, there's been a lot of progress in internet-accessible LLMs that are Google plus some language model. So, it's a shifting field.

Dan Jackson:

"Mansplaining engine" is one of the best descriptions of generative AI and ChatGPT that I've ever heard. I'm going to use that, but I will cite you.

Basem Aly:

No problem. I'll help you source.

Dan Jackson:

There's a question on legal education. How do you nudge legal education? I think Justin asked that question. How do you nudge legal education to require more technology courses?

Nudging legal education is a very difficult thing to do. The only real hammer, and I would love other educators, especially deans, associate deans and whatnot, the fancy folks on the call, is to get the American Bar Association accreditation standards to change, which takes eons to do, but can be done. If you just look at the requirement now for some experiential, but that was a 20-year effort, 20 or 30-year effort to get experiential learning required as part of... They might move a little bit faster these days, but that's my... Unless of course we can make the case, and I think we can make the case soon, that our employers are going to expect it.

I tell every student who's a first year when they talk about the courses that I teach in AI. I've got another one coming up this fall that we'll be teaching for the first time that's specific to AI uses and regulation and risks. It's going to be every employer's expectation, consistent with Basem's point on digital natives, that these folks, that new lawyers are going to know something about this and know how to educate the rest of the firm or the rest of the office about it. However unfair that is, that is going to be the expectation.

And I had one student who was in a co-op after taking my course, had the AI project on her resume. During the interview, they wanted to hear about. She showed them her prototype, they loved it. When she walked in on her first day, she was introduced to her team at a very large international organization, human rights organization, as an AI expert and she was put in charge. I know, she was horrified. She was put in charge of a team to start building and iterating on an AI tool. They loved what she was doing, and

now they're looking for funding to bring her on board full-time, postgraduate, as their AI expert.

Keith Porcaro:

Yeah, I think find-

Basem Aly:

That's the nature of disruption. Yeah. Oh, sorry. Go ahead, Keith.

Keith Porcaro:

No, no, I was just going to say I think finding small ways to foment demand is really important because we think about it from the teaching side and the curriculum-planning side, but students are super busy and they've got tons of things that they want to take, and there are lots of competing priorities for their attention.

So, one of the things that we tried with the algorithms class that I teach is trying to make it that it fits really easily into anyone's schedule. So, it's one credit, it's a big class, it's pass/fail, and that starts to get students interested elsewhere. So, we're probably going to do a session for the 1Ls on how to protect yourself online, which is not quite legal tech and it's not quite tech law, but it gets them to start thinking about how these things are going to affect them in everyday life, even if they don't choose the niche of tech law that we've cornered it into.

Donna Lee:

I'd like to pose Sateesh's question in the chat about clinical programs. A lot of law schools have programs that are based on a traditional legal services model, poverty courts, clients in housing court, family court, immigration court. But the huge difference is that our clinic students might have one client for the semester or they might be teamed to both represent one client as opposed to dealing with the volume that legal services providers have to deal with.

I'd be curious to know the panelists' thoughts about ways to potentially integrate AI into that kind of model of clinical teaching, or I don't know, whether you think a different model may be called for.

Drew Amerson:

I'll jump in here and answer the last two questions. So, one thing that's been great at our school is our academic dean has been a champion for us in leading and integrating generative AI into any course that wants to do it. So, we led a faculty retreat on available generative AI tools, and that was a lot of fun. It was also a little dangerous because I was taking sample exams that I pulled from professors and running it and getting answers

and then getting corrected by those professors. But one thing that I learned there is our tax clinic was using a tool that I didn't even know about, it's called Blue J, which is a generative AI tool for tax. So, I know that that clinic at least is going down that pathway. I'm not sure about others.

Alice Armitage:

Well, I think that, no, not yet. And part of it is because they're trying to get students into courtrooms and into actual positions of advocacy as opposed... I mean, yes, it's all access to justice focus, but I don't know that they're looking in terms of widening access to justice in general, but having more students go down that path when they graduate.

I do think there would be a huge role though for any clinical professor who wanted to do this to try to take a particular problem in the housing court or whatever, a particular form that needs to be filled out. And in San Francisco, you have an eviction notice. You have five days to respond, and so many people who get those don't even know what they need to do. But if there would be some way, and we do have a Right to Counsel program, but to automate that answer very easily and really all you have to do is put your name in and send it to the court and that is deemed an answer and then you have more time.

So, I do think that would be a great thing for students to do, but that's not what clinical, in the past, clinical courses have really focused on is that as their mission. It's more about getting the students to have some real experience with real clients.

