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 – 06
Lecture – 28

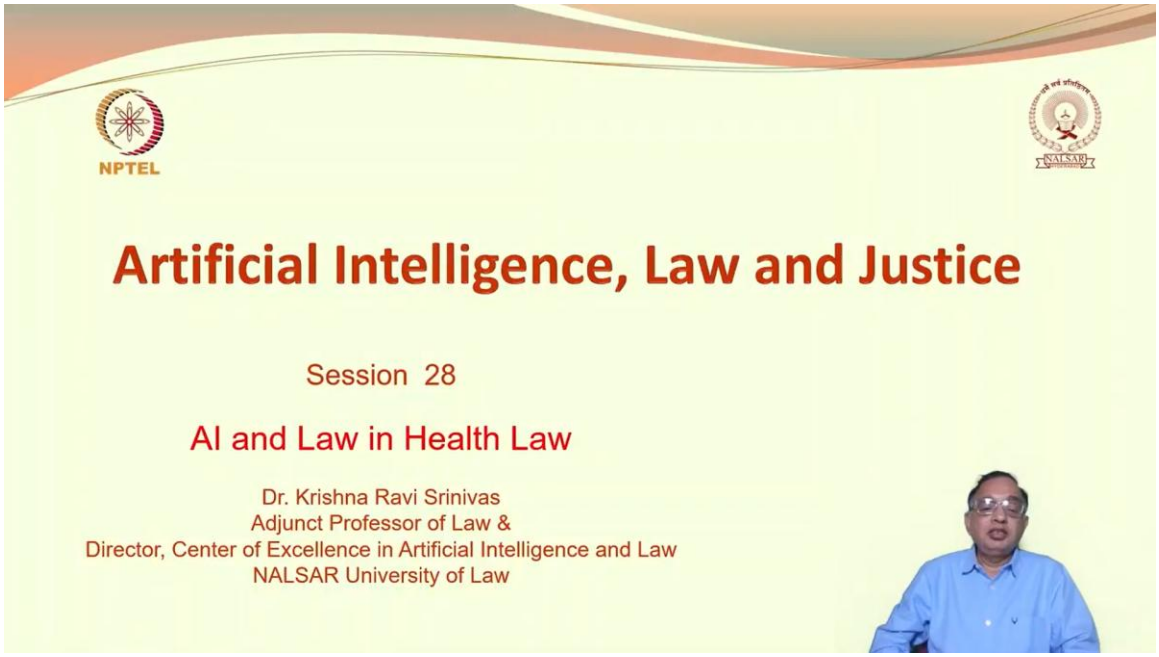

Artificial Intelligence, Law, and Justice: Session 28 - AI and Law in Health Law. Health law is a very important branch of law of which many of us may or may not be aware, but the dimension and scope of health law has been expanding due to technological developments, where things have become more complicated because technology not only makes many things possible but also raises a lot of questions that the law or the system is not prepared to answer. In other words, using technology looks very attractive, very good, and very fine, but it also raises many new challenges from a legal perspective, particularly in sensitive areas like health law.

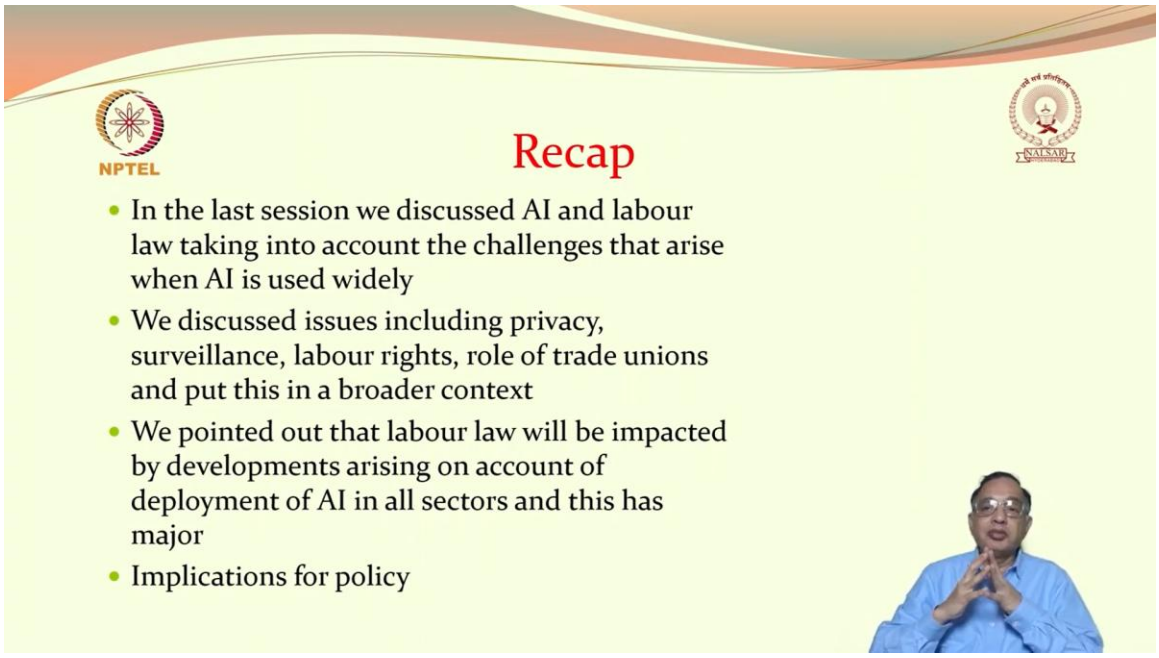

So, let us do a recap before we go into detail in this session. In the last session, we discussed AI and labour law, taking into account the challenges that arise when AI is used widely. We looked at the potential negative impacts, the potential positive aspects, and discussed in detail the need to strike a balance, as well as how labour law has to become agile to deal with the emerging challenges arising from the deployment of AI broadly in all sectors. We also looked into some specific stakeholders or groups, like people who are working remotely, gig workers, or people who are working on contract but are not necessarily working in a specified location or part of a larger workforce. And then we also looked at some of the things, whether retraining and upskilling alone will be potential solutions or whether we need to go beyond them, and how labour law has to come to grips with new developments like the extensive use of AI for surveillance or monitoring employees' behaviour.

We also discussed that since AI is being used extensively for all sorts of workers, including those who are doing a lot of remote work, there is a right to disconnect, and people should have the right not to be connected to the system for all 24 hours. So, then we discussed, how the right to reconnect or the right to disconnect is also linked with AI, because when AI is monitoring you, the moment you log off, AI should also be cut off; it should not continue to monitor you by having some systems inbuilt in the computer or the equipment. So, we discussed some of these things, and then we also mentioned that labour law would definitely need a major revision, taking into account all these developments. And we also discussed some of the policy-related issues; particularly, we mentioned in some detail the findings of the Economic Survey 2025 brought out by the Government of India. In this session, we will look at another important branch of law called health law.

