LIBRARY OF CONGRESS

+ + + + +

UNITED STATES COPYRIGHT OFFICE

+ + + + +

COPYRIGHT IN THE AGE OF ARTIFICIAL INTELLIGENCE

+ + + + +

WEDNESDAY
FEBRUARY 5, 2020

+ + + + +

The meeting convened in the Montpelier Room, 101 Independence Avenue SE, Washington, D.C. 20540 at 9:00 a.m.

PRESENT
MARIA STRONG, Acting Register of Copyrights and Director, U.S. Copyright Office
CATHERINE ZALLER ROWLAND, Associate Register of Copyrights and Director of Public Information and Education, U.S. Copyright Office
ROB KASUNIC, Associate Register of Copyrights and Director of Registration Policy and Practice, U.S. Copyright Office
JOHN ASHLEY, Chief, Visual Arts Division, U.S. Copyright Office
KATIE ALVAREZ, Counsel for Policy and International Affairs, U.S. Copyright Office
REGAN SMITH, General Counsel and Associate Register of Copyrights, U.S. Copyright Office
WHITNEY LEVANDUSKY, Attorney-Advisor, Office of Public Information and Education, U.S. Copyright Office

MARK GRAY, Attorney-Advisor, Office of the General Counsel, U.S. Copyright Office

ALSO PRESENT

FRANCIS GURRY, Director General, World Intellectual Property Organization (WIPO)
ANDREI IANCU, Under Secretary of Commerce for Intellectual Property, and Director, U.S Patent and Trademark Office
AHMED ELGAMMAL, Professor at the Department of Computer Science, Rutgers University, and Director of The Art & Artificial Intelligence Lab
ROS LYNCH, Director, Copyright & IP Enforcement, U.K. Intellectual Property Office (UKIPO)
ULRIKE TILL, Division of Artificial Intelligence Policy, WIPO
MICHELE WOODS, Director, Copyright Law Division, WIPO
SANDRA AISTARS, Clinical Professor and Senior Scholar and Director of Copyright Research and Policy of CPIP, Antonin Scalia Law School, George Mason University
ANDRES GUADAMUZ, Senior Lecturer in Intellectual Property Law, University of Sussex and Editor in Chief of the Journal of World Intellectual Property
JASON BOOG, West Coast correspondent for Publishers Weekly
KAYLA PAGE, Senior Counsel, Epic Games
MARY RASENBERGER, Executive Director, the Authors Guild and Authors Guild Foundation
MEREDITH ROSE, Policy Counsel, Public Knowledge
JOEL DOUEK, Cofounder of EccoVR, West Coast creative director and chief scientist for Man Made Music, and board member of the Society of Composers & Lyricists
E. MICHAEL HARRINGTON, Composer, Musician, Consultant, and Professor in Music Copyright and Intellectual Property Matters at Berklee Online

DAVID HUGHES, Chief Technology Officer, Recording Industry Association of America (RIAA)
ALEX MITCHELL, Founder and CEO, Boomy
AMANDA LEVENDOWSKI, Associate Professor of Law and founding Director of the Intellectual Property and Information Policy (iPIP) Clinic, Georgetown Law
MIRIAM VOGEL, Executive Director, EqualAI
JULIE BABAYAN, Senior Manager, Government Relations and Public Policy, Adobe
VANESSA BAILEY, Global Director of Intellectual Property Policy, Intel Corporation
MELODY DRUMMOND HANSEN, Partner and Chair, Automated & Connected Vehicles, O'Melveny & Myers LLP
SARAH HOWES, Director and Counsel, Government Affairs and Public Policy, SAG-AFTRA
IAN SLOTIN, SVP, Intellectual Property, NBCUniversal

C-O-N-T-E-N-T-S

C-O-N-T-E-N-T-S

P-R-O-C-E-E-D-I-N-G-S

9:09 a.m.

MS. ROWLAND: Thank you all for coming here today. We really appreciate it. We are really excited to be here with WIPO co-sponsoring this event on Copyright and the Age of Artificial Intelligence.

We did Part I in Geneva back in September, and here we have Part II where we are taking a deeper dive into all things copyright. We have a really exciting day planned and we have some incredible guests, like the Director General of WIPO, among many others.

Right now I'm going to invite Maria Strong, who is Acting Register of the Copyright Office, to introduce the Director General.

MS. STRONG: So, thank you, everybody, and welcome to the United States Copyright Office. It's our privilege to be hosting this continuing conversation on artificial intelligence.

I would like to thank all of our

colleagues at WIPO who are working on this, our colleagues in the Copyright Office, and our colleagues also in the Library of Congress who made this event possible.

I think there's going to be a lot of questions that are going to be raised during the day. I can't promise we'll have all the answers, but it's an important topic that crosses many in the copyright community.

With that, I would like to invite up to the stadium -- the podium -- it feels like a stadium in here -- Director General Francis Gurry. Thank you so much.

(Applause.)

MR. GURRY: Thank you very much, Maria.

Ladies and gentlemen, a very good morning to you all. It is a pleasure and a privilege to be here in Washington, the great city of Washington, for this event.

Let me start by thanking the United States Copyright Office, and the Library of

Congress also represented here. Thank you for all of your support, for reaching out to us, and for continuing this conversation here in Washington. We are very grateful for Ms. Maria Strong and all of her colleagues. Thank you very much indeed.

It's great to see -- I haven't seen him yet, but I know Andrei Iancu is going to be here this morning, the Under Secretary for Intellectual Property at the United States Patent and Trademark Office. And I'm pleased to see so many old friends as well.

Allow me to introduce two of my colleagues who you are going to see later in the morning. Michele Woods, known to many of you in the Copyright Office, of course, who is the Director of our Copyright Law Division. And Ulrike Till, who is newly Director of our Artificial Intelligence Policy Division.

So this is a very timely event, an extremely timely event, because, as we are all aware, we are increasingly seeing artificial intelligence applications deployed across the economy and in the creative industries. It's a little difficult to keep track of it always, but we know that there are many instances of this.

Just in a different field, I saw that last week, in the United Kingdom -- and I use the words of the proprietors, with inverted commas, of the exercise -- the first entirely invented, machine invented, molecule has gone into human clinical trials. So it is quite clear that machine creation and machine invention is with us. So it's good that we can discuss these issues.

And I think the existence or the appearance of all of these applications that we're seeing across the economy is raising, of course, many policy questions. We've seen lots of papers. We've seen, in some cases, some executive decisions, and we are seeing judicial decisions emerge. And that creates a situation, I think, in which we are facing two risks in the area.

I think the first risk for us in our reasonably small world of intellectual property is that artificial intelligence is, of course, a general purpose technology and it is being deployed across every sector, and has applications in every sector of the economy, as well as social applications and military applications, of course.

And I think in those circumstances, from a policy point of view, we face some confronting questions, and there is a risk that the copyright questions, or the property questions in general, get drowned out by the urgency and the importance of some of the other questions that we're seeing in the area of, for example, privacy or security or data integrity.

And I think, for all of us who care about copyright, it's very important that we assert the role and importance of property, which has, after all, historically been such -- whether real property or intellectual property -- such an extremely important institution for the organization of the market and the economy.

The second risk, I think, is a risk of incoherence. The policy responses we see emerging now, for those of us who are trained in and comfortable with the common law tradition, we're not so upset by the slow movement of the gradual evolution of a policy response, but that is not the whole world. And I think here in this instance we have two differences from other preceding situations. The first is, of course, the speed of technological development. I'm not sure we have the luxury for that slow evolution of a responsive policy response.

The second is interconnection and the likelihood that artificial intelligence will be all over production and distribution in the economy before we have a coherent international response which, as you all know, takes years and years and years.

At the same time we know that artificial intelligence capacity is at the center, or is perceived to be at the center, of competitor capacity and competitiveness in the contemporary economy.

I think there's a danger here that we will see that competition that is there spill over into regulation and that we will see regulatory, and I think we are already seeing, regulatory competition.

Those who have scale, and those who are the first movers, have the great advantage in an interconnected world when dealing with an issue of connectivity of imposing the global rule.

This risk of incoherence is something that I think we all need to be aware of. We are especially grateful, therefore, to the Copyright Office and the United States Government for leading this conversation and leading us in the direction of finding some solutions.

Let me say a word about the WIPO process of which this is part. We have turned deliberately to conversation because we would like to allay any fears that there is going to be any abrupt or sudden or precipitant move to establish international rulemaking in this area.

This is not the idea at all. We are simply wishing to share experience with the hope that dialogue on the basis of shared experience would inform national positions in a more harmonious way in this highly intense competitive environment where artificial intelligence is at the center.

It's very much a conversation, and I don't know about you, but I certainly don't have any answers in this area. I find it extremely difficult, extremely challenging and extremely complex. So let me go to those questions -- some of those questions.

I would like to make really just some comments, three sets of comments. Many of these issues are sub judice. They are the subject of litigation, so I will not be making terribly many suggestions but I think there are two general sets of questions that we need to address.

The first is the interaction, or the

impact of, artificial intelligence on our current copyright system and let's confine it to copyright here. The second, I think, is a much more challenging question which is whether we need to go further and the possibility of new approaches or new ways of dealing with this phenomenon of machine creation.

Dealing with the first, I think here I would just like to mention two -- I know there are many issues -- two issues that I think are at the center of the interaction of artificial intelligence and the deployment of artificial intelligence applications with the existing copyright system.

Here I think the first issue naturally is authorship. I was very interested to read coming over two decisions of the Chinese courts on this. One was the decision of the Beijing Internet Court, and the other is a decision of a provincial court in Guangzhou Province.

I'm not going to go through the exact details of the decisions but the decisions do, I think, bear reading because they have taken an approach to the question of authorship where I think that they are looking at the creative process in a very comprehensive manner.

They come to their decisions ultimately that a machine creation, an artificial intelligence creation, which is original is eligible for copyright protection because when you look at the entire and the comprehensive creative process, then you have a human being that is involved in this process and it's a question of identifying the dominant human being. I think that's a very interesting approach.

I think we're going to see a cleavage between those countries and systems that are going to answer this on the basis of the proximate inventor, the immediate inventor, looking at it this way, as opposed to a more comprehensive view of the whole creative process which might give you a human being and might give you an author for the purposes of copyright protection.

The other, I think, huge question that we face in the area of copyright, of course, is the use of copyright works in data feeds to algorithms for artificial intelligence creations. Here this is obviously an issue that is raised, if not directly, then certainly obliquely, in the Amazon Audible case. We need to wait for the court's approach and decision in that litigation. There are lots of views about it.

This is a very fundamental question, I think, because you have many government policies at extremely high levels encouraging the free flow of data; the free flow of data in the interest of the development of artificial intelligence capacity and its spread across the economy.

You see in Japanese legislation an exception with respect to this which I think is something that we need to be very careful about. Can we take the whole repertoire of music of a music producer and feed it to an algorithm without that being a copyright infringement.

Of course, we will get some guidance on this from the Amazon Audible litigation. We won't necessarily get an answer to the broad policy question but we will certainly get, I think, some answer.

I think a somewhat similar question arises with respect to the use of deep fake technology in audiovisual performances. I am told -- I am not an expert but I am told that it is possible to construct an actor's role without the actor playing the role on the basis of the data that is gleaned from the preceding performances, or appearances, or photographs, or sounds, or images of the actor concerned.

Now, this is a very fundamental question also and I think it's very related to the first question of infringement of copyright by the use of copyrighted works in data feeds to algorithms. I don't know the answer and whether this is something that is going to require permission.

Of course there are other legal

remedies that will come into play. Privacy might come into play. Defamation might come into play. The whole deep fake question is a huge question for society in many, many ways.

This brings us back, I think, to the point that we really have to in the copyright community impose the importance of the property question in the policy considerations in relation to all of the questions that we are now seeing.

Those are just some fairly random comments on, I think, this first set of questions, I know there are many others, about the interaction of the copyright -- of artificial intelligence and the impact of artificial intelligence on the copyright system.

I would like to go to the second -- say a few words on the second set of questions which is, and the way I would frame it is really as follows: When we look at the copyright system, as we all know, and there are many experts in the audience, it's a product really, or evolved out of, printing.

Printing actually was really the first industrial process many centuries before the industrial revolution in Europe, and many, many centuries in China before the industrial revolution.

It was the first industrial process because it was the first process of mass production that replaced human labor. The human labor of writing was replaced by an industrial process of mass production.

Of course we know that this eventually led to the consideration of a copyright system but at the heart of that transition, that huge transition, of course, was always human creation. What changed was printing. It was really the expression and distribution capacity with respect to human creation. It didn't change human creation.

I think here we are facing an entirely different situation. We're facing a situation in which machine learning has led to machine creation and machine invention. It's a very different thing from changing the methods of production and distribution. I think it's a very radical change that we need to give serious consideration to.

If you look at our copyright system, all over it are provisions relating really to human creation; moral rights -- integrity, attribution -- the length of copyright, the life of the author plus 70 years. Are we going to do the life of the machine and its offspring? All over our system is human creation.

I think we do need to give very serious consideration from a policy point of view in a quiet manner to whether or not we don't need two different systems. One system, which is there to reward and incentivize human creation and another that is there to reward and incentivize machine creation.

It sounds very dramatic but I think this is something that we do need to consider very profoundly. Of course it's going to raise civilizational questions really of the extent to which we want in our economy and society to favor machine creation or disfavor machine creation in relation to human -- in comparison to human creation.

I would say that is one of the new questions that is going to lie at the heart of our debate. You can imagine that you might have, for example, a system which is a property right in respective machine creations for 10 years or for 20 years -- I'm just drawing examples out of the hat -- which is modeled along the lines of the reality and use of machine creation in our economy as opposed to the reality and the use of human creation in our economy and in our society.

The biggest challenge as far as I can see -- besides the policy challenge of how we go about doing that -- the biggest challenge is going to be how will we ever know what is a machine creation and what is a human creation?