Dan Jackson:

I think Suffolk Law School did-

Basem Aly:

Just come to Columbia.

Dan Jackson:

Oh, go ahead.

Basem Aly:

I said, "Just come to Columbia." Visit our clinic.

Dan Jackson:

Sounds like it.

Donna Lee:

Yeah. And I love the idea of pairing that individual advocacy work with a broader project-based work that multiple students could engage in on access to justice topics. Thank you.

Elyse Diamond:

Alice, you've talked, too, a little bit about, leading up to this program though, about some of the work involved in creating a collaboration and presenting a problem, so there is that facet. I don't know if you guys want to speak to... Each of you described an interesting project where you took on something relatively manageable, I suppose, but could you, maybe, want to speak a little bit to that?

I think that perhaps what's a little bit intimidating in some respects to folks is it can be a lot to come up with, to be working with community partners and educated about an access to justice issue, and then present that in real time to your students. That can be a challenge in a clinic setting. But I know Alice, for example, you've talked about some of the challenges in putting together some of those projects.

Alice Armitage:

Well, I have to say that the hackathon I created... The first time I taught about homelessness, I taught the students the basics of it, but then I wanted them to interact with people with other disciplines. So, I created a live, public hackathon in San Francisco for a whole weekend, and that was very difficult. And I had a few professors before, before this, when they heard what I was doing, questioned my capabilities or my thought process in deciding to take that on. And it was incredibly stressful, I have to agree with that, but it was an amazing experience for the students because they were actually, in real time, talking to engineers and UX designers and stakeholders that had the problem themselves.

And it made them, every one of them, after the fact... In the beginning, they all said, "Oh, why are you making us do this? This is awful. We don't know anything. We're lawyers. We need to stay in our lane." And then, after that 48 hours, they were so excited by what they were able to achieve with their knowledge in combination with other people. Frankly, I think that's a problem in lawyering in general. We're so siloed and we think we do our own thing and we should stay in our lane, and I really like to push the students to not do that. And frankly, I think that generative AI is going to be a way that people will begin to realize that they can do other things besides stay in that lane that we all learned about in 1L year of law school.

Dan Jackson:

That's what clients want. Clients want lawyers who actually know how to work across disciplines, also picking up on Mark Palmer's helpful comment in there about other disciplines. Clients do not want lawyers who stay in their lane. They want lawyers who are going to poke holes and everything, see the big picture, who understand all the aspects of it.

The future of education right now out there from, I mean, any of us who are educators, and probably I feel sorry, given that we're in this particular breakout session, is interdisciplinarity. There's just so much potential. Not a lot of universities have figured out a budget model that actually will accommodate that, but I'm hoping that that will come soon. And it's just really exciting to see and I think there's a lot of potential for that in the law. There's a lot of folks coming into our lane in the law space and I welcome them in and I think we should be as deep into other lanes as possible.

Keith Porcaro:

Yeah, I think I've had a lot of really positive experiences teaching interdisciplinary classes. Duke has a bio and tech ethics program, and so I teach a class that is a mix of law students and public policy students and tech ethics students, and we do simulations. And so doing simulations that ask them to work together to, say, unpick issues that are developing in a software company, where they're looking through fake data and looking through fake emails and trying to learn about data issues there.

And, so, I think, I don't know, we don't tell all students that we're teaching them a new language and a new way of thinking. And I think exposing them to other disciplines and other structured ways of thinking is a great way to make them more aware of the strengths and limitations of what we're teaching.

Donna Lee:

Earlier this afternoon, in one of the sessions, one of the panelists spoke about AI and applying technologies with the right architecture and talked about guardrails. And, as someone who tends to be tech phobic, I wasn't exactly sure what the speaker was talking about. I don't know if you all feel like you could address that.

Alice Armitage:

Well, in terms of LLMs, I can tell you that when they're talking about guardrails, it could be in the case of hallucination, it could be in the case of memorization, but what they're saying is that they're adding different algorithms basically in the fine-tuning of the LLM or in retrieval-augmented generation capacity of the LLM to get it not to do things like hallucinate, to be biased, to undo some of the bias that is in the training dataset.

So, for instance, the New York Times, and when they sued OpenAI and Microsoft, had all sorts of examples of memorization where people prompting ChatGPT were able to get entire portions or maybe entire articles reproduced exactly or almost exactly. And, so, after that, if you put those same prompts in, now, that the New York Times used, you cannot get ChatGPT or GPT-4 to do that. It tells you, "Well, no, I can't do that. It won't be done."

