TRANSCRIPT OF AUDIO FILE:

REGULATING GENERATIVE AI

The text below represents an understanding of the words spoken. No guarantee of complete accuracy is expressed or implied, particularly regarding spellings of names and other unfamiliar or hard-to-hear words and phrases. (ph) or (sp?) indicate phonetics or best guesses. To verify important quotes, we recommend listening to the corresponding audio. Timestamps throughout the transcript facilitate locating the desired quote.

BEGIN TRANSCRIPT:

INTRODUCTION: Welcome to A Hard Look, the Administrative Law Review podcast from the Washington College of Law. We'll discuss how administrative law impacts your daily life, from regulatory actions by agencies and the litigation over them, to the balance of power among branches of the government. This is A Hard Look. [00:00:39]

ALEXANDER NAUM: Welcome to this episode of our legal podcast, where we'll delve into the challenges of regulating artificial intelligence, AI, and the copyright issues that may arise, particularly in the context of generative AI. As AI technologies continue to advance at a rapid pace, the legal landscape surrounding the regulation struggles to keep up. With the development of generative AI, which can automatically create original works such as music, art and literature, questions surrounding copyright ownership and infringement have become more pressing. In this episode, we'll explore the complexities of regulating AI, the legal frameworks that currently exist to address these challenges, and the potential implications for copyright law in the future. Join us as we navigate this fascinating and ever-evolving legal frontier. [00:01:28]

EVA PEDERSON: And that introduction was written by Chat GPT. Professor Carroll, could you tell? [00:01:36]

PROF. MICHAEL CARROLL: I could not. No, that was pretty good. [00:01:39]

ALEXANDER NAUM: It's wild to think how far AI technology has advanced, even to the point that it could be writing the script for your favorite academic podcast. I'm definitely joking. Chat GPT only wrote that opener. So let's hop back to humans. My name is Alexander Naum and I'm the senior technology editor for the Administrative Law Review. I'm glad to be hosting today's episode with our technology editor. [00:02:00]

EVA PEDERSON: Hello, it's Eva. [00:02:01]

ALEXANDER NAUM: And our guest, another incredible human, Professor Michael W. Carroll. Professor Carroll received his J.D. in 1996 at Georgetown University Law Center. Shortly after, Professor Carroll helped to found Creative Commons Incorporated, an international organization that encourages the free sharing of both creative and copyrighted works. Professor Carroll still

remains involved with the Creative Commons U.S. Project at American University Washington College of Law, where he's also professor and faculty director of the Program of Informational Justice and Intellectual Property. Professor Carroll additionally served on the National Research Council Board on Research Data and information from 2008 to 2013, and currently serves on the board of Public Library of Science and as an academic fellow at the Center for Democracy and Technology. I'd like to remind our listeners that these are the personal views that Professor Michael W. Carroll and do not reflect the views of his employers, clients, organizations or other individuals onto which these opinions can be imputed. Thank you for joining us today, Professor Carroll. [00:03:09]

PROF. MICHAEL CARROLL: Thank you for having me. [00:03:10]

ALEXANDER NAUM: Yeah, of course. Of course. So, yeah, let's start from the top. AI has become a hot button term and one that is commonly misused. What does artificial intelligence mean and how far has it advanced our society? [00:03:22]

PROF. MICHAEL CARROLL: Yeah, I agree with the premise. It is often misunderstood. Artificial intelligence it is not subject to a uniform definition. But the general idea is that, can we use computers to function in a way that mimics human intelligence, such that essentially the computer could actually quote unquote think independently. Within the field there is this sort of term of general artificial intelligence. That would be a machine that is generally able to operate intelligently like a human being. Or narrow artificial intelligence, which is a program that can operate like a human being within some domain. So Chat GPT and all of this generative AI is not general artificial intelligence, right? We can't ask it to drive a car. We can't ask it to do a variety of things that require intelligence that are outside of its programming. So as impressive as these new technologies are, they are still just examples of narrow artificial intelligence. [00:04:37]

