843

Wed, 6/23 8:32PM 1:02:12

SUMMARY KEYWORDS

work, ai, approach, data, application, people, ethical frameworks, system, practice, transparency, clients, practitioners, thinking, information, speaking, attorneys, technology, questions, point, slide

00:00

Hi, I'm James shear, I'm a partner with Baker hostettler. In New York, I co lead the emerging tech team, within our digital asset and Data Management Group. What that means is we've we've come up with a focus on emerging digital technologies, and certainly their application to client needs. And a lot of the focus there too, on both what the new technologies mean for clients and the work that they do, and of course, our support of our clients in the work that they do, as well as its intersection with the practice of law. So the ways in which these new technologies can help improve what we do as attorneys. And that's where what I'm gonna talk about today is this. I call it ethics AI in the world. So it's a long form presentation here is AI and algorithm development, and the new frontiers of professional responsibility. Long story short, what we're going to talk about is what, what you need to know about these new technologies as a practitioner, I want to give you enough tools to be dangerous, but really that that danger comes from asking the right questions, you do not need to be an expert here, but you no longer can afford my opinion, to, to be afraid of these technologies, to avoid them to pretend like they don't exist, or to push off that responsibility on to someone else. Now, it's certainly available to help us the responsibility here and bring somebody else into the fold. And we expect that and we'll talk about the rules of ethical responsibility later on, you can bring in the experts. But you know, I'd argue here that you need to know enough to know when you need help. And we're going to we're going to try to change some unknown unknowns, and unknown unknowns, to eliminate some blind spots, and to work through that as a practice. And along the way, I want to introduce some technology that you may or may not be familiar with. That's okay. You know, as I work through with my associates, and even with clients on these types of issues, these are new, they're they're certainly they're emerging, they're evolving as they go. And the more you hear it, the more it's going to breed some of that familiarity. And we want to get past this idea of a blackbox system and a computer or an algorithm or an approach for AI. That just seems as though it's it's, it's magic. And it's not Carson, I hope it if anyone tells you that it is then they're they're speaking out of school, but we want to get

you to the point where you're thinking about these things. And you're an active participant in those discussions in that question, asking process, reinvesting as you would

02:47

like a reporter, or you're doing cross examination, although I don't welcome cross examination for a conversational style. So outline, if you've got it on here on the deck, we're going to introduce you to AI and related technologies, we're going to talk about some types of machine learning algorithms. And to swell some of the surprise, a lot of what we're talking about here is machine learning. It's what they've called the space weak AI, as opposed to strong strong AI would be where you really do have replacements to human cognition, you've got these instances where you're replacing people, that's that's really not the case, here, we're talking about applied tools that are able to do certain things that that humans have done in the past, but can do them much more quickly, much more efficiently. And on a consistent basis. Although that consistency is a little bit different than what you might expect just coming out of a computer. We'll talk about an example machine learning algorithm, as well as the data sets we use to train AI. And again, I want to go through some types of machine learning algorithms and AI approaches, again, to invest some of that knowledge into the process so that you're familiar with it, so that when you hear these things in the future, it should ring that bell and say, Okay, I've got some experience with this. And in please do use this course, and this information as a jumping off point, to continue to further refine your skills and your knowledge in this area. And again, that's the the admonition I give to other people when they're assisting in this practice or working on projects with me, or helping with client advice, which is, don't be afraid to go and learn. certainly ask. And if you've got an instance in an opportunity to ask your clients or programmers or systems engineers or other people involved in the process questions. It's a great habit to get into. And part because you're involved in the process. If you're acting, it's much more likely you're an active listener, and you're learning as you go. You can get the knowledge and information from the people that you're asking. But also, if someone doesn't know the answers, usually as practitioners we can pick up on that pretty quickly. So it's a it's helpful to know if the people with whom you're speaking are the right resources. If they're not, of course, to move down down the line, I'll give you one of my tips here for for working on this practice, it's sometimes it's dangerous just to speak in a one on one environment, especially if you're speaking to the so called subject matter experts, they tend to be very, very confident. And sometimes that confidence may mask lack of knowledge. Sometimes it's more appropriate to get in a room with a couple different practitioners at the same time. And here I'm speaking of like the programmers to the engineers, and to ask these questions, because then they'll start to think more carefully about their audience, which would be their colleague or the other person in the room, you'll see some nuance introduced to that discussion. And that that should help you start to frame those questions. And you should come away with all of the answers in a discussion like that you're seeking, what's the next question? What's the next question? We're exhausting? What the particulars are here, so that when we're coming up with answers, we're

framing things as far as facts or evidence or in response to a regulatory inquiry that we've got. We've got a story and an explanation that's predicated on the facts. And that is going to be consistent than with the testimony if we get that, to that point, from the people with whom you're speaking. So that's how we build all of this up. This is foundational building blocks. Now, what's this in supportive? Well, we'll talk about some of the proposed legal and ethical frameworks that we're seeing beginning to emerge. And these go all the way up until the end of 2020. very much an emerging area. This is present day, we expect to with changes in administration that we may see some pivots and some changes. But this this is happening as we speak. There's hundreds, literally hundreds of different proposed AI ethical frameworks. Some of those obviously have a little bit more gravitas than others, but they're all in development. And that's part of the work that I'm doing right now, I'm working with several different groups on the development of their standards, including those that are applied to engineers, those that are applied to practitioners in support of the help that we're going to be giving. So of course, we're trying to help the helpers in some instances, whether that's on the programming side, the front end, and the practitioners on the back, will speak about AI transparency. And it's the first evolutions here in the context of the general data protection regulation, which was the EU standard. That's still in effect today. And I'll show you a chart and explain a little bit how some of these principles are coming from existing regulation, and how a lot of them try to trend toward where we expect regulation and regulatory enforcement to go in the future. Also, we'll talk briefly about use in the courts in the justice system, as well as considerations as we've seen them the case law, and I've got an example in the deck, this dimensionality reduction algorithm may speak to it, depending on how this discussion goes. And I'm speaking with myself, but imagining how I'd explained this and sitting down with an intelligent, informed, but interested in curious practitioner for how these things work. The objectives here, again, if you want to turn to that slide, slide three, defining machine learning and AI. For your understanding, we'll talk about the proposed frameworks, we'll examine its application and law, practice management, again, going back to different silos here that that are as, as you'd expect, going to bleed together, you've got your clients use of AI, and by extension, AI and technology out in the world. And then you as a practitioner,

