



Fox Rothschild Podcast

The Presumption of Innocence

Episode 42: AI in Criminal Justice: Opportunity or Opportunity for Misuse?

Featuring Matt Adams of Fox Rothschild and author Pramod Kunju

Adams: Hi, everyone, and welcome to "The Presumption of Innocence," a podcast brought to you by the White-Collar Criminal Defense & Regulatory Compliance practice at Fox Rothschild. Once again, I'm your host, Matt Adams.

AI is sort of the topic du jour in today's contemporary technology dialogues out there in the world. And to set the issue where it belongs, I have the good fortune, I guess, of having lived through the last technological renaissance in the 90s with the advent of the internet and the personal computer and email and everything that it brought to our lives. In the early 2000s we saw smartphones come on the scene. And really changed the way that data portability changed all of our lives. And now we, we are faced with the advent of artificial intelligence. And we've heard a lot of talk of artificial intelligence, but today we're really going to cut right to the heart of how it intersects with the criminal justice system.

And my guest today is Pramod Kunju. Pramod earned his nickname "Data Guru" for his contributions in the field of data analytics and AI. He pursued his passion for criminal justice by obtaining a master's in criminology from UC Irvine. He's a highly respected expert in the intersectionality between criminal justice and AI and a frequent speaker on those issues. He's also the author of "Artificial Intelligence (AI) in Criminal Justice: A Primer on Implications, Ethics, Policy."

Pramod, thank you so much for joining us on "The Presumption of Innocence." Welcome to the program.

Kunju: Thank you. Matt. Thank you for having me. And thank you for that great introduction. I also want to give you a shoutout for hosting this podcast. It's a very critical topic, something that's very close to my heart and something that's very important in our state, in our nation, so thank you, Matt, for hosting this.

Adams: So let's start broadly, from a 30,000-foot level. In your view, what is artificial intelligence? From your perspective as a data scientist who's now interspersing that expertise into the criminal justice arena, what is artificial intelligence precisely?

Kunju: Yeah, that's a great question. You ask five people what AI is and you're likely to get five different responses. And all of them are probably true.

Copyright © 2024 Fox Rothschild LLP. All Rights Reserved.

All content of this podcast is the property and copyright of Fox Rothschild LLP and may not be reproduced in any format without prior express permission. Contact marketing@foxrothschild.com for more information or to seek permission to reproduce content.



So, in my mind, having come from a data science and data analytics background: For me, AI is all about making decisions related to the future by looking at data and information that we have collected in the past. So that in a nutshell is what AI is all about. And then you get into nitty gritty, such as machine learning, where you're teaching a computer how to understand things, how to absorb the information and then how to respond to questions. For example, ChatGPT, that's becoming widely popular. So what it essentially does is fill in the gaps. So you ask a question and it's really filling in what it thinks comes next. So that's its way of responding to a question. And that's enabled under the covers, essentially, by what you said, data science and data analytics.

But in the case of ChatGPT, again, I'll throw some of the technical terms out there but, if you don't understand that completely, that's totally fine. So, there's a concept of large language models, or LLMs, that you may have heard if you're in the field. And essentially what it does is, you know, OpenAI, the company that built ChatGPT, has essentially scanned the entire internet over a period of time.

It takes a lot of effort, it takes a lot of money to do that. And ChatGPT has learned everything that's out there on the internet that's publicly available over a period of time. And then it's able to use that information to respond to any of the questions that you have. That's, again, just one use case of AI.

There is also things such as facial recognition -- and I'm sure we'll get into some of that from a criminal justice perspective -- where you can actually show an AI machine a lot of different pictures, for example, of cats and dogs. And you can ask a question, you can show AI a specific picture and ask what animal that is, and it's able to respond. So that's another aspect. So, again, to summarize it's learning over a period of time from past information to make reasonable predictions about the future.

Adams: Not perfect predictions, though, right?

Kunju: In no way perfect. And in fact, there are many examples of ChatGPT getting basic math calculations wrong because of the way the question is asked. So, yeah, it's nowhere near perfect.

There is this concept of AI which is sometimes referred to as general AI or general artificial intelligence, which is a nirvana of AI, where you almost have AI having human-like characteristics in terms of being able to interpret and analyze the question being asked. But we are nowhere close to that right now.