One other confusion is, a lot of what we're seeing now is a subset of the field of artificial intelligence, which is really what's called machine learning. You create some data and then you create a model based on that data. You train the model to give you an output. So let's say you want to train the quote unquote AI to be able to identify a cat or even produce a painting of a cat for you. You feed it a million pictures of a cat. It reads that data. It now understands or it acts like it understands what a cat is, and then it can generate cat images based on having been trained on that data. That's not really AI. That's machine learning. And that's what a lot of the stuff that gets called AI these days is really about. [00:05:28]

ALEXANDER NAUM: You mentioned Chat GPT, which I feel like a lot of our listeners just definitely have interacted with. Also another app called Lensa (ph), which a lot of listeners also likely have interacted with as well. What are these products? I know they're machine learning, but can you dive into just like these specific products and what they provide for users and how they generally work? [00:05:49]

PROF. MICHAEL CARROLL: Sure. So the technical term for these are LMM or large language model, and essentially what they are, are tools that are built on predictive analytics. So as impressive is the output of Chat GPT is, it doesn't quote unquote know what it's saying and it's not programmed to know what it's saying. Instead, it's looking at the patterns of data, primarily

textual data that it has been fed. And it's essentially saying, based on the word I just used, what would I predict the next word would be? So let's say we want Chat GPT to draft us a corporate report for the Securities and Exchange Commission. It's going to use the data of having had a bunch of corporate reports in its dataset to say: when the word is this particular word, the most likely next word would be this. And so it's really just predicting based on past usage of language. And it's impressive in the outputs because it really looks like the way a human might talk or write. But that's just because it's computing on a massive, massive set of data based on how humans have been writing in the past. [00:07:11]

In a weird way, the programming that goes into these tools is programming we've known about. This is effectively the same thing as autocomplete when you're trying to type in a search in your search engine, it's the same idea. It's just that these models are based on massive amounts of data using massive amounts of computational energy that is unprecedented. And that's why these tools are catching so much attention. [00:07:39]

EVA PEDERSON: Amounts of data that a human could never hold in one brain. I remember reading a story the other day of a woman whose dog was saved when she fed the blood test results from a veterinary appointment that couldn't be deciphered, and Chat GPT offered a couple of potential diagnoses. So just throwing that out there, what are the benefits of these products? And then why are they controversial? [00:08:06]

PROF. MICHAEL CARROLL: It's a great question. And I think we don't really know the answers to either side of the coin completely yet. I think the benefits are, how highly adaptable these tools are, that they are general purpose technologies, which means they don't specialize in science or math or English, or they can be used by people trying to solve all kinds of different problems. So kind of like a computer or kind of like the Internet. It's a general-purpose technology that can be used for lots of different reasons. And so we haven't seen all the different good and bad outcomes that can come from that.

What we're seeing in terms of the good, is it's a really good tool for things like idea generation, potentially outlining information, you can feed it information and it can format and answer back to you that helps you then take the next mental step yourself. What we're also seeing is that the tool is even more useful as you refine your own knowledge of how to use it. We already have power users who are learning that you don't just ask it one question, you go back to it over and over again as you refine the outcomes and it gets better and better as you do that. [00:09:21]

The downsides are, because it really doesn't know anything about the content that it's producing, it produces wildly, factually inaccurate statements that it can write in a voice that seems very confident and it can be very confidently very wrong. And it can also be very wrong in harmful ways, produce what a human would be held legally liable for, libel in defamation, because it will make factually incorrect statements that could harm a person's reputation. [00:09:53]

So I think we have to be careful with how we treat these outputs, and not ascribe intelligence to this. This is not intelligence. It is just a computer model that's producing an output. But it sure looks and feels like an intelligence. And that's one of the big challenges for society at large and for the law. [00:10:14]

EVA PEDERSON: That's right. And I think that Chat GPT is an example of a program that's tried to issue several disclaimers on its website to say, this is just a language model at this point. But you might still see in the future people relying more and more on Chat GPT, regardless of those disclaimers. So that's great. Pivoting to what's used to train the model. Are there examples of other programs that use licensed materials or potentially copyrighted materials to produce content for users? [00:10:46]