08:28

the ways in which AI may impact your practice. And, frankly, you have to know enough about this area, to understand its implications and the attractions involved. And then we'll talk about how this implicates the intersections with malpractice with disciplinary complaints, client dissatisfaction, it's an area that we have to know about so that we can service the clients and know what their needs are, while also understanding its application to what we're doing as practicing attorneys. And we'll see how this affects the education of attorneys going forward. Again, another quick tip here. As you know, if you've ever been stuck doing it, compliance work, or even like helpdesk type work, you know, you can often start at the source, which is the search engines are going to be very, very powerful tools here to find out what the current state of play is.

It's, it's great when a lot of these resources are being developed, as I said, a number of different standards across all these practice areas, because you can go to those for like one step or one stop shopping, as well as like that imprimatur of the organization that's compiling. It's a great place to get kind of the overall larger sense of how these technologies work. But if you really want to understand the the minutiae, and the grid, or to ask more pointed questions of a witness or of your client or in support of briefing or response to getting, you got to start digging, and you can dig using one of the big search engines because that's, that's where the professionals are going. A lot of machine learning application processes are now knew. And here's another idea to keep in mind, the math of all these, that that that math is for the very most part out there in the public. So this is math and approaches that have been developed through the government, they've been developed over the course of decades at this point, as far as specific or particular approaches. And just don't get afraid if you see an equation and we're going to throw one up here on the screen to or within the PowerPoint, so that you can again, get form familiar if so, how to read these things, and at least to ask questions about it. But when it comes to the application of these algorithms and approaches, you know, I've heard clients really push back when we've talked about the math and to say, Okay, well, what do we need to worry about here with this black pocket? Black Box algorithm? What do we have to be concerned about? When it comes to expressing what the math is? And clients will push back and say, That's not what's important to me? Certainly, the math is the vehicle that's going to get me there to my answers. But what's the true importance is the underlying data, the information that we're looking at, that we're concerning ourselves with, that we're building models and approaches off of, because that information is really the lifeblood of a lot of these organizations that are playing in this space, and or that they would expect to see if your, or your client are pulling in approaches to try to affect some kind of outcome. So I mentioned the GDPR. Before here, that's, of course, the European Framework looking at data protection there. One of the things that GDPR looks at is this automated processing and kind of defending what outcomes are there. You can't just say, Okay, well, I use this data, I put it into the black box, and I didn't really understand what happened, therefore, I'm not responsible for it, they're trying to push additional responsibility in there. And this starts with an understanding of the information going in the front end. And that goes to data privacy and other security provisions, thinking about how you're monetizing that data, and really what the effect and application are for that data across it. So that is the intro. And it's been quite an intro, I think, but we'll we'll dive into it. And again, speaking to some definitions, their AI, we're talking about this category of scientific and statistical techniques. It's computers, mimicking human behavior. But of course, doing it at a different speed, or sometimes at the same speed. And we'll see that we'll see this evolution a lot to the applications of this in your clients activities. And what you'll see in your, your day to day life, machine learning is a subset an application of the techniques. And it's allowing computers to learn from information. But I mean, learn here is is different, of course than what we do, at least in some ways. It's it's pattern recognition, it's getting to a place where you can then use these learnings extracted from the data or information sets in combination, you're using different techniques, and then approach

it in different decision making capabilities within a system. Finally, we've got deep learning here, subset of machine learning, where we're talking about multi layer neural networks, these go to further complexity of process. And again, a lot of this is introduction of technology for you to think about to say, Oh, well, someone mentioned a neural net. Well, a neural net is just this subset, then of deep learning,

13:12

again, a subset of machine learning, which is a subset of AI. But of course, it's weak AI, because it's not just replicating a person in a way that's indistinguishable from another from a person. Instead, we're looking at these applied techniques in their application against data. So if you turn to slide five, we're going to talk about some types of algorithms. Again, don't be afraid Don't just gloss over this slide. When I say logistic regression, or decision tree, or SVM, or KNN, or some of these gradient boosting algorithms. These are different sets and subsets of math and mathematical approaches, where you're evaluating data. Another idea to keep in mind here too, is that if AI is a science, but it's it's very much a work in progress science, and it's also I've heard practitioners say it's much more an art in some ways, because you're you're taking data, you're applying these different techniques, there's there's not just like, apply a apply, be apply, see, and you've got a recipe and it's always going to work that is nuanced that has got its issues, we'll talk about some of those and the ethical concerns about what your data still has within it. But not only does that have those issues, the techniques work differently. And it can, they can work differently depending on the order in which you use them. And that speaks to this layering approach of some of the different mathematical approaches. So you may have a data set, you may have a goal, you may want to try to categorize certain information in a way. And you say, Okay, well, I want to use this approach. Well, this approach doesn't get you there. Okay, well, I'll switch gears and try this other approach. Okay. Well, this approach didn't give me Well, what if I combine them What if I use linear regression first to come up with a pattern I'm trying to map to a curve and then I'm going to start using decision trees applied towards some of these These first parts, so you can stack these approaches. And you keep tweaking it and working with the math and the approaches, until you start to see results, that they get close to mimicking what you as a person would do. Now, of course, you think about that you say, Okay, well, that's just one thing, I couldn't just fit it in the system, the system knew what I wanted, and then acted upon it here, you know, you're, you're feeding together different approaches, setting them in line. And then you're seeing if it works on your subset of your sample set if it's getting to the result that you would have expected, or you'd like to see. Now, of course, the benefit here slash the danger is that once you come up with an approach, even if it's a layered approach like that, that starts to work, then you can apply to a much bigger data set, you can do this over and over again, it's that iteration, it can scale, suddenly, you've got a lot more efficacy, which means you have to be really, really careful on the front end, to do what you think you're doing. And one of the examples we've used before, you had a neural net, and approach that was being trained on pictures. And by the way, you guys are