Adams: And the subject of your scholarship, and certainly your book, is about this intersection between the criminal justice system and AI. And in episode 35 of "The Presumption of Innocence," we talked in detail with Mark McCreary, the Fox Rothschild Chief Artificial Intelligence & Information Security Officer, about some of the Department of Justice's policy pronouncements about seeking enhanced penalties for the misuse of AI as a intersection between AI and the criminal justice system. And certainly, that punitive component is out there, right? Anytime you've got a powerful technology, the opportunity for misuse is just as prevalent as the opportunity for positive, productive use. Is it not?



Kunju: Absolutely right. We just spoke about the fact that AI is not perfect. And you can also argue that humans are not perfect, right? So, our criminal justice system today, you know, we are always looking towards perfection, but you can argue that currently it's not perfect. But we are always striving to be perfect. And the same can be applied to AI.

And we can also get into ideas and concepts such as bias, which could be built into AI algorithms. And, you know, in the real world, again, people would argue that there is bias in the real world. And as we discussed, AI is using current and past data and information to make decisions. So, if there is already bias in the real world, which means that there is potentially bias in the data and information that AI is using to make decisions.

So, it's quite natural to expect that there could be bias in AI's decision-making. And as you know, AI is used in many fields other than criminal justice. But, I would argue that nothing is probably more important than its application in criminal justice. Because a lot of times it's, you know, life and death, literally life and death decisions that are being made in the justice system.

So then you talk about some of the other concepts is as explainable AI. So, what you're alluding to Matt, there is if you're using AI for decision-making, especially in the criminal justice arena, how do you make sure that it's explainable? You know, we spoke about how difficult it is to train AI, train the algorithms. You need large amounts of data, large amounts of machine power. And a lot of times these are black boxes. And even the creators of these AI applications may not know exactly how, you know, an AI application is making a decision. So it's very important for companies, especially in the criminal justice arena, who are building AI applications to make their applications as explainable as possible. You know, when it's making decisions regarding a person's life, you know, their term of imprisonment, their inclination to commit the crime again, what's known as recidivism in criminal justice.

So I think all of those decisions need to be carefully calibrated and, for the most part, explainable.

Adams: So when I read your book, right? And your book is self-titled "A Primer on Implications, Ethics and Policy." I view the position that you've taken a sort of a glass half-full view on AI's use in the criminal justice system. And I will self-identify as a glass-half-empty person when it comes to AI's use in the criminal justice system, because of those grave consequences that you just talked about. That life-or-death notion.

But before I challenge you on that notion, talk to us a little bit about some of the ways that artificial intelligence is currently being used in the criminal justice system and what you see as on the horizon for use in the criminal justice system, you know, within the foreseeable future.

And when I talk about the foreseeable future, I've never seen a technology in my lifetime evolve at the pace that AI has evolved. Even the internet, you know, we had dial-up and it took some permutations of ISPs before we got into what we all enjoy as that broadband internet connection that lets us seamlessly access the World Wide Web.



But just talk to me about the uses now, and what you see on the horizon as it pertains to the criminal justice system so that we can kind of dissect the pros and cons. Because I'm not convinced that AI is a 100% good thing for criminal justice.

Kunju: Yeah, I'm glad you asked that question, Matt. So, glass half empty or full? Yeah, that's always a relevant question to ask. The reason why I am optimistic, and I think it's more of a situation of the glass being half full, is because I'm comparing what AI can do today, and what it will be able to do in the future versus what we currently do in the criminal justice system, right? So, is AI perfect? No, but is our current criminal justice system, if you subtract AI out of it, is that perfect? I think the answer is no again. So, I'm looking at what are some of the positives or what are some of the benefits that AI application or algorithms can provide to the criminal justice system.

Now to answer a question about, you know, what's currently happening in the criminal justice system when it comes to AI. One way to look at the criminal justice system is break it down into three of its core pillars or dimensions. So, number one clearly is policing. And number two is prosecution. So policing is where you're responding to a crime. And prosecution is where typically the district attorney, or whoever is in charge for that region, prosecutes the crime. And the third pillar is clearly corrections, which is related to incarceration, the jail system as it's commonly known as.

So, if you drill down into each of those: when it comes to policing, you know, AI is definitely being used right now. For example, I'm in Southern California. LAPD has a predictive policing algorithm that they use to determine which areas to patrol, which areas are more likely to face higher crime versus some of the other zip codes. So that's definitely an application of AI that's currently being used.