PROF. MICHAEL CARROLL: Sure. I mean, I think because copyright basically applies to almost every kind of content, it seems to me that no training dataset would have zero copyrighted content in it. And yeah, I mean, we have to be careful about this technology. So there's a company called Clearview AI that basically just scraped a bunch of image data from Facebook, and it's now selling facial recognition technology to police departments by having scraped millions and millions of copyrighted photos without asking either the subject of the photo or the photographer for any permission. Now, I think they didn't need permission under copyright law because I think using copyrighted materials to train a model is an example of fair use. But I am very troubled by Clearview AI's use of those images because its facial recognition technology is inaccurate and it has gotten people arrested. There is a lot of angst among various copyright owners about their data going into these training sets without them having a say over that. I think right now, at least under the law in the U.S., that that is likely to be a contested issue. [00:12:02]

EVA PEDERSON: So somehow the privacy interest is a little bit more compelling than the intellectual property interest, what you're saying. I think there are a few states like Illinois, for example, that have their own statutes to protect the privacy interest but not the IP interest. Would there be any potential claim for creators whose works are used by these models? [00:12:23]

PROF. MICHAEL CARROLL: Well, we're going to find out. Because there are two big lawsuits that were recently filed against Stability AI, which creates the tool Stable Diffusion. One of those lawsuits was filed by Getty Images and the other was filed as a class action lawsuit on behalf of a number of creators of photographs and other visual art. And there the claim is that – so just to be clear about what I think the fair use piece of this is, it's the copying of copyrighted data on the input side. So if the only reason you're copying a copyrighted material, a piece of copyrighted material is to train the model, but then the output is not in any way similar to those copyrighted inputs. I think that's the fair use. [00:13:11]

What the claims are here, though, is that Stable Diffusion's output are in fact incorporating some of the protected expression of the images that were fed into it, so that the claim is not just about copying on the way in, but the copying that results on the way out. Getty Images also says that there's trademark harm because some of the outputs use the Getty Images brand, but in a distorted way. [00:13:38]

So we will see. And maybe what we're going to do is draw some fine lines around fair use. But I think that will be focused on the outputs of these models, not so much on the inputs. [00:13:49]

ALEXANDER NAUM: And for listeners who may not know what the Fair Use Doctrine is, can you describe what that means and when it's normally invoked? [00:13:58]

PROF. MICHAEL CARROLL: Sure. So copyright on its surface is fairly simple. If you want to bring a copyright claim, you really just have to say two things. You have to say, I am the creator of an original work, and an original work just requires you to have done something that required some choice about how to express yourself. So if you're a photographer, you chose the lighting, you chose the angle, you chose the zoom, how far zoomed you want to be. And those are your artistic choices. You get a copyright for those. And then the claim is that somebody else has come along and used your work in violation of your exclusive rights to make copies, to distribute copies, to make derivative works and so forth. [00:14:42]

If the copyright owner can show those two things, the analysis doesn't end because the user can then come back and say, Well, even if I did use your work, my use was a fair use. And fair use is a flexible balancing test where the courts look at four basic factors that boil down – I'm not going to give you the technical legal language unless you want me to, but the lay person's understanding is, you ask the user, Why are you using this? Are you using an amount that's appropriate for your purpose? How much original expression is actually in the thing you're using? So there's going to be a difference between my using a novel and my using maybe a table of data that has just a little bit of expression in how it's organized. And then the last question is, Is that use causing harm to the copyright owner economically? Because it's either a substitute or you're depriving them of some licensing revenue. [00:15:42]

ALEXANDER NAUM: And I they know that the Supreme Court is currently reviewing the grounds of fair use in a case known as Andy Warhol Foundation v. Goldsmith. What is this case about and does it have any implications towards AI models? [00:15:57]