whoever's watching this video, you've likely been participants in some of these trainings, if you've ever had to use this CAPTCHA, little boxes, if you're trying to access a webpage, or trying to get onto something that has a lot of traffic. And earlier on, they were Can you identify these these words, and you'd see like snippets from books, and that may have been part of one of the book scanning projects that were the rage, you know, a decade ago. Lately, you may say, Okay, well, can you identify a street sign? Can you identify a crosswalk? Can you identify bicycles or buses or mountains? Well, the idea there is you're part of a project that's trying to teach autonomous vehicles, how to drive, how to recognize all those things, because we're very, very good at that signal detection, certainly were good enough to allow us for the most part to drive. How can we make sure that the cameras and then the software within them are recognizing those same topics to do you can see, we're part of that process, we're tweaking the system, we're helping that. So there was one system where they were feeding a bunch of images in and they were trying to differentiate between a building and a

17:03
tank,

17:04
pretty good distinction to make. And if you see that in your caption, then I wonder what you're visiting. But here they are looking at that they came up with what seemed to be a pretty good recognition, until it turned out that they weren't hitting the right thing. And luckily, this is all done in a system as opposed to like actual field operations. But you had that practitioners here, they've been very confident that the system could recognize pictures of tanks. But it turned out that pictures of tanks were always taken with a clear sky behind and the algorithm was recognizing that as a salient element of the process and identifying the system as such, that wouldn't be so great if you've got a building, and it's a clear sky behind it. And suddenly, you no longer have a building that self resolving probably a problem probably, but not at all what the system was intended to do. If you turn to slide six, and if you've been flipping through the deck at this point, you probably stopped on this one, it is a visually arresting slide. This is these are some of the approaches or types of algorithms in visual depiction, these came from the noun project, which I highly recommend if you've ever looking for specific types of helpful icons. But you can see machine learning just trying to find our path, we're trying to find a way to get from point A to point B, where we'd expect it to Simple as that. And it's as complex as that.

18:18
And then you've got decision here, tree here. And and frankly, for the decision tree approach here, it's not much more complicated than you think. If you've spent time with flow diagrams, then you

know what a decision tree is, you've got a flow chart, you've got these nodes, you think of like the beginning circles, they're each representing a test on an attribute, you know, whether a coin flip comes up heads or tails, and then a branch there is the outcome. So the first one on the top, the top left there under the decision tree, okay, heads or tails, you go through the next iteration, same thing heads or tails, and you start to work through there. And the outcome is the decision after you've worked through all those attributes. Again, not complicated here, but complicated at scale. As I heard long ago, the difference between quantitatively complex and qualitatively complex, not necessarily qualitatively complex that is, we can understand it, but quantitatively complex because you can build this off of a myriad of different decisions here. Now, random forest, I like this one here. Again, don't overcomplicate this random forest comes from the use of a bunch of decision trees. And there's actually other ways in which you've got stubs or stumps thinking of these types of decision trees that are truncated or consolidated down to something. So a random forest may take a bunch of different decision trees, and then start to work through merging them to get a more accurate and stable prediction. And as I said, before, you think of these approaches, you're layering them on top of each other, you're trying to get to going back to the machine learning icon there, you're trying to find your way through so you can even think of this as Okay, well I try this approach and I get to this point, okay, that's not quite where I want to be. By layer another technique on top here, what will it be another one another one can see linear regression here. And linear regression should be for anybody who remembers their math classes, or who went back and did statistics or finance or any of those disciplines, you know, you're just trying to find a, you're trying to plot a line within your data to see if there's a trend to see if there's a pattern that emerges, this is the most basic of that. Now, support vector machine is pretty similar to linear regression, except you're looking for a plane, you've got a three dimensional attribute of all this information, you're trying to fit something within that pattern, and you're welcome for this great graphical representation, use my hands. But imagine that that cube of information that's that's that's scattered around and you've got the point, you're trying to find something that matches up there, that would explain the whole is, and if you've got that explanation, or that that power, then you might be able to apply it to new data sets, or similar data sets where you're trying to create the same type of outcome, or the same type of understanding. Finally, and I'll take a breath after this. You've got k means here, vector clustering. So instead of looking for a line, or a plane that fits instead, you say, Okay, well, my data is made up of all these disparate points, instead of finding one explanation, I want to find a variety of them. So maybe this grouping goes to this point. And this grouping over here, it goes to this one in this group and goes to that, I can start to distinguish characteristics. So instead of like, let's say, you've got all these different pictures that were assembled, who knows on my phone, you say I want to categorize them according to something like do linear regression. Or I can support vector to say, okay, all these pictures look like they're mammals, came into mind, say, Okay, well, I'm gonna start now to differentiate between what a cat is, and I've got these characteristics of a cat. The point here is in the long tail in this way, and the fact that they all look very angry. And then I can characterize dogs, then I can characterize my kids. And, of course, you see this when you're

using if you use an iPhone, or other phones that suggest images and say, Hey, we think this picture is associated with this person. And of course, that's the phone going through pulling different characteristics off the facial points of data, putting that into a model, extrapolating out if you have kids like I do as they grow older, and they still recognize it, and then pulling together that as a model, inferring certain things from the information available, and then creating that. So going on to the ML algorithm dimensionality reduction. So this is good, don't be afraid of the equation. But I want

