

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 – 04
Lecture – 17



Artificial Intelligence, Law and Justice

Session 17

Artificial Intelligence and Copyright -Part-IV

Dr. Krishna Ravi Srinivas
Adjunct Professor of Law &
Director, Center of Excellence in Artificial Intelligence and Law
NALSAR University of Law



Artificial Intelligence Law and Justice: Session 17; this is Part 4 of AI and Copyright. In this session, we will look into generative AI and copyright law, how generative AI is creating new issues, and also some problems for copyright and copyright law and practice. Basically, this is a continuation of what we discussed earlier; many of the points will be in continuation, but there are also many new points.

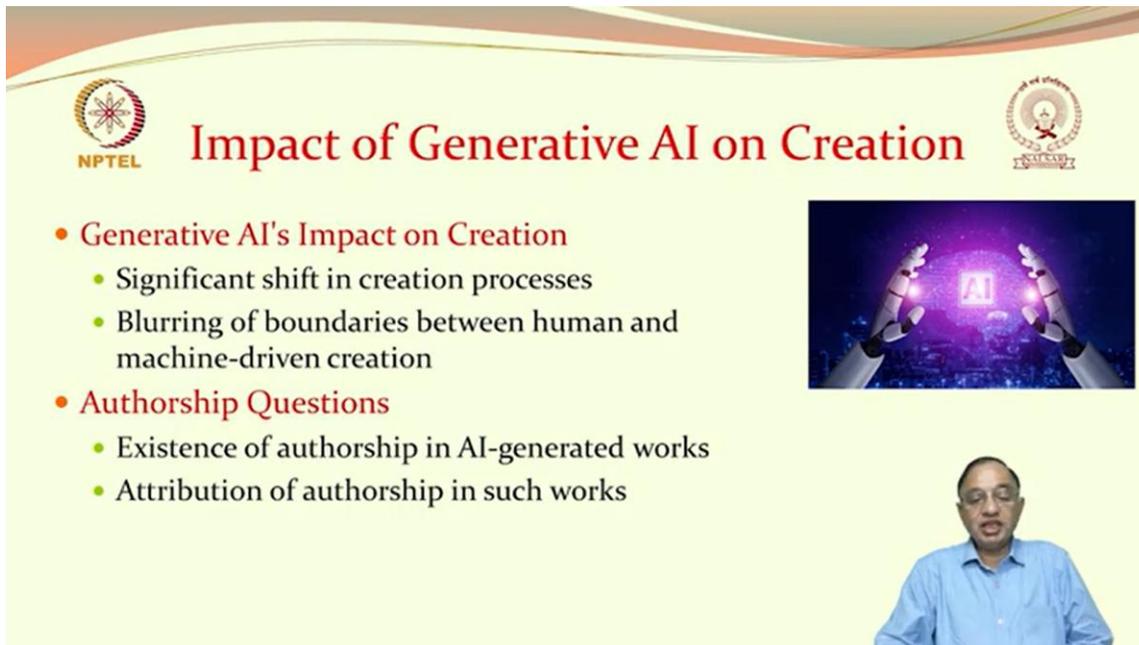


Recap

- We discussed issues including synthetic data
- We also highlighted issues on Text and Data Mining and AI
- We listed how this is being addressed and solutions that are suggested



To give a recap, in the last session, we discussed issues related to synthetic data. We also highlighted the issues in text and data mining and AI. Further, we listed how these are being addressed and what the suggested solutions are.



The slide features a yellow background with a decorative orange and white wave at the top. On the left is the NPTEL logo, and on the right is the logo of Anna University. The title "Impact of Generative AI on Creation" is centered in a large, bold, red font. Below the title, there are two main bullet points in red, each followed by two sub-bullets in black. To the right of the text is a square image showing two white robotic hands holding a glowing purple cube with the letters "AI" on it. At the bottom right of the slide is a small video inset showing a man in a light blue shirt speaking.

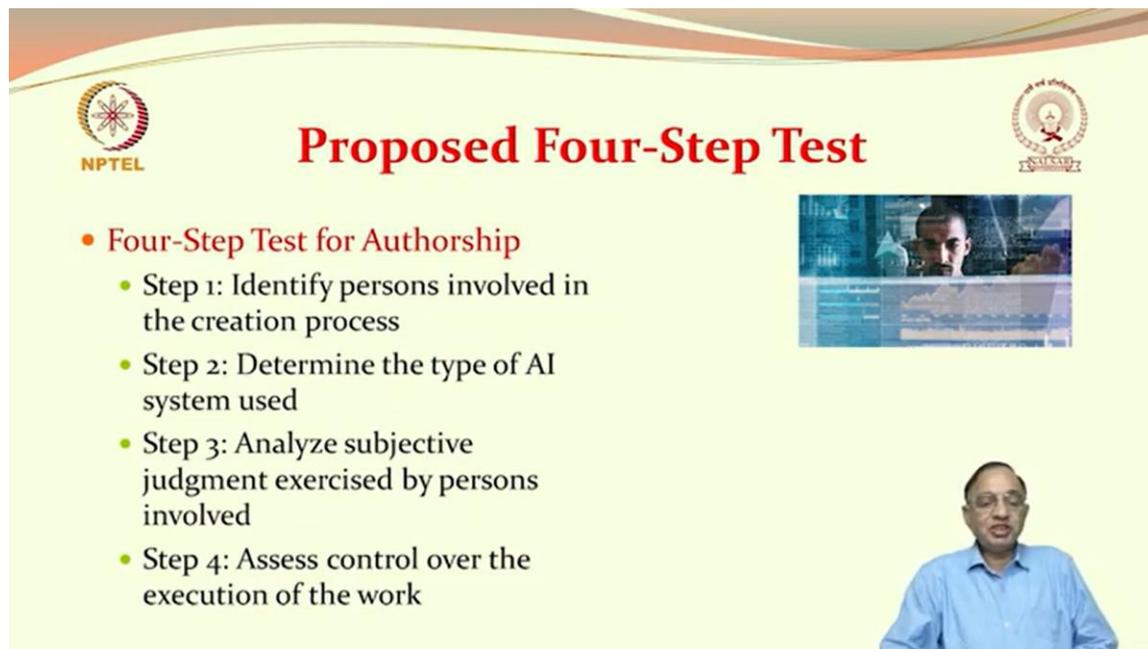
- **Generative AI's Impact on Creation**
 - Significant shift in creation processes
 - Blurring of boundaries between human and machine-driven creation
- **Authorship Questions**
 - Existence of authorship in AI-generated works
 - Attribution of authorship in such works

So, when we talk about generative AI's impact on creation, we all know that its core strength lies in its creative capabilities. So obviously, it has had some impact on creation. How big, how small, how large in the long term, we are not very sure. But it is creating a significant impact on the creative process. Because what it does is blur the boundaries between human and machine-driven creations. Human-driven creation encompasses sculptures, paintings, poems, and photographs—things made by humans; we know about creativity and the creative process. But in generative AI, the blurred boundaries between human and machine-driven creation are creating a whole lot of conceptual problems regarding who exactly is the creator: the machine, system, or AI, or is it human, and then what is their respective contribution to that. Applying specifically to copyright, the question arises: who can be attributed to or assigned the rights as the author? Normally, copyright hinges upon one fundamental concept called author and authorship. In the case of generative AI, since it is very creative and can enable creativity on its own, it can create many things, raising the question of where exactly this label should be fixed and to whom we should attribute it.