I don't have any answer to this but perhaps those of you who are experts in the area have thought about this and thought about the ways in which you would be able to make that distinction. I'm not sure that people are going to necessarily tell you, or indeed are telling you that this is a machine creation as opposed to a human creation. This, I think, is a real challenge for us.

Two final words if I may. The first is another major question that I think underlies many of the other considerations that we have with respect to artificial intelligence and that's data and the status of data. Of course, it's a policy confusion and a policy mix.

We've got all sorts of issues, policy issues, that come into play with data and the data economy and the increasing use of data throughout the economy. We have privacy issues, we have integrity issues, we have security issues, we have competition issues, and the list goes on. Everything is coming out and converging on data.

You see that all come together, for example, with respect to medical data where you have a market estimated at around about $240 billion United States dollars worldwide of medical data increasing rapidly. All of those questions come into play, whose data are they?

Property, I think, has an extremely important role here. Of course, the interaction of whatever you decide about policy in relation to data is going to have a major impact on the intellectual property system and on copyright.

If you start making policy pronouncements about the use of data and the flows of data then, of course, the intellectual property system is there saying that certain categories of data when they satisfy certain criteria are eligible for property rights.

I think data, besides machine creation, is another very profound question where we are a long way in the policy area from really knowing how we are going to be treating this.

My last remark would then be this. We have a lot of very profound questions here. We have at the same time a conjunction of awkward circumstances, I think, in the law because these questions relate to -- these questions relate to technologies that are deployed on a global scale. And if you consider the creative industries, the platforms, operations are global.

At the same time these technologies lie at the heart of competitiveness and competitiveness lies at the heart of international economic relations. And, at the same time, we see a faltering multilateral system with respect to rulemaking. There are many things that are happening well in the multilateral system but rulemaking is not one of them.

So we have, I think, the conditions of a storm here. I think we all need to be aware of that in this area. It's for this reason that we are extremely grateful to the United States Copyright Office for engaging in this way.

The United States Patent and Trademark Office -- Andrei Iancu just walked in -- we are very grateful for your support for an international conversation with respect to these extremely challenging questions. Thank you very much for this opportunity to say a few words this morning.

(Applause.)

MS. ROWLAND: Thank you so much, Director General.

Now we are going to hear from Maria Strong, Acting Register of the Copyright Office, to give a few remarks about the Copyright Office and artificial intelligence.

MS. STRONG: Thank you, everybody, for coming. Again, thank you Director General for your very insightful remarks to set up the stage for today's conversation and the continuing conversation that we'll have both here in the States and in other places around the world.

We are excited to begin and continue this conversation coming out of the Geneva event last September. The former Register of Copyrights Karyn Temple actually spoke and moderated the panel on copyright and today we continue that conversation.

Here at the Copyright Office we are keenly aware and interested in the intersection of copyright and AI. As the primary agency charged with administering our nation's copyright system, the Copyright Office has long understood the importance of keeping up-to-date with changing technologies.

We look at these evolutions from both a practical and a policy standpoint. Practically, we must adjust to the changing technology when we are examining works for registration because each application requires us to make a determination on whether a particular work is copyrightable.

In addition, we have long provided advice to Congress, the courts, and our intergovernmental colleagues on policy issues related to copyright and emerging technologies. This means that we must appreciate the complexities of how technology and copyright are developing both for individuals and companies, both at the national and global levels.

Thankfully, we have experience in adapting to new technologies. Each wave of technological advancements has brought new challenges. The first federal copyright act in 1790 protected only books, charts, and maps.

Over the years, and often after significant study, Congress has added a number of newly-developed types of works to the Copyright Act, all of which we have had to assess as part of our duties to administer the Act.

We at the office have handled a steady stream of evolving technologies since our founding in 1870. This is our 150th year anniversary. We'll be talking more about it then so it will be coming up to celebrate a little later. Some might criticize the speed, or lack thereof, of legislative change to reflect new technologies.

Yet, at the same time, others recognize the importance of letting the marketplace have some time to develop before attempting legislation. I think the Director General mentioned that sort of balance of time and thoughtfulness that needs to be taken into account for this complex system of change.

Now we are looking at a new technology, that of artificial intelligence. AI seems to mean different things to different people. Indeed, there remains a debate about its very definition. While many of the technologies that we're discussing today are new, the issues they raise are not. As the Director General said, we may not have answers for all these questions, and that actually could be a good thing.

The Copyright Office began thinking about some of these issues related to the intersection of AI and copyright back in the 1960s. Yes, 1960s. The Office's 1965 annual report addressed the problem inherent in machine-generated works, noting that the determination of a line between human and machine authorship would be a crucial question to establishing copyright authorship.

As we pointed out back then, finding the line between whether a computer is merely a tool or whether a computer has independently conceived and executed a work is very difficult. But, yet, this question is not so different from the question that we are already talking about here today.

Take, for example, photographs. When photography first appeared on the scene more than a century ago, there were debates about whether or not it would be protected at all. It was an issue eventually settled by the Supreme Court. Our office adjusted to that situation and today we register any photograph that displays a spark of creativity.

But there is a line as we saw in the famous monkey selfie. As our Compendium states, a photo taken by a monkey is generally not copyrightable regardless of how compelling that photo may be.

(Laughter.)

That is because the Supreme Court held in the 1879 Trade-Mark Cases decision the copyright law only protects "the fruits of intellectual labor" that "are founded on the creative powers of the mind." While the monkey selfie at first glance seems like an open and shut copyright matter, there are variables of complexity both in fact and in law that could mirror the copyright issues we see today in AI.

What could a photographer do to make a monkey's act of clicking the shutter something akin to a computer program, for example? That goes back to the 1884 Burrow-Giles v. Sarony decision where the court analyzed the copyrightability of photographs.

There is no question a copyright law allows for the use of tools in creation, but at what point does setting something into motion mean authorship? This is something photographs and AI have in common: where is the line of authorship? What about ownership issues?

So the past can help us to identify and sort the present and future questions about some of AI's complex issues. It may not give us all the answers but perhaps it provides us with a framework of a way to analyze things.

Today we are looking forward to hearing about how artificial intelligence is impacting different types of copyrighted works. Our first session will lay some groundwork on technological and legal copyright issues.

Our panels will delve into the details of specific kinds of works, such as music, visual art, literature, video games, digital avatars, and the impact of AI and copyright on consumer products like self-driving cars.

We will also hear about how important it is to identify an appropriate corpus of material for machine learning and the importance of being aware of bias that may arise in the use of algorithms.

So to wrap up, today is a conversation, and in conversations, sometimes there are more questions asked than answers given. With that, I welcome everyone to today's conversation. Thank you so much.

(Applause.)

At this point I have the pleasure of introducing Andrei Iancu, Under Secretary of Commerce for Intellectual Property and Director of the United States Patent and Trademark Office. Director Iancu has held this position since 2017 having left the private sector where his practice focused on IP litigation. He, too, was at the WIPO event last September in Geneva. We are very pleased he is able to join us today. Thank you so much.

(Applause.)

MR. IANCU: Thank you, Maria, for the opportunity to join you today. And Director General Gurry, so good to see you here in the United States. Obviously, we've seen each other in Geneva and elsewhere around the world. Actually, several times on AI issues. Here we are again.

Very pleased to join the Copyright Office and WIPO with whom we work so closely on a daily basis as they open this conference on an important and fascinating topic, as we have heard, and will hear a lot more throughout the day.

AI presents challenges for many forms of IP. Some of those challenges are similar from discipline to discipline. For instance, can an AI algorithm be an author? You've heard a little bit from Maria and I'm sure it will be a heavy subject today. Can an AI algorithm be an inventor when it comes to patents? For both, how much human contribution is needed before rights can be conferred? Any human contribution at all?

Some of these questions are both ethical and legal. Though the answers may not always be easy, it's vitally important that at least we ask the questions and engage in the discussions like this one today and like the ones we've had in Geneva and elsewhere around the world.

We appear to be on the threshold of having to confront these issues at the USPTO and obviously the Copyright Office and in courts throughout the United States and throughout the world. AI, though, this is nothing new. It has been part of the national and international industry, academic, and even kitchen table discussions for a long time.

Although some of the innovations that AI embodies like speech recognition, image recognition, search optimization, and even early neural networks, were conceptualized as early as the 1950s. Significant breakthroughs are more contemporary.

Today AI is becoming ubiquitous in our society. I was at the CES show just a few weeks ago in Las Vegas. It seemed like almost every booth from computers to automobiles to biological companies all were discussing AI. Some now maintain that AI algorithms can create on their own with minimal, if any, human interaction.

As the administration's lead agency on

intellectual property, the USPTO has been actively engaged on this topic of AI and IP policy. A year ago, for example, we gathered leading thinkers, policy makers, academics, and practitioners to discuss the growing capabilities and economic impacts of AI and implications for IP policy at the day-long conference similar to this at our headquarters in Alexandria.

Then last August we issued a request for comments on patenting artificial intelligence inventions. Two months later in October we followed up with a request for comments on intellectual property considerations for innovation more broadly including specific questions of the AI nexus with copyright law.

We've received many, many comments from the United States and internationally and we are now reviewing those comments that we have received and will be issuing the report hopefully in the coming months.

But as we go through these exercises, we should be careful. We should be careful not to jump to conclusions. Instead, we obviously need to be deliberate and have a steady hand. As Maria mentioned, we have faced many technological advances over time; increasing automation. This new advance may be just another step along the spectrum.

Our current policies may work just fine. On the other hand, they may need to be updated as AI could also present brand new issues which we are exploring now throughout the U.S. Government and here at this conference today.

This is why I am so grateful to be working with all of you, with the Copyright Office, with WIPO on these very important issues. It's great that you all are having this deep discussion today. I wish you a great conference and hope to see you around the policy circles in Washington, D.C. again soon. Thank you very much.

(Applause.)

MS. ROWLAND: Thank you very much, Director Iancu.

Next we are going to turn to a foundational discussion of AI and copyright with the international discussion at play. We are going to be talking about -- I'm getting it all confused. I'm sorry. Excuse me.

We are doing the foundational discussion of AI and technology and the law. I would like to welcome to the stage Rob Kasunic who is head of our Registration Policy and Practice and the Associate Register here at the Copyright Office.

And Professor Elgammal at Rutgers University and is involved in both computer science and art and AI issues. He runs a very interesting lab over there in Rutgers. Thank you so much.

MR. ELGAMMAL: Hello everybody. Thank you for inviting me. I will try to give you an overview about how AI is used in making art, in particular visual arts.

Let me start by this video.

(Video played.)

What you see is from HBO Silicon Valley show that was aired back in April 2018 about AI art and the world. The two guys are talking and one of them is asking who is the builder and the other guy answering this is work done by AI. A third guy is coming in saying its actually the first art made by AI ever sold at an auction in Sotheby.

That's fictional but AI work is real. That picture is now a reality six months later when Christie's actually auctioned a piece of AI art and it was sold for almost half a million dollars.

I'll give you a brief historical view of where we are now in the context of history making art. Art has been evolved from cave paintings to painting on canvas like Caravaggio and the Renaissance. Then in the 19th century came photograph which changed the way of making images.

Photography became an art form in the 20th century. Then came digital photography in the 20th century. There came digital manipulation of images like Photoshop and these software. In the '90s came graphic rendering with all its amazing abilities. Now we are entering a new age of the machine using AI can actually create an image, not just taking a photo like a camera.

However, the use of AI in making art is not new. It is as old as the AI itself, from the '50s. Here are two of the pioneers who have been experimenting with using AI in making artwork.

On the left is Harold Cohen, an artist who used AI in making works of art for a long time. On the right is Lillian Swartz, a graphics scientist who also experimented with AI. This is one example of Harold Cohen's work.

The difference here is this is what's called rule-based AI system where you actually have to write lots of rules about what to do, what AI is supposed to do. Here AI actually makes drawings with some flexibility or guidance.

Okay. Move to the last five years. A lot of buzz has been around File Transfer where basically you start with a photo like I'm using a photo in the top left. Then you can stylize it by any style you like. There are many apps available that you can do that.

Another version of that is Google Dream where basically you start with a photo and then put some objects on top of it. This comes from a machine for a different purpose of recognizing cats and dogs and other things so you end up with something that looks like Van Gogh but has lots of dogs over it.

People liked that. If you search Google Dream you can find so many things you can do with it. However, these two examples of File Transfer and Google Dream are transformative AI where they transfer one image to another adding a style.

It's not generative AI. Generative AI I'll talk about today which is really the important issue when it comes to today's topic.

So what is generative AI? This actually came about five years ago in a work by Goodfellow and others. It's called Generative Adversarial Network or GAN.

So what is GAN and how does it work? It's an AI algorithm that basically tries to generate images. We give it some data, for example, images of cats or images of flowers. You want to generate more of that. The way it works is you have two components, two players.

One is a generator, the one on the right there, who has no access to this data. It never sees the data in the whole process. The other player is actually a critic or a discriminator, a technical term, who actually has access to this data. In the case of flowers, it has access to images of flowers.

Basically the generator will start generating totally random images because it doesn't know anything about what a flower is. It passes the image to that critic and the critic will say, okay, this is not basically flowers and send it back to the generator. The generator then has to improve, think about what to improve in that image to make something that will please the critic.

After so many tries hopefully that generator will be able to do something that looks like a flower that can fool that critic. The critic also tries to get better and better in telling whether that's a fake image coming from the generator or an actual flower image.

Here is an example after a few adaptations this will show up and the form is developed more and more until it's something that reasonably looks like a flower. That's a fake flower in that case.

All right. So this is the process of the kind of evolution on how art is made using AI. These are examples of actual artwork by different artists using AI art like Mario Klingemann, Tom White, Robbie Barrat, and others.

Actually, if you look at any of these examples, you can now think of what are the data that came from behind them, like nude paintings, in one a fan in another. This has transformed the data into a new visual form.

What are the aesthetics exactly of this process? At the top here you see examples of what happens when we give it lots of images of classical portrait from Renaissance to 20th Century.