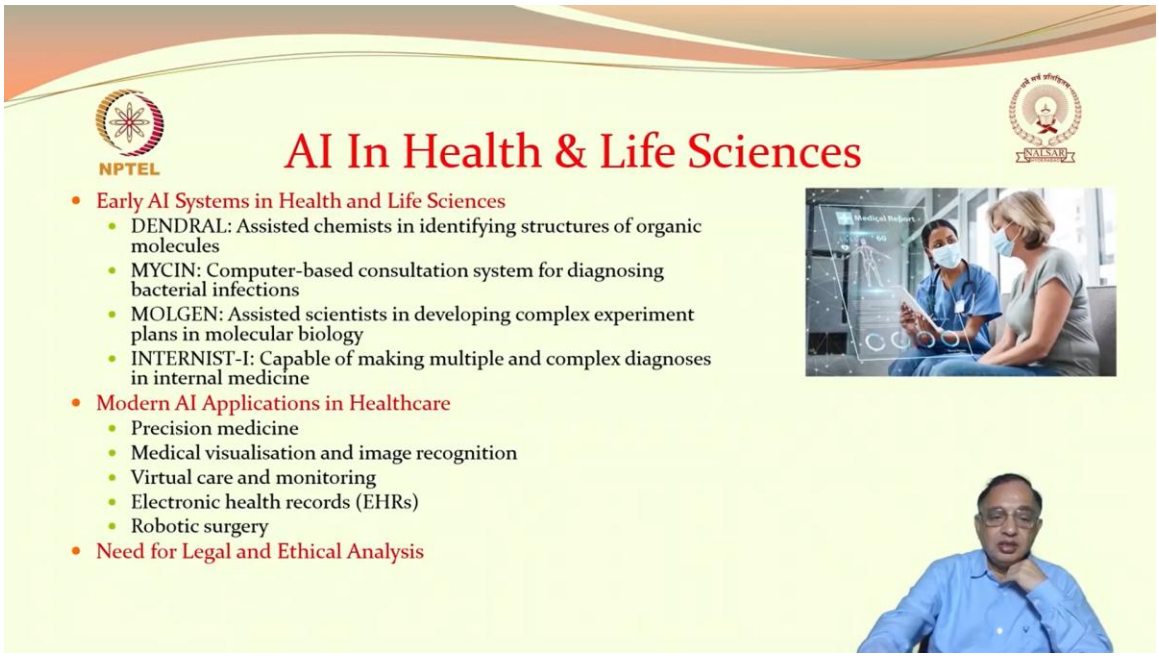

Now, AI in health and life sciences is not new. It has been there even earlier because some systems were used earlier, like DENDRAL for assisting chemists in identifying the structures of organic molecules, MYCIN for diagnosing bacterial infections, MOLGEN, and then INTERNIST-1; all these tools were available in a limited form to some selected physicians and hospitals. In the sense that AI was used, it was not widespread; it was used more like a diagnostic clinical tool than as something used on a day-to-day basis by many people. But today, AI in healthcare is used more widely than people would imagine because some of the smartwatches are also directly equipped with AI-based systems. And then when AI-based systems are connected to smartwatches, other wearables, and some of the devices that are embedded in a person's body either on account of treatment or for some monitoring, a lot of data gets gathered, a lot of information gets gathered and passed on. So, the question of health law having to deal with a whole lot of data governance-related issues is part of this larger issue of how health law should deal with data, not just data governance, but also the way data itself is impacting the health sector as such. So, the modern applications in healthcare are quite broad.

We are listing only some here, such as precision medicine, medical visualization and image recognition, virtual care and monitoring due to remote working, and remote hospitalization. One of the most important applications of AI is that it facilitates a lot of health diagnostics and health treatments that can be done on a remote basis using AI solutions, allowing people to be treated without the need for them to come all the way to hospitals located elsewhere. Then, digitization of health also means we need to rely more and more on electronic health records or digitized health information. Then, of course, robots are also performing surgery in many hospitals; this is too well known. But all these things need a robust legal and ethical analysis because, as of now, the legal system is not

fully equipped to handle many of these emerging concerns, particularly when it comes to data.

Additionally, there is no global organization on this, and there are no rules that are binding on all sectors of medicine with a similar understanding or basis, whether it is for data. So, what we are seeing is that rules and regulations are emerging, but in bits and pieces. So, there is no overarching health law or overarching health regulation that covers all the applications of AI being used in the health sector, including data governance. So, this fragmented, sketchy bits-and-pieces thing results in a lot of challenges not just for people who are dealing with the law, but also for policymakers, for people who are developing applications and devices, for patients, and for every stakeholder in the system. But then, this is not something that is unique to India. Other countries are also coming to grips with that, trying to understand it, and then develop some frameworks. WHO has come up with global ethical frameworks. ICMR has come up with some ethical guidelines. All of this is necessary but not sufficient in light of the developments.

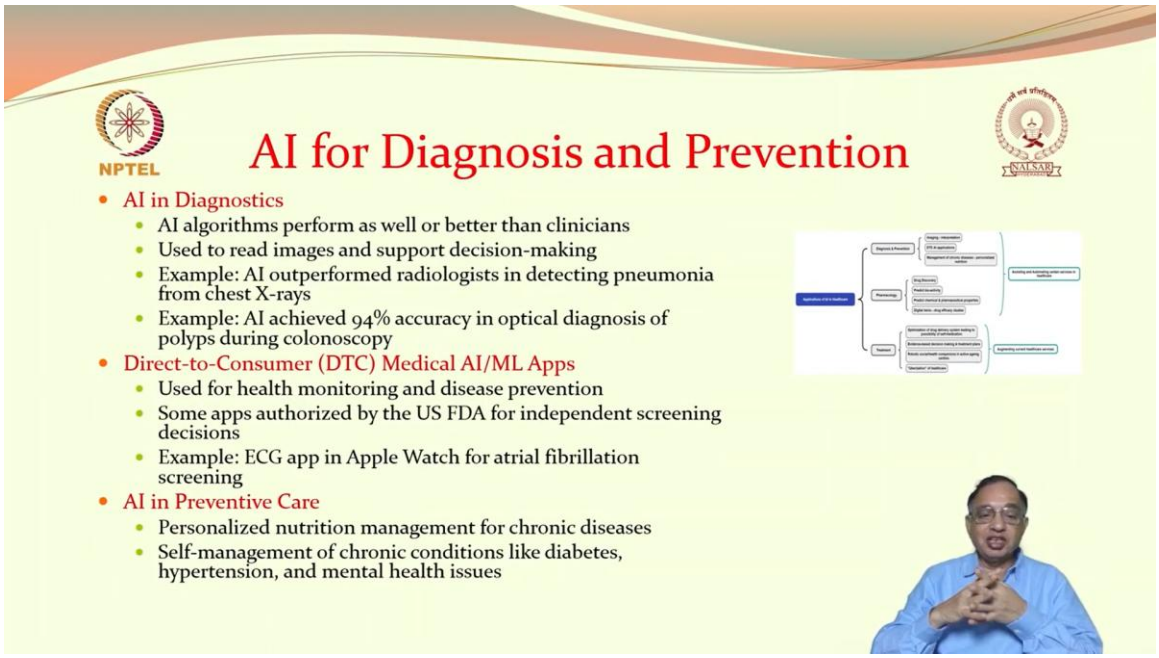

So, AI diagnosis and prevention can be used. Algorithms perform as well as or better than key clinicians. Algorithms can read images, can read X-rays, can read all sorts of images, and then make a diagnosis. They can also help in support decision-making. Some algorithms have a better reading capability or ability than radiologists, providing a quick diagnosis. Algorithms are also used not only in diagnostics but also in identifying which treatment or which sort of clinical solution will be more suited. So diagnostic is only one application, but direct-to-consumer (DTC) medical AI/ML applications are many because of the very idea that most of the variables, including smartwatches and other devices, have either AI inbuilt or are connected to AI systems.

The Apple Watch has been approved by the US FDA, and it has the mechanism to sense and alert the patient to whether there is a potential for arterial fibrillation. So, when the watch itself can alert, it could save lives, and in fact, it has saved lives in many instances where people can simply rush to the nearest hospital and get adequate treatment. So, when watches, devices, and a whole lot of other equipment are either embedded in the body partially or embedded closer to the person who is undergoing treatment or who is doing something, data transmission is only one aspect. But the other aspect is that this continuous monitoring and data collection also means that, whether people like it or not, they are being under surveillance. In the sense that you may say, "I agree to it; that's why I am doing it," but whether it is ethical or whether it is really necessary for so much monitoring through medical devices to happen to everyone 24/7 is questionable. Then in preventive care, AI has a huge role because it can help not only in managing chronic diseases, but it can also do a lot of things even in primary healthcare. For example, AI can be a good solution if we can identify people who are likely to become anaemic or who are likely to develop certain diseases well in advance. And then there are a lot of developments that are happening in this field, but since we are not dealing with AI in healthcare per se, we will not get into details here.

Now, AI in pharmacology, AI in drug discovery, screening, and digital twins in drug discovery are also well known. So, I will not go into the details here.