22:30

to put an equation out there. So you understand that this is how these things work. And I think as practitioners, certain of us, even if we're outside of IP, are going to have to start getting into the weeds more and more to understand, because you can query a system and you can audit on the back end. But you also have to look at what its reasoning is on the front end. You know, you'd normally do that by asking a person, what did you mean to do? What were the actions you took? How did you walk all this through? Now we're going to have to query in some ways, what it is that the approach was, we'd have to look at something like this. And maybe you'll get to the point where you have such facility that you can recognize where things are missing. And that that's a that's a great position of, certainly of interest, but of intellectual power to say, this is what I'm bringing to the table and helping understand this. They can remember being familiar enough to ask the questions. And the first question may be well, what are we trying to do here and within a dimensionality reduction. And, you know, I mentioned before on that slide, where you're trying to pull in different groupings. And if you think going back to slide six, like the K means approach where you're coming through, and you're plucking out different things, that something could be, you're creating all these dimensions, that you add a new information that can characterize it, it might not just be a cat, it could be a red cat, it could be a red cat with blue eyes, your red cat with blue eyes was thoughts. Well, that's all well and good. But you may not need all of that. And you may be in particular instances thinking as a practitioner, where you may not want all of that, if you've got an instance, where a client is using information to target home loans, and they've got all these sets of information. And they start using this approach where it's clustering data and making identifiers, if some of those identifiers of our of protected class information, and the system then wants to make decisions based on protected classes and their application, well, that may not be allowed. And if that's the case, then you have to start to dig a little bit deeper and say, Okay, can we reduce down the way in which this algorithm works? Can we remove certain dimensions. So there's math associated with that, but you can see where the guidance is going to be helpful. practitioners are going to have to work hand in hand, certainly with the setup of these things, but then in the protections on the back end to find out if they're working as appropriate, if they're working as designed. If there's anything else we need to tweak on that process. And some of the approaches here in the AI problem domains I've got here in this axelos space diagram on slide eight, again, familiarity with the ideas and the concepts and working through how all those

would look Read through it, get an idea. If you're more curious, start digging. But the the important part here, of course, is when you're introduced to new technology and supported the practice of your, your legal career, and or if you've got a client bringing something to the table, asking these questions to say, Okay, well, is there a component here that we have to think about? Is there something else that we have to be, if not worried about at least aware? That's a it's a good start to that discussion. And some of these are going to be fairly apparent if you're dealing with autonomous vehicles or autonomous systems. If you're dealing with an Alexa or Siri or Google Assistant. Or if you're dealing with a self driving car, you can see that you can you can make those connections right away and say, Okay, well, what's the implication here? How does this matter what you're doing? What do we have to be aware of and concerned with? And often, again, the secondary question there is, what's the underlying data? And how is that going to impact our analysis here? But the secondary question is, if it's a practice that might not necessarily call that out, but you have a sense that it's dealing with large volumes of information, and a lot of data coming in, then there is likely going to be an autonomous component to that application, and it's and the thinking you're gonna have to engage in if you're, you're collecting IP addresses through a web interface. Okay, well, even if you're using something called like a sim or a log aggregator, it's likely that that information is going to be addressed in an autonomous way first, before getting to the point where an alert goes to a person. Likewise, if you're collecting resumes online, or marketing data, the the deluge of information that's available and pushed through to people, it's just so incredible, you can't do that by person, you can't do that on an individual basis any longer. You just we just don't work that way. And apart, we've seen this, and in the advent of E discovery, using technology assisted review, there are certain tasks that are just going to be better placed in the hands of the system, until the point where human involvement is going to be warranted, and it's going to be evaluated. So a lot of this is going to be that combination, I just want you to think if you know you've got a ton of information coming in, how likely is it that some kind of autonomous system is involved in that practice? And then, once you determine that, or at least ask that question, well, what is the system? Who knows? If you're not speaking with someone that knows?

27:24

Who else might be involved in this discussion? Who else can you bring to the table, just to start to refine that point, you don't need to be perfect. And a lot of this that perfection isn't available. It is a fuzzy science, but you want to be knowledgeable. And certainly if you're asked by a Fact Finder, if you're asked by a judge, or by a regulator, you want to start to develop the story of what's happening, the explanation, a lot of this, and you'll see this in the the ethical frameworks, it goes to this transparency, explique ability explainability ability to relate or to create those analogies to explain what's happening. And that's that's where we see our value. As attorneys coming into play. It's well we've done the investigation, we've asked the questions we've audited and analyze the results, we've seen what the impacts are, this is how we can frame and define the behavior

and come up with explanations and defenses. In the very least a lot of it's not defense at this point. It's just responses to queries so that we can walk a regulator or someone through that process and explain what's happening. So I, I just spoke about the complexity of this system. And you can think about some of the data sets that are built into these processes. So for example, I, you've got financial institutions that have been playing this space a long time, they've had rich sets of data, they've been trying to make decisions, based on that data getting away from necessarily human factors. To whom do I give a loan? How long do I give it for? What's my interest rate? Yeah, what's the coupon on this? What do I need to set this at? How can I? How can I lessen my margins? How can I improve service? How can I improve the product? How can I come up with things that are new and creative in service to my clients? Okay, gracious view, certainly. But thinking about that. So you've got all of these different types of data sets on slide nine, thinking about transaction data, social network analysis, clickstream data, what people are doing online, a lot of those things to speak to emerging regulatory regimes like the California consumer Privacy Act, or ccpa. And it's assumed to be successor of cpra. And you've got some other US privacy regimes that are coming into effect right now to all of this data collection that makes it easy to or easier to aggregate and you're kind of defined process within the data can also go to start identifying people and that's the flip side here, more data we have, the more individualized the recommendations, the better the recommendations are going to be. But sometimes the creepier they are to the closer they get to understanding who we are and what we do if you set a high level or really digging in and suggesting things we hadn't even thought of, you know, you wonder how much that is, and and what the end results are going to be. They're about halfway through, according to the timing, introduced a lot about the algorithms and their approach, hopefully given you a little bit of background. So you feel more confident you can say, at the end of the day, you should you should know this right now. I, okay, we've got algorithms, we know they're applied to data, the data really matters. And that's what clients like to protect. We've got to be concerned if the data is related to personally identifiable information because that implicates people, or it doesn't have trade secret protections, or is it licensed? Or do we own the outputs? Like how do we think about all those things within the information as practitioners, that begins to get us to the application of the frameworks, we also want to introduce some of the concepts on the back end like okay, well, even if the process itself was fine, the way in which it was done a defensible process, you did the math the right way. You, you set your your schema, you've got knowledgeable people that can answer questions, things were documented, and we've got a clear chain of custody, we've got licenses, all of those things at the front end. And what are you trying to do? Is what you're trying to do aloud? Is it truly what you were trying to do? Is the system doing what it was intended to do? Are you double checking that? Are you making sure you've distinguish between buildings and tanks? especially depending on the application? And what is it that the the August thinking bodies out there, especially some of them control the practice that we do? What have they been thinking about? Well, I'll take you back to the SLR AI principles and a little bit of history on this one, it's the first one I introduce on slide 10. But this, this harkens back to concerns about biomedical research, and I believe was like, originally uses of RNA, of course,