We give it lots of classical portraits. One generation you see is kind of a failed attempt to make a portrait. It tried to make a portrait but it failed. From that failure come the aesthetics. I call it the aesthetics of machine failure. The machine didn't give a perfect portrait but gave that failed portrait.

That exactly reminds us of Francis Bacon portrait on the bottom with a difference. At the bottom, for instance, they intended to make this deformed portrait, while on the top she failed to make a portrait. From that comes an interesting portrait that an artist would like to put in an exhibition. That is to keep in mind of understanding the process here.

However, AI generation has evolved a lot in the last five years. At the right here you can see NVIDIA -- what is called fake faces where it generated lots of fake faces. By fake it means that it's not anybody the machine has seen in the training data.

It can actually look like somebody in reality because basically a human face is a combination of other faces anyway. This shows you how good is the rendering of these machines are now in creating realistic images.

On the left is an example of creating birds and flowers just from text. You give it a text of what you want to generate as a flower or a description of the bird you would like and the machine actually generates a good quality realistic bird which also can be a fake bird.

Where is artist in the process? So in this process of using GAN the artist would feed the machine some images like flowers and the machine would generate lots of other examples of flowers, not necessarily copying the data because the generator doesn't see the data, it just generates things from the same distribution.

What is the role of the artist here? The first thing is the artist actually pre-curates the inputs. The artists chooses what data to feed to the algorithm. You want to create something based on faces or flowers. What kinds of faces? What flowers? What paintings? What's the input? That's one of the roles.

Then the artist actually tweaks that. Most artists are not technical people so take an algorithm and change it a little bit or run it as it is or change some parameters. If the artists is technically savvy, he actually can change the code or write his own code. That's another role.

The third part is post-curation because - the machine would give thousands and thousands of possibilities and the artist chooses a few of them to show to the world exactly like a photographer would take one thousand photos and show ten or one in a show. The artist's role is pre-curation, tweaking, and post-curation. It's totally a conceptual art process where artists are involved and AI is a tool that can actually create something with the artist.

In October of 2018 Christie's sold that artwork on the left here. Here is the problem because that artwork was the result of a GAN machine that was trained by another artist. The other artist was Robbie Barrat who trained that algorithm.

He tweaked it and trained it on classical portraits and other authors take that algorithm and basically generated from that algorithm after being trained. Basically you're coming to the diagram here. The algorithm was created. It was already pre-curated and trained. He pushed a button, selected an image, and sell it.

Who is the author of that? That was the debate at the time. I have a whole article on art and that if you want to know more. There's a lot of issues about authorship of these kind of processes and the outcomes of these processes.

Another issue is that what you generate here is a very rich generation. You can - this video actually will, for example, show going through what the machine learned after trained on flowers. As you see it can really generate lots and lots of forms and flowers and intermediate forms and combination of forms. And that raise another issue. What if you train this kind of models and take away some of the results out with after being trained? Who owns that?

And for example, there are apps online like Artbreeder who actually do that. You can actually combine any images. This actually takes millions of images and train them and -- that include animals, objects, cars, many things. And train the machine to generate these kind of things. And once you do that you can actually navigate in that representation and generate lots of combination of things, amazing combination of things. For example, that video here will show an example where I combined - I just do this to combine a cat, a panda, and a hamster. And what you generate here are all possible combination of these three animals -- this doesn't exist in reality.

So if you go to this amazing website and you create your own image out of that, another person can actually create the same image accidentally or eventually. Then that is the authorship and ownership issues.

So in that particular process the artist really here is just basically the post-curation because the algorithm is already trained and has ability to generate infinite amount of images. And as an artist or as a user your role is really searching for something interesting to you to - out of there to show to the world. So that's a very limited role.

So moving to autonomous generation, can AI generate art by itself and can the role of human become minimal to that? So what if you take these GANs and fit it for example all art history, for example? What it will generate? That's an experiment we did at Rutgers.

So basically we took 500 years of art history and - images of art history, about 100,000 images and feed it through these GANs. By definition, by construction GANs will not make anything creative, as I showed. It will just basically try to create something that looks like the data you give it. So the outcome will be more like what you see in the bottom here where basically emulation of what you have seen in art before, and mostly it's failed emulation.

So how can you move from being generative to being creative? And that's basically our contribution in that area, which was called Creative Adversarial Network or CAN, where basically you would try to - pushing a little bit further in the autonomy and becoming more creative.

So what is that? We call that artist AICAN. So basically we based our algorithm here on a theory from psychology by Professor

Martindale who was at the University of Massachusetts, and the theories basically can be summarized as this.

Imagine you're an artist living in the late 19th Century. So Impressionism has already happened, and you already have seen lots of Impressionist work by many artists and Impressionists have really painted every possible facade of a building or a street or an amazing landscape.

As an artist you're already bored of that, as a new artist. And so basically if you keep doing the same kind of art, it was called habituation, you kind of get bored of it as a viewer. And as an artist, avant-garde artist basically you got bored very early about - from that, not like the general public.

So as an artist your role really is to come, create something that's really innovative to push against habituation with the least effort, because if you push too much, that can be totally innovative and can be shocking. And that was exactly what for example Picasso did in the Ladies of Avignon artwork that later sparks the cubism movement. So pushing innovation against habituation is right when you will drive the art forward.

So how can we implement a machine that do that? So what we did is we take these GANs and modify it in a way that we want to create something that's innovative. So if the machine creates something that's - repeats, for example, Renaissance or Baroque or Impressionism, it has to be (unintelligible), otherwise it's not going to be creative. That's how to push it forward.

In the same time if it generates something totally random that will not be copying any of these existing schools -- that's totally shocking. That's not really as interesting as art. So you have to put them here into a dilemma.

In one hand it has to follow the aesthetics, it has to follow the distribution of what art is. In the other hand it shouldn't generate anything that's identified as an existing art movement. And by doing that the machine started generating these amazing forms for us that was very really interesting. And when we make surveys and show these mixed up with actual art by human artists from recent contemporary art fairs, we're surprised that people thought that this art was made by a human artist 75 percent of the time compared to 85 percent of the time when they're shown art by abstract exhibitionist masters, and 48 percent of the time for actual art by artists from Art Basel 2016.

So that really put the machine at the level of generating visual images. It has certain interesting forms that you cannot tell whether they're done by a machine or a human. And what's the role of the human here? For me nothing other than developing the algorithm, the process and feeding all art history. So I have no control of what will be generated next.

All right. So let me finish here.

And also we will associate words like intentional and inspirational, communicative, having human visual structure with these arts.

Let me finish by this, a social validation of AI art in society. So the first one in the top left is an LA gallery show back in 2017 that show AI art. I think that was the first gallery show. Then the Christie's show in '18. Art - was shown in an art fair in Art Basel and SCOPE Art Fair. It was shown in galleries in New York and in Chelsea and another action by Sotheby later. The Barbican in London had made a show about the AI art. There was just an exhibition in China, in the National Museum of China in November -- expected one million people showing art and technology including AI art. And obviously many coverage in the media from top news coverage.

So I'll stop here for the time, and I'll start with the discussion.

(Applause.)

MR. KASUNIC: Good morning. That was

fascinating. Thank you.

Well, as the director general thoughtfully discussed earlier, we are at a time now where artificial intelligence is really forcing us to reconsider and question some of our foundational assumptions. And one of those would be just looking at the question of what is art?

I thought it would be helpful to look at some dictionary definitions of that just as context of where we now stand on that issue. The expression or application of human creative skill and imagination typically in visual form such as painting or sculpture producing works to be appreciated primarily for their beauty or emotional power. Works produced by human creative skill or imagination. And the conscious use of skill and creative imagination especially in the production of aesthetic objects.

So that is a - certainly art is not the question. There's a difference between what is art and what is copyrightable, but to get into the realm of what is copyrightable, we've had several Supreme Court cases over many years that have looked at this question, looked at the constitutional provision of works, of what constitute works of authorship and what constitutes originality. And in the Bleistein decision Justice Holmes did look at that question of what is originality and stated that it was - wrote that it was the personal reaction of an individual upon nature or expression of the author's unique personality as the key to satisfying the constitutional requirement of originality.

Similarly, when we get to the case of Burrow-Giles versus Sarony dealing with the photograph of Oscar Wilde that you see there, the question of authorship was raised and in - and the question of whether photographs could be copyrightable. So it also was addressing what constitutes a writing under the copyright clause.

And the court said that an author is the person who effectively or as near as can be the cause of the picture which is produced; that is, the person who has superintended the arrangement, who has actually formed the picture by putting the persons in position and arranging the place where the people are to be, the man who is the effective cause of that. The author is the man who really represents, creates, or gives effect to the idea, fancy, or imagination. And also said these views are of the nature of authorship and originality, intellectual creation and right to protection confirm what we have already said.

More recently the Supreme Court stated in the Feist decision versus Rural Telephone Service in an opinion written by Justice O'Connor that the sine qua non of copyright is originality, and to qualify for copyright protection a work must be original to the author. The court defined author in the constitutional sense to mean he to whom anything owes its origin, originator or maker.

While the word writings may be construed liberally - may be liberally construed as it has been to include original designs for engraving, prints, et cetera, and photographs, it is only such as are original and are founded in the creative powers of the mind. The writings which are to be protected are the fruits of intellectual labor embodied in the form of books, prints, engravings, and the like, and that again reinforcing that originality is a constitutional requirement.

So some of the questions raised by works created by artificial intelligence are: can machine learning produce original works or is the product of such software inherently reproductive, derivative, and/or the result of a system or process devoid of that personal reaction of an individual upon nature or expression that is devoid of the author's unique personality? And also it can be asked can a computer program be the author in a constitutional sense?

And part of that question then answers does Congress have the constitutional authority to provide copyright incentives to computer programs themselves as authors? Furthermore, if such authority does exist, should Congress exercise that? I think we heard earlier that this is a time to be deliberate and to not rush into answering some of these questions, but to see perhaps how the situation evolves.

Also, if authorship or ownership of AI output should be protected, is copyright the proper vehicle for such protection or perhaps is sui generis protection or some other form of intellectual property protection preferable to that?

Some of the types of uses. You've seen certainly a lot of examples, but just to more broadly look at how software may be used and a very high-level overview of software can be used in the creation of works is the computer program itself that is the tool to create other works. In that case the programmer of the computer program can unquestionably be the author of that computer program. So computer programs of artificial intelligence are protectable themselves.

Furthermore, software as a tool, as we learned from the CONTU Commission many years ago when computer programs were incorporated into the Copyright Act that at the time viewing software primarily as a tool for others to use. So the programmer may own the software itself, but the user of the software owns the works created with the software as a tool much like a photographer would own the copyright to the photograph rather than having the camera manufacturer as the owner of that photograph.

Software can also be used as a template, a sort of Mad Libs-like use where the software can complete the template. And this seems to be the case for use in such types of works such as factual news articles. Sports articles, weather articles can use this form of template in which certain facts fill out the remainder of that. And this may entail some authorship by the author of the computer program in terms of creating that template that is going to be used.

But the question is software as the creator of works based on machine learning or the output of random or process-driven expression is really I think the question that we're - questions that we're looking at today.

So where we currently stand in the United States Copyright Office is that the office will not register works produced by nature, animals, or plants. Examples that we've used in the Compendium of Copyright Office Practices include a mural painted by an elephant, a claim based on the appearance of an actual animal skin, a claim based on driftwood that has been shaped and smoothed by the ocean, or a claim based on cut marks, defects, or other qualities found in natural stone. Those are objects that cannot be copyrighted because the result is not the result of human authorship.

Similarly, as you heard, also another example is that a photograph taken by a monkey may not be protected by copyright. And we took no actual position in the actual case involving the monkey at issue in Naruto here. Although the 9th Circuit did decide that - although the court found that - the court's precedent required the court to conclude that the monkey's claim had standing under Article III of the United States Constitution, that the court did find that with respect to the copyright infringement claim that the monkey, as with all animals, since they are not human, lacked statutory standing under the Copyright Act and therefore could not - the court could not affirm the judgment of the district court, that there was no standing to bring a copyright infringement action.

It is a wonderful photograph though. Although there are some lingering questions about the degree of originality.

(Laughter.)

MR. KASUNIC: I think when you look at the comparison there and the pose by Naruto, there are some lingering questions.

Another example in the Compendium that

has been addressed is the Office will not register works produced by a machine or mechanical process that operates randomly or automatically without sufficient creative input or intervention from a human author in the resulting work.

And the example, the pictures that you see there were actually countertops that were produced by a mechanical process and for which there was no human that actually had sufficient input into what the resulting work was. So those were refused for registration, and generally those kind of processes will not be sufficient unless there can be some showing of sufficient human intervention or creative input into the actual result.

The question of process also raises questions with respect to the decision in Baker versus Selden, and ultimately what is now Section 102(b) of the Copyright Act, which states that in no case does copyright protection for an original work of authorship extend to any idea, procedure, process, system, method of operation, concept, principle, or discovery regardless of the form in which it's described, explained, illustrated, or embodied in such work.

So in many of these cases it does sound like the algorithms or the software that creates artificial intelligence could fall under the category of being a process itself that is creating these works -- in some cases.

There's also some question that I think needs to be considered at least in terms of - typically in the merger question. It would be not wanting to allow a - the way that an idea can be expressed in only a limited number of ways to allow that - those limited number of ways to basically protect the idea itself.

And I think that when we're looking at the question of the computational power of computers as opposed to the human mind and human ability that there are some concerns as well about whether computers could fix so many variations of expression in a very short amount of time that there is basically no room left for human expression or original human expression from that. And at least one court had said that copyright should not be viewed as a game of chess in which the public can be checkmated. So it is a lingering question.

There -- as was discussed earlier, and I think is very obvious, is that the question of is a non-human work distinguishable from human authored works? And along with that, are machine-generated works creative? I think the answer to the first question: is a non-human work distinguishable? In many cases that answer is no. But I think that is not necessarily the relevant question.