When it comes to prosecution, again, you know, looking at our legal system, there's a lot of backlog. Because we do not have enough staff in our attorneys' offices to analyze and come up with the appropriate sentencing or decisions even, before sentencing the decisions that need to be made. And AI can accelerate a lot of that. It's able to process as we saw earlier, it's able to process information at a much larger scale, much more quickly than humans can. So, that's an area where AI can certainly provide beneficial in the prosecutory process within the legal framework.

Now, the third one is corrections. Now, when it comes to corrections -- and again, if you look at, for example, a state like California, and the number of people incarcerated, unfortunately, you know, our jail systems at this time cannot even handle the number of people being sentenced for imprisonment. And how do the staff manage all of that population? How do you make sure that there is no violence as you see very often in prisons and, you know, incarceration facilities? With the evolution of AI, with AI being able to, even face recognition and all of the analytics it can do within the video domain, you know, recent staff is able to monitor and intervene in advance of any incidents that could potentially happen.

So those are just, three simple examples. And it can also look at, you know, the past data of incarcerated population and determine who is likely to be a threat to some of the other population within the correctional facility. So all those are applications that are currently being used.



Adams: Before we go on to where you see it might go. Again, I admittedly, identify as a glass-half - empty person when it comes to the concept of AI and policing, prosecution and corrections because decidedly absent from those three pillars is that defense perspective, right?

And when you talk about policing, my question to you -- and I want to go through each of these -- when we talk about policing, what about the bias? What about the bias that the data has for over-policed populations? There's lots of data out there that says communities of color are over policed at staggering rates in this country. And if that is the case, aren't we just continuing to perpetuate that cycle if the very data that's in the system, in the algorithm being relied upon by the artificial intelligence is already so flawed and biased in the first instance? Can we correct that data for bias so that the AI is not as skewed and flawed as the human condition has already brought us to?

Kunju: Absolutely, great call-out, Matt. And that's the concept I was referring to earlier in terms of AI being explainable. And in fact, we can and we must correct biases.

I would even take an example, which is, you know, less life and death, which is financial decisions in a bank, right? So somebody goes for a mortgage application. And a lot of times now AI is making those initial decisions. And to your point, it can be biased against, you know, certain communities and populations. And that needs to be corrected. So it's really a responsibility of the companies that build these systems to look for biases proactively and not take it as a given or an accepted fact, simply because bias exists today in real life.

You almost need to look for biases in an algorithm when you're going through testing. Because every application goes through testing. So that's the first phase, using test data. And there, you'll be able to determine if there are biases, just looking at the outputs of some of the testing that you do.

Adams: How would you see, for example, racial disparities in policing? If the concept is, AI can bring us into the neighborhoods where there's crime, how do you correct the algorithm, eliminate the bias for the already documented historical patterns of racial disparities in policing against what are patterns of white-collar crime? And how do you put all that into one algorithm and expect it to spit out something that is free of bias? So how do you correct for that?

Kunju: Yeah, so that can get quite technical. But in a layman's language, once you run your existing data through the algorithms and the models, for example, you know, it may come up with two or three zip codes, which require extra patrolling. And then, as a responsible AI creator, you have to look at those zip codes and look at the population and see, to your example, if there are instances where, you know, demographics are skewed towards a certain part of the population.

So, an AI creator tester will be able to see that. There is queues, there is distribution which doesn't look right. And they will be able to question the results of the algorithm. And they should. Because, as you said, we don't want to copy and paste the existing biases into... so that's where human --

Adams: That's exactly it. You just captured it in a nutshell. We cannot copy and paste the existing biases. That, there's no better way to put it.



Kunju: Thank you. And it's not easy. Because, you know, AI cannot think for itself, right? That is, right now it cannot. So, it depends on the information you provide. And in some cases, you know, you have to make sure that, for example, in Facebook, they say that you can do microtargeting. And what that means is that applying the right filters, I can actually send an ad directly to Matt. You know, it's that.

Adams: I get them all the time.

Kunju: Yeah. You can go to the zip code demographics. It can actually customize to an individual out of, you know, billions of people potentially on Facebook.