PROF. MICHAEL CARROLL: On the answer to the second question, it might. We have to see what the Court says. The Court is answering a very specific question. So the first question under fair use is, is the so-called purpose and character of the use. That is, tell me user why you are using this copyrighted work. And the Court has said, if we can say that the use is quote unquote transformative, then that's the kind of use that is likely to be fair. [00:16:29]

Now I have to put air quotes around transformative. That's in the law what we call a term of art. A transformative use does not require that you actually make a derivative work or change the underlying copyrighted work. You can also use it for a new purpose. So if I take ten pages out of a novel and I bring it into a classroom to do a critical reading exercise with the students, that would count as a transformative use because I'm not using the work for the pleasure of reading the novel, I'm using it for this different educational purpose. So whether you sort of recontextualize the original work or add some new meaning or message to it, that can also count as a transformative work under the sort of Court-made interpretation. In Andy Warhol, what happened is, Andy Warhol used a photograph of the musical artist known as Prince in order to make a pop art sort of portrait of Prince. And the photographer Lynn Goldsmith is suing for the use of that photograph. She licensed to use to Vanity Fair for an artist to use the photograph as a reference to make a single artwork which could be published in a single volume. But Andy Warhol went beyond that and actually made a series of portraits. And so that's the basis of the lawsuit. [00:17:58]

Ans the specific legal question the Court is going to answer is, how do we interpret when a use counts is transformative if it's also creating what we would call a derivative work from an original work. How do we sort of put those two things together? So we'll see what the Court says. But what it will say is something about transformative use means X. And in the AI context that may have some bearing on what we were talking about before, whether the outputs of these models and the similarities would count as an unauthorized derivative work, or would they count as a transformative work that is treated as a fair use? [00:18:37]

EVA PEDERSON: I wonder if there might be an additional legal element of, is there human effort required in the transformation? Is writing a Chat GPT prompt or maybe selecting an image enough human effort to qualify? And I won't give you that one. I will ask you the flip side question, which is what have courts said about AI as a creator itself? [00:19:00]

PROF. MICHAEL CARROLL: Right. So here we have fairly clear answers. We have a Ninth Circuit case and we have a Copyright Office, both of whom are saying that copyright belongs to, quote unquote, authors under the law. And an author has to be a natural person. And I guess whether that's a matter of constitutional law or statutory law, the Ninth Circuit answered it in the question of, if you remember the case about the quote unquote monkey selfie, the photographer who, a macaque monkey got a hold of the phone, took a selfie, and then he sought to register copyright for it. And the Copyright Office said, no, the monkey is not an author under the law. And similarly recently, AI-generated or machine-learning generated outputs were brought to the Copyright Office and saying, Please register my claim, the author is the AI. The Copyright Office says, I'm not going to register it. AI can't be an author. [00:19:58]

But even to the point you just made, there's a really interesting registration that's happening right now where the author uses a machine learning model to create a graphic novel over the weekend, and then sort of refines the use of that AI. So there's some human inputs, copyrightable human inputs, and then there's some AI-generated outputs. And so there's a question about how you can register that work; can you register the whole thing, or only the human piece? So lots of fun, interesting copyright questions to be answered as we go along. [00:20:36]

ALEXANDER NAUM: So if the courts to determine, this falls outside of the grounds of fair use, can there be a balance between respecting creators' copyrighted material while still allowing AI innovation? I know that you were one of the founders of Creative Commons Incorporated. Can you talk more about what Creative Commons is for our listeners who may not know and can you see it being used as a tool to achieve this balance? [00:21:01]

PROF. MICHAEL CARROLL: Sure. So I have to take you back to the early days of broadband Internet when we didn't even have social media and people were doing a lot of self-expression through blogging mainly, and people were copying and pasting from each other's blogs the way that we now share on social media. And a group of us lawyers, mostly, and technologists, were worried that all of this stuff was happening under the shadow of copyright law, without any guidance as to how copyright was going to apply to this. And so when we created the organization Creative Commons, we created a series of standardized licenses that you can attach to your copyrighted work to tell people, Hey, I have a copyright, but I want to share it with you and I'm willing to share it with you under these conditions. [00:21:53]

