Course Name – Artificial Intelligence, Law and Justice
Professor Name – Dr. Krishna Ravi Srinivas
Department Name – Center of Excellence in Artificial Intelligence and Law
Institute Name – NALSAR University of Law
Week – 03
Lecture – 13

Artificial Intelligence, Law, and Justice; session 13, second session of this topic. Algorithmic justice is the topic.

Let us do a quick recap of what we discussed in the last class. In the last class, we introduced the idea of algorithmic justice and then discussed why this concept has become very interesting and has also evinced a lot of interest among the people. We also

gave a definition of the origins of this and discussed how the idea of algorithmic justice is viewed not only in a negative sense but also in a positive sense, as it is much more than just a criticism. It is also a tool or an approach that will try to separate certain things, make the companies understand certain matters which otherwise they would have ignored, and is also something that pertains to the use of algorithmic decision-making in public services.

So, drawing upon the work of Prof. Karen and others, the problem of using artificial intelligence as well as automated algorithmic decision-making in the criminal justice system was discussed extensively. We also discussed what the implications are for such uses in human rights and how the algorithmic justice thing may not be what we actually need to look at in criminal justice principle, in the sense that algorithmic justice demands that we adhere to criminal justice principles scrupulously when we are using ADM. To put it another way, we also need to look at algorithmic justice as a critique, as something that could provide a solution, and also as a concept that would be useful given the widespread use of algorithmic decision-making across sectors in different applications, particularly in various sectors like health, police, the criminal justice system, and elsewhere. In this class, we will take that discussion forward and then see what the other important points in this discussion are.

Professor Simon Chesterman has identified three challenges posed by opacity: we have discussed opacity, the black box nature, and the lack of transparency. He says these can result in encouraging or failing to discourage inferior decisions because the problem is that there is no potential for human oversight and accountability. That itself is something that could lead the system to come up with inferior decisions, and certain decisions that we normally wouldn't allow in the regular context. Criminal justice systems or any such systems prohibit any discrimination based on certain identity markers explicitly. But the problem with ADM is that such impossible decisions may be allowed on account of the inherent bias in the data or on account of the bias in the algorithmic training itself. So,

what we normally do, or what normally the judges, the courts, the criminal justice system, or the administrative law do not allow or permit, could be possible through the back door when we use ADM for such purposes. In the sense that they may create a discrimination-based ruling that otherwise would not have been possible or admissible at all. Then it is also possible that when a court gives a decision, it is really legal, it is really legitimate, it can be challenged, it can be questioned, it can be countered, argued; it can go all the way up to the highest courts. But here the problem is that the process itself, when it gives the answer, is important for us to know.

But when that process comes up with answers that cannot be defended fully or that cannot be understood fully, or which could be classified as illegitimate decisions, we are in a serious dilemma or we are in serious trouble. So, the process, due process, the fair process, or as we said, the court should treat both parties equally and at an equal distance. That may not be possible in algorithmic decision-making. One more point he makes is that, look, the opacity may not be fully 100%. We may need some novel explanations for this. But some decisions may still not be explained at all. And for some decisions, they should not rely on machines to make them. Some decisions should always be beyond the scope of machines to decide upon. So, when we look at opacity and decision-making algorithms through these ideas, we find that opacity is not a mere technical problem that can be solved through better algorithms and better design systems. Rather, opacity should be looked at in a different way than what happens because of this opacity.

When the opacity itself results in some decisions that could be termed patently illegal or decisions that a normal court would not have given. Or a decision that is there, but then it could be explained, or it is something that can be really understood or explained by anyone. So, the idea of using opacity as a mask and then saying opacity is a technical thing; we can cover all that, we can remove this, and the problem would be solved is wrong. And we need to understand one more thing. If we talk about opacity and then transparency as totally two different things, and if we say that 10% opacity is okay, 95% transparency is okay, and 5% opacity will always be present in a system because not everything can be explained. Systems may have a margin of error of 5 to 10%. That should be permissible because the system is very efficient and it can handle many things that humans cannot handle. So, this is a trade-off. This trade-off should be allowed, particularly in criminal justice systems and in public welfare systems. This is not something that is about efficiency versus cost. This is much more than that, and we will see that a little later.