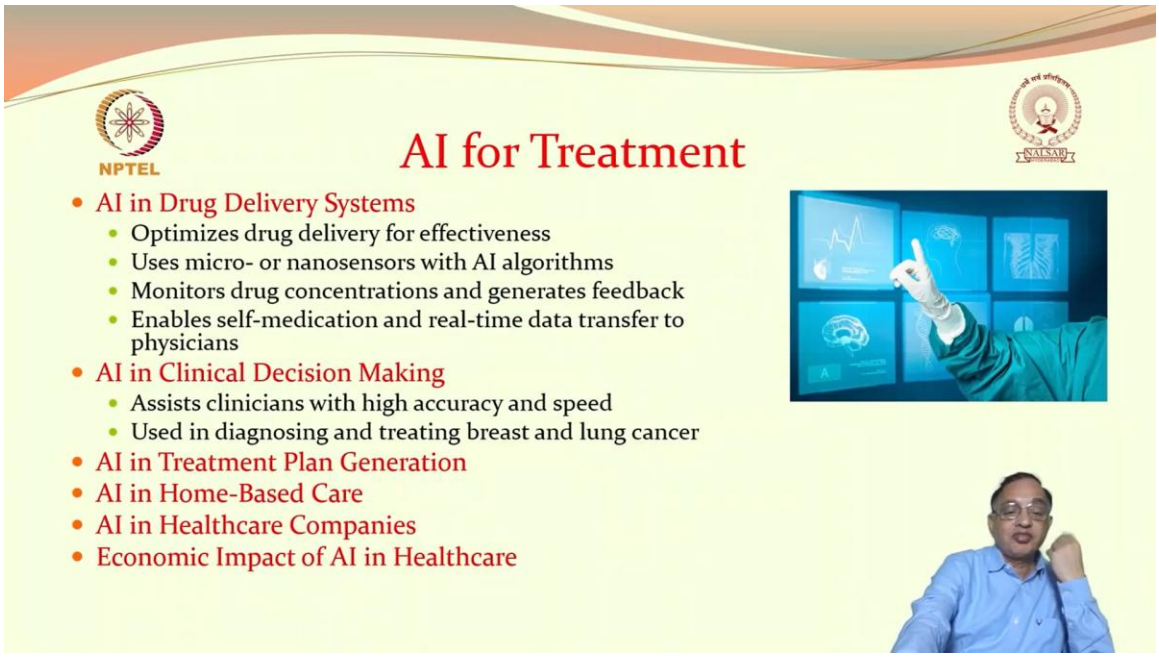

Then, for AI treatment, we discussed some aspects here.

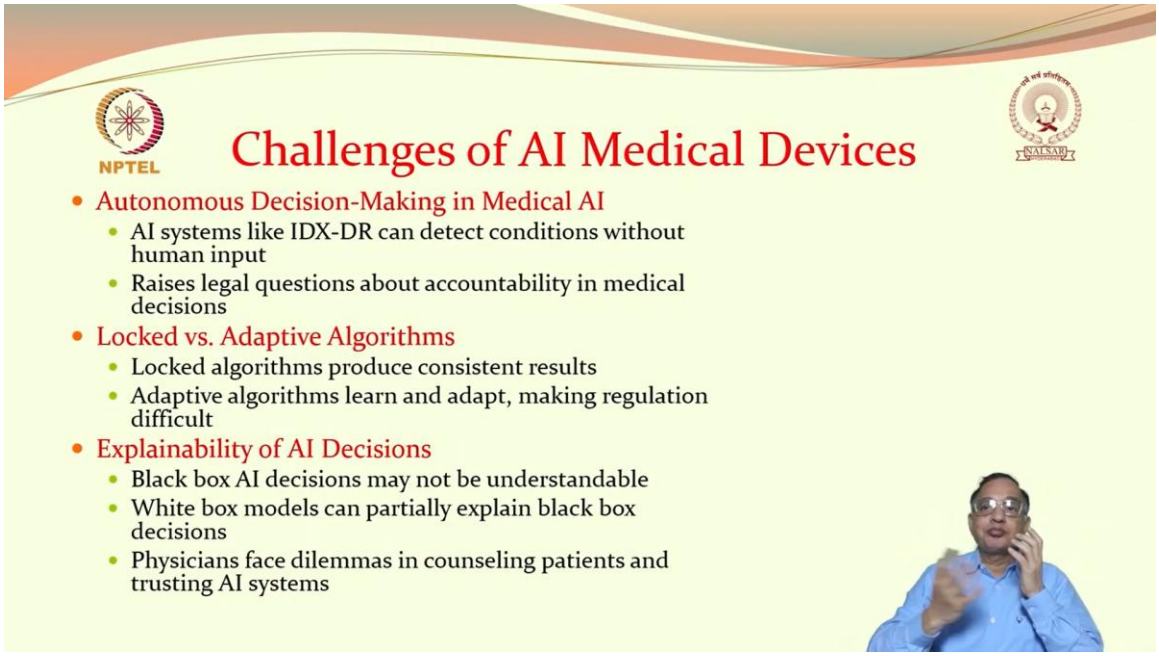

So, broadly speaking, AI is bringing about a transformation of the health ecosystem—not just diagnostics, not just clinical services, not just data. It's a transformative ecosystem. Transformation is happening. AI is resulting in a paradigm shift. But this paradigm shift is occurring at a juncture where the health system itself is undergoing lots and lots of digitisation. Many hospitals now have electronic health records. And then electronic health records are also part of the national health mission in many countries, including India. And then electronic health records are also part of the larger system where insurance companies can demand access to electronic health records; not only insurance

companies, even employers in some instances can have access to or can demand access to electronic health records.

So, the digitization of the health system, which is again happening at a rapid pace, is being combined with the AI revolution in healthcare. So, this opens up lots and lots of interesting possibilities and then lots and lots of challenges as well. But when it comes to medical devices, there are some specific things in which AI is part of the solution but can also be part of a larger problem. But some of the AI-based systems, including the medical devices, can diagnose or detect conditions without human input in the sense that they can know, as I said, the person is likely to experience arterial fibrillation within the next 15 minutes. It can warn, but then you need not do anything. The information is gathered, data is collected, and then based on algorithmic thinking and pattern recognition, it tells you that within 15 minutes you have to rush. It raises an alarm. So, this raises questions about accountability in medical decisions. Prima facie, when everything goes wrong, the question of accountability and responsibility will become a huge question mark. When everything goes fine, people will say everything is fine; we need not question the system.

Some of these systems are built in such a way that they connect data through algorithms that are locked, providing consistent health results, whereas there are some algorithms that learn and adapt, making regulation difficult in the sense that you really do not know what exactly they are learning, what exactly they are adapting to, and what sort of data governance framework is in place as part of the larger information-sharing network. Then, when we bring in more and more AI into diagnostics as a decision-making tool, the question is: where do humans fit in? How many of these decisions are taken by health AI black boxes, which may not be understandable, and what we need in explainable AI in terms of health care is something that is again very emerging. So, there are lots of uncertainties, lots of queries, lots of doubts, and then a lot that we do not even know what we should know in the first place in many of these cases. So, we need to move towards white boxes when it comes to AI in health. But white boxes and black boxes alone are not sufficient.

We need to have a much better system where explainable AI should be made mandatory in some circumstances. And then the problem here is that often when people indulge in self-diagnostics or self-testing using AI systems, they come to some conclusions, rightly or wrongly. So, when patients tell them, after examining, that the AI systems may not be accurate or may not be right, or that their health conditions are such that they should discard that, then the question of credibility arises. For physicians, they really struggle hard to convince such patients who think that AI is the best solution for their healthcare needs. It's a problem for physicians as well. Because physicians can explain their diagnosis, what they can look into, then on that basis what they have understood. But then AI wouldn't be able to do that. But still, people may have better faith in AI because AI is machine-based.