So, when you are looking at AI and data that you're feeding into AI, you also need to look at a more aggregate level to eliminate some of the biases. So that's one of the techniques that you can use. And again, there are several techniques that you can use to detect and eliminate biases. And a lot of them are very technical, but as much as AI has advanced, the methodologies to detect and eliminate biases, fortunately, have also advanced. So there is no longer an excuse for a company to say that we just fed AI real data and this is the algorithm that they're using based on the real data. They cannot hide behind that excuse anymore.

And I think at some point, we'll talk about legislation, and this is one of the areas where usually I would argue against, you know, government intervening with, our creative process, especially in a very, still a very green field like AI, where innovation and creativity are going to be crucial to bring out all of its benefits. But when it comes to bias and making sure that you encourage, or in some cases -- we're talking about criminal justice -- penalize companies for not looking at biases that they could have caught in advance before releasing those algorithms for practical use.

Adams: When we move to the second pillar, that pillar of prosecution. Noticeably absent from your description of prosecution was the defense, the rights of the defendant. Which in our system are actually held above the rights of the prosecution. So, how can AI be useful to a defendant? You mentioned backlog. And certainly AI can process information. But that really is a view of the criminal justice system as a, sort of a conveyor belt of, sort of detached from the human element. Where does AI come in for preserving the rights of the defendant, which are in our system above the rights of the prosecution?

Kunju: Absolutely, well said. You know, innocent until proven guilty. That's the cornerstone of our criminal justice system. So, when it comes to defendants, again, they can leverage the same or similar AI capabilities that the prosecution has access to. You know, one of the things that I talk about in my book when it comes to future applications of AI is, for example, reconstructing a crime scene, you know, in a 3D format by looking at all the information that's collected, two dimensionally, right, from a crime scene. And a defendant's attorney can certainly, if the defendant is indeed innocent, use those techniques to prove his innocence or to justify an alibi that person has.



So, again, on the defendant side, all these technologies that the prosecution has access to are also available to the defendant. So that levels the playing field, so to speak. So that's one way to look at it.

Adams: But does it really? Because there's a couple of really high-profile examples of technologies that are made available to prosecutors. Which, under our constitutional principles in this country, Brady and its progeny, are supposed to be disclosed to the defense. And when I talk about these technologies that use, in some way, AI, I'm talking about ShotSpotter. I'm talking about DNA-mixture algorithms that use AI to try to take a complicated sample with multiple DNA and break them apart.

When defense lawyers push, these companies who create these algorithms, create these technologies have agreements with governments that say, if you are forced to turn over your intellectual property you agree that you will dismiss the case before you will do it. Yet the government still used these technologies. Yet the governments still upend the presumption of innocence by taking these technologies and using them in cases where they know full well that when a good defense lawyer pushes for the code, the source code, the understanding of the science behind it -- or the so-called science, I would argue some of it's pseudoscience by the standards of evidence, right?

This takes AI to a new level. We're not just talking about administrative convenience. We're not just talking about eliminating backlogs by letting cases get processed. We're not talking about taking data and correcting it for bias. We are talking about the wholesale use of AI to accuse and prosecute individuals, human beings. And then there being an acknowledgement that if the defense pushes for how that was come to be, that they'll just assume dismiss the case rather than disclose that source code. How do we reconcile that with our criminal justice system?

Kunju: Yeah, I mean, great topic. I don't think there is black and white or a straight answer. It's a very difficult question. Because even leaving AI aside, you know, even before AI, there are again, classic cases where the prosecution and/or defense wanted data from iPhone, right? Something that you would think that, you know, it's easily available. But Apple as a company oftentimes would push back unless it's a national security issue. So --

Adams: In the case of a very high-profile incident of terrorism, they still pushed back

Kunju: Yeah, because Apple has promised its customers that the data is always private. So now we're getting into constitutional rights of each citizen versus what the government can demand from a company. So, we're getting into... and I'm sure there will soon be a case in front of Supreme Court where they can make the ultimate judgment on which way the company should act in that situation.

But going back to your original question about AI and the fact that, you know, the source code, or going back to the explainability concept, that's simply not available for the defendant. But I would say that they should use all available discovery demands at their disposal. So, right now, the defendant has full rights to seek discovery of how prosecution arrived at a certain conclusion. But



again, they cannot influence the private companies to -- I don't think at this time, at least -- give up their source code, as you're saying.

Adams: Would you agree with me, though, that if the private company does not give up their source code that if the charges are predicated on this type of technology, they should be dismissed?