Algorithmic Law

- How do we understand algorithms and particularly when they take decisions. Editors of European Journal of Risk Research point out that algorithms often work as a functional equivalent of a legal rule, similar to a structure of command and consequence. Further we need to study the legal rules as well mechanisms embedded in them as part of new emerging discipline algorithmic law.

- They further make some important observations, after reviewing the literature regulation of algorithms they ask "Today, we should also evaluate the contemporary experience of algorithmic decision making and set standards for computer engineers to eventually pass a test in programming code that could resemble the evidence-based, justifiable, reasonable and prudential activities of a legal decision-maker. In terms of the performance that we would expect of AI in an imitation game, can machines provide useful predictions for decision-making systems in the legal domain? ". They propose something that is similar to Turning Test for it.

- They also point out there is a counterview that only human beings can do decision making as they have empathy and that is a core human feature essential for lawyering and for Judges.
- Artificial Intelligence Risks and Algorithmic Regulation Pedro Rubim Borges Fortes1, Pablo Marcello Baquero2,* and David Restrepo Amariles2 European Journal of Risk Regulation (2022), 13, 357-372

Algorithmic law is a new concept because algorithmic regulation itself, as an idea, is not even 20 or 25 years old. So how do we understand algorithms, particularly when they make decisions? In the sense that, what do we make of them? The editors of a European Journal of Risk Research pointed out in a special issue that I am citing here. That algorithms often work as a functional equivalent of a legal rule, similar to a structure of a command and consequence, in the sense that, as we have seen in numerous earlier examples, like a decision tree, we can also reduce the legal rules to a decision tree. So, their argument, which we have made earlier, is that algorithms work as a functional equivalent of a legal rule—not the normal or identical equivalent, but functional in the sense that what the legal rule does, the algorithm also does.

Maybe the legal rule is written in a way we understand, although not all of us will be familiar with legalese. Algorithms are coded; they do their calculations in a different way, but the rules are the same, or the rules that run on the algorithmic decision-making systems and the corresponding rules that are being run or used could be the same. So, the command-consequence structure is being maintained. We really need to study the legal rules as well as the mechanisms embedded in them because a new discipline is emerging as algorithmic law. Algorithmic law is a sub-discipline of law that looks at algorithms not as mere algorithms but as part of the legal code.

Although they have borrowed quite a good number of ideas from Lawrence Lessig and his particular work on code as law, as well as his work on understanding legal principles versus code, and how this is important in the cyberspace context where the code operates like a gatekeeping mechanism, or functions like an administrator, or decides this way or that way for you. So, based on that scholarship, they have made substantial arguments in this paper, which I think anyone who is interested in algorithmic law and regulation should go through. So, if we have to study legal rules embedded in the system as part of emerging algorithmic law, it is very important and interesting. So, the second point they are making is that after reviewing the literature on the regulation of algorithms, which has

been present for the past 15 years, we should also evaluate the contemporary experience of algorithmic decision-making and set standards for computer engineers to eventually pass a test programming code that could resemble the evidence-based, justifiable, reasonable, and prudent activities of a legal decision-maker. Put in other way: we are demanding something very exacting, something very difficult from the people who code and build this system.

So, what are we saying? We are saying, "Fine, we want to evaluate your algorithmic decision-making and then set standards," so that you pass. We are going to prepare some tests for you to undergo, or the system should undergo, and then we can really say whether the system is as good as a person who makes decisions based on the law, or if it comes anywhere near the legal decision-making process. This is equivalent to what Turing proposed decades ago to test whether the person at the other end or the system at the other end is intelligent or not. So, who would expect that in an AI imitation game, machines can provide useful predictions for decision-making in the legal domain? They propose something that is similar to the Turing test for it. The paper has a very interesting description of it. And then they create a very interesting view. It is always assumed that humans can make decisions as they have empathy, and there are four human features that are essential for lawyering and judging; thus, this view is acknowledged. But don't forget that there is also a counter view that says humans have this specific capacity to become lawyers, to become judges because they have empathy. But there is also a counterview to this in the sense that why can't algorithms be judges, why can't ADMs be judges and advocates, or for that matter, why can't AI systems act as advocates and act as judges?