Artificial intelligence can create creative works, but the question is whether the creativity is the output of artificial intelligence software; it may be considered authorship, without sufficient creative input or intervention from a human author in the resulting work.

Another question is whether AI creative output itself is derivative requiring authorization of the works used in the course of machine learning. And again, this was something discussed by the director general in terms of the data or the works that form the basis of machine learning. And so there are significant questions about access to those works and authorization for those works.

It's also important to consider that not all creative works are copyrightable. Certainly under the Copyright Act, under federal law, a creative unfixed work is not protected by federal copyright law, and a creative work that is not within congressionally-designated copyrightable subject matter is also not protected by copyright.

That picture there represents a claim that we received for a genetically-altered fish in which we did not find a Section 102(a) category that was sufficiently - in which that fell and so could not be registered.

Should a creative fixed work created by a non-human be within the scope of federal copyright law? That's a question.

Another consideration with that question, and it may seem unfair - I think there are similarities to this quote from Feist. It may seem unfair that with these works being created, why should not somebody be considered the author of that work, or also be considered - is it - should not somebody be considered the owner of that work and be able to reap the rewards from that work?

Well, a similar question was raised with respect to the concept of sweat of the brow and industrious labor, that the money, expense and time and effort that went into creating certain factual works. And the court said it may seem unfair that much of the fruit of the compiler's labor may be used by others without compensation. As Justice Brennan has correctly observed however this is not some unforeseen byproduct of the statutory scheme. It is rather the essence of copyright and a constitutional requirement. The primary objective of copyright is not to reward the labor of authors, but to promote the progress of science and the useful art.

So a significant question is does artificial intelligence require copyrights incentives? Do we need to or - and again, if it does require some kind of incentive, is copyright the proper vehicle? Thank you.

(Applause.)

MS. ROWLAND: Thank you so much Rob and Professor Elgammal.

Now we're going to turn to our international panel to talk about the developments in the international sphere, and we have some very interesting panelists here today.

It's going to be moderated by Maria Strong, and we are going to have Michele Woods who is at WIPO, but also an alum from here, and Dr. Till, who's going to also join us. We've heard earlier that she is the newly-appointed

Director of AI. And I believe we're also going to have Ros Lynch, who is the head of the U.K. IPO to talk about this.

So they're going to talk about all of the different issues that are arising in the world and in their countries. Thank you very much.

MS. STRONG: All right. Thank you, everybody, as we're getting settled here. Again we want to welcome everyone. As the Director General said, this is a question that is facing all national governments, but it's also a global issue. And so this is a panel where the international gets to meet the domestic.

And so the way we're going to do this is I'd first like to start off with the national experience from Dr. Ros Lynch, who is the Director of Copyright and IP Enforcement at the U.K. Intellectual Property Office. She will be discussing some of the challenges she has experienced in her country. And she had the benefit of working with WIPO last summer also

preparing a two-day presentation on AI.
Then we'll go to - after that we'll go to Dr. Ulrike Till, who's new to WIPO. I think this may be your first month maybe of work? So she'll be able to discuss some of the work of her new division.
And then we'll turn to Michele Woods who can also talk about additional WIPO initiatives at the international level.
And we hope to have time to have some Q&A between the panelists on some of the issues that are pressing.
So with that, Dr. Lynch?
DR. LYNCH: Good morning, ladies and gentlemen. And thank you, Maria, for correcting my job title. Whereas I would like to be the head of the U.K. Intellectual Property Office, I am not.
(Laughter.)
DR. LYNCH: I'm only responsible for copyright and IP enforcement, but I've also taken on the responsibility as the senior responsible

officer for everything we do on AI.

So I'm really very grateful for this opportunity to be here just to share with you some of the things that we've been thinking about and doing within the U.K.

For the U.K. the government has set AI as a priority, and it set an ambition to make the U.K. a global center for AI and data-driven innovation. And in the Intellectual Property Office we've been thinking for the past year about what our role is in relation to helping the government deliver on that ambition.

And we've been working on - basically on two streams. So we've been doing quite a bit of operational work, looking at the use of AI in terms of trademarks and patent registration, but that's not what I'm going to focus on because we think - we feel that the most challenging, the most interesting bits, of the AI field to fix on is the big questions around the policy around what happens in terms of trademarks, copyright, et cetera, policy and legislation. And that was why in June last year we held an international conference to look - also in cooperation with WIPO to explore the economic, commercial, and legal implications of AI for IP.

We didn't focus exclusively on copyright, but covered all the rights, but it became very clear from the two days of discussion that some of the biggest challenges posed by AI lie with the copyright framework. And I think Director Gurry already posed some of those questions to you.

And he mentioned authorship, and I think it's come up in every conversation so far. And we do actually feel and agree that AI is actually challenging some of the very settled notions of what authorship is, and it's calling into question some - probably the very foundation on which copyright is based, which is human authorship.

We in the U.K. thought that we could answer some of these questions around who is the author if - well, whether or not there is an author involved and who would that author be if an AI produced works in art, whether it's art, music, et cetera, and especially if there's very little human input into that work.

So we have in U.K. law a computer-generated works provision which is not something that you find in many other countries around the world. I think there are only about three other countries or so that have such a provision. And in our provision, computer-generated works are defined as works generated by a computer in circumstances where there is no human author.

This was proposed in 1987 and the provision, if - reading the - I certainly wasn't in the IP Office then, but reading the debates around it, it said that the provision was specifically included to deal with the advent of artificial intelligence.

I was quite shocked that they - the U.K. Parliament was actually thinking about AI, but - at that stage, but I - in my mind what we know of AI now, I don't think that they could have conceived of where the technology would have gone and what it would be capable of doing now, let alone what it will be capable of doing in 20 years' time.

For our provision, computer-generated works provision, the author is the person who made the necessary arrangements for the creation of the work. So this could be the operator of the AI. It could be the computer programmer who wrote the algorithm, whatever. It has never really been defined. It's just been left as the person who made the necessary arrangements.

Protection of these works lasts for 50 years from the date the work was made, so it has quite a long period of protection within our law. However, there's a huge tension for us between the concept of computer-generated works as defined in law and the concept of originality.

To be protected, computer-generated works must be original. The law appears to assume that originality is a human quality only, so in the U.K. and in Europe an original work is one which is the author's own intellectual creation. The concept has been further developed by the courts, and considerations include whether the author has made free and creative choices or reflected their own personal touch. So again, it's very, very human-defined.

And for us it's not immediately obvious how current AI-generated works could fit with this definition, and the tension makes it unclear to us how a work that is generated without a human author can actually be protected under the computer-generated provision within our law.

So we've got something in law, but actually we don't know what it means, and we don't know how we could use it in this context. So that's for us an added headache because we thought we were being very forward-thinking 20-plus years ago, but actually it's just complicated the situation even more.

So looking beyond the specifics of

U.K. law, I just want to touch on a few other things, and this is something which has already been raised. There's an interesting global debate to be had around the role of copyright protection as an economic incentive when it comes to AI-generated works.

So copyright is generally considered somewhere between an economic tool that incentivizes and rewards creativity, but also a natural right of authors to protect their creative works as expressions of their own personalities.

There is a question then around whether AI development needs incentives, and I think that this was just raised. And if so, what is the best vehicle to achieve this? Is copyright protection a necessary incentive? Could there be another mechanism? And again, I know that Rob just raised that. Does AI actually need incentivizing?

We seem to be having - from what we've heard, AI is being used in so many contexts. More and more works are being produced every day. It's being used across the economy. So is it doing okay, or do we actually need to do something to actually incentivize more creative works?

We've been asked do we need a new right, which is something we have to consider. We've also been asked to look at our exceptions to copyright, and should we have an exception for works created by AI? But also what is the role for licensing in this space? These are all questions which we are still trying to come up with answers to.

If it does make economic sense to protect AI-generated works, we need to ask when the copyright should subsist. So what type of works would qualify for the copyright protection, who should own the copyright, which we've had some of those conversations already this morning, and then how long, which again I think Director Gurry already alluded to.

The speed at which an AI can create

creative works, whatever they are, whether they be music or paintings, et cetera, just doesn't feel like 50 years is the right length of time to give it because these works are constantly coming out. They're being adapted. So we do need to have a conversation about if we're going to protect AI works, for how long do we need to protect them?

And I appreciate what Director Gurry said earlier about the length of time it takes to agree anything globally, but I do feel that we need to have a global conversation, and we do need to have some kind of global standards, or whatever, in this space. I'm not talking about a big treaty, but we do need to have some kind of agreement, some kind of shared understanding of not just what the problems are, certainly how to define it, and what we need to do in this space.

I also wanted to just briefly touch on infringement and liability. I know there are huge questions around whether or not the AI is making copies. Is it keeping copies? Where are these copies stored? If it has permission for some of the data that's being fed in, who gave permission for that data to be used? If so, if there is infringement happening, who is liable?

Is there a case for also considering secondary infringement of copyright? Would that lie with the institution where the computer is based or with the individual who wrote the algorithm, et cetera? Those are all things which we need to consider.

And my final point is one which is not normally talked about in relation to copyright, but I think it's incredibly important that we do not lose this from the discussion on the relationship between AI and copyright, and that is the ethical use of AI in this space, in the creative process. I do really think that this is hugely important, and it shouldn't just be left to the medical field or to something else. We need to actually look at the ethics of doing this within the creative space.

So just to conclude, in the U.K. we

are still thinking about these questions, and we're still trying to work on answers to them. We are shortly going to launch our own call for views going into - we're going to do it across all the rights, not just copyrights, but we've asked - we will be asking quite detailed questions, some of which I've mentioned, some have been mentioned earlier around data and computation, et cetera. And following that we will be publishing a government response giving answers to these questions where we have them.

But we do recognize the importance, as I said earlier, of having a shared understanding of the issues, and I think also having consistency in some of our answers, even though we may have some kind of domestic flexibility. So we look forward to continued work with WIPO and also with the U.S. Copyright Office as we try to grapple with this. Thank you very much.

(Applause.)

MS. STRONG: Thank you. Thank you very much, Dr. Lynch. We too look forward to working with you as you grapple with, as you said, your nagging headache. I think that might be coming across the pond to us as well.

And with that I turn to Dr. Till.

DR. TILL: Good morning. So I am the director of the brand new AI Policy Division at WIPO, and today marks three weeks in the role and three weeks for the new division. In some ways I feel like you should be standing up here and I should be sitting down there, because I certainly have more questions and no answers at all.

When I was preparing for this talk I was thinking I've only been in this job for three weeks. What can I do to add value to the discussion?

What I thought I would do is take you a little bit through the WIPO process, the conversation on AI and IP and where I came by my background. I come from private practice. I'm German-qualified and U.K.-qualified. I'll give you a little bit of an international perspective on why some of these questions might have different answers or different perceptions in different countries and also why I think having an international discussion and a wider discussion of all the issues is so crucially important. And that feeds in why I think the WIPO conversation is so crucially important in raising awareness of different angles to the different questions.

The growth of AI. I'm not telling you anything new. It's everywhere, and it's expanding rapidly. I'm not sure how many of you have seen the WIPO Technology Trends report that came out at the end of last year. It's not copyright; it's patent, but it looks at more than 130,000 patent applications.

It's one of the first studies to analyze the trends. It's a big volume, but some of the things are fascinating in it. All of them are, but what I took away is there's a real shift that you can see from theory to actual inventions. So there's patent application now, rather than just publications.

Machine learning is mentioned in more than a third of those applications, and really quite telling for me is that it is across a whole range of industries. It's from agriculture to education, from transport to IT technology. So it really hinges home that it concerns all industries. It's a rapidly evolving field.

With the field that is rapidly evolving, as we can tell even in the discussion today, is a multitude of discussions that are exploding. The -- discussions are national level, certainly of governments. A lot of governments have AI strategy papers or policy in place. Not very many of those mention IP at the moment, but that discussion is certainly starting.

There are four here nationally where issues are raised in the U.K. and the U.S. The patent offices around the world, the IP5, have got a cooperation to look at how they can use AI in the administration of IP, but also they've started amending the examining guidelines to deal with issues of AI.

When you look into the academic papers and legal publications as I've done over the last three weeks, it is overwhelming. There are so many of them. There's so many views and so many discussions. And the discussions unsurprisingly go across all IPRs. It's from copyright to trademarks to patents to confidentiality to rights of data.

For me personally, where it comes down to policy question is really: is the current IPR system fit-for-purpose? Does it fit the new AI world or does it need adjustments in order to strike a fair balance between protecting the investment into AI and the fostering of technology going forward and innovation? And that really for me is at the heart of all of the discussions.

And really this is where the WIPO conversation on AI and IP comes in. Having all of those discussion points, having all of those opinions, having the question asked and having them asked internationally in different forums is a fantastic way of helping us all understand the issues, not so much have the answers, but understand the issues that we ought to be looking at.

WIPO is giving a forum to that. As you know, many of you were part of the discussion before I even started; the conversation started in September last year in Geneva. It is at the point at the moment where WIPO has published a draft issues paper. It's on our website. If you haven't seen it, have a look at the website. Should get a little bit more intuitive going forward; changing that at the moment.

The call for comments is until February 14th. At the moment it simply is asking the question have we got the questions right? Are the things we're thinking about the things you are thinking about? Do we need to change that? Are they on the right emphasis?

The idea is very much that we will, having received the comments on February 14th, revise the draft issue paper, and that will guide the agenda for the next conversation that will happen in Geneva on May 11 to 12, and a second conversation planned for the end of this year Q3, Q4.

One of the things that is quite close to my heart is, as I said, we're at the stage of asking the questions. Are we getting the questions right? And getting the questions right for me is really important and comes a little bit from my background.

I spent very many years as a litigator, and I know what happens when systems get changed and they get changed in unforeseeable ways. That always happens when you tweak systems, but I think the more time we actually spend having the discussion to get the questions right, the more we can minimize unforeseen consequences when the systems get changed.

That is of course - as the Director General also said, it's a race for AI. The field is moving very fast, and there is a tension between spending enough time on getting the questions right and seeking answers to those questions.