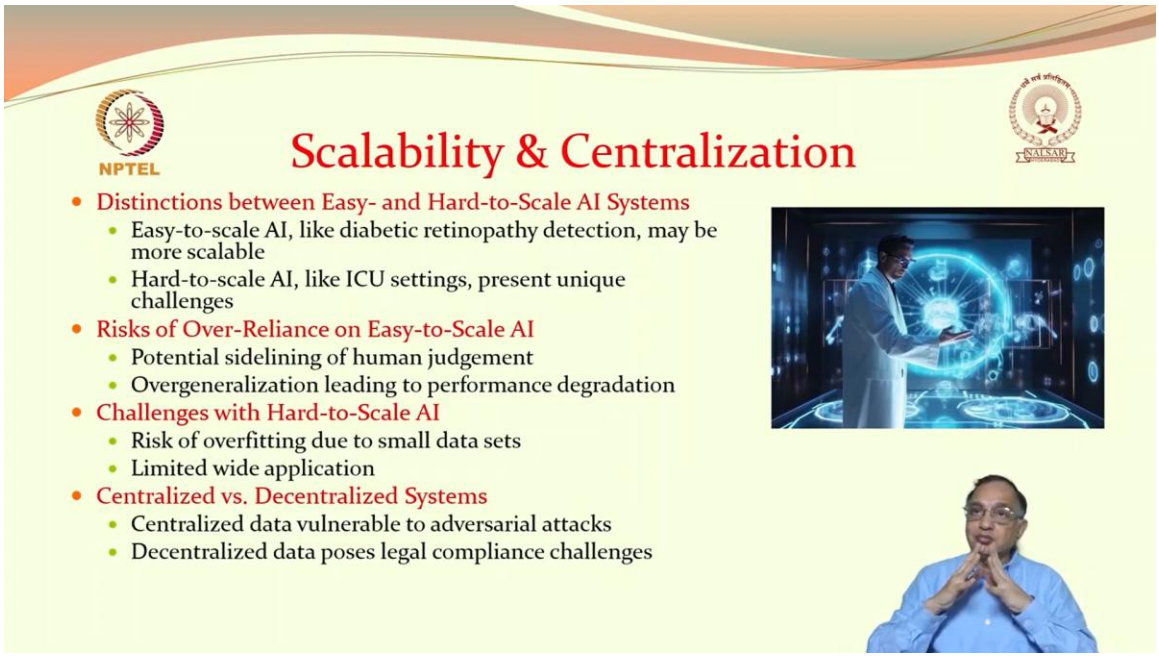

And then some of the things like scalability and centralization are major issues in health AI. Because scalability is about how to scale an application that is tested with a limited population or tested with limited data sets into a larger application. What are the ethical and legal challenges that arise there? Then some of the applications that are to be used in critical, life-threatening, or life-saving settings like the ICU need to be evaluated for how well they have been tested, their reliability, and the legal liabilities and responsibilities involved. This is a very grey area as of now. Relying too much on AI could become a problem because if you are going to automate AI decision-making at some stage, and AI is going to take over partially or fully the role of the clinician or physician, the question is where human judgment will fit in the overall health system when AI is a major driving force. And then this overgeneralization or overreliance on AI may also result in some sort of degradation of human labour, human work, and human understanding and appreciation. Then the problem of scaling with AI could also be a problem, as what is derived from small data sets or small data samples may not be the one that will be truly workable with a larger population.

However, when they try to fit it incorrectly or forcefully and then attempt to reach some decisions or conclusions—either preliminary, partially, or fully—it raises a lot of questions. So, the question again is that if hospitals are going to maintain electronic health records, then it becomes a cybersecurity issue. But what happens when there is a national-level database of health records? And if there is going to be a national-level database of health records or EHR, it could result in cybersecurity risks; it could result in systems getting attacked either to hack the system, to steal data, or for some other ulterior purpose. But health data governance may again be something that forbids decentralized data governance or decentralized storage, or rather may favour large data governance,

particularly in a centralized manner. How health law should deal with these issues is a major topic.

The problem here is that some of the issues we are talking about are not something we have seen before, particularly in health law. So how do we apply them? And then we cannot apply a single approach to various AI-related legal issues because of the context, the accountability, the responsibility, the decision-making process, and the role of AI in that varies. So, we cannot compare AI-based decision-making in health with driverless cars or robots. They are in a different category. This is a different category. We cannot extrapolate either the rights, the responsibilities, or the accountabilities. So medical AI needs a specialized regulatory governance framework in which health law is a key component. And then the complexity of the evolving medical AI systems, which also includes data governance, the governance of various other AI systems, and robotic systems that are being put together, means that the complexity is going to increase in the days to come. Particularly when AI is part of a larger datafication of the healthcare system, where robots will play some role, AI-based systems will play some role, and in diagnostics and treatment decisions, AI will either assist or be a major decision maker. So, the complexity of the AI-based medical system raises many challenges and questions for health law.

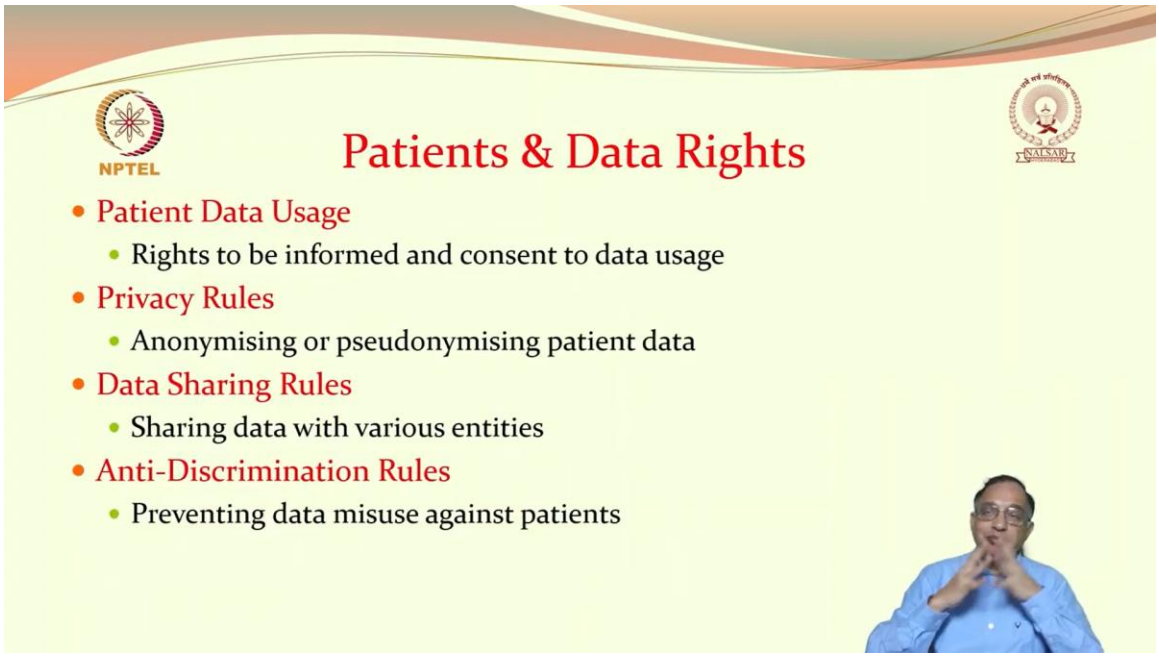

Then obviously, patients have data rights. Data governance is a major issue; privacy is only one aspect of that. Another aspect is how people use data, how not only people get access, but then what happens when people's health records are shared or benefits are derived without any consent. Additionally, when patients' health data is being anonymized or pseudo-anonymized and then shared, what sort of conditions should be imposed there? Whether the patient has the right to know or to extend something further, and whether the patient has the right to have some health data deleted. If I say my father is dead. He died yesterday, but he did not want his data records to be maintained in any central database or any hospital, so this is something like the right to be forgotten or the right to have your data deleted when you are not alive. Therefore, the data rights component, or its scope, is expanding because the applications are also expanding, leading to new issues and new challenges, including the right to be forgotten.

So, for developers and AI decision-makers, lots and lots of new queries and new demands arise. So, they need to develop new datasets. They need to come with grips on the health data, with its own unique features, its own risks, as well as the need for better-equipped, less biased, more reliable algorithms and data. And more importantly, the decision-making process or the AI-based system should be something that tries to understand how the physician will interact, how the physician will make a decision, or what sort of health parameters and health data a physician will look at prima facie to come to any decision. So, a physician should also have the right to override any AI decision if the physician feels that the AI decision is prima facie wrong or is not based on a bona fide understanding of the data. So, some of these things can also be discussed in terms of tort law, where the question of righting the wrong in terms of giving compensation or coming up with a broad tort liability scheme could be addressed. So, integrating some of the tort law principles into AI-related health law and then deriving some responsibility, compensation, and liability perspectives from that is feasible. But, as we said, as of now, the health law itself is undergoing changes in bits and pieces in different countries. We will take a look at this shortly in this session. So, we know some principles, we have a lot of cases that have been tried elsewhere, a lot of prior examples, but how to really fit them into health law when the context of AI is a real challenge. Because some of the ideas, like tort, may not be exactly translatable when it comes to dealing with black boxes or systems where accountability cannot be pinpointed to a single person or institution.