Because in the literature, there are many arguments that support that AI can really function as a judge, and there are also counterarguments. In fact, this debate has been going on for quite some time as to whether AI can replace a judge or be part and parcel of a judicial system where there is AI courts run by AI systems. Now in the Agentic AI age, or when we are seeing the emergence of Agentic AI, particularly AI agents that are

becoming more and more sophisticated and can handle different tasks, as well as handle a lot of query-based replying and understanding, it is possible that an AI agent can function as a quasi-judge who will do the first round of vetting, run the first round of analysis of evidence, and then do the first round of understanding the arguments before coming up with a preliminary judgment. So technically, with their increasing capacity to reason, argue, assemble evidence, collate data, and analyse them, today's AI agents may emerge as quasi-judges or judges, helping human judges come up with decisions, whether AI emerges as a full-fledged judge or not. So, this possibility, which was not there earlier, has arisen now on account of the developments in technology. So maybe in the next few years, we may have quasi-AI advocates, quasi-AI agents acting as quasi advocates and quasi judges, or as members of tribunals in the sense that they can replicate those tasks if they are given enough data and training. Coming back to this, they also point out one thing: there is an idea called smart law, which unfortunately we won't cover in this course.

Scientific mathematical algorithmic law is again shaped by risk and technology. In the sense that smart law tries to bring the rigor of computing algorithms, mathematics, specifically statistics, and science into law. So smart law and the ideas that AIs can be judges, and that a good portion of the judicial system can be taught as code or reduced to code, or that a good portion of the legal system itself can be translated and used as code or the ideas that repeatedly tell us that this is not simply a technocentric view. Let us also keep that in mind. This is not simply a technocentric view.

This is a view that finds the parallels or the equivalence between law and code. And then the way law operates, and then the way code operates, it is also possible that the legal process itself could be made more within codes, more scientific, more proof-oriented, and then run by statistical formulae; also, an AI system can handle it much more efficiently. In the sense that they can do regression analysis, they can do a whole lot of statistical reasoning. So, the idea of smart law or AI judges is not merely a technical perspective; it is a perspective that many people argue will improve the law, enhance the legal process, and make law and the legal process much more scientific, much more accurate, and much more effective. Convincing because it is based on science, it is not based on mere legal norms, legal values, human rights, or anything; so, making law more "scientific" is the idea behind smart law.

The technology today is available for it, so the question here is whether we should go for such choices, even if it is technically possible, or if it is morally acceptable. There are two or three views on this. One that we can use it to aid the judges, but not beyond that. Another view is that, given the capacity and the way technology develops, we should not eliminate the possibility of them becoming judges or doing judicial functions. There is also a third argument that could be placed between these two, beyond these arguments: in certain circumstances, we can use AI judges but with human intervention in the sense that there should be a human who oversees that, and this will also help us address many issues that the legal system or the justice system today is not able to address more efficiently and much more economically.

So how we use AI systems as judges is also something that is related to algorithmic justice because our fundamental problem here is that we have been arguing that

algorithmic justice demands something like justice. It is not something that calls for technically sweet justice. One way that the meditators have argued is that some activities are considered to be essentially human, and that our society will value the presence of humans in the loop and the decision-maker. So, AI and robots may carry out the services and activities we no longer want to perform. In this, we should carefully examine whether we prefer to be judged by human intelligence or by AI.

So, this is what the argument is that they make. Algorithmic law and regulation challenge everyone to rethink power, democracy, regulation, and institutional design. This is inevitable because algorithmic justice precisely engages with these issues. Now coming back to the point they have raised, they again cite a legal scholar, Balkin. In warfare, is it better to use robots and automated systems run by robots, or is it better to use humans? The problem here is that a society may feel that it is better to use humans because no human will lose their life, or it is better to use robots and such systems for the simple reason that human life is precious. And then a whole lot of accountabilities will come. And then if 1000 soldiers get killed or 2000 soldiers get maimed in a battle or in a war, lots and lots of questions will arise. Families will complain. People will be very upset that we lost so many lives. But if we are going to use robots for that, that problem is eliminated. In the sense that, at least as of now, we don't consider robots to be humans or equivalent to sentient beings. So, they build upon this logic, and then they argue that algorithmic law and regulations challenge everything to rethink power, democracy, regulation, and institutional design.