And there are six flavors. We don't have to go into the details. But the most open is, you can do anything you want with this as long as you give me credit. And the most restrictive is, you can copy-paste this but you can't make derivative works and you can't make a commercial use and you need to give me credit. [00:22:10]

And then there's, we toggle in some other options in between. And so yes, people do go to the repositories of Creative Commons license content for all kinds of things. If you go on Flickr, you can find millions and millions of CC licensed images and I know a lot of corporate people use it for their PowerPoint presentations because then they don't worry about copyright because the license tells them they're allowed to use it. I know a lot of these AI model generators are sucking up as much Creative Commons license content as they can. Notice that Wikipedia is licensed under a Creative Commons license. So if they're using Wikipedia, which they all are, then they are by definition using CC content.

There are some concerns among some creators who put their works out under a CC license, really thinking they were sharing this with other humans, and are less thrilled with the Clearview AI type uses of their images. But as a copyright matter, the license permits that, even if its intention with a sort of sharing spirit behind the reason for the sharing. [00:23:22]

ALEXANDER NAUM: Yeah that's definitely very interesting. And I didn't even think about the creators that may have wanted to license it to Creative Commons, but didn't even think that it could be used in other ways that they didn't intend. I also wanted to discuss some of the broader ethical dilemmas posed by AI. Do you see these programs advancing to the point that they eliminate entire sectors of work, especially in the market for creators of intellectual property? Are we already there? [00:23:48]

PROF. MICHAEL CARROLL: Well, there's a lot of speculation about this issue. Certainly there's a lot of people who are very concerned. I think the thing to think about is, yeah, I'm sure there will be some forms of substitution. What computers are very good at is doing anything that can be patterned as a set of rules or as a set of steps. So anyone who has a job that feels kind of formulaic is probably at risk. Because if you can actually reduce the task to a formula, that's what computers are really good at doing. And so we've already seen like in the field of journalism that stories about the weather, stories about a sport like a professional game or something where there's a lot of data that gets fed in – computers are good at working with data – and the words that we use to describe the weather or the sporting event are very standardized and kind of patterned. We're already seeing machine learning being used to generate stories that are getting published with almost no human oversight at all. So journalists are worried about the capacity of that to get stronger. [00:25:03]

I think we're also going to see within the creative fields, again, think about a field of music like electronic dance music. In theory, right, a D.J. comes up with a flow but it wouldn't be that hard to write an algorithm that sounds like what a deejay would do at a club. But let me just ask you, would it be the same experience for the audience member if somebody just went up to the front of the stage and pushed play? I don't think so, right? [00:25:39]

I think there's more angst among the creative community than there needs to be. But I'm not going to deny that there are certain kinds of standardized, copyrighted outputs that AI may end up becoming the generator where people don't really care too much about the quality as long as it does the job. Like background music in a dentist's office, kind of, or images at a hotel, I just need a nice picture of a pretty flower on the wall. Yeah, I think there's a risk that the artists who currently get paid for the outputs like that maybe may be at risk. But every economist I've read on this issue, and I've read many, always emphasizes that this is what one economist one time called creative destruction. You know, nobody makes money manufacturing buggy whips anymore because the automobile came along. But these new technologies now create new jobs. So already law firms are hiring people to formulate queries for Chat GPT, like, do you want to be a master query formulator? That's not a job yet, or it's about to become a job. So I can't predict the future, but I think we're going to see some job losses, some job creation, and I don't know what the net will be. [00:27:00]

ALEXANDER NAUM: Yeah, that's definitely always the issue of progress, you know. But maybe it's a future that we all can work with. Another type of generative AI, colloquially known as deepfake technology, has the ability to superimpose the likeness of another person onto videos without their permission. Are these programs regulated in any way? [00:27:20]