But there was an interesting example with Google. So, when they released Gemini a few months ago, they had an image generator with it. And, so, they had very carefully created guardrails to prevent if someone is looking for an image of a doctor and it's always a white male, but they didn't quite craft it well enough so that if someone went in and asked for the images of the popes over the years, they would come up with Chinese women and African American tribal chiefs as examples of popes, and one famous one that went around the internet that someone asked for an image of George Washington and it was a Black man.

And they're right to be trying to fight the bias, but it's very hard to build those guardrails in a way that they get an accurate functioning of the LLM, and yet also we're finding the bias and the discrimination that underlies all the datasets.

Basem Aly:

I think it's very early days, but I think before very long, there'll be an emerging small language model ecosystem and agent ecosystems. So already right now, you can go on GitHub and download a framework that will let you have multiple agents chat with each other and delegate jobs, et cetera. I think before too long what you'll have are very specialized, focused language models, that their job is to double-check and verify facts or try and overcome stereotypes or do other things like that.

Essentially, the way LLMs are designed, they're designed side by side with GANs, a generative adversarial network. That's how LLMs sound, that's how they can produce fluid language. So, I think very, very soon, it'll be easier and easier to deploy these types of adversarial agents or GANs or other things like that that will verify those outputs and build up those types of guardrails.

But it doesn't obscure the fundamental fact that at this point in time, we have a tendency to anthropomorphize LLMs. So, we tend to think of them as oracles, when in fact, they're muses. They're something that will reflect back on us or give us some type of symbolic output that matches our expectations and tries to essentially please us. So, you might tell an LLM that the world is flat, "The scientists have just discovered that the world is flat," and then it'll proceed and discuss with you based on the idea that the world is flat, et cetera, because it's trying to please you. It's basically serving you words, which is what you're asking for.

So, I think over time, people will, I think, be more comfortable with the idea that they're more muses than Oracles. I think it would help if people engage with them in low-stakes, fun activities, and don't publish a BuzzFeed article about how my LLM planned my vacation or did my meal planning. Nobody needs to read that. But I think the more people are exposed and play around with these things, the more they realize their limitations.

But ultimately, I think one of the most intriguing scenarios for LLMs is in terms of being an interlocutor negotiator, a sounding board, a shoulder to cry on. On Reddit, it's really hot to have erotic chatbots. So, these are things that are filling a type of need that people have, and oftentimes it's really a need to just speak with someone who is understanding of you, who doesn't judge you, who reflects back your own biases and tells you you're a good person, all of that.

So, I think it'll be a little bit of time till we get over that and sort of develop those types of tools and defenses that will help us navigate through that.

Alice Armitage:

I would just like to chime in and agree with you, and also say that even now though, you can find online, and I think in Claude's prompt library, it's a long prompt, but it asks the LLM to act as a tutor. And I have the students do that and interact with it, and it actually, it does what you said. It supports them, it helps them, it doesn't give them answers, but it nudges them in ways to get better answers, to think about something from a different perspective.

And I think that is a wonderful... Especially in law school, we don't have that much time, or usually, to give individual assessment or to sit down with somebody and point out exactly which part of their thinking went wrong and the LLM, when properly prompted, can do it already. But I think you're right. Eventually, there'll just be a tutor LLM that does that for you. And I think that's a wonderful thing. And as people who teach, I think we should try to embrace that and understand its limitations of course, but also embrace it when it can provide a service to the students that we can't.

Dan Jackson:

Yeah. For those of us who are running these innovation labs or design shops where we're trying to get students to do all the ideation and the brainstorming and the co-design and collaborative design, all this stuff, it's very difficult to do that alone. It's like I used to come up, before LLMs, I used to come up with these weird exercises that tried to recreate what it would be like to brainstorm with another person, but it's just yourself because law students are busy all day, like 10:00 PM they get home, at that time to start working on some stuff.

And we have had tremendous success in our classes with using those tools as brainstorming partners, as co-creation partners, even as product testers, user testers. You give the chatbot a direction of exactly what type of entity it is, here's who you are, what your story is. "Now I'm going to upload this document, give me feedback on it." And it's pretty successful in that regard.

Keith Porcaro:

I think the one thing I might say is we talk about how it might be 90% of the way there, that last 10 percent's a real devil. We heard about it with self-driving cars, "They're almost there, as long as they can get better at taking naked left turns and driving in the snow and not running over bicyclists and on and on." It's that last little bit that tends to be really difficult. And the other bit that comes out of that is that we design these tools and then we expect humans to take over when they fail, and it turns out that we're not very good at that either.