There is an existence of authorship in AI-generated work. Attribution of authorship in such works is also possible. But there are a lot of questions. In the sense that if I write a novel using AI, I can acknowledge that AI has been used in writing this novel. But as we have seen, AI cannot claim copyright on its own. I can claim copyright on my own even if I acknowledge that I wrote this novel using an AI system, an AI tool, or whatever tool it is. But neither the persons who provided me with the tool nor the ones who developed the tool can claim authorship because I merely used the tool. On the other hand, by using AI generative tools or generative AI, I can create works endlessly. I need to do some prompt engineering; I need to give it some hints; I need to upload certain files; I need to

ask it to do certain things, and then it will keep generating. So, in one sense, if I can use Generative AI very creatively, I can keep producing output after output, and I can still claim authorship.

But copyright does not depend on the quality of the work. Copyright depends on "the authorship," and then it needs to be established who can be proven or who can claim to be the author. So even for low-quality works that are generated by AI or through generative AI, humans can claim copyright. Of course, it all depends on the market. In the sense that you may generate 100 books with generative AI, but none of them may survive in the market. Nobody may buy it, or they may just become yet another thing done out of curiosity. So, in attributing authorship in such works, one can claim that generative AI created them, but I am still the author. But even if the copyright office accepts it, even if it is there, the proliferation of such works is something that is bound to happen, and it is already happening. AI-generated works are increasing day by day. But they are not exactly the ones for which the AI gets credit for copyright. So, the paradox is that AI may generate thousands of works, but the person who used AI will become the author.



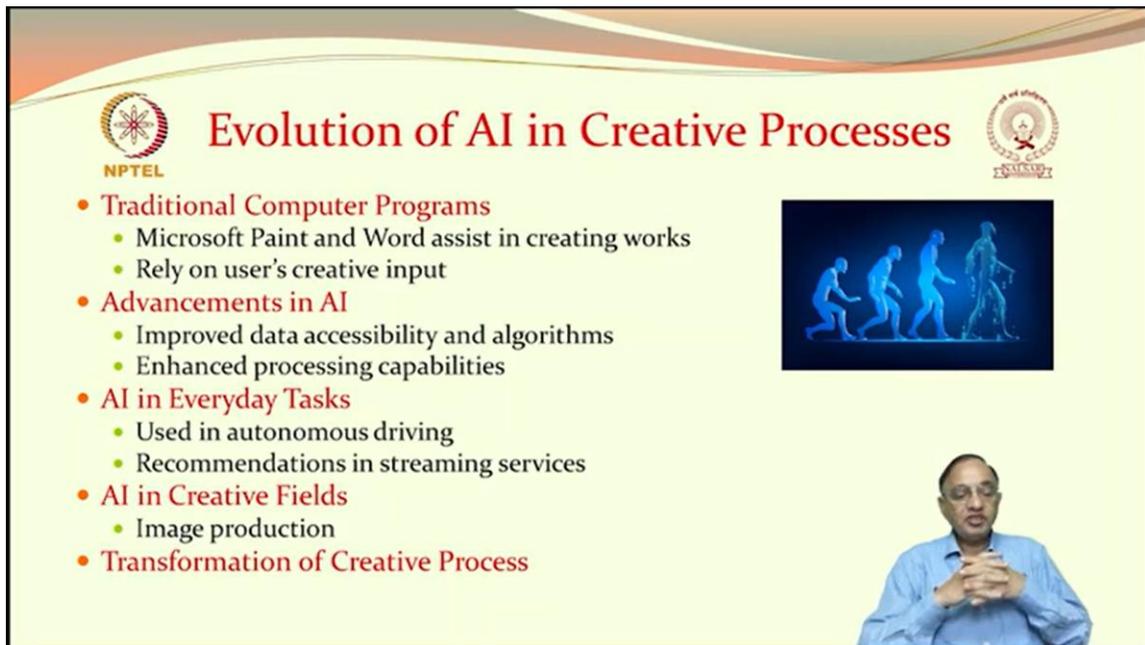
Proposed Four-Step Test

- **Four-Step Test for Authorship**
 - Step 1: Identify persons involved in the creation process
 - Step 2: Determine the type of AI system used
 - Step 3: Analyze subjective judgment exercised by persons involved
 - Step 4: Assess control over the execution of the work

But to avoid this, a test has been proposed: whether AI is the author or who the author is; the question involves a four-step test. The four-step test for authorship is to identify the persons involved in the creation process and then determine the types of AI systems used. Now identify the person, which means it could be a single author or a group. Determine the type of AI system used, as this will help us understand the generative capacity and the creative capacity of the system. Did I use a very elementary AI without much creative capacity, or did I use a sophisticated one? It will help. Then, was the subjective element exercised by the individuals involved? What exactly is a subjective element? To what extent did they edit the processes? To what extent did they help the AI systems? And to what extent are their involvement and contributions in the book and in the work? For example, I wrote a book with the assistance of AI that is 200 pages long, but I didn't

publish it as such. After getting the book output from AI, I start extensively reworking it, revising it, adding paragraphs and text, rewriting the full chapters, and then shuffling the chapters. So, my judgment, my work, and my assessment are also a part of it. I am not simply using what it generated.

And then who is able to access control over the execution of work? Is it I, or is it someone else, or is it the system in the sense that the system itself can simply write, write, write and then produce? So, in that case, who is able to access the control over the execution? Suppose the process is that the system generates, but then I intervene, review, pause the system to re-correct, delete many pages, write again, and then give an output. I am able to exercise control, so this four-step test will tell us who can claim authorship or to whom the authorship can be given.