Based on these principles and these arguments, one way to look at potential solutions is this. We need to differentiate what exactly we are looking at between the ADM and the AI systems. Many AI-based systems in sectors like health enable AI to diagnose, suggest solutions, and make recommendations. AI systems can identify the need for cataract surgery and may recommend it. But it is left to the humans to decide, although experts can agree with the findings of the AI system. Our AI system here is more like a

recommender position; it is not the one that acts as a final arbitrator, nor is it the one that will perform the operation unless you know you are going to combine that with a robot that could perform surgery. The argument is that which is again very interesting and perhaps also somewhat convincing to some extent. In some sectors, AI expertise can be as good as that of a human being. So, there is no harm in using AI, for example, in health. If the AI's expertise is as good as that of humans, why not use it? If the AI can evaluate students' papers like any other professor, provided it is given in a form it can understand, it can then evaluate, what is wrong with doing it. Since wherever possible, if AI can reach this level of expertise and domain knowledge, and it is being trained in it, and it can be proven or used in a non-controversial way, we can go for it. Similarly, in Law and Justice, the argument is that AI can provide a recent solution, a finding, or justify its order. In the sense that as systems become more sophisticated, some problems, like opacity, transparency, or labelling them as black boxes, will be reduced, and systems will also be able to explain better, reason better, and come up with 100% explanations like any other human would.

But can that per se justify that we should employ them as judges, particularly in the way that we look at algorithmic justice, or should we take the view that AI can be an advocate, or AI can be a judge? It is not permissible, but we can use AI advocates and AI judges as proxies in some contexts. It can be argued that, as I said, these things, such as opacity, can be addressed. So, the counterargument goes that many of the concerns raised regarding the algorithmic justice discourse and narrative can be countered by better systems that will be less discriminatory, much more technically sophisticated, less biased, and as good as, if not better than, human decision-making. So, how to use AI and how not to use AI can also be reduced to a level of technical efficiency and then to technical potential.

Solutions and Approaches

- Can systems be designed in such a way that that users, designed can collaborate and develop them based on lived realities of persons. According to Siddharth DeSouza absence of careful legal design can result more exclusion and margialization (S de Souza, "The Spread of Legal Tech Solutionism and the Need for Legal Design" European Journal of Risk Regulation)
- Algorithmic audits can be another solution.
- Ethics committees to evaluate algorithms can be a good solution to address ethical issues
- Bringing in Responsible AI and explainable AI concepts to practice will also be useful.
- But we need clarity on what we want to address or solve and what are the trade offs we are willing to agree to.

Then comes another approach or another way of thinking that people are also now discussing in terms of inclusivity, bringing in inclusion right from the beginning. There

are ideas like inclusive design and inclusive development of software, and then bringing users and developers together. I mean these ideas are present in other disciplines and applications, such as participatory plant breeding, user-led innovation, and innovation development, where the users also play a major part. So, Siddharth DeSouza, whose work we have already cited and discussed in the context of the use of AI by the Supreme Court, comes up with a very interesting proposition that users and then developers can collaborate on the systems based upon the lived realities of the persons. So, our major argument is that you are trying to reduce people to statistical numbers; you are trying to desegregate and then reduce them to data, numbers, potential, and possibility. The potential or possibility of this person committing a crime is 27.3%. But still, he cannot be given bail on account of these things. Or the probability of this person jumping bail and then committing the crime again is 80% based on his previous behaviour and on other parameters. So, we are saying that if people's lived experiences can be brought in, how do we bring them in, and what sort of values we assign to them is a different question. That is a question of not just technology but also of a bit of social science understanding and methodology. He says that careful legal design can address many of these problems.