PROF. MICHAEL CARROLL: Well, so in general the law does not regulate technology per se. What the law tends to do is regulate people's uses of technology. So I think deepfake technology is a technology to be concerned about because it creates such realistic outputs that are fake. But only the skilled observer can tell that. And so the nightmare scenario that everyone has in mind is a deepfake president who then orders a nuclear strike. And that's like the nightmare scenario for deepfakes. But we're seeing all kinds of other harmful uses of the technology, to sexualize people without their consent, do things that would harm people's reputation. So I think if there's going to be regulation it will be on these kinds of uses of deepfake technology. And unfortunately, it's going to be hard to figure out how to scope a regulation around those uses. [00:28:27]

EVA PEDERSON: Sometimes, as we alluded to earlier, artificial intelligence is so advanced that it's hard to tell what's human created and what's artificially created. Some argue that the danger of AI comes from the fact that the human user is unaware of the AI's influence. What consequences could hidden AI bias have for human artists? [00:28:46]

PROF. MICHAEL CARROLL: It's funny you should ask. I actually wrote an article about that. I wrote an article for a symposium called The Rules of Engagement. That piece is really about looking at this problem on social media. So if you're a social media content creator, your goal is to go viral. How are you going to go viral? Well, you need the algorithm, the recommendation algorithm, to pump up your material. So as an artist, you're the middleman, if you will, or the middle person, is no longer necessarily a record executive or a movie producer who you need to get through as the gatekeeper. Instead, we're using artificial intelligence or the machine learning as the automated gatekeeper. And what we're seeing is that artists then try to game that system. They try to figure out what the gatekeeper wants in order to get their content to go viral. And as a result, it's feeding back into to the kinds of content that they create to figure out what the

algorithm is going to like. So indirectly, the algorithm is influencing the human creative decisions. [00:30:00]

And we have seen cases where content creators have been the subject of algorithmic bias. And there were a group of Black artists on YouTube who felt that their work was being suppressed by the algorithm and filed a complaint. So because how the algorithm decides what to promote and how strongly to promote it is not shared and is not transparent. We can't detect the bias other than by looking at the patterns of its output, and that's hard to monitor. [00:30:34]

So yeah, I'm very worried about that question and I would like to see greater regulation and more transparency requirements around anyone who's relying on these tools for decision making, whether it's in employment or national security, public safety. There's just too much reliance on these black box models that that have bias baked into them, and we know that. So we really need to have some tools to regulate that and fix it. [00:31:02]

EVA PEDERSON: Speaking of disclosure and transparency, the European Union and then more recently the Biden administration's Office of Science and Technology Policy have released respective bills of rights for AI. Can you talk about this and what it means for the future of AI? [00:31:18]

PROF. MICHAEL CARROLL: Sure. So I think this is an example of where policymakers feel like it's too soon to be setting in stone actual hard law, where there are legal obligations that immediately get created, because everything is moving so quickly and it's easy to regulate yesterday's problem.

So one thing that you can do is to articulate your regulatory principles. What are the values that you want the law to be promoting? And what are the sort of principles of behavior or conduct that you think ought to align with how you promote that? And one of the early areas where this approach was effective is in the area of information privacy. So back in 1973, when the U.S. government was going to digitize a bunch of its records or start relying on computer systems to keep its records, something called the FIPPs, the Fair Information Practice Principles, were articulated, which are at a very high level of generality about how a system of records should be run and what the subjects of the data that is being collected and used, what rights should they have to know that there's data about them in the system to be able to see that data, to be able to correct errors. Those are the kinds of principles that that – and those principles end up informing and largely becoming law in the EU's privacy law, the General Data Protection Regulation. And I see these Bill of Rights for AI as a version of that. [00:33:00]

This is not yet law. But these would be the principles that would inform what law is likely to look like in the future. We're lawyers, so it's like an issue-spotting exercise. We've gone through, here are all the legal issues that we've spotted that are likely to need regulation. And here are the principles that we're going to rely on in thinking about what that regulation should be accomplishing, and even thinking a little bit about how it should accomplish it. [00:33:28]