we have looked forward into the future near in seeing CRISPR and its application for editing genome sequence. But back in the day, they were looking at this and saying, Okay, well, we don't want scientists floating around the human genome, doing different things like that. They're specifically oriented toward

32:12
human experimentation, especially irreversible, especially the type of experimentation that we just don't know where the outcome is going to be. Definitely parallels here with AI. Clear. And that's that's where they were thinking of this callback. Now, there's original ethical or soft law framework approaches, we're trying to figure out ways that how do we address an area that's so complicated, that they expectation that our legislatures or legislator, our legislatures, are our bodies of law setting people, they're not going to get this you don't have trained necessarily trained AI scientists sitting there. So we understand that then not only is it tough to understand this at a point where you can really give good feedback that's not going to have that Cobra effect, or the unintended consequences where you grab it by a table and tail and it bites yet,

33:06
how do we

33:08
navigate that, so that trying to put this in the hands of people understand what they're doing, but also appreciate that this changes so rapidly, it's not just the Moore's Law of transistors getting half as big or twice as powerful, every 18 months. This is, things change at breakneck speed, you've got these approaches where you're trying to put things into practice, and it is a fuzzy understanding. So we don't have a perfect understanding of this. And it's also the the mindset of Okay, well ask forgiveness and move quick and break things. Because these things scale, you want to be the first one there as a client, as a, as a company, as a technology. And you can push it out at scale, given all of the platforms that we have. And this is publishable things, things can become viral, things can explode overnight. And suddenly you have a new paradigm that you have to be concerned with. Let's, you know, pretend 20 years ago who thought that you could have a self publishing mechanism that could reach almost every corner of the world. But that's where we're at. And of course, we've only built things to scale more over the interim. So starting with the similar, they'd looked at soft law kind of curving behavior. And they had tied it, the original one to publication like nature and some of the bigger journals, they'd said to scientists, we're really playing in the space. If you don't abide by this, we're gonna freeze you out. And I think some of the present day ones too. And if you looked at the future life Institute, I mean, really a lot of luminaries kurtzweil was there if you know him from his Google work or even his work within

acoustics, Elon Musk was another one. It's some of the people that you look at playing in this space. They were there for these discussions and the hope was get some of these rules and concerns in front of the engineers and the people designing the systems at the front end. So that a they understand the concerns and be they're invested in the process. They think of these as their ethical constraints, almost like the admittedly too simplistic or the laws of robotics that that came once upon a time within science fiction. So Asimov looking at this and saying, well, we expect that these technologies are coming, what can we do to kind of, to at least couch them terms that still work with us as people, because that's who we are, in our service, these projects, or these, these processes are supposed to be working. So similar here, it looks at safety, failure, transparency, judicial transparency. So again, fail your transparency, think of if the system doesn't work, how can you go back through and audit or at least see what happened? judicial transparency? Well, if you've got people who are in charge, if you've got rules about these things, how can they look at these and say, Okay, well, what happened here? If it had an ill effect was that Ill intended? Was it negligent? Is it something that we can fix if the genies out of the bottle, you know, what, what additional protections or recourse Do we have. And then, of course, a focus on human values, personal privacy, liberty, and privacy. So privacy here is, it's been a fundamental concern of AI application all the way through. And, and that's something we're going to see revisited over and over again. So that offers another ethical point here, to the extent that these issues are raised in your practice. I've heard it said that the best systems come from this behavioural economists point of view, which is,

36:40

if you really want to know what people are going to do, it's not the classical economists view, it's not asking them, it's looking at what they did, their past behaviors are really, really, really good predictors of what they're going to do. Well, that's fantastic. And if you want to sell a lot of product, you probably want to know what people want to do, or people are going to do. However, if you start to dig in too deeply, as to what people are going to do, you're going to find out what people have done. And that that starts to get a little bit trickier when it comes to the privacy concerns and implications. The US is catching up, but around the world got a lot of direct views toward privacy. And then of course, you've got exceptions, depending on the application of the algorithm, for example, in in Europe, and certain GDPR and other countries, you've got different protections afforded to medical care and clinical data, you might think, Oh, it's always gonna be more protected. Even in the United States, we recognize HIPAA, protections for data and limited data sets and kind of this cascade through of how data is protected. Again, mentioning a lot of terms, you don't need to know all of this, I want to gear you into thinking about it so that when you hear these points again, you're like, Oh, yeah, that rings a bell. What more do I need to know about that? Oh, if I've heard the data collection of the points, well, if I'm thinking about the data, then I've got to think about its use. Going back to clinical data, you've got additional protections on the front end, we really want to be careful about how we're collecting clinical data and how