The slide features a title 'Evolution of AI in Creative Processes' in red text. On the left, there are two logos: NPTEL and the Indian Institute of Technology (IIT) logo. The main content is a bulleted list:

- **Traditional Computer Programs**
 - Microsoft Paint and Word assist in creating works
 - Rely on user's creative input
- **Advancements in AI**
 - Improved data accessibility and algorithms
 - Enhanced processing capabilities
- **AI in Everyday Tasks**
 - Used in autonomous driving
 - Recommendations in streaming services
- **AI in Creative Fields**
 - Image production
- **Transformation of Creative Process**

To the right of the list is a small image showing the evolution of man from an ape-like creature to a modern human. In the bottom right corner, there is a small video inset of a man in a blue shirt speaking.

But to go into this, let us also look at the evolution of AI in the creative processes. Computers have traditionally been used by many painters and artists, starting from Microsoft's Paint and Word, to assist in creative works. They all rely on creative input from the author in the sense that MS Word or Microsoft Paint is more or less driven by us; the system helps us write. But then it does not contribute to generating ideas, giving ideas, or executing things. They are helping us in the sense that we are using them, and then we are painting; we are creating work. So, the creative input and the creative labour come from the person using it. But with the advancement of AI, improved data accessibility, and algorithms, AI's processing capabilities have enormously increased. And now AI is also part of so many things in life, ranging from autonomous driving and recommendations in streaming services to many other things you have already seen; AI is part of that.

So, AI, is it really creative? What exactly do we mean by creativity when AI recommends a streaming service to us? These are big questions because we think that recommender systems are not creative. But someone will say they are also creative because they organize something, decide something, and tell you creatively what you should look at.

And then AI in creative fields such as image production, photography, and filmmaking is also possible. So, the transformation of the creative process through AI has come as part of the larger transformative process, the creative process that started with computers. But we should also keep one thing in mind. Long before the advent of Microsoft Paint and Word, computers were used for creative purposes in different ways. This shows that even in the earlier generations of computers, some sort of creativity was possible. So, it is not that all the creativity in computers started only in the last 30 years.



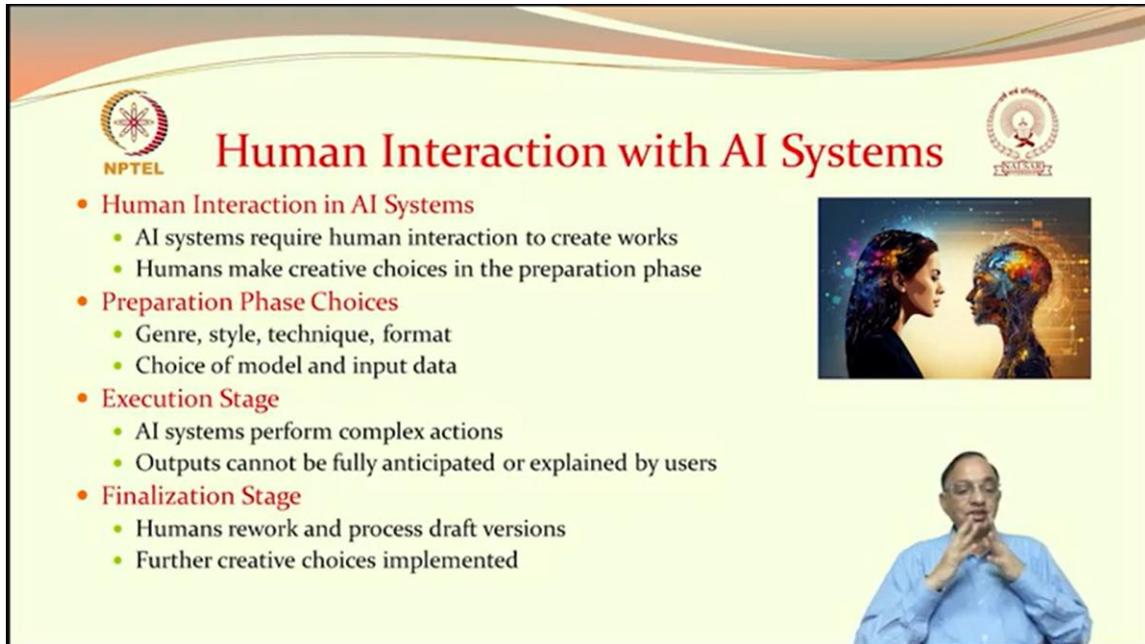
The slide features the NPTEL logo on the left and a circular emblem on the right. The title 'Legal Challenges in Copyright' is centered at the top in red. Below the title, there are three main bullet points, each with sub-bullets. To the right of the text is a small image of a person with a glowing aura. In the bottom right corner, there is a small inset video of a man in a blue shirt speaking.

- **High Autonomy of AI Programs**
 - Raises questions about human intervention in authorship
- **Determining Authorship**
 - Essential for preventing unauthorized use
 - Legal action can be taken if authorship is established
- **Publication of AI-Generated Works**
 - Applies to both artists and everyday users

But the problem with copyright, particularly in the AI context, is that AI programs are very autonomous in the sense that they operate on their own. So, when AI writes a novel, the question is whether we say that there is a need for human intervention, and what the level of human intervention is. But we need to determine authorship for many reasons. Who exactly can be identified as an author, attributed as an author, and given the copyright? Two, authorship is necessary, and it has to be proven or accepted if there is going to be a case of copyright infringement of a previous work or some other issue where it is found that the book is obscene, creates violence, could result in violence, spreads hatred, or speaks derogatorily about women. The problem is that authorship is not a mere copyright question alone; authorship is also a legal question if certain actions are to be taken against the author or if the author is to be prosecuted for something. Authorship should be clearly established, and the court should be provided with proof that it is his content. But when it comes to the publication of AI-generated works, it applies to both artists and ordinary users.

Like my students who use AI to make PowerPoints, it tells them that this was created using it. So, are they the real authors, or is the system the real author? That's a big question. On the other hand, many researchers, students, and people use AI very creatively. But not everyone goes for copyright, and not everyone takes authorship straight away and then says, "I am the author; this is my work," because artists, along with many of those who use AI for creative purposes, only selectively seek copyright

registration. So, AI-generated works could be any number, but then the copyright by default will be applicable to anyone who has generated it; however, those who stake a claim to copyright or those who apply for copyright, and then those who get it registered in their name, are not many compared to the output generated through AI.



The slide is titled "Human Interaction with AI Systems" in a large, bold, red font. It features the NPTEL logo on the top left and the IIT Bombay logo on the top right. The main content is a bulleted list of stages and choices. In the bottom right corner, there is a video inset showing a man in a blue shirt speaking with his hands clasped.