ALEXANDER NAUM: Yeah, definitely. I mean, you brought up the GDPR, which is a big example. I mean, it's influenced even how companies even operate within the U.S., even though

it is an EU regulation. But hypothetically speaking, if companies completely disregard the AI Bill of Rights guidance and let their AI programs just run completely rampant, what federal agencies would step in, and how do you imagine them regulating these programs? [00:33:52]

PROF. MICHAEL CARROLL: It's a great question. Again, I don't think they're going to regulate the programs per se, but they will regulate the way the creators and the users of these programs. And I think it would be unwise for a company that wants to be in the game for a long term to be completely ignoring this because it is a form of what you might call soft law. It is saying these are the kinds of things we're concerned about. And basically if you want to invite faster regulation, go ahead and violate these principles, and then you'll see the legal system kick into gear more quickly. [00:34:28]

Certainly I would expect the Federal Trade Commission to be involved right now. That's the agency that comes the closest to dealing with consumer privacy. And I think sort of any kind of disclosure requirements could easily be found in the FTC's authority to regulate deceptive and unfair business practices. I would expect the Consumer Financial Protection Bureau maybe to have a role if finance is in any way involved. Potentially the FCC, depending on how the technology is being used. And then I guess, again, if we if we start talking money, then the various Securities and Exchange Commission, maybe the Commodity Futures Trading Commission, the people who are looking at Bitcoin could easily have a role in looking at these technologies. It all would sort of depend on what's being done with them. [00:35:28]

ALEXANDER NAUM: Yeah. I mean, I know these are still the early days, but in this new and ever evolving frontier of AI technology, do you see even the potential future need for Congress to step in to enact new laws or even potentially create new federal agencies in the absence of our current federal agencies that are better equipped to handle issues emerging from AI technology? [00:35:50]

PROF. MICHAEL CARROLL: Well I'm sure we will get legislation at some point, but Congress is generally a follower and not a leader when it comes to technology regulation. And usually you get regulation when there's enough stability in the system that we can feel some confidence that if we put down law today, it won't be obsolete tomorrow. And often there will be stakeholders that have sort of organized themselves to be able to successfully push legislation through. And I don't think we're there yet. I think figuring out who the stakeholders are, do we have a lobbying organization for the AI industry? Maybe, but not quite. And whether we need new agencies, that feels like a reach at the moment. Although again, you could imagine a circumstance arising where it feels like our existing agencies are too siloed. [00:36:48]

But I would see that's pretty far down the road. I think the earliest kinds of regulation will be disclosure requirements, the requirements that you disclose that you used an AI in this decision, that you used an AI to produce some kind of output so that people on the receiving end can sort of evaluate that. Or if it's in the context of a decision, take that decision to court and challenge it. But because disclosure is an easy one on the regulatory side, substantive regulation about what you can and can't do with the technology, that's a lot harder to draw lines and easy to game if you get it wrong. So I don't see those rules emerging anytime soon. [00:37:34]

ALEXANDER NAUM: I did want to give you the opportunity just to say any parting comments that you may want to make for our listeners. [00:37:41]

PROF. MICHAEL CARROLL: Well, I would just say these were all very good questions. And to the future and present lawyers, you know, law is also a lagging industry. We tend not to be leaders in embracing technology, although we know we need to. And I think this is one where, whether you like it or not, these technologies are going to infiltrate and become part of legal practice as well. And so I think especially for younger lawyers, you are well advised to become familiar with how to use these technologies as problem-solving tools in the future practice of law. [00:38:18]

ALEXANDER NAUM: Well, I want to thank you so much, Professor Carroll, for your substantial contributions to today's episode. And for helping us to understand this very new and complex subject. I would also like to thank the American Bar Association's Administrative Law Section, the Administrative Law Review, and of course, our podcast's own Eva Pedersen for her support in creating this episode.

If you're new to our show and enjoyed the episode, give the episode a like and be sure to follow and share our podcast with your colleagues, friends and family. Thank you and you'll hear from as soon as we discuss other topics impacting administrative law. [00:38:51]

END TRANSCRIPT