Kunju: I can give you my opinion on that and, you know, that may not be the legal standard right now in our country.

But I could give my opinion that, we are talking about procedural justice here, right? So that's the cornerstone of a criminal justice system. The system is what's called procedural justice, where you are assuming fairness for a defendant, especially because, again, innocent until proven guilty. So I would opine that that's a fair standard, what you're suggesting Matt, but it may not be a legal standard right now.

Adams: Well, there are examples. I will submit that there are very many examples where those technologies were not made available for the defense. In one of the most high-profile examples, it's the StingRay technology, which is the interception of electronic communications. There are some high-profile examples where national security has been used to justify the use of these devices where they... not just any data being transmitted through mobile channels, but frankly, all of it, including all of the metadata about the substance of the communication. And they're being utilized in these national security events that designate as such and arguably warrant this type of technology being employed. I have my reservations about whether that is justified. But at the end of the day, this technology has been used time and again by governments, knowing full well that they have a contract with the private company that developed it that says, you'll never turn over our source code. You'll never let this black box get opened by some defendant. And you will, if you have to, if a court orders you to, you'll dismiss the case. Isn't that a Pandora's box?

Kunju: Yeah, I think, I always look at it, and I'm sure you do too, you know, from all perspectives. So, as a professional, as I mentioned, I come from a data analytics background. I run my own data analytics company. And I can understand where some of these large companies are coming from. Because they have invested potentially billions of dollars into making their algorithm. Potentially have a patent on it. And they're always concerned about, you know, what if my source code gets out there? What is preventing somebody else from maybe spinning off another company overnight, right? That it took me 20 years to build and, you know, yeah, a billion dollars to build. So, so there is always that. And there's constitutional protection, as you know, which is a totally different realm of the justice system, you know, the patenting and on trademark and all of those aspects that come into play.

And again, Matt, when you ask me questions about AI, I usually try to -- or anybody else -- I usually try to step back and look at the world without AI, right? So, for example, in the state of California, at least, you could have an audio recording of an evidence to a crime that happened. But it's not even admissible in court because you did not take permission of the parties involved, right? So you have it



right there and on your phone or another device. So, I think there is interplay of various elements of the justice system.

But again, you know, my personal opinion is that you always want to be the presumption of innocence, that's the title of your podcast, and I'm completely on board with that. In fact, I'm actually talking right now to victim rights organization where I'm looking for volunteering opportunities. And that's something that's very close to my heart. There are instances of, I think there's an organization out there -- and I forget the name, I should remember -- which actually advocates for people who are already incarcerated and there are again high-profile stories about people who are proven innocent after committing to a prison for several years or even --

Adams: We just had somebody from the Innocence Project on this program, just a couple of episodes ago, talking about people have been put to death in this country because of mistakes. And to me, as we sort of tease out the paradigm with respect to AI, at least in its current form, right? Where in my view, I guess I remain unmoved that I should take this glass-half-full approach because all I see is it being wrought with potential issues. And we haven't even gotten to the worst of the issues that I see yet, and we're going to get there in a few minutes because I want to get your perspective on it.

But if we can distill down all this skepticism in my mind and balance it against some of the positive views that you have for the use of AI in criminal justice, there's one word that comes to mind and it's guardrails. It's that AI may provide some benefit. There may be some application of AI that is useful. But, would you agree with me that unfettered, uncontrolled, government hands-off, no laws, wild west AI is potentially devastating?

Kunju: Absolutely. No question about that. And that's what I was also alluding to when I spoke about legislation and what needs to be in place, even though I'm a proponent of creativity, innovation when it comes to new technologies. But especially in a field like criminal justice before a police department adopts any application, it has to go through several tests. And one of them is clearly if the company has made an attempt and documented it to make sure that the biases are eliminated. I guess that's really number one. And I don't think anybody will disagree with that, especially in a criminal justice environment.

But again, going back to half full or half empty. You know, you look at the scale at which AI can, for example, "Minority Report." I mean, that movie came out many, many years ago and I mean, that's an example of the potential of AI right? Where, in that case, the application, I don't think they call it AI, but the application that we could predict when a murder would happen and who would be committing that crime.