- **Human Interaction in AI Systems**
 - AI systems require human interaction to create works
 - Humans make creative choices in the preparation phase
- **Preparation Phase Choices**
 - Genre, style, technique, format
 - Choice of model and input data
- **Execution Stage**
 - AI systems perform complex actions
 - Outputs cannot be fully anticipated or explained by users
- **Finalization Stage**
 - Humans rework and process draft versions
 - Further creative choices implemented

When we look at it, there is also the question of how we interact with AI systems and who decides what is creative. AI systems need humans to work with them. In the sense that they can do it on their own, but when we intervene, when we talk to them, or when we give input comments, or when we fix the parameters, or, say, read the novels of Fyodor Dostoevsky, read the novels of Jane Austen, read the novels of Chetan Bhagat, now, after reading all this, you write a novel on your own. If I say that, AI will produce, but the creative choice is being made by whom? If the humans make the creative choice and then tell AI, "Look, this is the plot, this is the narrative, and this should imitate the style of Jane Austen, or it should imitate the style of Rowling," I should tell it. So, the creative choice I make in the process is whether I am intervening in it or simply giving a clue to the AI and then starting to work on it.

So, there are a couple of things, like genre, style, technique, and format, and then the choice of the model's input data. So, choices have to be made in the preparation phase itself because if I want to choose or test, I can take three or four generative AI models and then assign the same task to them. I can assign the same task, saying you read all the works of Fyodor Dostoevsky and then come up with something like "The Idiot" in 300 pages; if I assign the same task to different AI systems currently available in the market, the output from them will be very different because, although they are AI systems, although they are LLMs, and although they are fundamentally trained to write, they will come up with totally different works, so the choice of model input data also makes a huge difference. But they perform a lot of complex actions. But as we say and as we have seen many times earlier, if you ask AI why it wrote in this style or why its character committed murder at this stage and escaped, or why its character decided not to take any

position and then committed suicide, AI wouldn't be able to explain this, nor would AI, when it is writing, have been able to anticipate this decision by the character and then go forward as to what should happen next in the sense that AI generates the output without "understanding" or "the narrative capacity to explain." So, outputs are there. If I anticipate something from the AI, that it has read all the books of Fyodor Dostoevsky and is going to give me something like "The Idiot," it may turn out that the book is titled "The Idiot," but then the output could be totally different. It could be a totally different story altogether, not at all similar to "The Idiot." So again: if AI is the author, it should not be able to explain the narrative like an author who says, "I chose this; this is the reason why this character appeared on this stage and then, logically and rationally, this character disappeared from the stage only to reappear later." Now authors have their own narration; authors have their own rationale for each character they can explain, but AI can't. However, when humans rework and repurpose the draft versions and then re-edit, the creative choices are already being implemented there, in the sense that humans take over the creative process at some point, and then they make it for better or worse.



The slide features the NPTEL logo on the left and the Indian Institute of Technology logo on the right. The title is in red. The list of steps is in red and black. There are two images: one of a person's head with a digital brain and another of a man in a blue shirt with his hands clasped.

Implementation of the Four-Step Test

- **Identify Humans Involved**
 - Differentiating between developers and third parties (users)
- **Determine AI System Used**
 - Classify systems into partially, highly, and fully autonomous
- **Consider Subjective Judgment**
 - Free and creative choices by the user
- **Control Over Execution**
 - Ultimate step in the evaluation process

But if we apply the human test, then we need to differentiate between the developers and third parties, and we also need to develop the idea that systems can be categorized into partially, highly, and fully autonomous, because not all AI systems are fully autonomous; some systems are partially autonomous, as they are of the earlier generation, while others are highly autonomous but not fully, and the latest ones are fully autonomous. So, depending on the system's autonomous capacity, the tendency for its authorship rights or authorship claims may be questioned. And more importantly, we can even say that the creativity of those who are partially autonomous is less than that of those who are fully autonomous. But again, it is a question of whether the partially autonomous system itself is sufficient to produce something and claim authorship, or should we say that only the fully autonomous ones should be given authorship? This, again, is not a matter of debate; it is not a matter of contention. It is a matter that delves into the fundamental question of authorship and then into fundamental issues in creativity.

But if the user or the person who "ordered or commissioned the output" makes many free and creative choices in the sense that he or she doesn't fully accept the output, does a lot of editing, shuffles it around, adds certain things, changes the narrative, changes the tone, or changes the active and passive voice of the characters in many places, and then puts everything together, edits it, and tries to come up with a coherent novel. So, in that case, the user or the author makes a lot of subjective judgments, which also involve a creative process. So, the ultimate thing is control over execution. Who controls the execution? Here is the author or the one who commissioned the study. But a highly autonomous system need not even be told by humans or by "the author" to function. A highly autonomous AI can simply generate output given the preconditions, the instructions, or what we would call prompts. So, if you implement the four-step test, we do get some answers, but those answers may not be very clear.



Legal Certainty and Future Adaptations

- **Bridging the Gap Between Traditional Authorship and AI**
 - Promotes legal certainty
- **Developers as Authors**
 - When working with partly or highly autonomous systems
- **Third Party as Authors**
 - When using AI systems created by developers
 - Must work with partly autonomous systems or edit highly autonomous system outputs originally
- **AI-Generated Work**
 - Initial AI-generated work not considered human authorship
 - Edited work may be considered human authorship
- **Fully Autonomous Systems**
- **Case-by-Case Examination**



The problem here is that while the legal serendipity exists, we need to think about the big gap between traditional authorship and AI. We need to promote legal certainty so that developers or authors, when they are working partially or with highly autonomous systems, can decide the logic, the rationale, and set the algorithms that determine how AI will work. So, they know the work that will come out of AI; they have seen the narrative, they have seen how the system is going to work, and they have an idea about it. So, does it mean that developers would have anticipated all the outputs, or can we say that developers know that in the initial stages itself, if the AI is going to write a poem, it will be based only on Emily Dickinson and not anyone else? No, it's very difficult to say. So, developers as authors, yes, feasible, but can they be full authors? They can at once say that they are the ones who set the tone or who set the path, but then they are not the ones who drove the process. So, developers are not the people who drive the process; they enable the process to go forward. Third parties can become authors when they use AI systems. Developers, as authors, must work with partially autonomous outputs or edit the originality of highly autonomous outputs so that they can claim authorship.

But initial AI-generated work is not considered human authorship. Edited work may be considered human authorship because human involvement is substantial. It not only changes the text, content, and direction, but it also involves the human taking control as editor and then taking full responsibility for the output as an author; so here, fully autonomous things, when they are in place, make things very complicated. For example, a fully autonomous one may write a very sophisticated text or a very fantastic thing, but That text in the end may not be something that a typical novel or a traditional novel might include or read. That text may be good, brilliant, and well-written; it might have been able to imitate the style of some authors, but its literary standard or literary quality may not be great.