But still, in the absence of careful legal design, more exclusion and marginalization will occur. Or, to put it in other words, legal design can alleviate or address some of these concerns raised again and again by people who talk about algorithmic justice. But that needs to be done with care and that needs to be done with caution. Another solution that has been in vogue, which many companies are doing, and which in fact many of the consulting firms that develop and then deploy AI systems also do, is to develop algorithmic audits. Having an algorithmic audit is also a good practice because it will tell you whether your systems are compatible with some selected AI ethics codes or whether they fulfil some AI ethics norms. These things we will see in detail in the sessions on AI Ethics. So algorithmic audits can be another solution because they are not something you do in the algorithm for the first time and then forget. Algorithmic audits are discussed in the continuous sense that as the system is improved or updated, we should conduct them. No, that's not the way. Algorithmic audits can be conducted where the outputs can be compared, and if there are inherent biases that are reflected, they can be addressed by auditing them and then by trying to recode or redevelop the system.

Another solution that has been proposed is that we need to bring in ethics committees to evaluate algorithms, and this is a good solution because ethics committees understand these problems in a very acute way, as ethics is not a very dry domain. And then there are different schools of ethics, different thoughts on ethics where people talk about normative ethics; you have Aristotelian ethics, you have pragmatic ethics, and then these different schools of ethics also have different perspectives on understanding human suffering and the impact on humans by different systems. And particularly speaking of ethics, some disciplines like bioethics can really open up many things for us to understand better. So, ethics committees, which are normally there as part of conducting clinical trials or experiments in hospitals or for any large-scale experiments involving humans, can also be used to deal with algorithms and their inherent biases; they can assess them based on ethical principles and parameters. Then another idea that has been brought in is bringing responsible AI and explainable AI.

We have already discussed this solution, but we will discuss both these concepts in much greater detail in the subsequent classes. But the fundamental question in all these things is what we want to address or solve and what trade-offs we are going to agree to. So, what do we want to address? And what we want to address is whether it is a mere question of opacity in technical terms or if we are talking about a larger solution in the sense that an opacity of 15% or 10% is okay, provided some of the other factors related to bias and discrimination are fully addressed, or if a black box is 95% transparent and 5% black box, or if there is a trade-off. An ADM that is proven to be more biased and seems to have more discrimination, if that is the major weakness, could be addressed and technically improved. There could still be minor other problems in the sense that there could still be some discrimination and some bias.

Can that be taken as a good example that a much-improved system is better than a previous system? So, what are the trade-offs we need to keep in mind? In the sense that these are some of the solutions if there is an idea that AI systems can be deployed. If they can be deployed when they are improved. So, the normative principles, if they can be brought into the AI systems, and then if ethics committees can evaluate them based on ethical parameters, and then if people can come together, as Siddharth DeSouza has pointed out, with a legal design carefully taking into account the lived experiences of the stakeholders coming together. Teaching the people who are trained in algorithms and then software and then engineering to bring in human elements to generate ideas about how they can assign values or come up with better solutions. Can that result in meeting some of the needs of algorithmic justice? The answer is wide open; we really do not know because the topic of algorithmic justice itself is quite new, as we said, and some of these debates have been happening only for the past 15 or 20 years in a much more concrete way. Although you might find that discussions about using AI in the legal system, particularly as judges, are much older in the sense that they have been thought of as a solution even earlier.

Good Intentions and Reality

- It is presumed that AI based systems can offer some services such as legal guidance if not advice and they can be used to provide information to citizens in multiple ways. But according to **Blank and Osofsky** the reality is different. Based on an extensive study they point out "However, the use of automated legal guidance also comes with important costs. Critically, we show how, precisely because of its perceived strengths, especially when compared to the often messy, or ambiguous, formal law, automated legal guidance can obscure what the binding legal provisions actually are" They also point out while such guidance may be given by tools when legal counsels argue on behalf of governments may not give the same advice or may provide something different. Thus those who can afford can get better legal advice. So while AI systems can provide legal assistance , that addresses only partially the issue of getting the right legal advice.
- Joshua D. Blank and Leigh Osofsky 2025 Automated Agencies Cambridge University Press