And you look at, you know, some of the concepts around always prevention is better than cure, right? That also applies in the criminal justice system. So, a trivial example of, if you are able to patrol an area which is prone to crime, and you could save lives, right? So we're talking about what the benefits are. You could save lives, literally, by making sure that your police officers are patrolling the right places at the right time. You could save lives, you know, in a large soccer stadium with



potentially thousands and, you know, even more depending on the venue and the event and you have bad actors. And AI is able to use facial recognition to tell you that there is a bad actor, there are bad actors in that crowded environment. And you could save thousands of lives.

So those are all benefits that you need to weigh alongside. But I also agree in a criminal justice system even one defendant falsely accused or sentenced for a crime is one too many, right? You almost have to get it right 100%. But are we doing it today? Probably not without AI being in place.

Adams: My mind goes to immediate panic and a sort of dystopian "1984," George Orwell-type of view when it comes to the use of this stuff. But I don't know that we're going to resolve that today, so to speak.

But when it comes to guardrails. Are we there at this point right now where Congress has to act? Where, as a country, as a nation, it seems like we can't get anything done. But as a data scientist, somebody who, frankly, knows enough about this technology to know how powerful it could be, and frankly, how it could be misused. Are we at the point where we need action immediately? I mean, it seems to me that this technology is evolving at such a pace that the answer has to be yes.

Kunju: I agree with you on that, Matt. Certainly, we are at a point -- I know that there was a presidential declaration from the United States calling out some of the guardrails, as you're talking about, that need to be in place. That's more of, more along the lines of an executive order. But I do see a real legislation coming up very soon to Congress and the Senate.

And, you know, I think there are elements there that everybody can agree on, especially when it comes to an arena like criminal justice. I know that Europe is probably ahead of the U.S., where they actually have legislation in place to monitor and address some of the issues that we just spoke about.

Adams: If you were advising Congress today, where would you tell them to start when it came to those guardrails, specific to the criminal justice system, that you believe in order to curb abuses must be imposed upon AI today.

Kunju: I think for me, the number one consideration will be explainability within the AI applications that are used in the criminal justice system and being able to ensure that every decision made by AI. As difficult as it might be, right? But that's the responsibility of the companies that produces applications.

We are talking about, certainly talking on one side about companies investing heavily in terms of labor and resources, financial. But also on the other side, looking at these companies profiting out of that for many years to come and --

Adams: If explainability, you know, tell the audience a real-life example of what that would look like. Take, for example, ChatGPT -- which by now, everybody pretty much knows what that is, right? You ask it a question. It spits out a response. As you said, at the outset, it has learned the internet. It has learned the universe of that information that's out there. And we know that information is flawed. We have so many examples in our society. And if we're not going to just copy and paste that bad



information and create this hyperloop of disinformation that is... frankly, the potential exists. How would you mandate the explainability, from a technical perspective, of ChatGPT?

Kunju: With ChatGPT, I think everybody recognizes that it's a model, that a computer has learned over time. But, but specifically speaking about explainability in criminal justice system, you know, some of the examples. We already talked about predictive policing, and if you're assigning patrols to a certain zip code, you better be able to explain, as a chief of police of a city or a sheriff in a county, why you made that decision. And then even question the creators of the application that enabled you to make the decision. Hopefully, you're asking these questions before you adopt that application. That's going to be key, right? So, before you sign up for a predictive policing application, you have to ask those questions in advance. You have to see it in action before you say, yeah, we are going to adopt that.

And another example would be, you know, facial recognition, right? So, if you're choosing one individual out of thousands of people, and AI is predicting that that could be a bad actor, you have to be able to explain why, right? It shouldn't just be based on the demographics or the race of that person. So that's where the explanatory concept needs to be in place when it comes to technologies like that.

Adams: Does explainability start with an inherent recognition that this is an imperfect science?

Kunju: It does. And you could also argue, you know, somebody patrolling the area. Again, I tried to look at the world without AI, right? Before AI there was patrolling. And you know, how does a policeman or policewoman or officer choose to pull up one individual versus another, right? So what's the decision-making process there? And they should be able to explain that, right? We're not even talking about AI. And I think the same standard should apply to an AI application making the decision.

Adams: Well, we are certainly only scratching the surface, Pramod, on some of the most critical issues of our time, I would argue to say. But I want to end on one point. And I want to ask you a very direct question: What keeps you up at night about AI? Against the backdrop of your, book, which, admittedly, I think presents AI with the benefits and the positive benefits. And I'm not here to say there are none. I'm not that half-empty kind of guy. But I am still extraordinarily nervous about the misuse of AI and some of the problems that AI could present.