Because we know that machines, at least as of now, do not have the capacity to think like us, imitate us, or write like us, or at least, if not us, at least like the great authors in literature. On the other hand, in the future, this could become totally different. Autonomous systems are becoming more sophisticated. They may have the capacity to edit, revisit, rework, and revise. In that case, can we say they are authors, so they should be given authorship? That, again, is a very debatable point because we cannot predict how the systems will evolve. But even if you go case by case, it becomes a little easier in one sense; however, it is complicated because exactly proving some of the claims related to authorship and control is not all that simple.

The slide features the NPTEL logo on the left and the IIT Bombay logo on the right. The title is centered at the top in red. Below the title is a bulleted list of points, each with a red header and green sub-bullets. In the bottom right corner, there is a video inset showing a man in a blue shirt speaking with his hands clasped.

Fair Outcomes and Case-by-Case Examination

- **Consistency with EU Copyright System**
 - Results align with the EU copyright framework
 - Ensures fair outcomes for creators
- **General-Purpose AI Systems**
 - Works created without significant effort cannot be attributed to users
 - Further editing to imbue originality is required for attribution
- **Attribution to Developers**
 - Works created using AI systems by developers can be attributed to them
- **Case of Mario Klingemann**
 - Creates artworks through AI system development and utilization
 - Reasonable to attribute resulting works to him as the author

One way is FAIR outcomes and case-by-case examination results; it aligns with the EU copyright framework. The EU copyright framework ensures fair outcomes for creators. This is consistent with the four-step test. But works created without significant effort cannot be attributed to the user. Because if the author does not contribute, or if the author simply gives a prompt and then vanishes, then he cannot or should not be given authorship because his contribution is not great; it is very minimal. At the same time, if the author does a lot of editing, rewrites, and reworks to imbue originality, then it is possible that authorship could be considered. But it is also possible that the attribution

could be given to developers who develop the system if it is shown or proven that the system has been used solely to create literary works like novels, poems, short stories, novellas, or dramatic works. In the case of Mario Klingemann, he created artwork through the development and utilization of AI systems. So, it was reasonable to attribute the resulting works to him as the author because he did system development and utilization, and he was able to prove that he was an author, as attributing the works to him passed the four-step test.



The slide features a yellow background with a decorative orange and white wave at the top. On the left is the NPTEL logo, and on the right is the Indian Institute of Technology Bombay logo. The title 'Ongoing Legal Refinement' is centered in red. A list of three main bullet points is on the left, each with a sub-bullet. An inset image shows a sequence of human evolution from an ape to a glowing digital figure. A video inset at the bottom right shows a man in a blue shirt speaking.

Ongoing Legal Refinement

- **Difficulty in Ascertaining AI Authorship**
 - Complexity in identifying the creator of AI-generated works
- **Future Evolution of AI Technologies**
 - Continuous advancements in AI
- **Need for Ongoing Legal Refinement**
 - Adapting laws to keep up with AI developments

On the other hand, ascertaining AI authorship is problematic as continuous advancements occur; certain things may change, but the need for ongoing legal refinement is also necessary because adapting laws to keep up with AI development is a complex, challenging, and difficult task, not just in copyright but also in various spheres, including patenting, which we will see soon.



How Generative AI Models are Trained

- **Generative AI Models and Deep Learning**
 - Generative AI models like OpenAI's GPT series use deep learning.
 - Deep learning involves neural networks with multiple layers of nodes.
- **Functioning of Deep Learning**
 - Information is processed at different levels of complexity through deep layers.
 - Early layers identify simple patterns; subsequent layers recognize complex patterns.
- **Learning from Large Data Sets**
 - AI models learn from large quantities of data.
 - Patterns and rules identified during training are used to create new content.
- **Data Retention and Memorization**
 - AI models encode patterns as numerical parameters, not storing entire datasets.
 - Generative AI models can sometimes recreate identical or near-identical copies of training data.



So, after all these discussions, we also need to understand, at least in the context of copyright, how generative AI models are trained so that they are able to give such outputs. Generative models like OpenAI's GPT use deep learning. Deep learning involves neural networks with multiple layers of nodes. And then, deep learning information is processed at different levels of complexity through deep layers. Each layer identifies simple patterns, while subsequent layers recognize complex patterns. And then AI models learn from a large amount of data. Further, patterns and rules are identified during training and used to create new content. So, AI creates content from content, learning from what it learns. And then AI models encode patterns or numerical parameters; they don't store the entire datasets. And then, generative AI models can sometimes recreate identical or near-identical copies of training data. So, when the model is being trained, it can create an entire or almost identical sample from the training data. It is also possible that when we want to generate synthetic data, we can use some trained data as a model and ask it to generate synthetic data, specifying that the synthetic data should resemble 70 percent or 90 percent of the original data. We can also give some conditions for generating synthetic data. That is a different topic, but I am saying that it is possible, when we think of generative AI, that even synthetic data can be developed in that way.

AI and Copyright -Part-IV

NPTEL

Example of GPT-3 training data

- **Massive Quantities of Training Data**
 - Generative AI models require extensive data for training
 - Example: GPT-3 by OpenAI
- **Sources of Training Data**
 - Common Crawl dataset
 - WebText2 dataset
 - Two datasets of books
 - English-language Wikipedia
- **Fine-Tuning with Curated Datasets**
 - Further training on smaller, curated datasets
 - Refines model capabilities for specific domains




So, what is the GPT-3 training data? It is a generative model requiring extensive data from OpenAI. We have the Common Crawl dataset, the WebText2 dataset, two datasets, and books from the English language Wikipedia. So, it was again fine-tuned with curated datasets; this is the training data needed for the GPT-3 version.

NPTEL

AI Developers' Increasing Secrecy

- **Shift in Transparency**
 - AI developers have become more secretive about training data
 - Detailed explanations have been replaced by single sentence descriptions
- **Examples of Reduced Transparency**
 - OpenAI disclosed main sources for GPT-3
 - GPT-4's data described vaguely as a mix of public and licensed data
- **Motivations for Secrecy**
 - AI developers have not provided detailed reasons
 - OpenAI cited competitive landscape and safety concerns
 - Sharing details could facilitate replication of models
 - Detailed information could aid malicious actors



But here the problem is that developers are increasingly worried about people knowing about their data sets, how the training was conducted, who trained it, and what exactly the whole lot of training was based upon copyrighted material and non-copyrighted material used without authorization. Because I could have asked for and then received copyright authorization from the copyright holder to train the model. But then I could have used it for different purposes, or I could have used it in excess of what I asked for, or in excess of what rights were granted to me. And there is also the problem that copyright holders

may try to sue if they know that the training process used copyrighted work. So, what do the developers do? Developers do not tell us exactly which ones were used in the training data. So, they have become very secretive. Instead of giving detailed explanations, they provide misleading explanations or a few descriptions.