But we also need to look at some things with a pinch of salt. A normal solution or a normal approach is that there is a gap between the public and the government, and the public often is not able to approach government offices, service providers, and others because of a whole lot of other problems, including people not being available, overcrowding, or insufficient human resources within the government. So, government departments are using more and more digital tools, including chatbots and do-it-yourself tools for the people who can really use them to address some of their concerns. And of course, there are a whole lot of ways that governments try to automate their own operations. And if these can be expanded to legal services, in the sense that I am sitting somewhere 300 kilometres from the capital city where the high court is situated. I need some legal clarity on whether my understanding of this legal issue is correct because I want to file a petition in the nearest court to me. Can I have an AI-based system that provides legal services? It provides legal guidance; it could be an interactive AI bot; it could be a web-based decision-making tool that can also help me in the sense that we can simulate; an AI tool can also simulate in the sense that I input questions, it gives me answers, then we simulate a game or we simulate a situation to arrive at what could be the potential solution or what could be the potential way out. So, there are so many possibilities.

It can also be a good source of information. But a recent study by two professors in the U.S., Blank and Osofsky, indicates, and this is a very interesting book which I would recommend anyone to go through, they said the reality is different because they found that the promises made or the potentials assumed—that more and more use of AI tools or more and more use of such digital technologies in legal systems, particularly in government systems, is going to make things friendly—is not making things easy for people. Often, that translation is not 100%, or there is a huge gap between the potential that is expressed and what is actually realized. Here they point out something that is very pragmatic and is all the more relevant for countries like India.

However, the use of automated legal guidance also comes with important costs. Critically, we show how precisely because of the perceived strengths when compared to the often messy and ambiguous formal law, automated legal guidance can obscure what the binding legal provisions actually are. In the sense that the legal provisions can be smoothened and then it can give a very short, crisp answer that could be valid and acceptable, it will not go into the complexities, or what the legal assistance system provides through an AI system can offer something similar to a solution that looks perfect and understandable. But then what is not said by the system or what does not caution you is far more important. To put it in another way, as I said, someone approaches this AI-based tool for some assistance. It can come up with a solution; it can say something very specific and then say, "You do it." But it is also possible that it is trained in such a way that it identifies only one solution, or it is directed towards choosing that as the best solution for the person who asks this question. In contrast, when I go to an advocate or a legal counsel, that legal counsel will have a broader perspective and understanding, and will try to explain the various options available. This process will take more time; if you go to court by applying and invoking these provisions to ask for this relief, it is feasible, but the probability is not very high. On the other hand, in this process, we invoke these provisions, and then if we use this act or if we totally go to a different forum instead of going to court, if I go to a tribunal, we may get a better judgment or it could be faster.

So, an advocate can open up all the options, taking into account my priorities and concerns, while ensuring they are legally sound. But their experience with their case studies shows them this is not something that is happening with the legal AI agents because the legal AI agents give them crisp, quick answers, but which is not the true and full picture. Now, I am fighting a case against the government. I have filed a case against the government. So, the government is recommended by the government leader or the public prosecutor. I, as a person, if I rely on the advice given by a legal tool that the government itself has nominated or made available, can I really expect that the public pleader, the government pleader, or the public prosecutor will 100% use the same arguments or will 100% agree with what is said? No. It could be totally different because the pleader, the public prosecutor, or the government advocate is a person who will look at other possibilities, including counterarguments, who will also look at various other precedents, binding precedents, cases where what was argued and with what clarity, what should not be touched upon, or where the case law can be used and cited with confidence. But a legal service tool may not be very efficient. More importantly, it may not even happen that the argument put forth by the legal tool is the same one that the advocate representing the government will present in court; both could be very different, and both could be totally different as well. So, their problem, they say, is that the real-world life experience of people in courts, tribunals, and advocacy is something, whereas the advice given by the legal tools, AI-based tools, is something that is totally different.