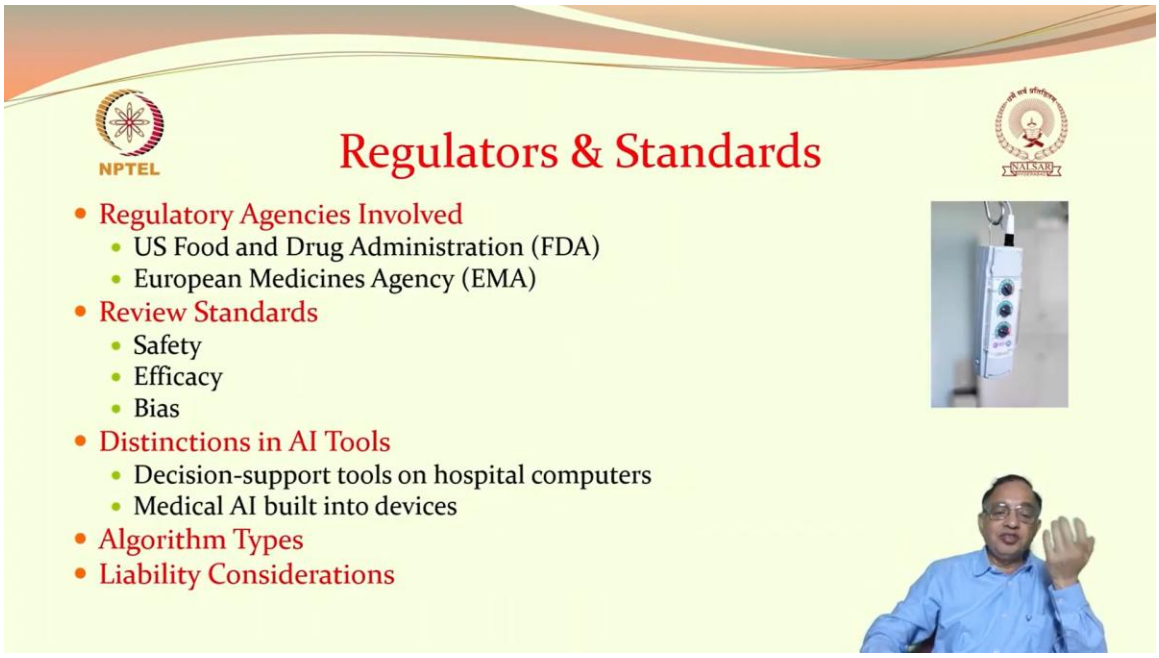

Then there is the question of standards. A lot of regulatory agencies are working on it, and these agencies are developing standards for AI decision tools and AI medical devices. They also consider the whole idea of what sort of algorithms are being developed. And of course, when it comes to regulation, liability considerations are part and parcel of that. A lot of things happen, particularly in the context of regulatory agencies, as they also deal with medical devices, AI-enabled devices, and a whole lot of new approaches, including AI systems entering standard medical healthcare, which is something very novel.

And then we have issues in the hospitals. Hospitals have to deal with a whole lot of new

challenges, including electronic medical records, data governance, sharing of data, and how they should use medical AI reliably, what sort of precautions they should take, and what sort of legal regulations they should adopt; more importantly, they must ensure that the ethical guidelines and norms of the regular health hospital system are not violated in part or full when they deploy AI, so this is an issue of critical importance for health law as well.

Then, as we said, what happens when a physician relies on AI but then also becomes liable? Can a physician take recourse and say, "I relied on AI, so I am not fully liable because ultimately the decision-making was left in the hands of AI, and I only followed the way the system asked me to make the decision"? Then informed consent: when should the patient be informed? Should patients be informed that they should get all the details about AI? Suppose the patient says that I don't want to give informed consent to an AI-based system; I want only a human-based system. So, what sort of exemptions and what sort of limitations should be used? And then, in the case of AI, certain things could not be interpreted, or some things could not be explained at all. So, what are the ethical points we need to consider? And then can we simply say, since AI decisions cannot be interpreted and cannot be explained, we are not liable for it? So, if the hospitals cannot simply tell the patient the diagnosis, the principle behind it, or if the patient dies on account of some medical procedure that was recommended by AI, the whole question is where AI's responsibility stops and starts. Where does the hospital's responsibility stop and start? So, this is a very grey area, but we need to come to grips with it and deal with it as we evolve.

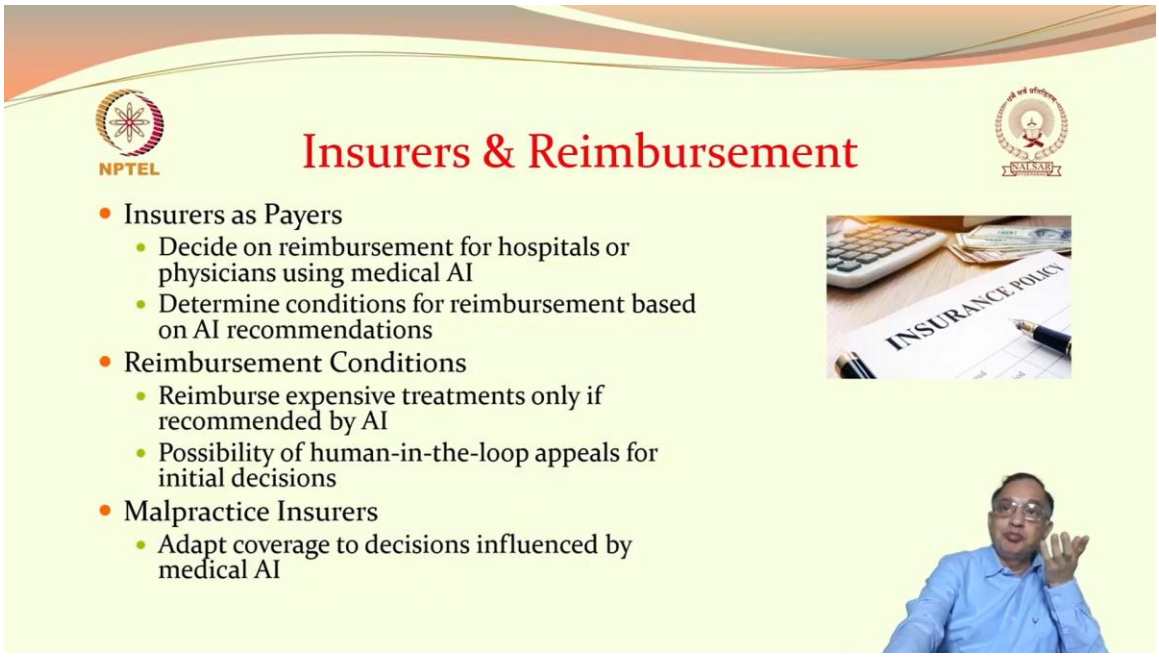

Of course, insurance is a big deal. Health insurance has become more or less inseparable. So, insurance companies may also feel that AI systems' data or the way AI makes decisions should be made available to them, or they can demand that explainable AI models in the healthcare sector be made available to them, or they should also have a say in developing them.

Then we have this whole question of algorithmic decision-making, which again comes in the context of health algorithms being based on problems in data collection, problems in bias in the algorithm itself, or in the data itself. Then there are the AI-related IP matters which we discussed in detail in the previous sessions. So, the policy solutions for

algorithmic discrimination that we discussed earlier can be transposed here to a great extent in the sense that, through algorithmic decision-making, we found out what the major issues could be and how to rectify them. Some of these points that we discussed earlier are also applicable to health law. But then again, health law is a very specialized field that needs to draw from examples in other sectors where AI has been deployed and then try to adopt it as well as adapt it.

Data privacy differs from country to country. In the USA, it is covered under HIPAA. So that means that when data privacy norms differ, and the AI system that is stationed or being used somewhere accesses data from a different country where the data protection norms are totally different, a whole lot of data flow-related governance issues arise. Different countries then deal with it because there are some agreements, trade agreements, that allow certain trade flows in data. There are certain regulations like GDPR and the EU AI Act that are very strict and come with a lot of liabilities and responsibilities. In general, in the absence of any specific restriction, it is presumed that data can be shared or can flow across borders. But some of these developments are too complex in the sense that the rules are not fully clear. Interpreting them again is a very grey area, and how do we identify and de-identify AI strategies? How do we look into data privacy in health research, because health research is again something that is using more and more AI? What rules of clinical tests, what rules of health research ethics, and what laws should be modified to accommodate AI? A lot of things are happening; ICMR guidelines are also getting revised, and the WHO is also revising some of these applicable guidelines. So, things are happening, but whether the pace is adequate is a big question.