So GPT-3 had more transparency, but in GPT-4, that transparency is lower. And then we have seen why they are not going for extensive transparency and why they prefer secrecy. There are very good commercial reasons for this. And then there are also some reasons, such as if information is available in full, commercial actors can use it to exploit the system. Commercial actors can fool the system and then try to do certain things by hacking it. So, this increasing secrecy has both elements. One reason it could be done is for safety and to avert risk in the future. It could also be done for commercial purposes and to avoid claims that there was a misappropriation or infringement of copyrighted works.



The slide features the NPTEL logo on the left and the Indian Institute of Technology (IIT) logo on the right. The title 'Arguments Against Transparency' is centered in a large, bold, red font. Below the title, there are three main bullet points, each with a red circular icon and a sub-bullet list:

- **Merits and Objections of AI Training Data Secrecy**
 - Preventing rivals from replicating technology without investing resources
 - Absence of transparency could facilitate anti-competitive practices
- **Concerns Over Dangerous AI Technology**
 - Justifies measures to hinder market entry of competitors
 - Restricting access to training data may not prevent dangerous AI tools
- **Potential Harm from Lack of Transparency**
 - Difficult for regulators to identify harmful or discriminatory behavior
 - AI companies have vested interest in preventing data release

In the bottom right corner, there is a small video inset showing a man in a blue shirt speaking.

So, there are two things: arguments for transparency and arguments against transparency; both are somewhat valid.




Full Access Approach

- **Benefits of Training Data Transparency**
 - Multiple advantages for AI development
 - Arguments against transparency are weak
- **Achieving Transparency**
 - Various approaches provide different levels of information
- **Full Access Approach**
 - Highest degree of transparency
 - Datasets made fully publicly accessible
 - Third parties can verify training data
 - Possible if AI developers retain data post-training
 - Some developers already provide full dataset access online




However, the full access approach has multiple benefits because it helps everyone. Transparency is good for all, and more important data sets being publicly available allows others who want to build LLMs or models to use them, of course after getting all the valid permissions from others. Additionally, third parties can independently verify training data. This independent verification of training data is very useful if we want to know whether the training data is adequately representative, whether it contains any bias or discriminatory data, and whether the training data is given in such a way that it is fully representational, error-free, and of high quality. So, it is also possible that if the data sets are made available and full access is given to those who want to test them, then certain things can be said to be transparent. However, providing full data set access online has problems in the sense that somebody can hack it or use it with malicious intentions.

So, full and transparent access is very desirable. But who should be given full access again is a debatable question. For example, I can say that I can give full access only to trusted public authorities or trusted institutions that I trust in case of litigation, or full access can be given with some conditions or under some restrictive conditions, and then full access can also have non-disclosure, non-binding agreements. For example, after giving full access, I can say that the person who has received full access will not reveal what contents were used, will not reveal where the contents were obtained, and will not use any sensitive information about the data.




Gated Access Approach

- **Logistical Challenges of Full Access**
 - Hosting a large repository of training data is difficult
 - Managing millions of individual works is complex
- **Copyright Issues**
 - Rightsholders may object to their works being freely available
 - Licensing deals do not cover unrestricted access
- **Personal Data Concerns**
 - Datasets may contain sensitive personal information
- **Gated Access as a Solution**
 - Users can request access to restricted parts of datasets
 - Access is granted upon verification of credentials
- **Logistical Issues with Gated Access**



On the other hand, another approach is to give gated access. Gated access is necessary if you have too many volumes of trained data, as managing that is difficult, so right holders may object to their works being freely available. Licensing deals do not cover unrestricted access because when the licensing deal is given to an AI developer, it is granted for a specific purpose. That does not mean that the AI developer has the right to put the data online again or give access to a third party. So, when the data access given by the copyrighted person to the AI developer does not ensure that the system will provide the output or that the AI developer will make the material available in the public domain for all, it could again be restricted. In the case of data sets with personal information, where the copyright owner has granted access for training purposes, a gated approach with a string of conditions is preferable. But still, users can say that they want to check; a gated approach is fine. I have every reason to suspect that you have used 10 books written by me, you have used 15 articles written by me, you have used some of the text from the books I edited, and then I can list out and say I want to only cross-check.

I want to check whether you have relied on them, whether you have used them for training, and what sort of attribution was given. I can also ask the developer to clarify: can the developer guarantee that the attributions to me or the views attributed to me in the output or during the training data are correct? That's also important because if I write 10 books and then my views are attributed to X, Y, and Z, and my views are not attributed to me but somebody else's views are attributed to me for some reason or another, then I have every reason to not only be worried, but I can also sue. In the sense that this wrong attribution is a question I can question not only as an author but also as a contributor to knowledge. So, there are a whole lot of things, but there are also a lot of logistical issues, like giving access, monitoring access, and ensuring that those who have access do not abuse it, because someone can claim that they have a very genuine interest in having access, but that could turn out to be with some malicious intention. So even the gated access approach, although there are myths, has many problems.




Rightsholders' Perspective On Transparency

- **Transparency for Rightsholders**
 - Helps establish if a work appears in a dataset
 - Requires clear explanation of data sources
- **Training Data Summaries**
 - Should identify individual works
 - Explain sources like existing datasets or internet domains
- **Assessing Likelihood of Use**
 - Rightsholders need to assess if their works were used
- **Challenges in Copyright Infringement Claims**
 - Unauthorized use identification is not straightforward
 - Varies significantly between jurisdictions



So, the transparency from the rights holders' perspective is very important. And then they have a stake over there, but there are also issues that are not easy to resolve. So, if we can respect the right holders' perspective on transparency, we should also understand that even if the developers want transparency, there are some very clear-cut valid reasons that do not favour transparency.