we're using it. But on the back end, it's so important to human The, the promotion and the advancement of our medical practices and approaches are so important that unlike other data sources, and sets where someone might ask to be removed from a data set, as you can do through GDPR. And through the application of the ccpa, and California, for California consumers. In that case, it's so important, you can't take the data out. So there are some bells that are unreasonable even under statutory regimes. So thinking carefully about that the movements of data, the use of the data, wherever you can take it out, whether you're able to take it out. And remember, this is not a denigration, especially for those of you who are engineers, and also practitioners, a lot of the engineers you're going to encounter are very, very concerned with, can we do it? Are we able to do this? Is it possible? Can I stick this other thing in? Or Wouldn't it be cool to do this within my device? Why not collect all of this information? I mean, I've already got the devices, I've got all these different instrumentation, opportunities, why not do that? You know, our ability and our training to come in and say, Okay, well, may you do that? Is it allowed? And if so, how is it allowed? And what constraints and curves Do you have to put around it? That's part of our discussion and considerations here. And again, moving on to the tenants of partnership, and I still on 10. We're looking here at where, again, the scientists and the engineers are taking feedback into the process and designing that we want this to be designed for us not as a matter of, Oh, well, what happens if we do this thing? It's should we and can we do it? Now we've looked at different applications and how in turn to slide 11. Here, the House of Lords Remember to seek the Lord climate about this, where they're they're trying to look at in developing these things. Okay, well, it's not just in service to people, but also we understand it's going to be part of our economy going forward. Hopefully we'll take the burden off Certain people for certain types of tasks, but we know since it's going to scale and it's going to start working autonomously, we're gonna have to think carefully about its application, how it's going to work. Likewise, the Montreal declaration, we start to move around the world, as you can see different concerns here. We've got the Singapore personal data protection commission, looking at the model artificial intelligence. Now, this is an area in which if you're practicing in areas that touch on the use of artificial intelligence, you see these big data sets, IP protections, data security protections, data privacy protections, in particular, we're starting to see how these things are coming together. So for Singapore, for example, their their data privacy office, their commissioners, their administrators,

40:45
they look to

40:46
some of the rules within the data privacy regimes where they've got enforcement powers, but then they're looking at how they can actually enforce those, they start to go to these other tenants. So if you're ever in a position where you're trying to figure out what's happening, you

should think of some of these standards and approaches is being the middle ground right now, for application of other laws that touch upon them. So GDPR looks at this autonomous processing, you know, how do we look to that? Well, the EU has a number of standards that it's trying to promulgate, it's working toward that. But there's all these interim measures. And from our vantage and our ethical considerations about how to practice in this space, you're looking for explanatory power, you're looking for essentially, some way in which you can replicate a kind of audit process on the back end. So if you're familiar with audit practices, and how those work, you're looking at different controls associated with audits, we're not may say, here's my framework. Here's what I'm trying to do. Here are the things against which I'm measuring is are my metrics, and this is the information I want to collect. Well, if you're doing legal work here in this space, a lot of times you're walking a client through in saying, Okay, well, I want to know the story, I'm going to be the reporter here, and I want to find out more information, I want to speak to people who are knowledgeable, I want to collect my 36 or other subject matter expert witnesses on the front end, so I know what they're going to say. But also, if we're looking at whether this isn't a reasonable practice, well, what am I going to measure that against and some of the soft law standards is or ethical standards that are promulgated in law yet are as good as we get right now. And and having a standard, many of them overlap within the van or eular diagram that shows all the overlaps. There's a lot of stuff that's similar. A lot of these frameworks, especially finding one that's specific to your your industry, or your clients, geography or their area of practice, or subject matter. That's how you're going to walk through and get to where you probably need to be for tomorrow, until the law starts catching up at a higher level. the I triple E, global initiative here on autonomous Intelligent Systems, this is still ongoing and full disclosure, I'm part of the I triple E and some of the working group work that they're doing. But of course, here we're speaking to the engineers and their application of these technologies and what they what they can, should and may do, Australia jumping down under to their framework. And again, you can see this generation of net benefits. We want this to be in service of humanity really doing no harm. You can harken back again to those azimoff rules of robot behavior, regulatory and legal compliance, certainly privacy, protection, fairness, transparency, and explainability. accountability. Transparency is a tough one. And I'm going to return to it after we go through some of these these frameworks. It is an area of particular interest of mine, during which I've published and presented a number of times, just because I think that that's where we as practitioners are going to get more fully involved. We're going to help, I would hope in the design of certain systems where you've got clients concerned about the the long term effects, but really, it's going to be that explique ability of the system and guided transparency of the process, both of the math, the data underneath it and the effects that we have, as well as the thinking behind it are really the direction that we have. That's where attorneys in the same way that you're presenting facts in the way in which it's appropriate to the court, you're presenting your legal arguments, you're crafting that story within the bounds of your profession. That's a lot of what I expect attorneys are going to be doing within the practice of AI and this this kind of autonomous systems and advanced algorithms as well. We're going to be telling the story within the bounds that set forth, we're going to be tying the

story against different regulatory frameworks. And certainly the soft law or ethical frameworks to and we have to be knowledgeable enough to make that story makes sense. We have to connect those dots just as we saw within the original machine learning diagram on the preceding slides. jumping to the OECD recommend the nations on the Council of AI and the G 20. ministerial statements, again, trying to apply these worldwide to effect kind of a consistent approach that's that's where we are going.