And, so, I think the caution that I might invite... Two things. One is, I think leaning into a future or imagining a future where that last 10% gap never actually closes, where they remain good for some things but not good for others, and what does that mean for how we can use them and how we can use them responsibly? And the other is, I think, to some of the other points that have been made, it's an invitation for us to really consider whether or not any of the pedagogy that we've been using over the last century or so is actually working.

But one of the interesting things that comes out of the early human-computer interaction literature is chatbots as an interaction mechanism are actually terrible at helping people learn, for some of the reasons that Basem talked about, is they don't reveal their underpinnings. So, there's some question of, okay, what kind of evidence do we need to show learning advancement with this type of interaction model? And it raises a question of how can we look at the way that we teach the law and what evidence do we have that's the case method that we've developed because it seemed like a good idea 100 years ago from somebody at Harvard is actually working?

And, so, I have a suspicion that it's not going to be a large language model, is not going to take over 1L classes, but my hope is that it's going to invite us to think about how we can help students develop these foundational reasoning and knowledge skills that we care so much about, and that will make them responsible professionals.

Donna Lee:

Keith, that's such a wonderful transition to what I hoped we would end the session with, which is maybe takeaways or ideas that our panelists, experts, and experienced tech users and teachers have that we can share back with folks in the other concurrent

sessions. So, if you guys want to drop your pearls of wisdom now, that would be fantastic.

Elyse Diamond:

Yeah, it sounds like Keith got us started with that, so if other folks want to chime in-

Dan Jackson:

I know, right? That was beautiful. I was going to say that I can't-

Elyse Diamond:

... with lessons learned.

Dan Jackson:

I guess I would just say that I think it's incumbent upon all of us as educators to get law students genuinely invested in the direction of artificial intelligence, both in society and in the law, and that we must not just hand this over exclusively to the technologists and commercial interests. Lawyers have to be in the mix as co-creators.

Basem Aly:

I would say that I think it's time to redefine AI. It should actually mean augmented intelligence. It should mean me plus the computer equals augmented intelligence. And it's very, very good at structured output. So, aside from just playing with it and producing low-stakes content, I would advise you to be playful with it, to play a game, to try and generate code. I think one of the most unheralded aspects of LLMs is that it's really the ultimate no-code engine. It's opening the door for people who have absolutely no programming chops to produce workable code and then to have the LLM explain to them the code and then to modify it. That's a tremendous advance right there.

But from my understanding, I think a lot of students and a lot of lawyers would just be happy to basically have an LLM consistently number paragraphs in a long document. A lot of people have trouble with that. So, I think the more you can interact with it, the more you can demystify it, the more you can be playful with it, the better. I wouldn't go so far as to say to try and break it, but that's always helpful.

When I'm troubleshooting technology, I like to try and find its breaking point or what you can do to bump up against those guardrails. A lot of times, you can ask it to reveal its assumptions. You might also ask it to, after it gives you an answer, you can ask it to interrogate that answer and to find the flaws in that answer. You can have the LLM be reflective, but that will never be a substitute for you being reflective, for you exercising

judgment, and for you to wrestle with the quality of the output and to try and get it to really follow your orders.

Elyse Diamond:

Drew or Alice, any last, a lesson learned, or what you want to leave folks with who are thinking about or already involved in diving into this with their students?

Drew Amerson:

Look, I would just echo what Basem was just saying. There's no substitute for playing with these things, and there's a really low barrier of entry. You can get your students on this immediately and have them play with it. And I do tell my students to try to break it, to find those guardrails to push. And typically, they're into it. They have fun once they understand what they're trying to do. So, no substitute for actually getting there and getting your hands dirty on these.

Alice Armitage:

I guess my final takeaway would be all of what has been said before, but also go a little farther and embrace generative AI. First, maybe just in your personal life. You can take a picture, get it on your phone. You can take a picture of the food in your refrigerator and ask it to come up with recipes that you can make for dinner that night, or a picture of your closet and ask it how to better organize your clothes. And it's actually really good at that.

So, use it for what it's good at and it's going to get better and better. GPT-5 is coming out pretty soon, and all the rumors are that it's going to be a paradigm shift, and I think everything's going to change, our personal lives as well as how we practice law, and many of the doctrinal topics of law will also be impacted.

So, I don't think we gain anything by trying to ignore it. I'm not saying use it for everything but get comfortable. I agree. The bottom line is get comfortable with it. This is in our lives and it's not going anywhere.

Donna Lee:

Thank you so much for such an interesting and generative conversation. I am going to share the, sorry, I don't want to end with such a boring thing, so I'm going to share the CLE code, which is AJC2-041124, AJC2-041124, and then maybe we can show our appreciation to the panelists. Thank you so much for an amazing conversation.