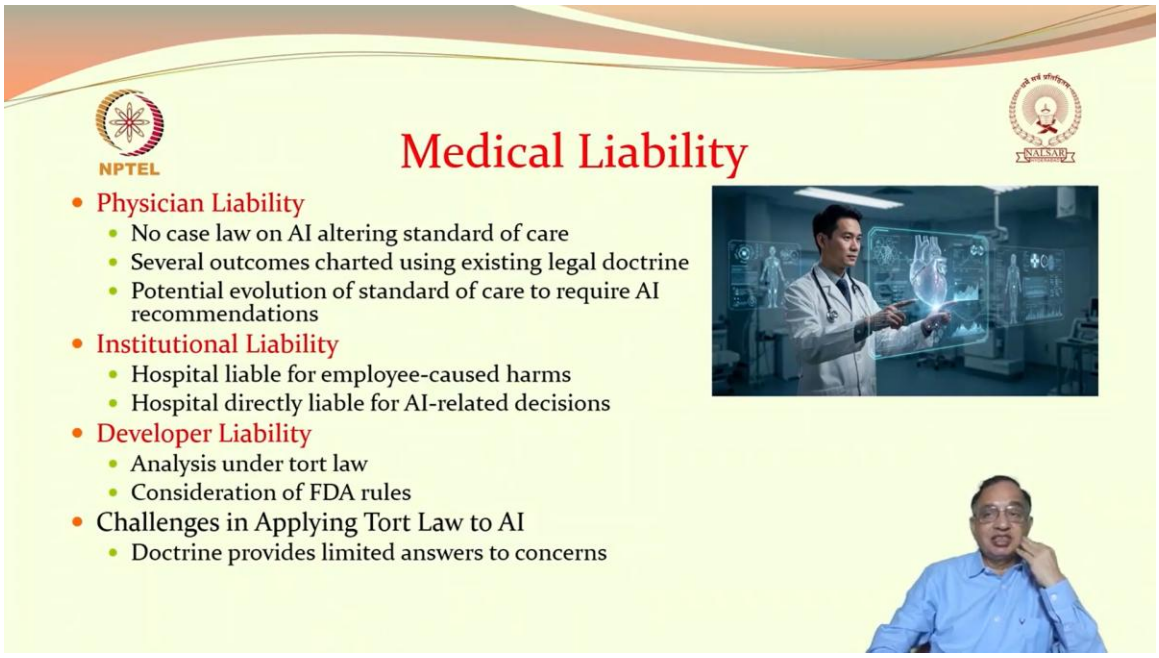

Then the medical liability we discussed could be physician's, institutional, or the developer's - who is really liable, and can we transpose tort law to AI-based decision-making in health law? A lot of things are not clear over there. Because health law is facing many unique challenges from AI.

Then some of the international organizations, such as WHO and OECD, are major players here. The Council of Europe, through its various human rights and health data-related legislations and policies, along with the European Union and the United Nations through its various agencies, including UNESCO, contribute to this effort. Additionally, there are many organizations like PATH, WHO's regional offices, and private

foundations, including the Gates Foundation, as well as major hospitals and universities, which are also significant hubs for medical innovation in AI. Everyone is working on a whole lot of cross-sectoral issues, specifically those related to health law and AI. But then a lot of things are also getting fragmented, and a lot of things are also being done in different directions with a different understanding of AI and medical law. So, the pace at which things are happening is enormous. But whether they are really converging or not is unclear. When the convergence or divergence will happen is something we have no clarity on. However, what could arise or emerge after maybe four or five years is that we may see a lot of good practices, standard operating procedures, guidelines, and then standard laws to deal with many of the problems we face due to data governance and the application of AI in health law.

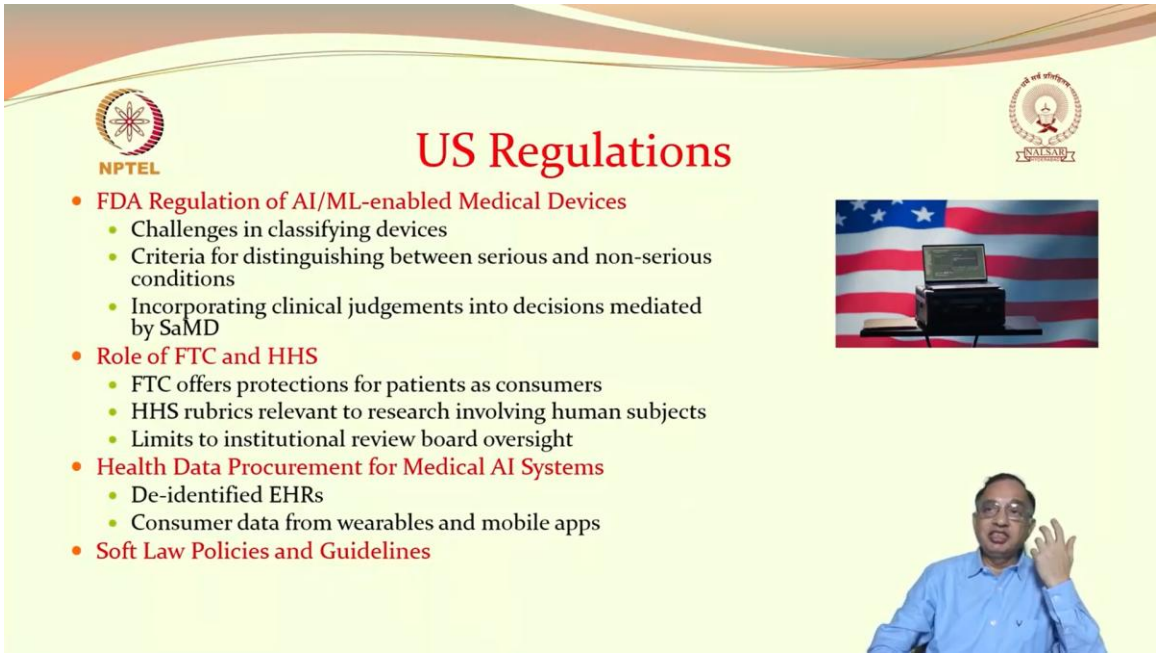

US regulations: the FDA is the major regulator that also regulates artificial intelligence and machine learning devices, medical devices, and then the Federal Trade Commission (FTC), which looks into the monopoly, import, and export of drugs, as well as many other aspects that are part of the larger trade framework. And then, health data procurement for AI systems is again a major issue. Then there are a lot of soft law policies that the health system components, such as universities and major clinics, adopt.

The UK has its own governance mechanism, but it is not fully developed. So, it is relying more on soft law. But the fragmented health regulatory governance is once again a major problem there.

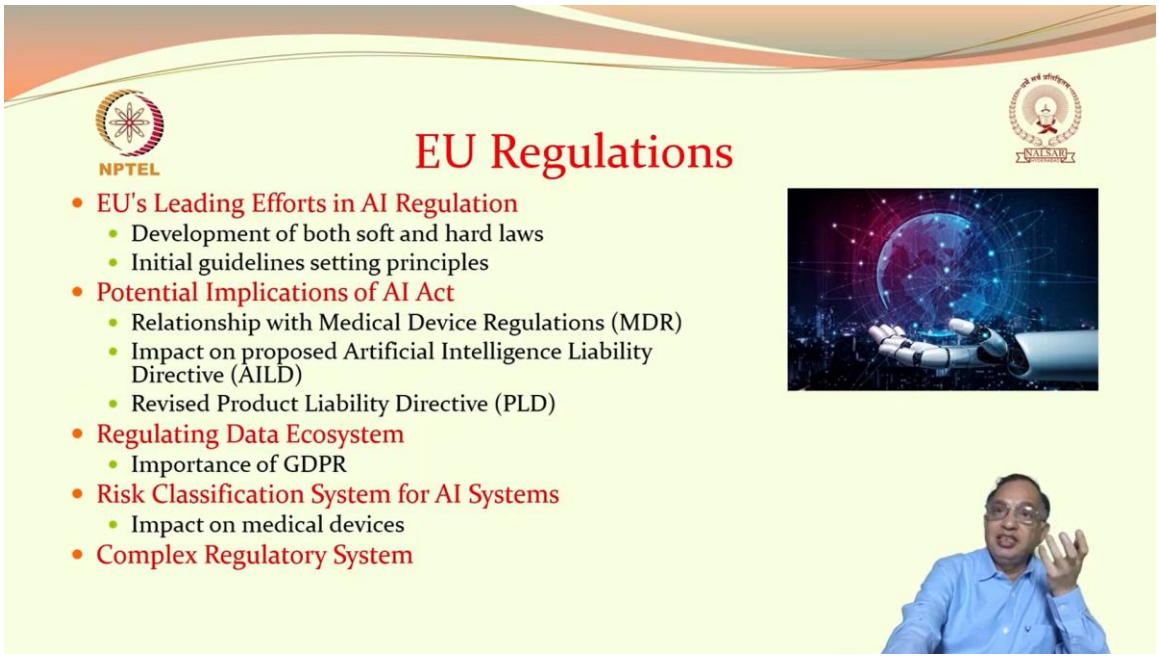

So, the EU is leading in AI regulation, and part of the AI regulation also significantly governs health data through GDPR. But then the EU AI Act has a lot of implications for medical device regulation. It also has implications for automated AI liability and also has a lot of implications for product liability. But then EU AI Act has come into effect. Certain things are not clear in the sense that the broader EU AI Act has a lot of things that are very specifically mentioned; full implementation of the EU AI Act, when it happens

over a period, will reveal to what extent health law is getting impacted by the EU AI Act. But mind you, such a broad EU AI Act, which covers risk, data, and liability, is something that can draw a lot of inspiration for people who want to amend health law to come to terms with AI developments.