44:58

We expect that these things We're just going to start to tether together in certain ways. And, and in the same way that these systems can scale, we'd expect that as the framework starts to knit together, you get different groups joining together to say, this is how we're going to approach this, this makes sense, then the regulators are probably gonna start to meet in the middle. You can't just unring the bell, you know, I imagine the design of the internet and certain approaches to it would have been different if we would have let regulators design them. But of course, the internet happened and then regulations started to catch up based on what we see. And that's still evolving very much the same expectation here in the area of AI. And we've got different principles of responsible AI and ABA resolution 112. So hearkening back to August 2019, a very different time for everyone. But, of course, looking at the usage of AI, this is a very, very short document, I recommend you take a quick look at it easy to find, really, it's these duties, duties of attorneys to be competent, to communicate, to be confidential, certainly, and speaking to IP, and to supervise. So this goes to the use of E discovery as well, it goes to the use of any other of these new technologies that are put into place that obviously deal with the practice of law, both in support of clients, but also to what it is that we're doing in early AI iterations here are going to be seen in helping review contracts or documentations. And or approaches of those, those kinds. And then also within these areas where we as practitioners see a ton of information. So again, bigger contract review projects, bigger ediscovery projects, and litigation or regulatory response, doing due diligence across a deal. Those are areas where there's this deluge of information, we want to perform certain types of analysis across the information using these technologies, the next step, but you can't just rely on somebody else's representation, you have to know enough that if the work is being performed at your direction, or under your license, that you know what's going on, think about what explanations you'd have to offer up if something didn't go quite according to plan. Well, if it did go quite according to plan, what was the plan? What was it predicated on? What did you understand, we just started to put that into place. And continuing on here, you've got the European Parliament coming up with its own ethical framework for AI in beginning of September this year 2020. And again, looking at transparency, focus on bias, unfairness and discrimination. We'll see a lot of that in the US where we're trying to look at how protected classes work. Since we don't really have the the regulations associated specifically with AI yet we're trying to get to the end result where you've got the FTC or the FCC or different state Attorney General's who are t state attorneys general, were coming in and saying, These are the laws on the

books, you're using that technology, it's not meeting the standards. And we want to find out more about what you're doing. So existing enforcement practices, how will you be able to explain how will you be able to turn over information? How are you able to distinguish what you should and what you don't need to turn over? And how does it all work. And we've got the OMB here, where they had put forward some final guidance on the regulation of AI. And again, you could It's not that I missed numbers within the slides. By the way, I'm, I pulled out certain principles, but I kept them numbered as the principles presented in materials. So if you want to go back and find any of these, they're fairly easy to find. Among the, as I said before, literally hundreds of different ethical frameworks associated with artificial intelligence that are available right now. So you can check out the OMB guidance, take a look through here you can see the scientific integrity and information quality. That speaks to yet another point, again, trying to assemble all these different points to think about in this practice, Ai, data quality is paramount to and, and the old axiom garbage in garbage out, certainly works within an artificial intelligence approach. Again, we can stack different algorithmic approaches on top of something to try to mitigate bad data, but you're never going to do better than good data. And when I said that was the lifeblood of certain client applications, I meant it that's where they see their value. Some of them could care less if you know what algorithmic approach it, if you've got this key value, if you've you've done this SVM, or you're just using these types of plain analyses, fine. Everybody does that. Where the distinction and difference is the the type of information you hold the use, you've got that information, how good is it? If it's not, you know, you don't have people who are hundreds of years old within the system, or were born in you know, 1789 or everybody lives in Algeria like your you because they defaulted to one system. The the type of information you have underlying there speaks to the integrity and quality of the information. And that may be a sticking point, if you've got a regulator or even effect finder later on, asking, Well, how was this set up? who looked at the sources of information, how were they evaluated? How did you know it was a good data set, especially if you're impacting individual people? So flexibility here in approach, fairness and discrimination, again, thinking to protected classes, that's where we think a lot of the applications for enforcement are going to be on the back end. And we've seen that within financial services already, they're well aware of this area, it's just the newer poles where a client may be using again, as I mentioned this before, an automated system to pull in resumes, you may not be thinking about the application to these things, but the system, if it's building its own model, according to your prior practices, it may interpret those as preferences, and you may start seeing things appear in the system, or you may never see the results you would have seen otherwise, that can represent a problem, especially your auditing on the back end, and you're seeing some serious implications of protected classes. Likewise, disclosure and transparency, transparency, key one lines cross all of these, although I was cherry picking to make sure I pulled those principles out, as well as safety and security. Now, going back to the European Parliament's ethical framework for AI, they did provide a table within that, and I've got it on page 17, you can see here where they've got main legal issues. And they will actually talk about where there's EU legislation that's applicable and provided. So they're, they're going to go through and show you where the AI

applies and how that works, how that works. So if you're doing cross border practice, or your European practitioner, you've got some of that as part of your portfolio, you can think about that. But I'd argue a lot of us work for multinationals. They have concerns about this to not only the data transfers and their implication, but then the use of that information. So it's a two pronged issue. You know, it's the data quality and the the location from which the data is coming. Are you affording the right protections?

51:38

If you're using an easy if you're using that data in the systems, then are you using it the right way? Are you using according to regulation? And how are you going to build a defend that on the back end. So then, I've got a nice little slide here. Remember, we still do have to have heart within all of this, if you want a look at slide 18. Slide 19, this is the application of AI framework considerations. This was a, a piece done in 2019, where they were comparing AI, really a number of different principles, a very small subset, but some of the more important ones, I'd argue that have been really, I they have helped industry direct and focus. And they pulled into really five categories of principles here. beneficence, this well being it's supposed to be for the good of humanity, and not necessarily just profit driven. So there, there are concerns there that we've just got to think about longer term, non maleficence. You're not hurting people, but you're also respecting privacy and security, again, privacy, and the flipside of transparency. It is an interesting coin here, when you think about these systems, then autonomy, obviously you don't want these things to run amok and just do what they're doing on their own. What is the oversight? What's the off button? How can you work through these things? And if you do have a process, how can you unwind it, then we've got justice and explique ability. So explique, ability transparency, I'll mention this at the end a little bit more, again, one of my my favorite pieces here, but the thinking about like what the explanation is going to be and how deep in the weeds you go. Okay. And then. So that transparency, then is for at least eight of the other 12 regulations on slide 20, including a couple that are directly relevant here in the United States under the equal credit Opportunity Act and the Fair Credit Reporting Act. So again, using existing frameworks to apply against these autonomous systems and expecting consistent results with them. You may be going for the same result or a same defensible result. But you're using a different technology, you have to think carefully about what that means. And I mentioned that GDPR, and you saw that system, we've got a little bit more explanation for that on slide 21. And then, of course, we've seen AI starting to be used in the US court systems. If you look at slide 22. I I'd mentioned before the ABA resolution very much encourage you to take a quick look at that. It's very, very, very short. But we've seen the use of these technologies and sentencing state v. Loomis was a very, very important one, where they were using this compass tool to figure out really what the report the parole risk is. So assuming factors, putting them in building that system off of the rates of recidivism within the prisoner population or those those contexts. That's tough. That's tough to look at Liberty, apply essentially a back a blackbox solution here to say okay, what is the result? How well can we dig in