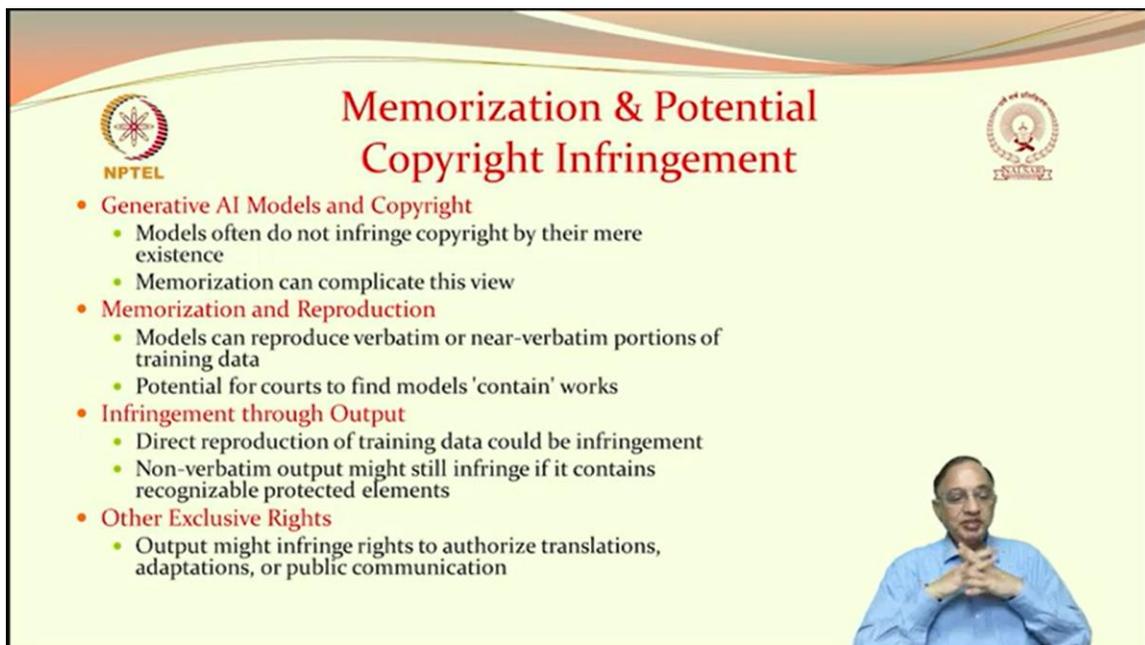

Copyright Issues With AI Training Data

- **Copyright Protection of Training Content**
 - Most generative AI models use content protected by copyright law
 - Copyright arises automatically for works meeting minimal requirements
 - Protection lasts at least the life of the author plus 50 years
- **Alternatives to Copyrighted Works**
 - Public domain works are outdated, mostly from before the 1950s
 - Creative Commons (CC) licensed content is limited
 - CC licenses may impose 'Share Alike' obligations on derivative works
- **Potential of Synthetic Data**
 - AI-generated data could replace real-world data
 - Experts are divided on the feasibility and timeline of synthetic data




The copyright issues with AI training data are many, and as I said, the question of synthetic data is also there, but people are not sure about the feasibility and timeline of synthetic data. We need to be very clear here: when AI training data consists of a lot of public domain works, which may not be outdated, some of the public domain works are very valid and have not only longevity in the sense that they are being repeatedly read and used, but they also still have popularity.

People do buy them, use them, and everyone talks about them if they are classics. So what is protected by copyright law and what is not protected by copyright law in the training data has to be sorted out by the developer so that the developer knows if there is going to be a contingent plan, if there is going to be an injunction when being sued, and who will be able to justify that he has used only those in the public domain; not all material that he has used is protected by copyright. At the same time, the Creative Commons license and what sort of attribution is being given there need to be looked into, not just the Creative Commons license but any such license. Suppose the license forbids any further commercial use, or if the license says that commercial use is permitted provided you share the additions or share the innovations made there on similar terms with others. Then that would become really problematic because no developer would prefer doing it.



The slide features a title in red text: "Memorization & Potential Copyright Infringement". On the left is the NPTEL logo, and on the right is the logo of the Indian Institute of Technology (IIT) Kharagpur. The main content is a bulleted list:

- **Generative AI Models and Copyright**
 - Models often do not infringe copyright by their mere existence
 - Memorization can complicate this view
- **Memorization and Reproduction**
 - Models can reproduce verbatim or near-verbatim portions of training data
 - Potential for courts to find models 'contain' works
- **Infringement through Output**
 - Direct reproduction of training data could be infringement
 - Non-verbatim output might still infringe if it contains recognizable protected elements
- **Other Exclusive Rights**
 - Output might infringe rights to authorize translations, adaptations, or public communication

In the bottom right corner, there is a small video inset showing a man in a blue shirt speaking with his hands clasped.

Now the problem is that there are two things we need to understand. When I read and memorize, it is a mental process. It is something that my mind does. It is not equivalent to digital copying, photographic copying, or Xerox copying; it is not equivalent to "I read the text here, I write it here, I read it here," or "I could type on the computer." My memorization could be partial, could be full, or I may memorize just four sentences of a 20-sentence poem because I like those four sentences. But when computers read, they also remember. This is something that they also see. So, whether that could become an infringement or not is a very debatable point because they do not have copyright by their existence. Memorization does, but memorization again is not something equivalent to reproduction. So, if we talk about reproduction in digital format, exact reproduction, then memorization may not fully meet that. So direct reproduction of training data could be infringement, but non-verbatim also is still infringement if it contains recognizable protected elements, so somewhere we need to segregate in the final output what the recognizable protected elements are and what is not. But the exclusive rights in a copyright also include the right to adapt and the right to translate for public communication. So, these rights could be violated.




Reproduction Right & Copyright Exceptions

- **Reproduction of Training Data**
 - Training data must be reproduced at least once during the training process
 - Considered a prima facie infringement of the right of reproduction
- **Debate on Copyright Protection**
 - Some argue temporary electronic reproduction and 'non-expressive' uses should not be protected
 - If accepted, most AI training copying would fall outside copyright protection
- **Current Assumptions and Legal Stance**
 - Assumption that reproductions during AI training are copyright-relevant
 - Permission from rights holder required unless an exception applies



But again, there is not much clarity in it because the legal copyright regime has yet to come to terms with the whole idea of AI-generated outputs and the way AI is using them.




Challenges in clearing rights for training data

- **Difficulty in Clearing Rights for AI Training Data**
 - Extremely large number of works involved
 - High transaction costs for identifying and negotiating with rightsholders
- **Expense of Paying Licence Fees**
 - Prohibitive costs associated with each work
- **AI Developers' Approach**
 - Using content without identifying or seeking permission
 - Basis for ongoing training data litigation



But clearing rights for training data is also very difficult because who will pay a license fee to them? So, the AI developers prefer to use content without identifying or seeking permission because then they can say that if you prove that I have used it, then we can talk about paying you some money or paying you some sort of remuneration.