Without giving anything away, a lot of the comments on the draft issue paper are racing ahead to trying to answer the issues, but my call at the moment is: help us define the issues so we can better answer them, but we'll continue the conversation.

Coming to copyright. And really for me, when I was thinking about what I was going to say, copyright and AI falls into three buckets. There is of course the really important question on ownership and authorship. The works created with AI, does copyright apply? What are the implications of copyright?

But I think copyright for me implicates every part of the AI creation system. When you think of the algorithm that creates the works or does the invention, there are jurisdictions out there that will not patent protect software or mathematical method per se, so is copyright the answer to actually protecting the AI underlying all of that?

And then there's the input. I mean the reason or one of the reasons AI is taking off is twofold. First of them is the increase in computation power, but also it's the increase of available data feeding into the AI systems. So what do we do with the new dawn of data? Does the data need protecting? Is copyright the right thing to do? Copyright impacts all parts of AI. For me, it's the input that goes into the AI. The AI does its thing. And the output. And really looking at all three areas is going to be absolutely crucial and has a lot of questions attached to it.

This is just whistle-stop tour for those that have not seen the draft issue paper. Here are the big headlines of the questions asked with regard to copyright in the issue paper. Unsurprisingly authorship and ownership is a big one. And really the question there is: can you think of it in legislation; and what does legislation say; and do you need a human author in the legislation -- should be changed?

If you take a step back, really in two very broad-stroke buckets, there are two ways of thinking of copyright. One of them is an economic way of saying what's worth copying is worth protecting. The other one is seeing copyright very much as an extension of personality with moral rights. And then having that discussion, depending which view you take on copyright, might give you very different questions and very different answers to that, but a lot of the questions that we have at the moment center around that.

Infringement and exceptions. I think everyone's already mentioned it, but again infringement and exceptions are very different depending which countries you look at, from fair use to private copying exceptions, from data mining and really what does it mean for copyright and AI if you start tweaking those exceptions internationally?

Deep fakes are on everyone's mind. What should copyright do? Should they be excepted from copyright? Copyright or a right per se is another question. The big question of data. Should data be protected? And bias. I think we are going to have a session on bias later, so put that aside. But they're all buckets of questions currently raised in the draft issue paper.

And I will finish with a plea for your help. Help us get the questions right. The deadline is February 14th. It's not very long, but if you can help us define those questions, give your views on it, our commitment is to revise the draft paper and issue the revised paper at the end of March. And then I will look forward to seeing as many of you as possible for the second conversation in Geneva on May 11 and 12. Thank you.

(Applause.)

MS. STRONG: Thank you, Dr. Till.

Michele has the final presentation

today on additional work of WIPO. And I'd like to leave time for one question to the panel. Thank you.

MS. WOODS: So I'm happy to try to be very quick. And first though I just want to say how lovely it is to be invited back and to see so many friends here through different phases of my copyright career, including that continuing now, and also to be on this panel with distinguished women leaders in copyright.

I'm not sure if Francis is aware, but his office told me to make sure the Copyright Office complies with the requirements of his policy on gender equality, and particularly on the role of women. And I was able to say: Copyright Office, Copyright in D.C., no problem.

(Laughter.)

MS. WOODS: So it's great to be here with this distinguished panel.

So a lot has been said already. I'm going to be very quick here and just point out another side to the work at WIPO with respect to artificial intelligence, and that is that while asking all these questions about policy and working with all of you and asking you for comments so that we define the important questions, we are also working with AI ourselves to come up with useful systems and tools that can be used both to enhance the functions and processes at WIPO, but also to be shared with international organizations, with IP offices, including copyright offices, and that work is very active.

There are a number of very interesting and useful tools now, and the development is really continuing at a fast pace. We have the Advanced Technology Applications Center taking the lead on this.

I'm just going to mention what some of the areas are. This isn't my area of specialty, and in any case we have other copyright-specific

issues to talk about, but I did want to note that we have, for example, the WIPO Translate cutting edge translation tool for documents. Right now it's mainly being used in patents, but it extends language coverage.

This is a very important area for patent examination as many of you I'm sure would be aware. And so this machine translation tool is really cutting edge and is being shared not only with other international organizations, but with IP offices around the world.

Then we have a Global Brands Database. This uses image search, or, is an image search service within the Global Brand Database. This allows trademark owners to identify visually similar trademarks and other brand information using AI tools.

Also another technology that's being shared, classification, automatic patent classification. This is a little different, but is also providing this information in order to help patent filers and examiners examine patent applications. And indexing tools for patent area once again. So as you can see, the area is really very active.

There's a new tool that's been announced, which is a new digital system for WIPO meeting records. And this, instead of text-to-speech is speech-to-text technology. And our AI colleagues have been using our AI tools to build this system.

The hope is that instead of having resource-intensive verbatim reports, and those of you who have been to our meetings have seen those reports; very thick, very labor-intensive to prepare, we'll be able to have automatically-generated and translated speech-to-text transcripts of our meetings in all six U.N. languages using these tools. That'll be complemented by an audio-visual system to provide digitally-indexed and searchable conference records.

This is a project very dear to my heart as in my position I have to oversee the preparation of these reports for the Copyright Committee, the SCCR. And this takes a huge amount of time. Big budget implication.

And it would be great if we had this. I was very unhappy with the colleagues when they announced this on my birthday and my committee was not in the pilot. But I'm very hopeful that in another year or so we will also be taking advantage. And in fact we're allowed to use the tool in parallel with our current system, which is very, very interesting and useful in terms of seeing how it works and how we might use it in the future.

Very quickly, we're working at WIPO with IP offices in a number of different areas. A lot of them have been mentioned. I'll just note that, for example, one way AI can be useful, and is being experimented with in a number of offices, is help desk services, automatic replies to clients. This can be very useful for copyright offices.

Machine translation, linguistic tools and terminology, data analysis. There's a lot of potential here. I know we'll be hearing more from some of you about how some of these tools are being used either in your offices or in your industries.

I wanted to just finish with a little bit of a snapshot and one that's dear to my heart and those of you who've worked on the Marrakesh Treaty, and that is in the area of accessibility and the potential for AI.

So I'm sure many of you know that at WIPO we have the Accessible Books Consortium, the ABC, and it operates the Global Book Service. There are almost 500,000 works in the catalog now that are available for sharing across borders through the Marrakesh Treaty. And another big function of the ABC is capacity building, helping countries, particularly developing countries, to prepare materials, usually educational materials, in local languages in accessible formats.

One of the big issues that has arisen in a lot of educational materials is using images or making images accessible. And up to now that's been a very labor-intensive and usually human-intensive work, but there's a lot of research going on now on image captioning using computer vision. So this is a technology that the ABC is looking at very carefully and could be very useful in a broad sense for accessibility, something we're very excited about.

It's not entirely ready yet. When you have a possibility that the computer still says that people are various types of animals, you can imagine that that could lead to a lot of problematic situations, and we wouldn't want to put the persons using the tools into a situation where they're commenting as if the people in photos are animals and creating problems for them. So it's not ready yet, but there's a lot of really exciting work going on.

A lot of it's being driven by companies with social media platforms to increase inclusiveness of those platforms so users don't just skip over images and photographs, for example, but are able to engage with those materials. There are of course a lot of economic implications that are not related directly just to accessibility, but the lens we're looking at this from right now accessibility is one where this is very important.

It also provides a case study of quite a few of the data and ethics issues that are being looked at in the WIPO AI conversation. So the machines need to be trained with large sets of images for machine learning. Are those copyrighted images? If we say, well, we should only use images that aren't in copyright, does that lead to questions of bias? Does that inadvertently lead to unrepresentative data?

There are privacy concerns with respect to the data sets and also the use of the tools. There are liability concerns if there is a mistake, a problem, and it leads to some kind of situation that generates liability. Who's responsible for that?

All of these are very important questions, but nevertheless the exciting thing here is that AI also offers tools that will help us to deliver on the promise of trying to in this case provide accessible format works, published works, but that could also have much broader implications for accessibility, for inclusiveness. So this is another area where we're very excited about this work and happy to partner with all of you. Thank you.

(Applause.)

MS. STRONG: Thank you very much, Michele.

That leaves us time for a question I'd like to throw out to the panel and also for perhaps later discussion on other panels, and it's this: It seems like much of the discussion we have had is with respect to what is the policy, and we all agree I think on the importance of trying to get to the understanding level that Ros talked about. But it seems like it focuses on two aspects. One is the promise of the technology to address large societal problems, or it could be the threat that's going to be caused by unregulated or undefined problems.

So how do we as policy makers and participants of the copyright community address this seeming polarization in both promise and threat?

And then another question, should you choose to accept it, is - I'd like to follow on what the Director General mentioned, is the role of competition. How can licensing play here and in the copyright community, both for creation and for use?

Take one, both, or anything you'd like to say. We'll just go down the line, maybe starting with Ros.

DR. LYNCH: Oh, okay.

(Laughter.)

DR. LYNCH: I think every new technology brings promises and threats, and I think that where we need to get to is to have a very good understanding of what those are. And I think even from the conversation we've had this morning, we know that there's not even a clear definition of what artificial intelligence is.

So we need to get to a point where we understand both sides to find some kind of balance, and I think that's where we always try to end up, with a balance.

MS. WOODS: Well, tough to say anything after - yes, of course we always want a balance and obviously agree with that. And so I don't think that I want to talk in terms of threats, but in terms of making sure that we look at all the opportunities.

So I mentioned accessibility and inclusiveness. Let's not leave that behind in what I think are right now mainly commercial developments. And that's fine. That's a good thing. And usually throwing - thrown off from that can be some very good developments in other areas as well, but keeping in mind social inclusiveness, whether that's developing countries or individuals who have a need, there's a huge opportunity here for this technology.

And I'd also say let's be careful not to get so worried about the threats, the liability, the definitional issues, the competition concerns, that we don't fully explore the opportunities that AI brings.

DR. TILL: From my view, a lot of the answers, what we're doing here is engaging with the topic. I think we all know that we don't have the answers. I think the only wrong thing to do would be not to ask the questions. So I think for me the balance is while the conversation might not be clear, while we wish we'd have the answers to it, actually feeling that uncertainty but having the discussion will go a long way there.

MS. STRONG: No, I agree. And I thank you all for your contributions today and for not only answering the question, but asking the many questions. We look forward to continuing this conversation as we welcome everyone to participate in the WIPO survey.

And I really appreciate, Dr. Till, your statement that the copyright seems to be not at this - at the center, but throughout the entire discussion of AI, and I think that's really important for everyone in the intellectual property community to understand writ large.

Thank you, everybody. I think we're going to be taking a short break. Thank you so much.

(Applause.)

(Whereupon, the above-entitled matter went off the record at 11:00 a.m. and resumed at 11:21 a.m.)

MR. ASHLEY: Ladies, gentlemen, I think we have successfully remedied some technology problems. So we are ready to resume. So if you would, please, join us. I'm glad to hear the buzz in the room and I know that it was all about this morning's brilliant and stimulating presentations you heard.

This panel we have this afternoon -- or this morning -- is very appropriate. We have three presenters who combine lots of legal knowledge, lots of information about -- technical information about artificial intelligence, and at the same time can contribute to today's discussion a general and very much needed information about how artists themselves are using artificial intelligence, and responding to the ethical, the commercial and the economic issues posed by artificial intelligence and other issues.

So you've met Ahmed. And Ahmed at the appropriate time will present some responses to some of the earlier issues you heard this morning. He has some new points that he wants to make about artificial intelligence, especially how he reacts to them and how he observes that authors and artists are reacting to them. Sandra Aistars is a principal at the Antonin Scalia School of Law. She runs a clinic there. That clinic is friend to -- supporters of assistance to individual artists, teaches them law, helps them preserve their rights, among other things. So she has much information about how the individual artists are using AI, reacting to it in much the way that Ahmed does.

A third presenter is with us by telephone. Thank goodness for technology, right? He is a hostage of many of the passport issues that are raging around the world these days and so could not be with us. But Andres Guadamuz is a professor at Sussex. In addition, he is an artist in his own right who has both exhibited and presented in exhibits that are artificial intelligence related. And he's a researcher. So a geek in several ways, as Ahmed is and to some extent as Sandra and I are. And so we will proceed with each of the three presenters taking roughly, but hopefully not more than about 10 minutes to make their individual presentations. We will begin first with Andres, then Ahmed and then Sandra. We hope that we reserve a bit of time at the end for some inter-panel, as well as -- conversation as well as conversations and responses to questions and comments that you have. So with that, let's hope the technology brings in Andres so that he can make his presentation.

MR. GUADAMUZ: Okay, thanks. Can you hear me?

(Simultaneous speaking.)

MR. GUADAMUZ: Okay, excellent. So we should be on the first slide now. Thank you very much for the kind invitation. Apologies for not being there in person. I will try to be brief to let the other speakers have an interesting conversation and be able to foster some questions. First -- can we go to the second slide? Okay. It may be a little bit trite say that this subject has become extremely interesting -- you have already had a very fruitful morning, from what I hear -- and we may have caught a lot of people by surprise. We were just recovering from simian copyright and all sorts of questions about Naruto when the issue of artificial intelligence became even more trendy.

For those of us that are looking at this since almost the beginning, the options for protection have always been very clear. First there is the answer, none of these works should be protected under copyright. They should all be in the public domain because there is no originality and there is no creativity. Only humans can create copyright, therefore all of these works should be in the public domain.

The other option for those countries with -- that have some form of registration is not to register the works. And I will let U.S. experts talk about these specifically. An interesting third option is on the table. This option is to recognize that these works deserve some form of protection and therefore either recognize that with copyright, or with some type of sui generis right. This is not, obviously, AI rights, or we are going to allocate rights to artificial intelligence. But this would be something akin to the EU database right.

All right, can we go to the third slide? The main proponents of the public domain option look at Europe as an example of where things are headed to think about this as not protection, or it being in the public domain.