there to see what the information should show us? In a case of liberty like literal, you know, life or confinement. You might expect greater transparency, you want to know exactly how this thing works. Even if understand Our circumstances, it's just been kind of done by gut feel according to how you think everything should work. But in other areas, how much does that matter? And we've done some work on transparency and to say, okay, there really is kind of this spectrum approach. And depending on the results of it, and the type of information used, and whether we want to keep people from gaming the proverbial system, then maybe we don't explain everything. So it really can depend on it. And that takes us right to this in entrata. Case, and the District of Utah, where they're looking at the production of tar metrics and statistics. So that's technology assisted review, where one side, of course, is asking for information from the other in the context of the discovery, the other says, Okay, well, I'm going to search through the information here, but I'm not going to have attorneys eyes, rest on each and every document, we're talking about, you know, hundreds of 1000s, or millions of pages or documents itself, we're going to use a computer system that's going to help us determine what's relevant, what's responsive, what's privileged, etc, and then turn over the results of that process. Well, some of that's going to body attorney, Attorney thought and application, some of its going to kind of divide into a proprietary approach and different decisions that are made. Ultimately, those productions were made under the license, and they're subject to professional responsibility requirements. And some courts have said, Okay, well, the parties can agree to exchange that stuff. And we'd expect that's the case, others here, you're not going to compel the production of the metrics and statistics,

56:27

you have to show that something was wrong in an instance like that to start to, on on unbox the box or approach that was used. So thinking carefully about that as it goes. And then we had a couple other considerations in case law, looking at via call intelligence and safety, and just where we see AI coming in as part of the case term construction, we've got some, so admittedly, pro se matters, that that we're seeing it starting to applicate are starting to apply here. And then the last one, I thought was very interesting, this force verse, Facebook to look at messages, whether or not you're using different technologies. I'm about allegations that you've got systems that are set up to try to bring you content that's interesting to you, you know, if that's a movie recommendation, or a song recommendation, or you like this picture of puppies, as my phone would indicate, then you might like these other pictures of puppies? Well, that's all well and good, unless you're trying to disseminate information that's deleterious or you're trying to propagate something that's going to be harmful, whether it's outside the bounds of free speech, if you're trying to, in this case, perhaps recruit people into these militant organizations or, you know, evolve and try to help help support terrorist activities. There are bridges too far. And I think that's where the practice is starting to define what that looks like. and integrate it into the practice as a whole, like what our responsibilities as attorneys are, and how those work. So again, to continue to beat that drum, you've got a number of different responsibilities that are associated with this as, as this

practice evolves, they speak to asking the reporter type questions, which are entirely warranted right now, this area is very much still in development, there aren't a lot of hard and fast answers. There's not a lot of ossification, within the practice, you really are looking for things that are happening, for the very first time decisions are being made as we go. And we'd expect that they're going to continue to so to the extent you can get involved in that practice on the front end, and you can help issue spot, it's going to be a wonderful, wonderful opportunity, slash responsibility that it's going to be difficult to get out of, likewise, and thinking about how you tell the story. Remember using some of these ethical frameworks as an application to essentially create, like an audit framework controls against which you're going to test the practices, see if you have answers. If you have answers, do you have sufficient information to support those? And then how would you compile that information, keeping in mind intellectual property protections, as well as proprietary interest as well as Attorney advice, as well as other types of privacy? and general protected applications here for security practices? How do you how do you coordinate all of that off in assembling a response that you'd have to a regulator or to an opposing party, or even presenting something to a Fact Finder? And then as far as that, that that XSplit explanatory power and the manner in which you do it? I mean, I think it is likability and transparency that hopefully is permeated this presentation. You know, we're looking here and saying it's not just give them everything. That's that's not what's going to be warranted and available in the practice. It is not in the practice of law. at trial, you have to look at things that are responsive, that are relevant that that speak to the matter at hand, but likewise on what to expect and we're starting to see develop here is kind of a stepped approach. To transparency, where you're saying, Okay, well, it depends on the audience, it depends on what's warranted, it depends on what's asked for, we want to give explanations that are specific to the questions that are presented, what's the topic at hand, we want to speak to the audience that's available. So just as this presentation, believe it or not, is supposed to be a high level presentation, an introduction to some of these concepts and digging a little bit on these stacked approaches. You know, we're looking here and saying, okay, to whom are we speaking, a lot of the transparency available within some of these different frameworks is looking at the the person who's using these technologies. So if you're being impacted by an AI approach, and you are submitting your resume, and it's important for you to get that job, then you may have an explanatory power that's sufficient to tell you regardless of what job you're applying for, what the system does,

1:00:50

but it's not going to be such a transparent approach that it would tell you exactly how the system works to say, oh, by the way, if you really want to get this job, then you should reformat your resume in this way. And you should add these keywords. And you should tweak this part of it, because that's going to game the system that will work its way through, and you're guaranteed to get to the next point. You can see it's this balance here. And well, that's a pretty simple example, we're going to see that kind of permeate throughout the practice. And I expect that's where a lot

of the attorney work in the future is going to be coming from and in the same way that we think about these things in their application of client practice. Then, of course, we're thinking about, however, we're using them in support of our practice of law, because of course, in the same way that a paralegal is working at attorney direction, and investigator baby working at attorney direction. Likewise, work processing likewise, docketing. We have to know how those things work to make sure that they work. We have to do the same here. And that's what the ABA speaks to. So that is my time. And I hope you've enjoyed this presentation. There's a lot here. And if you've got other questions, my contact information is within the presentation. I'm not the only one working in this space. Of course, there's a lot and we expect that people to some degree or another are going to continue to develop these areas of interest and experience, if not expertise in the coming years.