The ground reality is something that a tool may not be able to really grasp and give us advice, or the advice given by the tool may be legally sound and perfect, but the ground reality of the public prosecutor or a pleader coming up with a totally new set of arguments cannot be wished away. Another thing they point out is that often what happens in such circumstances is that if an AI tool is the only option available to me, I am too poor; I will use it, and I will rely on it. But those who are much more financially

secure or who can explore other options will not simply relate to him. They will try to guess the best option; they will go to more than one advocate, or they will consult different people for this. So, when an AI system can really provide legal assistance, that is perfect, they may not be able to give the right and exact legal advice we need. But the problem here that we need to address is access to courts, and solving that through AI boards or an AI legal system, which provides better access. This is only one part of the problem. The other part of the problem is that people are too poor, or they don't have the capacity and mobility to go, which is a different problem. Now we cannot substitute technology as a sweet solution without considering the other problem, which is poverty, or the lack of financial services or resources to go to court. So, the lack of one cannot be 100% substituted by another one and say we have found a fantastic solution.

If we were stuck somewhere 300 kilometres away, and then could use the AI system. You can do remote counselling through an AI system. So, this sort of thing, where reality and good intentions are different, is something we need to be very careful about, and the reason this makes an important point in the context of algorithmic justice is that algorithmic justice demands full justice; it does not believe in finding speedy technical substitutes and then promoting them as good solutions that will give us algorithmic justice.

NPTEL

Final word?

- "As AI's increasing use in the judiciary makes codified justice more appealing in other contexts, its downsides are likely to be reproduced, too. Problems analogous to the ones discussed above will likely arise. And the basic menu of responses, with all their limitations, is also likely to recur. Finding a path forward will require attention not only to technology and law, but also to technology's impact on conceptions of justice, in both its human and artificially intelligent forms "

- Developing Artificially Intelligent **Justice Richard M. Re & Alicia Solow-Niederman** 22 STAN. TECH. L. REV. 242 (2019)

Now, there is a very interesting observation that was made five years ago in the Stanford Technology Law Review. This is an interesting article that people who are interested in law, justice, and technology should really read. After exposing the way artificial intelligence and justice have been talked about in the literature, the pros and cons of that, and how different schools of thought have discussed it, their final words are this: "As AI's increasing use in the judiciary makes codified judiciary more appealing in other countries, its downsides are likely to be reproduced too." In the sense that when America codifies justice through AI, and when some other countries repeat the same thing, the downsides will always be there. "Problems analogous to the ones discussed above will

likely arise. And the basic menu of responses with all their limitations is also likely to recur." So, "finding a path forward will require attention not only to technology and law but also to technology's impact on conceptions of justice," which we have discussed in the context of ADM and in the context of AI in relation to the rule of law "in both human and artificially intelligent forms." In other words, what exactly is the technological impact on conceptions of justice in human intelligence, and then artificial intelligence is very important. Or this is the final word which cautions us that we should not take things for granted, or we may develop artificially intelligent justice, which may be possible, but still, this is something we cannot wish away.


In another context from a different country, this is a report that could be an eye-opener for many of us. It was the House of Lords committee report of 2022. They appointed a committee that looked into the application of new technologies in the justice system. They did a thorough job of getting stakeholders' views, looked at different technologies including facial recognition technologies, and then they had a lot of things to say about algorithms, but due to a lack of time and other factors, I am not fully reproducing them. In fact, one thing they said is that a country does not have any algorithmic registers in the sense that there are so many algorithms used by different public agencies, by municipalities, by police, and by various other entities, but there is no central registry where you can say that this agency has used this algorithm, and then this is developed by them, and this is exactly what it does. So, in the absence of an algorithmic registry, the spread of algorithms across different entities creates a situation where the users, stakeholders, or those who use these systems are not communicating with one another. And then another point that they made is that each agency wants to come up with a technical solution in terms of algorithmic decision-making. So it goes for bidding, it goes for calls, it goes for people to develop and then deploy.

But then these are often done without full understanding, without the need to develop standards, or without the need for assessment, and often these are also done without the

proper technical assessment of them in the sense that when public service utilities and public service authorities go for some of these technical solutions, including algorithm-based decision making, they do not pay adequate attention to the things they should pay attention to. Instead, what happens is that they either rely too much on technology, or they think that technology can provide us with a solution that will be much more efficient. But the observation that they have made is something that cannot be wished away because it is very valid for all countries, including India. When deployed within the justice system, AI systems have serious implications for a person's human rights and liberties. At what point can someone be imprisoned on the basis of technology that cannot be explained? Informed scrutiny is, therefore, essential to ensure that any new tools in this sphere are safe, necessary, proportionate, and effective.