In particular, I think what keeps me up at night is the notion that. AI could create fake evidence. Could create issues that are completely manufactured, and there would be little way for us to actually prove that, or disprove the authenticity of that evidence. I was cutting my teeth as a young lawyer as the e-discovery age really came to be. And I remember trying cases in the early part of my career where government lawyers were trying to authenticate primitive forms of communication, like audio recordings, emails even. And things of that stripe. And I remember being successful at pushing back on the authentication of those pieces of evidence. Because we had a bricks-and-mortar analog to compare those electronic, technologically driven things to.



This is a new frontier. This is something that there genuinely is not a bricks-and-mortar analog to. The email has the letter. The recording, you know, has primitive versions of it that were done and magnetic tape rather than in digital format. And there's metadata that secures and offers indicia of reliability.

But with AI, I fear that we're going to get into a position where we don't know what is reality? And that keeps me up at night. What about you?

Kunju: Yeah, well, first of all, Matt, thank you for the work you did there and you did exactly the right thing. And I would also argue that some of the examples that you just highlighted are really illegal acts, right? Even within the current legal framework, manufacturing evidence. I mean, those are illegal actions. And we have to hold --

Adams: Except there's bricks-and-mortar ways to demonstrate that right? We have people that can analyze metadata when it comes to an audio recording. We have people that can determine that a videotape has been spliced together. We have DNA experts who can testify about the composition of the alleles on a DNA strand such that they can compare it and give a predictive probability that that is who we say it is. How do we do that with AI?

Kunju: Yeah, and that's currently happening. Even trivial examples, such as, you know, students are using ChatGPT to write their essays. And now there are tools out there which is almost counter AI. It can tell you what percentage of an essay was written by AI, right? So think about that.

And you can extend that to any industry or even criminal justice where ... Deep fakes, right? So, if somebody is faking a crime scene, or somebody is faking somebody's voice or appearance and there are counter applications to the effect that can detect that, you know, those are being manufactured or artificially generated. So I think there is that.

Adams: Are they reliable?

Kunju: They're getting better, right, as we speak. And, and again, I look at the criminal justice system and the people in the criminal justice system again, going back to glass being half full, as people with good intent. And that has been my experience, just working with, for example, the District Attorney of Orange County's Office, or, you know, the Irvine Police Department chief. And I think you have to assume, just like you have the presumption of innocence on the defendant side, you have to also assume positive intent when it comes to people. And a lot of them are there because they want to serve our country, and they want to serve our communities.

But, you know, are there bad actors? Yeah, there could potentially be. But, you know, that's a presumption on the prosecutorial side and the policing side that I go into looking at some of these scenarios.

So, with that said. In the unlikely event that somebody manufactures evidence, I hope and I'm confident that there will be enough counter applications in place, ready to investigate those claims and ready to prove them false, if they are actually false. I don't know if they're there yet, but as much



as AI is advancing, there is also a parallel stream, which is actually detecting if something was done by AI, if something is fake. And we have to look at, you know, both sides of the coin, so to speak.

Adams: Well, Pramod, this has been a really fascinating, at times terrifying, conversation. I think as the technology does begin to evolve even further, I'd ask for one commitment: Come back on the program once we get some more development in this area so that we can talk about it some more. Because I don't think that we have even scratched the surface yet.

And, while I, again, will not say that I am an empty glass, I am certainly not a half-full glass either. But, it's been a fascinating dialogue. Your book is "Artificial Intelligence in Criminal Justice: A Primer on Implications, Ethics, Policy." And we've just spent the better part of the last hour talking about implications, ethics and policy. And I think it's a conversation that really must continue. I know I sit on our New Jersey Supreme Court's task force on AI and the conversations that have been had in that environment with people from the legal community, judges, people from the academic community and the data science community, have really been fascinating.

I'm glad we were able to, Pramod, share that with our audience today. Keep up the good work on this subject. And I look forward to hopefully talking to you again real soon as we have something else to talk about as it relates to this rapidly unfolding evolution of a really new frontier and technology.

That's all the time we have here today on "The Presumption of Innocence." I'm your host, Matt Adams. Until next time, take care. Bye-bye.