Then, of course, there is Singapore's approach, which we will pass through.

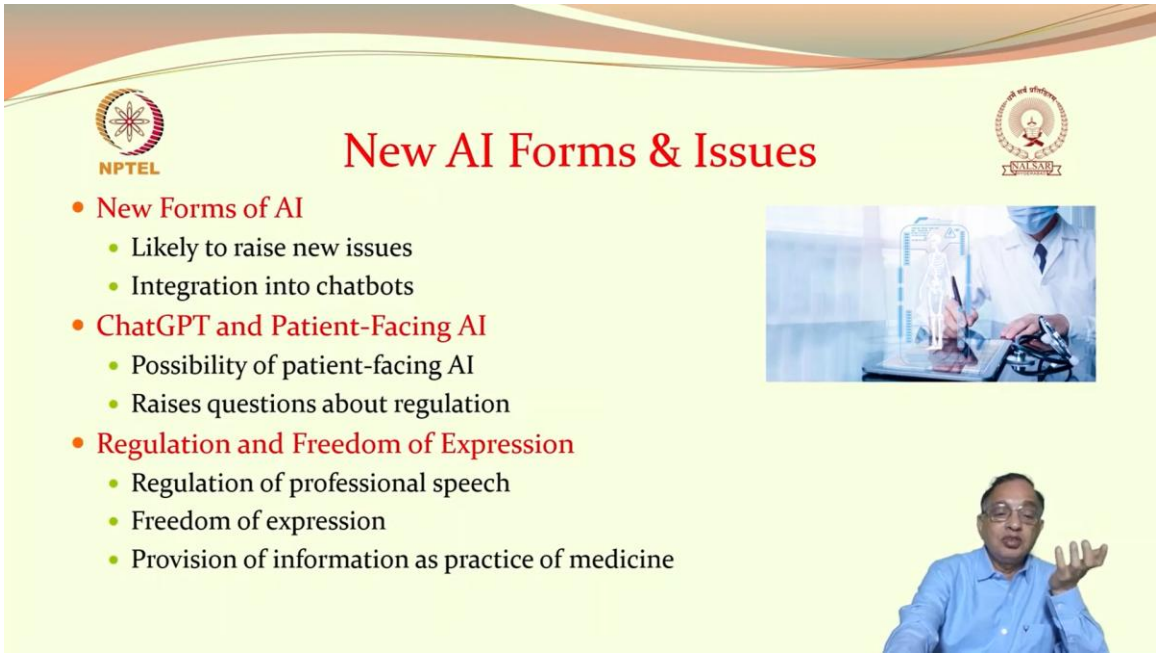

So new AI forms like Generative AI and developments in AI also raise new issues. We saw that when patients use AI tools like ChatGPT, Copilot, or any other tools, they are also not regulated. Of course, ChatGPT and Copilot have come with a lot of disclaimers. But when they give health advice or when people rely on their health advice and then

make wrong decisions, who will be responsible? This is again a major issue. So, when AI advances or when AI is becoming so advanced with new LLMs and new models coming up with much more capability, who exactly will rein in some of the health-related concerns, or who will address them?

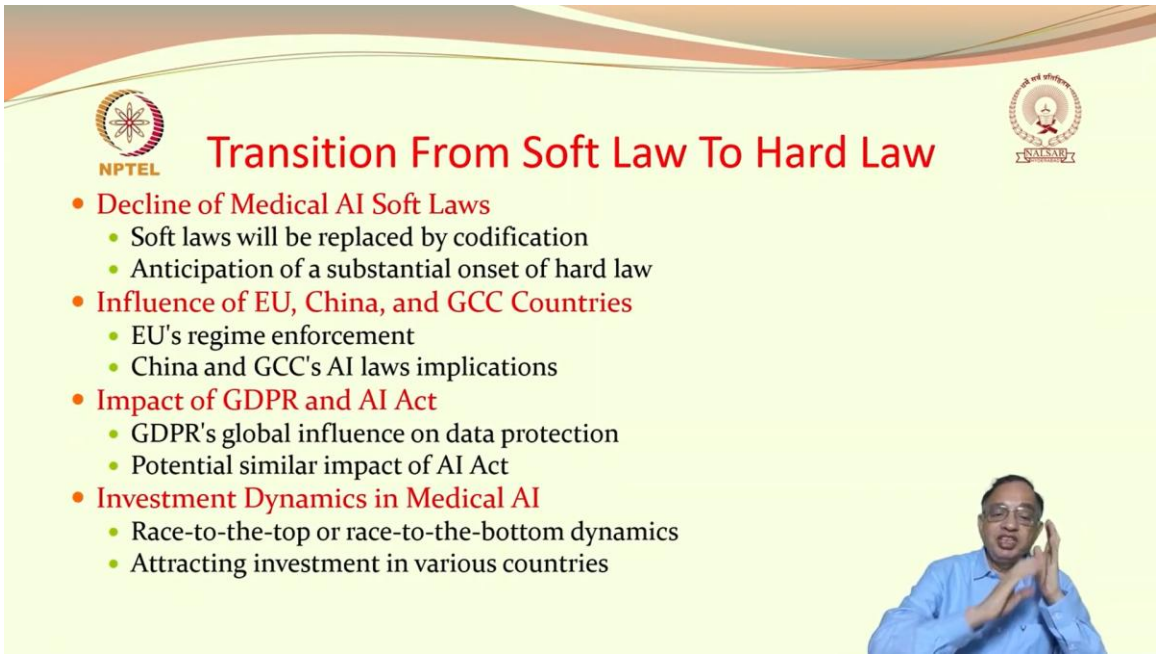

We see the need for soft law, it is obvious, but then soft laws alone won't be sufficient. Right now, soft law plays a major role in the sense that companies, institutions, hospitals, universities, and health regulators understand and abide by some guidelines and notifications, or by well-adopted standards. That is good, but sooner or later, they will be moving towards codification, either as hard law or as part of a larger framework. What countries like the EU and China are doing is very important because the EU is looking at things in one way, while China is also addressing the same issues but from its own perspective. So obviously, I discussed GDPR, the EU AI Act, and the enormous investment that flows into medical AI. So, the medical AI investment also directs innovation in specific directions, and those specific sectors that receive significant venture capital investments will be the ones where medical AI will be deployed more. So those sectors may need more regulation in the future, or health law needs to be adapted more to the sectors where a lot of investment as well as innovation happens.

So, medical AI can be thought of as something very special. There could be special regulations and special laws with their own regulatory authorities, as we have regulatory authorities for different sectors like banking and insurance. So, there can be a medical authority, or there can be an authority for medical AI dealing with the specificities of medical AI, but also interacting with the larger ecosystem in health. But whether such a thing would be required or whether countries would adopt a hybrid approach where the traditional health system for governance also takes into account the medical AI issues not through a specialized agency but as part of a larger health law governance framework is unclear. These are the things for which clarity is lacking, but a lot of developments are happening.

So, in global development, some countries are doing very well; they are trying to come to grips with overcoming fragmentation and move from soft law to hard law. Some countries, like the US, are rather proactive in the sense that their regulation and governance are also in tune, if not one step ahead, of the development. The European Union is doing a lot, but in many smaller countries, where AI is not even in place in a significant way, nothing much is happening. In India, we have some regulatory framework that is emerging on account of the Information Technology Act, the guidelines by ICMR, and the upcoming rules with the full implementation of the DPDP Act. But we don't have a regulatory framework or regulatory system for medical AI, so we don't have a medical AI law. But some of the things that would arise on account of AI being used will have a lot of implications for consumer law. It will also have a lot of implications, particularly for data governance. So how do we address these things, particularly in the context of medical devices, diagnostic treatments, AI-based diagnostic decision-making, or larger issues that need attention?