This again is the fundamental principle on which any justice system will work. In the sense that it will not unduly punish a person. The proportionate test, the necessity test, and the effective test, or the benchmarks that judges apply. In the sense that you will not put a person in jail for 20 years for a small crime. That person may be asked to pay a fine, or they may be imprisoned for 2 to 3 months. And what they see is that this sort of scrutiny is not happening. What they have found or what they have uncovered is a new wild west. You all know how a new wild west operates, how people operate in that, where new technologies are developing in a phase that public awareness and government regulation have not kept pace with. This is something that we also know as a pacing problem, which we will see later, particularly in the context of AI. So, technology rules; the advent of a new technology justice system is an interesting report that also discusses algorithmic audits and how some of these concerns should be addressed by public authorities. So that the deployment of these technologies does not result in a new Wild West.
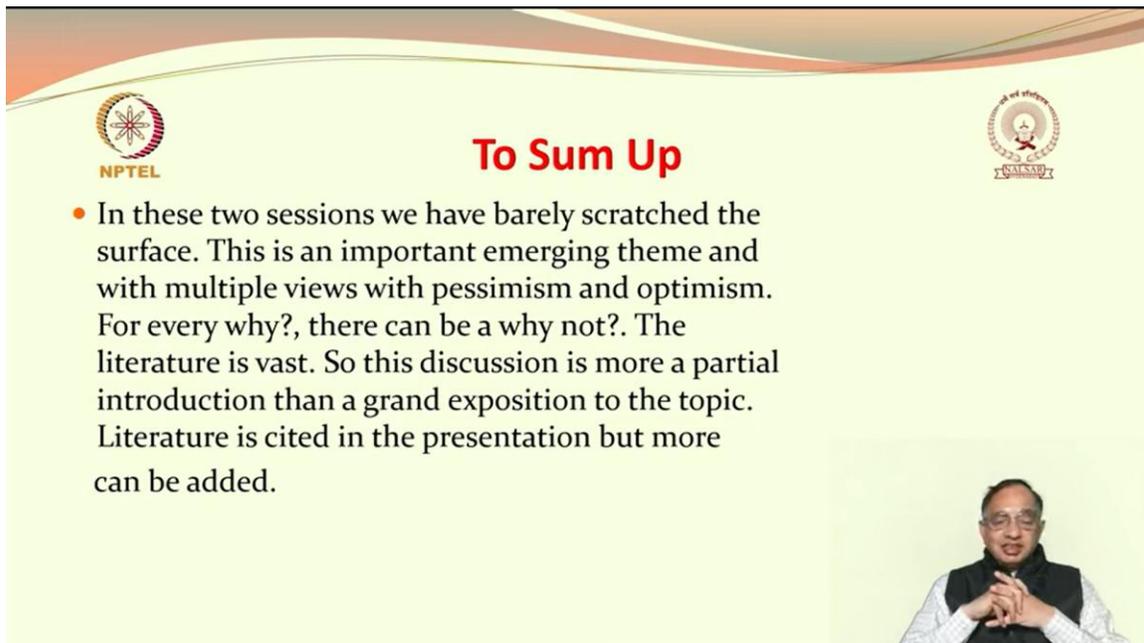
To sum up, in these two sessions on algorithmic justice, we have barely scratched the surface. This is an important emerging theme with multiple views of pessimism and

optimism. That is obvious; it could be expected. So, for every "Why?" of using ADM, there is an equivalent "Why not?" of using ADM. The literature is vast and expanding. So, this discussion should be seen more as a partial introduction to the topic than as a grand exposition of the topic. And the literature cited in the presentation has been given, but there is more that can be added. So, this is how we are winding up the two sessions on algorithmic justice.

The next session is going to be about copyright; we will devote four sessions to copyright for AI. Since AI and copyright are very hot and emerging topics, we will address threadbare some of the key emerging issues as well as some of the fundamental questions that again and again pop up when we talk about AI and copyright. Thank you.