The requirement for originality in European law is that a work can only be protected if it's the author's own intellectual creation reflecting his or her personality. This is well established in the directives and also in case law. We've had several cases that have elaborated this several -- this element.

So also, what conveys this idea of the author's own intellectual creation tends to be a combination of composition, it may be originality, et cetera. It is very unclear if some parameters and algorithms would be enough for -- for this. And so, for example, if you have -- if your parameter -- if you're an artist and you are selecting things that should go into your work, or you have lots of outputs and you select one or two that look very good. Is that enough to convey the author's own intellectual creation, or is it not?

So we have those questions. And can we go to the next one, slide 4? UK law has not been the same as the rest of Europe. And you can insert your Brexit joke here.

(Laughter.)

MR. GUADAMUZ: Of course, there has been quite a -- quite a big change between -- or a divergence between UK law and European law. And this actually predates Brexit. What has happened with the UK is something very interesting. We have had in UK law, in Section 93, paragraph 3 of the Copyright Designs Act, the CDPA, this very interesting formulation.

In the case of a literary, dramatic, musical or artistic work which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken. So you can see first that we have this idea of computer-generated work. Now there may be a question of whether or not this is enough to cover things like artificial intelligence. People like me think that it is.

And the other question would be, who should this go to? Is that the programmer? Is that the user? And these questions to open both Parliament discussion before the law was passed, and the only one case that we have leave me to believe that we think that this should cover the person that made it, also for the work to be created -- therefore the user, not the programmer.

And we go to the next one. That's slide five. Something really interesting has happened just this year. Earlier this year China surprised everyone by becoming the first jurisdiction to rule in this question. While it doesn't pertain to artistic works, I have included it here because it is still very relevant and very interesting. A court in the Chinese city of Shenzhen has decided that an article that was written by an artificial intelligence program has copyright protection. Now this is from Tencent, the big Chinese tech giant, they have something called Dreamwriter, and this is a machine learning program that writes about half a million articles per year in subjects that range from sports to technology news to financial news to sports -- sorry, to some basic news. You can also insert the joke here about how sports is actually not very difficult to write about.

Now, the court agreed that the article has copyright. Sorry, the -- a competitor published an article from Tencent. So Tencent sued for copyright infringement. The defense was that this was in the public domain because it's not protected by copyright, this was created by a machine. And the court actually said that this article had copyright, and that it was original. Therefore Tencent is the owner and it is the first work that we know that has been allocated copyright. Okay, can we go to the next one?

Our -- the opinion of artists in my experience has been almost to ignore copyright. I've been talking about this subject for quite a while and I've met quite a lot of artists in Europe that are dealing with this subject, and that are creating things with artificial intelligence. And they tend to be really not preoccupied whatsoever with the question of copyright, originality or copyright ownership. They're more concerned about whether or not they will be infringing works when they are feeding the artistic intentions of the machine learning algorithms, and whether or not doing this is actually going to be infringing copyright.

So the question is not so much the one of ownership and authorship, but one of infringement. And that's, I think, where probably we're going to be headed next. Can we go to the -- slide seven, concluding? I would like to leave you with a more philosophical question and that is, why do we protect artistic creations in the first place? From helmets to sofas to pictures to red buses, the courts in the UK have struggled with the definition of art. So we have very interesting case law, for example, dealing with whether or not the storm trooper helmet pictured here is a sculpture. Or whether a sofa is an artistic craft. So the courts have been struggling with this. But lately, what seems to be the question -- the answer, is that the author's own perception, or intention most importantly, is what is needed for something to be artistic. A helmet cannot be artistic -- this cannot be a sculpture, because it is just a prop. It has a utility. It's meant to do something.

Now why do I bring this to this question? Let's think for a second if we can have things that have a very low threshold of originality protected by copyright. Then maybe we can stop thinking about the author and concentrating on the author so much and start thinking instead of originality. Maybe it's a question of originality in the intellectual creation sense, in the European sense that something has to be an intellectual creation. If the work would be original by means of at least a modicum of intellectual creation from the start, maybe in the shape, of choices of pictures, selection of outputs, and even some sort of programming that goes into -- into the program, why shouldn't we give copyright to this work? As long as the process is not entirely random and mechanic, then there is a good reason to think that some sort of protection should be allocated.

Think of, for example, the Next Rembrandt. If you know the case, there was a -- a team in Amsterdam that created this amazing painting that was a machine learning version of what a Rembrandt portrait should look like. The -- the programmers and the researchers and these put a lot of work -- a lot of work, a lot of thought into the creation of this painting. So that to me is enough to convey originality in the sense that this an intellectual creation that reflects personality of the author.

And maybe the -- the final execution is not -- wasn't created by -- by one artist. But I think that there is enough to give copyright to someone. And I think that that's where we should probably be heading. Thanks very much. Now you can go to the final slide. Thank you.

(Applause.)

MR. ELGAMMAL: Okay, hello again. I'm not going to, of course, repeat what I said earlier. I just want to follow up the discussion from where it was left now. I want to talk about what artists do with AI and amazing -- and a lot of things is being done using AI in the last three years. In particular we created a platform for artists to -- to allow them to experiment with AI. And the motivation is artists in general are usually not technology savvy. Most artists don't know how to write code or change code or even run code. If you go to one of the code sites that can find codes, for example GitHub, you can find 27,000 versions of GAN, that I talked about this morning. So an artist who is just starting there cannot even know the difference between them and what to write and what to use.

So we create a platform that allowed artists to create art will experiment with this kind of technology. And I want to show what artists have been doing and different ways of interacting with this. Even some schools, like NYU Art School, have been using this now.

So this is an example of an artwork that the artist -- his name is Devin Gharakhanian -- used photos of Charlie Chaplin, fitted to AI algorithm, and generated this abstracted surrealist version, and painted it -- so just painted the outcome. So he and others are using the algorithm and finding, selecting an outcome and painting it. A very standard way. Another example of an artist, is Marco (phonetic) from New York, who had basically fed the -- the AI his own photos and fed it also his own style of artwork, how he makes art. Engineered this stylization of photos based on his own style, and then printed it and then painted over it. So you see -- so basically this work involves choosing the art, feeding through the machine, with the -- manipulating the input and output and hand painting it.

Here's other example of an artist --

her name is Anne Spalter -- she actually was an artist on -- who for a long time, actually had been doing digital arts. And she used the platform and she was amazed by what it can generate. But she wasn't happy with that -- with the quality. There was a -- it was very small, actually, one inch by one inch basically. And she -- actually painted based on the idea that the AI give it to her. So this was all painted by -- by -- on canvas. And she was thinking that for the first time in 25 years I will go back and paint it using -- myself, using tools because I had been digital artist. So the AI inspired her and her ideas of what she wants to paint.

So the AI is an idea generation, totally different. This is again some of her work. And finally, an -- Carla Gannis who is a professor at NYU who is actually using also these kinds of tools for many amazing ideas, try to convey a concept.

For example, she made a competition between herself as a human artist and avatar artist, it's called C.A.R.L.A G.A.N. Funny, her name starts with Gan. So in this competition, basically, she post -- posted works by herself and post works by -- AI-generated works based on her own inputs. And after a while you could not tell which one is the real person and which one is the avatar person of the artist.

And she did amazing, for example, virtual reality experiences that use AI in the making. So this is just an -- some examples of what artists would do. So, as you see, basically artists are using AI now as a tool. However, it's a very different kind of tool than -- because, all this AI technology comes around, artists using it. From the invention of oil paints to print making to photography, always artists take notice and use this technology.

What's different here is the tool for the first time becoming -- has some element of creativity, it can surprise you even as an artist from -- and that can lead you to a new direction. And here is the -- the issues that we are discussing today.

And in terms of authorship, there are different versions of using this process. Either you can -- you can be technology savvy and have control of the inward, the algorithm and the outcome, where you control the whole process and -- and this way -- that it deserves wider protection. Because if you develop an algorithm or a post that can generate your own art and it can generate an infinite amount of art, there's no sense of copyrighting or asking for copyrighting a single output of that. Because this machine, basically, you can generate an infinite amount forever.

And at the same time it is very hard, as we describe here also earlier, to patent this process because it's an algorithm and a software, which is very hard to patent. So we come into a point where it is very hard to protect the outcome and very hard to protect the process. And that means some -- some -- different ways of looking at -- at things.

And the other way of doing things, if you are not technology-savvy and just using something that is already made, and you're just pressing a button and generating more or searching what can be generated and finding something or combining things to generate in something that you like. There is some creative process here also, but here it's more into what a photographer would use a camera and create something of their -- of their own.

Anyway, I -- this is just some of the -- the feedback I was giving over the discussion that has been in the morning. One last thing I want to discuss also, open the question about the use of data in the process. Suppose you use data by a living artist, or somebody who is protected. Take Warhol, for example, art and feed it to the machine to create something. What is the situation here? I basically see there is no difference in that from a living artist who actually go to a museum or take out a Warhol book and look at all the photos and digest them and make a creation at the end. Because the AI at the end, we are creating based on data, but in the same time, as long as that creation is not a direct derivative and it's transformative enough, it's -- for me I see it as no different than a human digesting that body of work and making a creation. So I don't see a reason why to treat the machine, or the AI generation -- authors using AI generation here differently. So this is just my comment about this. Thank you.

(Applause.)

MS. AISTARS: Should I just jump in? Okay. Well thank you again for inviting me to present here and for the stimulating conversation we've had so far. As was mentioned, I run a clinic at George Mason University Law School's -- Scalia Law, and it is a program through which my students and I represent individual artists and small businesses in the arts, advise them on various copyright issues, help them protect their rights and so forth.

Through that work and through simply

the family I've grown up in and the community I have interacted in, I've had the opportunity to speak about AI issues and just arts issues in general with a variety of artists and a variety of fields and a variety of generations. And I have to say that, you know, there are many issues they are reacting to and considering, much like we have been raising today. But also much like today, they have more questions than answers. So I'm glad that I am not the only one starting my presentation with more questions than answers.

The Next Rembrandt Project was mentioned by Andres. And I thought it might be good to play the video so that everybody is starting from a common understanding of at least one of the ways that machine learning and AI is being used. So, if we could hit play on that.

(Video played.)

MS. AISTARS: Thank you. So let me ask folks in the audience, how many of you think that this is brilliant? How many of you think it's troubling?

(Laughter.)

MS. AISTARS: Seems like an even split. How many of you think it's brilliant and troubling?

(Laughter.)

MS. AISTARS: All right, well, let me tell you that you're not so different from the arts community on this. And to me, you know, as in many discussions of fair use, it's really how the work is being used that creates that feeling of, oh, this is brilliant or this is troubling, for the artist. So this particular example, I would say -- because it's based on Rembrandt works which are all in the public domain, and that's all that the AI was fed -- it may not raise many copyright issues at all. But think about if we were talking about the next Warhol. Or you know, identify any current living artist who's -- who's selling their work.

That seems to raise quite a few more issues. There are various ways that we can deal with the ingestion of works. And I think that may be the place to start the discussion because that's certainly the place where an artist whose work may be being used to teach the AI starts. And so there are various possible policy answers.

You could consider it to be fair use because you're never actually reproducing the work that has been ingested publicly. And so maybe you would draw an analogy to the Google Books case. That may be troubling to some degree because the Google Books case, I would say, presents a completely different ultimate scenario which -- you know, arguably the court found benefitted the artist at the end by making the work, you know, find-able and purchase-able. And I view the Google Books decision as a narrow decision which, you know, requires a variety of factors to be in place before we would consider it fair use. But that's certainly one -- one option you could take.

Another possibility on the other end would be everything needs to be licensed that's fed into the work. You know, that's -- that's also not an unreasonable position. And then I guess there's a position in the middle which is it's fair use, so long that the work that's produced in the end is automatically public domain -- which I think is what Andres was suggesting.

But clearly, if you are the artist and the next Warhol is what comes out at the end, isn't that actually the worst of all possibilities for you? Because the end result is competing with your -- your creative work. So even something that seems like a potentially reasonable policy analysis on a commercial level doesn't necessarily, you know, play out that way from the artist's perspective.

But let's look for a moment at the resulting artwork itself and how we feel about that artwork. And I just -- I have to read you this. And once I read it, you'll see why. So this is a UK art critic, Jonathan Jones, and he wrote about this project in The Guardian. These are his words, not mine.

"What a horrible tasteless, insensitive and soulless travesty of all that is creative in human nature. What a vile product of our strange times when the best brains dedicate themselves to the stupidest challenges, when technology is used for things it should never be used for and everybody feels obliged to applaud the heartless results because we so revere everything digital."

So my reaction to that was, tell us what you really think.

(Laughter.)

MS. AISTARS: He went on, and you might appreciate this, you know -- explaining why he feels this way. And he said, you know, art only has meaning if it comes as a historical record of the artist's encounter with people and with beliefs and with the anguishes of the time, and that great art is not just a set of mannerisms that should be digitized.

He also then suggested that the AI should go to bed with Rembrandt's lover first before trying to replicate Rembrandt's work. And that also the AI should experience plague, poverty and old age. So --

(Laughter.)

MS. AISTARS: So you can see that these issues are not going to be any less controversial than any copyright issue we have ever considered. So, but just to be serious for a moment -- not that that wasn't serious -- but I think what the critic here is raising is really the intense moral rights issues that artists feel around the works. And, you know, as -- as you from the EU and UK know more than I -- you know, there are more than 60 countries internationally that recognize moral rights of attribution, of disavowal, of publication, of withdrawal from market and of modification. And even here, in the United States, at least 9 states that have some form of protection for moral rights. And that comes most often in the realm of visual arts. So this is a great place to think about it, but also, I think, explains the outrage that some in the community express.

I'd caution though, even where there are protections here in the states, those apply to works that are produced in a single copy, or in the case of photography in an exhibition print. And so, certainly the end results of this product likely would not qualify, if you're going to be able to run off however many next Rembrandts or next Warhols on your 3D printer as you'd like.