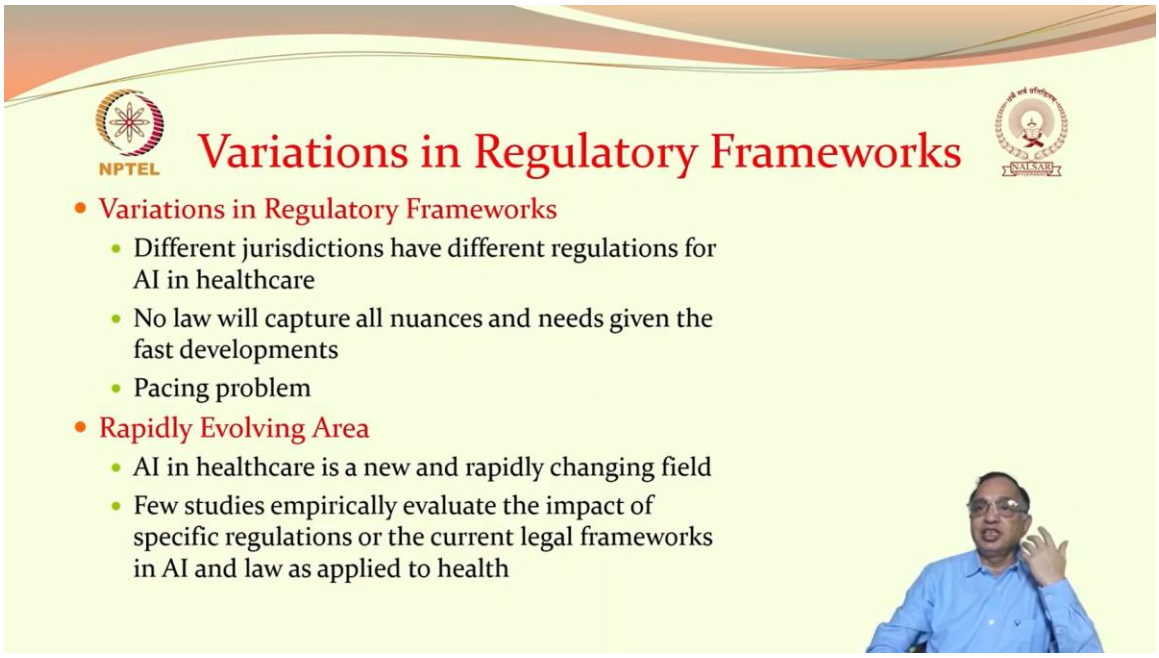

Then, as we said, huge variations exist in the regulatory frameworks, and no law can really come to grips with the developments in technology and be comprehensive enough to tackle all the issues. Law needs to be agile; it needs to adapt itself and keep pace with technological development. Therefore, health law itself would need a rework or a total revisit when AI becomes a major driving tool in health—not just AI, but also the whole datafication of the health ecosystem. If we consider AI as a part of datafication, that is happening in the broader context of the datafication of society. Thus, health law challenges can be understood as part of a larger societal problem of coming to an understanding and dealing with the developments in AI. So, that is the level of challenge we have to face.

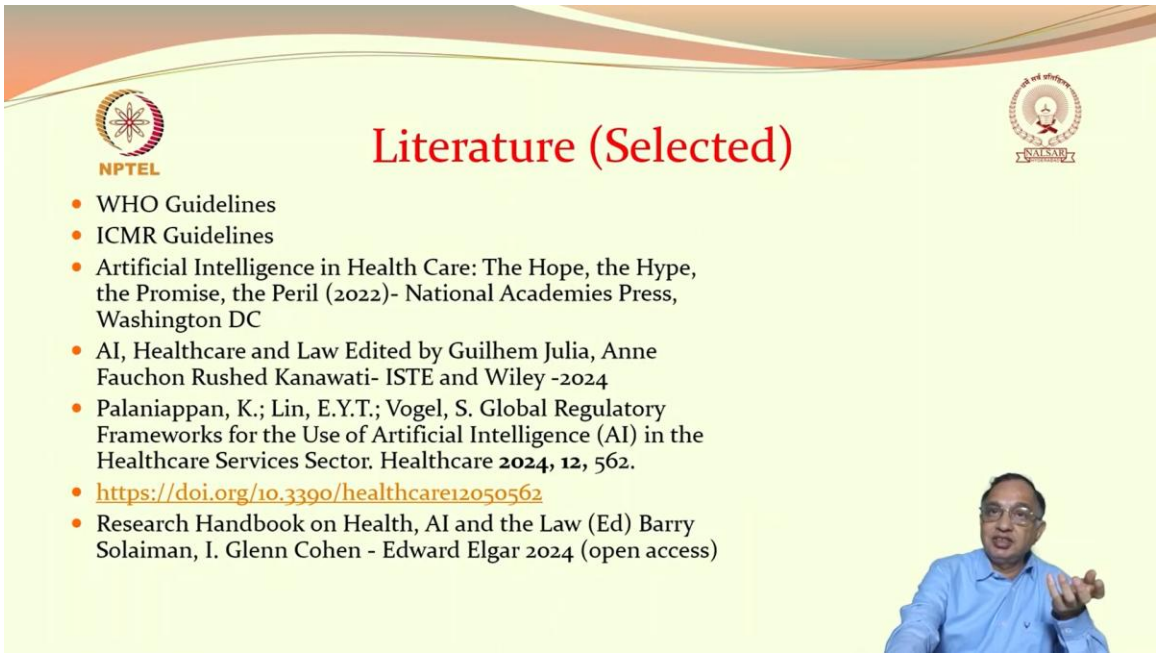

So, I am citing some literature here; of course, the literature I have cited is not exhaustive. I have not even scratched the surface here, but then the topic is huge, emerging, and complex, so we have just barely touched upon some important aspects.

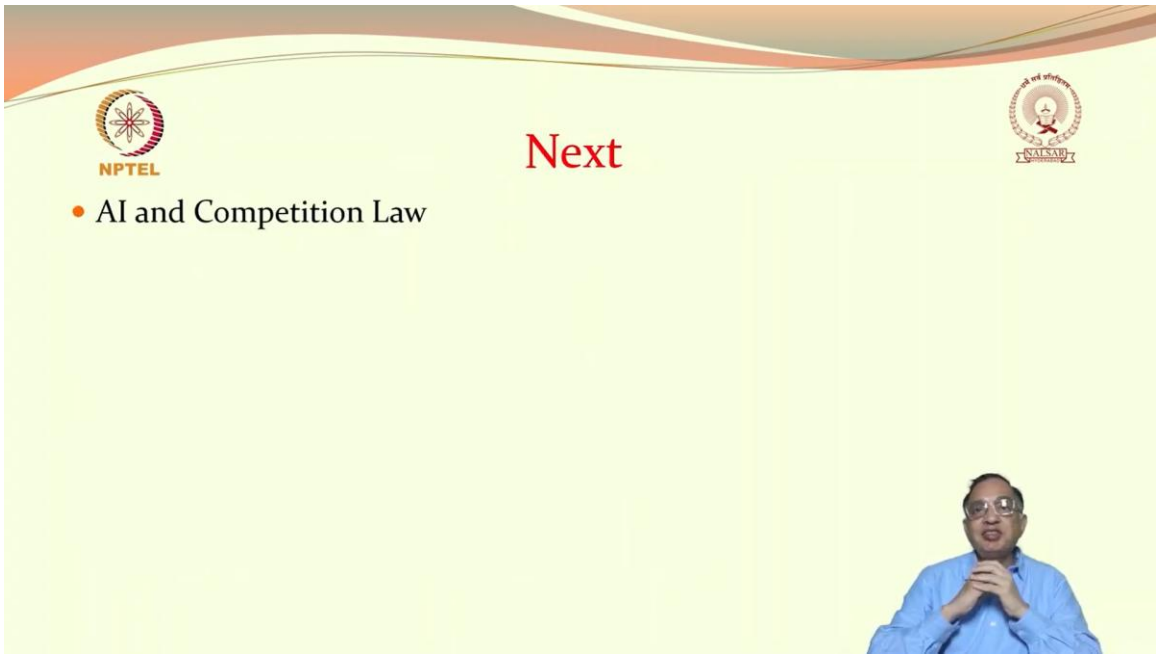

In the next class, we will deal with another part of law or another discipline in law that is AI and competition law. We dealt with IP in the context of patents and copyrights; competition law also has a significant IP component, but then competition law goes beyond that because competition law ultimately is the regulation of the market. It is also the law that ensures that monopolies do not deprive consumers, monopolies do not cheat countries, or monopolies do not override societies. So, AI and competition law are again

emerging as a challenging area, and in our class or session, we will touch upon the key components of that and why competition law challenges through AI are something that competition law will have to constantly keep in step with as AI systems evolve. Thank you.