Varied Impact of Training Data Transparency



- **Varied Impact of Copyright Law**
 - Legal systems differ in their treatment of copyrighted materials in training data
 - Some require explicit permission, others do not
- **Transparency Requirements**
 - Impact varies between jurisdictions
 - US fair use vs. EU Article 4(3) CDSM exception
- **Pro-Rights holder Advocacy**
 - Assumption of similar outcomes across jurisdictions
 - Flawed reasoning due to differing legal contexts



So, a lot of litigation takes place here. The varied impact of training data transparency on copyright is very important. So, U.S. fair use and EU Article 4.3 of the CDSM are very different. So, the point here is that since the legal system itself is not really clear on what fair use is or what is meant by the fair use exception in the U.S. context, a whole lot of problems will arise in the context of AI tools developed in one country, adapted in another country, reused, or made better in yet another country. So problematic assumptions will also be present because if I am going to use the fair use doctrine of the European Union in India, it won't work. The Indian one is much broader; it is also much wider, and it is based on our reading of the copyright, where certain things are totally permitted for non-commercial academic purposes. The US fair use restriction or right is very limited and restrictive. So, if I transpose one over the other in different contexts and in different countries, there will be chaos. But there is no universal definition, nor is there a universal standard for fair use. So, this means that AI developers will have to be sensitive to the copyright exceptions and then exemptions in different jurisdictions.




Training Data Transparency in the AI Act

- **Impact of Transparency Requirements**
 - Determined by existing copyright laws
 - Limitations without revising those laws
- **Transparency Provisions in AI Act**
 - Facilitate opt-out mechanism in Article 4(3) of CDSM Directive
 - Provide information on sources of training data
 - Unlikely to improve rightsholders' position due to flaws in opt-out
- **Article 53(i)(c) Requirements**
 - Providers must comply with Union law on copyright and related rights
 - Identify and comply with reservation of rights in Article 4(3) of Directive (EU) 2019/790
- **Article 53(i)(d) Requirements**
 - Providers must make publicly available a detailed summary of training content



The AI Act of the European Union tries to address this in detail, and I have listed it here. I am not going to explain in detail because we are not going to discuss the AI Act in detail.




Benefits & Limitations of Transparency Requirements

- **Benefits of Training Data Transparency**
 - Clear advantages in understanding data usage
 - Promotes accountability and trust
- **Dependency on Local Copyright Law**
 - Varied outcomes across different jurisdictions
 - Local laws significantly influence the impact
- **Challenges with Copyrighted Materials**
 - Complex issues remain unresolved
 - Transparency alone is not a complete solution



So, there are some benefits and limitations of transparency requirements that we have already described.



Need For Broader Policy Measures



- **Importance of Training Data Transparency**
 - Advantages of transparency requirements
- **Need for Policy Focus Beyond Transparency**
 - Amending copyright law to balance interests
 - Varied approaches based on jurisdiction
- **Protecting Rights holders vs. Promoting AI Development**
 - Further measures for rights holders in some cases
 - Ensuring copyright law does not hinder AI in others
- **Engagement with Key Stakeholders**
 - Determining effective policies for specific contexts
- **Frequent Reassessment of Policies**




So, what we ultimately need are broader policy measures; we need to balance the rights of holders versus promoting AI development, and there needs to be an effective policy stakeholder engagement process. Amending copyright law in a rush will not work. Some sort of varied approaches should be taken, and then more importantly, depending upon the jurisdictions, there could be some tailor-made exemptions. For example, if it is India, applying the same broad fair use doctrine, lock, stock, and barrel will not and should not be a good policy option because the original intention of the broad fair use was not to develop AI systems and was certainly not to develop them for commercial purposes; it was done for the specific purpose of promoting academic scientific research and then for educational purposes. So, extending that to AI systems in the broader context in India will be something that does injustice to the primary objective of fair use in India. Similarly, we need to go into each country and then trace the roots, identify the cases, identify the rationale behind the exemptions, and then see how the whole idea was developed and arrived at. So, unless we contextualize, we cannot think in terms of broader policy measures or exemptions. Of course, Europeans have come up with certain things that may or may not be suitable in other contexts, but only through a comparative analysis of the Fair Use Doctrine, particularly when it is applied to AI, can we come up with some tangible, relatable, rational policies.




Future of copyright law in the context of AI

- **Challenges posed by Generative AI**
 - Comparison with past technological advances can help in crafting policies but some issues are novel
 - Pacing problem in technology and regulation as example
 - Not just copyright – IP and AI – IP for AI ?
- **Need for Adaptation**
 - Copyright law must adapt to new technologies
 - Importance of decisive action by policymakers
- **Crafting Effective Policy**
 - Difficulty in regulating new technology during early stages
 - Too difficult later -
 - Importance of understanding societal impacts
- **Issue of Entrenchment and**
 - Rapid integration of generative AI by major tech companies
- **Training Data Transparency**



So, there are a lot of challenges, and we really need to think about many things regarding the future of copyright law in the context of AI.




Balancing Innovation & IP Protection

- **Balancing Innovation and Intellectual Property Protection**
 - Exploring the impact of generative AI on copyright laws
 - Examining the need for updated copyright exceptions
- **Generative AI and Copyright**
 - Understanding the challenges posed by generative AI
 - Identifying potential solutions for copyright protection
- **Future Directions**
 - Proposing new frameworks for copyright exceptions
 - Ensuring a balance between innovation and protection



And then the whole question of balancing innovation and IP protection is the big elephant. It is in the patents. It is there as well in copyright. It is there whenever we try to link any innovation with IP rights, particularly in the case of technologies like AI.



Next

- Session 18 and 19 will be on AI and Patents



So, in the next session, we will talk about AI and patents. In these four sessions, we have touched upon the broad issues of copyright and AI. We have also gone extensively into many topics, but let me caution you that we have simply scratched the surface. We have not gotten into extensive detail on many complex issues for the simple reason that it is beyond the scope of this course. True, it is also uncertainty; illegal uncertainty being prevalent, it is also not fair at this point in time when offering a course like this to put certain things as conclusions or conclusive evidence for one policy or against another, or even to say that most probably in the future that copyright policy will move in this direction. Even to guess and then state such a thing may not be fully correct because policy options are evolving; they are not fully exhausted, and what the EU does or what other countries will do will also change. Things may change depending on how countries want to promote AI and how copyright stakeholders want to negotiate. So, these are complex matters, but trying to understand even the complex matters in a session like four is a necessary beginning but not sufficient to get an advanced idea of the developments happening in copyright and AI, so we will devote the next two sessions to AI and patents.