So the last thing I will say, so that we don't run totally out of time, is that there are other common law issues that we want to consider too, as we are considering these works. You know, any rights that might have been created by contract, rights that arise by tort or fraud, so-called passing-off concepts. And the one question after looking at this Next Rembrandt Project that I had myself was, is this just a, you know, brilliant forgery machine? Because if you listen to how the project was brought about and how the AI was instructed, that's actually what art forgers in the analogue world do. And, you know, while we may recognize that the resulting work is a work of art, people have always had very complicated relations to forged works or works that are created in a style of an artist.

And so I would ask not just what's the impact on the artist, but what's the impact on the art-buying public? Are there trademark-like concerns? Are people being swindled? And I think we will have very many interesting discussions around that. And I'll stop there so that we have some time to actually talk.

MR. ASHLEY: At least eight or nine minutes for general conversation -- including with you, in our audience, if you have. But let me, as I usually do, play devil's advocate here. Because there should always be a devil in the room. Let's assume just for the sake of discussion that Rembrandt survived today and his works are protectable, all 10 zillion of them. And you raise an interesting moral rights slash ethical issue, slash forgery issue. If my data mining is only to do the type of studying of all of his currently protect-able, in the hypothetical, works. I am not competing with those works. I am not copying and reproducing those works, at least, not reproducing and distributing them publicly. I'm merely for the process of research, doing what computer programmers do all the time with others' computer -- copyrighted programs, which is borrowing content, using it to advance the state of the art. In this case, I happen to be using it to produce -- advance the state of the art of visual artists. How is that scenario different from what the data miners do in the Rembrandt case?

MS. AISTARS: So you're not actually --

MR. ASHLEY: I haven't gotten to the --

MS. AISTARS: -- printing off a new --

MR. ASHLEY: No, I haven't gotten to that point. It was just -- just --

MS. AISTARS: It is actually a good question, and one that I was going to mention but didn't have time to. Another application of AI has been to teach AI machines to distinguish forgeries. And so in those instances, at least the one example that I am aware of, they weren't feeding the entire work. But they were breaking -- ultimately, they fed the entire work, but they're breaking it up into tiles which are fed individually in to inform the AI machine. And that, for whatever reason, makes it more likely that the AI will be able to detect a forgery because they are looking at so much more, you know, data I suppose. You probably have a much better reason and explanation why it works.

But, I mean, there I think as a matter of copyright law you'd look to things like Google Books and say are you actually ever putting in output? Are you harming the market in any way? You know, how much of the work are you using? Is it just a -- you know -- but -- so I would feel more comfortable with that sort of a use than if you're at the end of the day printing off, you know, the new Warhol, as I said -- or the new whatever, you know, today's, you know, most popular artist is. The new Banksy, perhaps.

MR. ASHLEY: This is really -- this is really -- this is really Ahmed's rhetorical at the end of his presentation just now where he really challenged us to think about how data mining in the hypothetical I just created -- scanning, uploading, saving is how I'm defining it, data mining -- is any different than the artist actually data mining by living and interacting with encyclopedias, websites, the living world. Pulling from that his sense of Rembrandt's style and then mimicking it in the way that artists do -- or some photographer.

Okay, so you have some point that I am sure -- that Andreas wants a piece of this as well.

MR. ELGAMMAL: Yes, thank you for bringing this Rembrandt example. I think we have actually two different examples, the Rembrandt case and the artwork I have shown using the artists. So the Rembrandt case is brilliant in making forged art, basically using tools available today to make some art in the style of another art. Basically, that's not known as -- to be art, to start with, if you do that. If you do something in the style of Van Gogh, that's not art. That's basically maybe -- that might -- think of it as decorative art for putting in your room, but not really art.

So an artist doesn't do that. A real artist tries to do something creative and new and - a new idea. And the example I have shown -- I have shown today was artists using AI. All of them are really brilliant in the way they're creating new things --

(Simultaneous speaking.)

MR. ASHLEY: But I am not -- I am more endangered by -- by that. That you would refer to creating something which is artistic as not art.

(Simultaneous speaking.)

MR. ELGAMMAL: -- as art, but as --

MR. ASHLEY: -- so that is a more -- that is a more dangerous proposition to copyright than it seems to me AI is.

MR. ELGAMMAL: Yes. As I mentioned in my talk earlier today, what will drive progression of art is really innovation. So if you now or even hundred years ago try to copy Van Gogh, say, or Monet that already had been done before, you wouldn't be considered an artist. You wouldn't consider something -- whether to make some -- some beautiful images in the style of something in the past you can -- you can put in your room. But that wouldn't put you in a gallery or a museum anymore. That's not what the art world considers art.

(Simultaneous speaking.)

MR. ASHLEY: Okay, well would it be --

MR. ELGAMMAL: -- that's not what -- I mean, you cannot copyright that in such ways because it's not -- with no innovation or creativity behind it.

MR. ASHLEY: But that's a high standard. We -- we -- copyright protects trash too.

(Laughter.)

MR. ASHLEY: Right? Very, very low standard. And the reason -- the reason for that low standard is a very good one. It is to avoid the aesthetical and, it seems to me, the moralistic judgments of deciding that's not good enough art, or innovative enough art, or whatever. So my point about trash is to preserve the principle that it seems to me are -- now, real quickly -- I am supposed to be wrapping up.

Real quickly, I wanted to -- and again, a friendly devil's advocate challenge here. And to you, as well Andres, if you want. The -- the sort of categorical treatment of the product of AI, at least in the Rembrandt, as really good forgery. They forged a style. Mr. Rembrandt did not own his style. He may have created one, but he alone did not own it. And that is true of all of the innovative people who created categories. That is an essential proposition of copyright. That I can take your style -- I can research and reverse engineer to figure out your style -- even if by AI. And then I can use your style to produce things. That runs up against the moral rights issues that you were talking about.

MS. AISTARS: Right, so that's why I actually -- and I say this in a lot of instances -- I don't think copyright needs to address every single issue. And I don't really view this -- this Rembrandt image as a copyright issue so much as a either, you know, consumer fraud or, you know, passing off or -- you know, contracts issue. And I -- I would urge copyright not to try to take on too much. Because we will get into these arguments, whether, a la Blurred Lines, or -- or, as you were pointing out, you know, is this art? Is this not art? I personally don't want judges making that decision, either for me or for my clients.

MR. ASHLEY: All right, we will let

Andres have the last word. And then that will be it.

MR. GUADAMUZ: From a -- I guess from an international perspective, there is no question that this is -- for example, this would not be infringement if we were thinking about that. We would -- if Rembrandt was alive, even, I don't think that there is enough there to -- to say he has this. You were saying he has no ownership over his own -- his own style. There is no such thing as a Rembrandt-ness copyright.

But from another perspective, perhaps that we -- we are not going to see this brought to court in things like art, I think. It is already -- it is happening and it is going to happen more and more with low-level music, a very large maker of artificial intelligence music, for example, Jukedeck, just got purchased by TikTok. You can probably see where this is going. This is going to free music for people to -- to put in their videos instead of having some problems with copyright.

This is going to have -- to going to the fourth point with really low-level type of art. Maybe computer art, computer-generated art that is -- can be used as a background in a game, or even a movie or something. If we are not going to protect this, everything is going to go into public domain. Then the artists, musicians, journalists, are going to be competing with free works. And this is where I think we should think really hard about where we are going.

I'm not concerned that this is going to become an ownership issue. But definitely a -- or an infringement issue, but definitely it's going to be a -- litigated, and it's already being litigated.

MR. ASHLEY: Well thank you for that. And the Copyright Office and WIPO, like to thank Sandra, Ahmed and Andres for their very thoughtful participation. Thank you.

(Applause.)

MS. ROWLAND: Thank you so much. That's such a fascinating panel discussion.

We're saying goodbye to Andres, I believe. And we're here to take on the next panel, which is about the world of other works and AI. And we're going to be talking about interesting things, like literature and video games. And so I would like to welcome to the stage the panel for the world of other works.

MS. ALVAREZ: Hello everyone. So welcome to AI and creating a world of other works. Also the panel right before lunch, so it will be a good one. Obviously, AI is being used to create a wide variety of works. We just heard about visual art. After lunch we are going to hear more about music. So now we are going to hear from people who work in some other areas, specifically video games and literary works. I am going to have everyone give their own presentation and talk about AI and how their own works tie into all of this. And at the end, if we have some time, then we're going to have some Q&A. So I will just go through the lineup. Not everyone is quite in order, but Kayla is going to be first to talk. Kayla Page, she is Senior Counsel for Epic Games. Then we're going to have Jason Boog, who is the West Coast Correspondent for Publisher's Weekly. Then Mary Rasenberger, who is Executive Director of the Author's Guild. And then Meredith Rose, who is Policy Counsel for Public Knowledge. And Kayla is up.

MS. PAGE: All right. Hello, so Kayla, Senior Counsel, Epic Games -- small video game company that no one has ever heard of and definitely isn't in the press all the time. So Epic is kind of at an interesting intersection of AI because it is both a game developer and publisher with that little-known title, Fortnite, that everyone knows. And also a game engine distributor. So in addition to actually creating immersive game play experiences and all of the AI that you may utilize to ensure you're enriching game play or matchmaking players effectively, they also support, via the architecture, the game engine, the actual build-out of this.

So we're in a kind of unique position

to watch how these items can play out from the programmer's standpoint of the AI, all the way into the final creative output that lands in the game. And just a little bit of background, I don't know how familiar everyone is with software engines, but -- so games run on a game engine. You need an architecture below you that allows for efficient game design. So you need a rendering engine, right, for your 2D and 3D graphics. You need audio. You need physics. And you need AI at this point.

I think that it's become a hot topic lately, but really video games and AI have been just inextricably -- that's a word -- linked since the 1950s when the first computer program was trained to play chess. And we have been pitting AI against human experts to try to, you know, learn from that for forever. You go back -- even Pac-Man has the fundamental essentials of what is now environment-based AI. It is going to detect where you are. It's going to predict where you're likely to go. And based on those two inputs, decide on the creative-ish decision of whether it is going to chase you.

And it's simplistic, but it's actually a really, really effective model. And so I like the model because for me, AI is a tool. Like, video game companies are absolutely creators. They provide the creative direction. They provide the narrative structure. They provide all of the inputs that, like, eventually that AI is going to utilize as a tool to create these outputs.

But -- it's an increasingly sophisticated tool, but the bones are all the same. And so I kind of analogize it often to artists -- visual artists in a video game company. They're using -- they're probably using Photoshop, right? That's a tool. And if you give a creative direction to an artist and every -- to 10 artists, and 10 artists Photoshop it, you're going to get 10 different results. And so it's driven very much by the creative decisions that are being made inside of the company.

So what I've enjoyed watching is the evolution. And I think a lot of the discussion here is, like, okay, well at what -- what percentage -- at what point are we going to say that there's been so much machine intervention that we've lost that human authorship that -- that has the -- typical, you know, copyright protections attached to it. And it's definitely evolving. You know, early on it was really static rules.

It was a deterministic environment. So you're playing chess, right? There are only so many moves in chess. The board state is known. And it's -- it's very meticulously written code at that point. And when I say code, it's really just if-then statements. So if the player moves the knight to B3, you are going to react in this way.

And now that -- that doesn't really work anymore, right? These predictable experiences are gone. We have massively -- you know, MMOs that have tons and tons of players all around the world. The environment isn't static. It changes. And even if the environment has been coded at a base level to support certain physics, player interaction is going to change that in a second. I -- it never -- I have, you know, nieces and nephews that are like 14 and 15 years old. And I can't predict that they're going to kick the cat in that game.

And so when that cat hits the building, which it shouldn't -- this is not my game. I don't have anything to do with this hypothetical game -- nor am I in an way advocating -- sorry PETA.

(Laughter.)

MS. PAGE: You have to consider, like, the physics of that -- the physics of what we now support in games. Of smoke. Of light. Of ray-tracing that actually enables us to have these really realistic engines that are capable of doing things like we see in Pixar movies and giving real -- real-time, real kind of humanity to these virtual worlds.

And so -- well, I look at it now and we've moved from this sort of static model to very much needing to utilize AI to ensure that things are functioning in game as they should be. And that's just from a game play perspective. From an engine perspective, we support much, much more. But you end up in a -- in a situation where you simply could not efficiently, effectively, cost-effectively code for every possible contingency when you allow people to interact in a virtual environment.

And so, you know, some of the things that we see significant development in are things like adaptive game play. So you suck at video games. You're like, I have never picked up a controller. I don't have any idea what to do with this. And if you enter a video game, and you are playing with these -- with the AI, with these bots that are non-player characters -- NPCs, if they just kick your ass every time, it's not fun. It's not fun at all. And you have to hit this sweet spot, right, between this is too easy and I am bored and I don't want to play and this isn't enjoyable -- to I'm just being embarrassed.

And so you get adaptive game play that allows players to either match at commensurate skill levels, or an AI that's capable of kind of engaging with your game play style and giving you a personalized experience so that you're not miserable the whole time. I heard somebody on an earlier panel mention accessibility. That's huge too. Natural language processing in AI is amazing. And, you know, I don't know if anyone is familiar with the numbers, but there are a staggering number of people with disabilities that play games because it's a really easy place to connect socially. It's like, our online environments aren't fake online life anymore, right? They're part of our real life.

So that sort of natural dialogue that happens is really well supported by natural language processing. We're not seeing a ton of that yet. It's pretty sophisticated. But there are many movers and shakers and -- in Microsoft and in Amazon that are also trying to figure out how we can harness some of that to remove, like, toxic players in chat. Get them out of there if they're harassing people, or if they're just being inappropriate.

Similarly with immersive game play, you want -- you want non-player characters that actually feel like they're not -- you know, they're not walking into a wall 16 times and you're just trying to interact with them. And it gives you this option to have, like, really believable and realistic body motion -- or dialogue options. Or they could react to you emotionally in a way that they just couldn't do that prior.

And in addition to procedural content, which is just anything in the environment that needs to render real-time as you walk around. So that's kind of a -- the high-level of those. I guess what I find interesting is the way that this is coded. Because I think that often AI is either highly, highly simplified in a way that doesn't make it seem -- it doesn't really let us get into the meat of the copyright issue. Or it's just this black box of, like, it's a neural network and we have no idea how this works. But it's actually -- it's pretty interesting to watch it from the engine side when you go to code it.

So if you look at Epic's Unreal Engine 4, it has a lot of AI that it supports. And one of the things that you can do is take -- let's take a very simple creature in a video game. We'll do an animal, right, because there's no dialogue or anything. So you have a cat in a video game. And you need this cat to behave in a certain way. You -- you have a creative expression. This is going to be a mean tom cat who -- devil may care -- whatever.

And you need to code certain things for his behavior. You have -- you have to have action tasks that exist. So here's your root AI, right? And then, you have a behavior that says the cat is going to groom itself. Okay. That's a pretty easy path. But you need things to make it seem lifelike. It can't just -- it can't be a constant 20 seconds, and then it grooms itself again. It's not very engaging. So you put conditions on it. Cat grooms itself. Cat also has a startle response, right? And you have to marry those two concepts as related. You have to say, okay, cat would not be grooming itself. Like, it's not lifting its leg and playing the cello if it's startled, because someone has thrown a bottle at its head. I am really mean to cats in this. I didn't actually prepare this at all.

(Laughter.)

MS. PAGE: I love cats. I have a cat. And so you have all these interrelated tasks that without AI running real-time -- because you think about the -- the 3-dimensional space we occupy as human beings. The environment we have to interact with. Body placement, movement, the way we react to things, the way that that reaction changes when it's acted upon by another force.

And you have to assign these all priorities, or -- or exceptions, or they're conditional. And so, without AI to actually help facilitate those sorts of -- of real-time, like, machine-learning feedback of, okay, we need to make this more and more realistic -- you just don't get the same sort of immersive narrative experiences.

(Applause.)

MR. BOOG: So I am going to -- I am a journalist and an author. I am going to speak about how writers use artificial intelligence right now. But I wanted to start by reading just a really short excerpt from some -- it's a collection of essays.

"I had asked my parents for advice. They were at odds over whether or not I should go to the upcoming poetry competition. I grumbled at first, but I started to take action anyways. I asked my father for some of his old writing supplies, and he happily lent them to me. From that point forward, I took to spending my days in the studio perfecting my craft. However, as the weeks passed by it became increasingly difficult to stay focused on my goal. I turned to alcohol and illegal drugs. And I found that instead of relaxing and reflecting, my thoughts grew more chaotic -- my actions, more destructive."

"But when the judges arrived to hear that poem, written by me, I was left with a bittersweet smile on my face. I was proud of myself for taking the challenge. And I wanted to continue this legacy and become a renowned poet."

I didn't write that. That was the absolute -- that was my moment of narrative singularity. Last year I worked for two months to create an AI that could actually tell a small little story. And that -- after reading hundreds and hundreds of pages of output, nonsense, all sorts of crazy stuff -- that was the first time that my little AI told me a story. So if you go to the first slide. I just put the text up there and you can see, I broke it into the three-act structure that you see in film. So you have act one when you introduce the conflict.

And act two, you have this young poet struggling with his craft. He even gets into drugs and alcohol, struggling to finish his -- to reach his goal of going to the poetry contest. And then you have act three when he or she actually goes to the poetry competition and delivers a poem.

And I cannot tell you how happy I was when I discovered that. And it was -- it was really magical. And I am not a coder. I worked with something called GPT-2, which is an open -- a model released by OpenAI last year. It's -- it was scraped from 8 million different webpages. And this language model -- basically, when you feed it input, it will try to guess what the next part of that story is. So what I did is I fed this -- this really super-powerful AI a whole bunch of very short stories from this section of Reddit called writing prompts. And it's where people write very short stories -- much like this -- and I just gave it thousands and thousands of examples of those and then turned it on and set it loose. And then I had to read hundreds of pages to find something like this. But when I found it, it was really magical.

So that's one example of how people can work with AI. Here is another -- this is Janelle Shane. She is an author who works a lot with different AIs and bots and things. And she took the names of spaceships from Iain Banks novels. I don't know if anyone reads Iain Banks novels, but they're amazing. And the -- one of the best things are these self-aware, AI-powered spaceships that fly around. And they all have really funny names like Prosthetic Conscience, and things like that.

And so she fed those names -- hundreds of those names to an AI, and it started to generate new fake names. And you can see some of them are up there. They're pretty hilarious: "Happy to Groom Any Animals You Want," -- which would probably be good in our video game -- "Surprise Surprise," "And That's That!" And then she puts it up there. And thousands -- hundreds and thousands of people love to read these.

She's got a very popular blog, and she wrote a book recently as well. But her work is really wonderful. And that's one example of what writers are doing.

Here's another one. This is from Allison Parrish who teaches -- she's a poet. She teaches at NYU. So she took a whole 3 million lines of poetry scraped from Project Gutenberg. And she put those into a collection and gave it a -- turned it loose on -- with a natural language processing AI. And that started to sort these millions of poetry lines into what they meant. It tried to come up with the best interpretation of what they meant.

And then she took descriptions of movies from Wikipedia and fed those into it. And so, what the computer did is it substituted for a synopsis of Star Wars -- like the -- the crawl that you see at the beginning of Star Wars where they say "it's a period of civil war," she fed those sentences into it and it gave you a line of poetry -- the line of poetry that corresponded most beautifully to that line. So you have this really excellent Star Sonnet: "What field of civil war whose lightnings make a terror of all Space. An army when its king has fled: for alms of memory with the after time, are slowly borne to earth, with a dirge of cries."

And so you come up with something really unique, special and -- I've never seen anything like it in my life. So -- let's go to the next one. And this is Robin Sloan who is a novelist who has two really super-powerful GPUs in his house. And he runs the same technology that I did for my thing, but he can do it with his own GPUs, which is pretty amazing. And he took 100 different fantasy novels; fed them into that same engine; and then gave the names of 1,000 of his subscribers to that AI; and then generated a personalized, customized fantasy story for 1,000 of his readers. And here is just one example. And you can see, it's -- it's a really -- it's a little nonsensical, but it sounds like a quest. It sounds like this fantasy story written about you. And he sent that to 1,000 of his email subscribers. And people shared it all over the place. It was really special.

And then, I work with Film Threat magazine. And this is AI generated non-fiction. So using another engine, we fed hundreds of thousands of movie reviews into, kind of, the short, sort of user summaries and things that people leave on movies. Fed those to an AI, and then we came up with -- it's actually -- we invented kind of -- we call them artificial intelligence reviewer. And so we started sharing movie reviews from -- artificial intelligence reviewer. And they're really special. And here I am just going to read just a little bit of it.

This one is about -- I recommend this movie -- it's Pain and Glory, Pedro Almodóvar's new movie. And this is what he said -- this is what our AI reviewer said - "Pain and Glory, a movie by Spanish director Pedro Almodóvar that's filled with beautiful moments and some of the most intense, beautiful scenes I've seen in a long time. The last act is nearly non-stop action and big explosions. It's epic, awesome and bloody awesome. Feels like a cross between the Matrix, Matrix Reloaded and Indiana Jones and the Temple of Doom." So -- I've seen it, but I don't want to spoil it for you. But --

(Laughter.)

MR. BOOG: And then finally, this is called AI Dungeon. And this is an interactive game that you can play. You can download it on your phone for free. This -- this college kid named Nick Walton built it using GPT-2. And it's basically -- it's loaded with hundreds and hundreds of -- it was trained on hundreds and hundreds of story games where you say, "I would like to go west," and see what happens. And you go -- it goes, you walk west. You find a treasure box. And you say, "open treasure box," you open the treasure box and it says, "inside the treasure box is a trap." You just died. It's that sort of storytelling thing.

But this is AI-generated. And it will literally never end. This is just an example of what it was. But this -- this app will literally feed you story until your phone dies. And it's really magical.

(Laughter.)

MR. BOOG: Or you die -- yes. And then finally, I just -- about that Rembrandt thing that we watched, the thing that makes me the most upset about it is it's all smoke and mirrors. Like, that was a highly produced video by a bank. And it's like, they try to pretend like they're doing something really sophisticated, and they're not. Like, anyone in this room could learn over the next month or two how to do what those people are doing. And and it's just out there. And I think it's important for you to learn about it. I think it's important for kids to learn about it. So I introduced my kids to GPT-2. And we took one of the stories that that the story engine had made. And then we used the GAN, that same technology that we saw in the art thing, to generate computer images and to illustrate the story. And we picked them ourselves. And I showed them how the GAN works, and I showed them how GPT-2 works. And we actually made these little zines. These little physical things that we can share with people in the world. And I think that's really important. Because I want my kids to be comfortable with that. And I want everybody to just know that you can handle this stuff yourself. And it's coming. And there's a world where this stuff is going to be in everyone's hands. And I think that's important.

So if you come up to me afterwards, I'll give you -- I brought a few copies of these, in lieu of cards. But thank you.

(Applause.)

MS. RASENBERGER: Okay, thanks. Well, I am going to be a little more boring. I am going to talk about copyright, actually. As you might guess. But you know, before I do, I just want to say, you know, artificial intelligence is bringing a lot of really cool new tools to writing and to the writing profession. And we're all already using some of that. Spell check. Google completes your emails now for you if you want them to. I am looking forward to artificial intelligence that -- where I could just, like, dump all my ideas into it, and then it organizes it. So -- I am sure that's coming.

But I do -- because Jason spoke about GPT-2, I do want to mention another example of somebody using GPT-2. John Seabrook -- I don't know how many of you read his article in The New Yorker, The Next Word. He did an experiment where he took 20 years or so of New Yorker articles and fed them into GPT-2 and then gave it a prompt to see if it could -- the experiment was to see if it could write a New Yorker-style article. And the prompt was an article about Ernest Hemingway.

So I just want to read you what the AI writer did. As John said, the resulting language had imitated the cadence and narrative rhythms of

The New Yorker. You know, I said, when I read it I was like "God, this is a New Yorker piece. This is amazing." But let me read you the first couple sentences. It's -- it describes the -- the AI writer describes "walking up to Hemingway's gate in Cuba and seeing a dog who had been a common visitor before the war gallop up a path to the main building with a tiny cow standing by her side. There is a puddle of red gravy in the front yard."

So -- you can see -- it's like -- these are very good sentences, but there's a lot wrong here.

(Laughter.)

MS. PAGE: Where's the part about the cat?

MS. RASENBERGER: Yes -- I'm sure, if you go further, there's some cats because Hemingway did have about 30 cats when he lived in Cuba. But I want -- you know, it's like a dog has been a common visitor before the war? The dog is galloping? The tiny cow is standing next to the dog that's galloping -- it's like hard to figure that out. The cow is tiny? Anyway. And then, of course, there's the puddle of red gravy in the front yard -- which, you can see how that might happen, but --

(Laughter.)

MS. RASENBERGER: So anyway. I do, though, want to focus on a couple of, what I see are, really leading issues in copyright regarding AI. The digital age, as most of you who are in the copyright industries are very aware, has already been really hard on creators. On, you know -- authors' incomes are down over 40 percent. The mean income of writers today is -- from their writing -- is around $20,000. I know it's been hard on songwriters, photographers. But it's nothing like what AI is going to bring. And what I want to talk about today is the fact that there are some -- that our current law does not adequately address a couple issues. And that if we don't proactively address them, we are going to see the creative industries further decimated.

After all this fun, I hate to be so negative. But there you go. So these are the two issues. One is the -- allowing the unlicensed use of copyright works to train AI under fair use, an issue that's been mentioned by a couple people already. And the second is the current volitional conduct in secondary liability cases where it's very possible that in the future we will have AI infringing, and there's no one to sue, there's no one to hold liable.

So first, when you talk about the -- the ingestion issue. So you know, as we've already heard, for AI to generate -- an AI machine to generate new works -- and I will focus on literary works -- it has to train itself on large volumes of existing literary works. Right? So that's how -- that's how you -- and we saw this with -- with the new Rembrandt.

So it's copying those works. Now, the output, though, is not infringing. The New Rembrandt -- and you know, infringement is an analysis of does this work -- is it substantially similar to this work? The outputs are really -- they're -- in a sense, this is a very broad, sweeping way of describing them, but they're mash-ups of a lot of different works that they've been fed. And, you know, based on rules and parameters that they're given. But like the New Rembrandt, If Rembrandts were still in copyright, I don't think that you could say that the New Rembrandt necessarily infringed any true Rembrandt.

So -- now the -- we have to look at the actual -- the initial copying. What some -- cases are called intermediate copying, and whether this is fair use. And I think there's some who argue that all of this should be fair use. This is -- you know, this is intermediate copying. Look at the existing case law, which Sandra spoke a little bit about.

I mean, this goes back to the older case in copying works in order to basically reverse engineer them was fair use. Google Books we see -- the Court found that the end product was not infringing, it was just snippets. That was fair use in terms of what was human readable, but there were millions of copies made. Many copies of millions of different works made in order to get to that end product.

And in that case the -- the -- and as well as the HathiTrust case, the -- these -- they treat this mass reading as though it weren't real reading. That somehow because computers are doing it, it doesn't count. And companies, when they use computers to copy, you know, that's -- it's -- they're given -- they're given a pass.

So -- and you know an issue is that there is -- there's value in those works, even if humans aren't reading them there's value in them. Obviously I don't agree that it should always be fair use, because what about when you're ingesting works to create competing works? And we can definitely see a future -- it's not that far off -- where a -- maybe it's a more advanced version of GPT-2, ingests a lot of romance books, though in fact, the -- all the romance novels have already been ingested. And you can write new romance novels. Romance novels tend to follow a particular formula.

And I'm actually not sure this isn't already happening. Kindle Unlimited has -- it's a -- basically a subscription service which has a lot of self-published books in it. And many of them are genre books, particularly a lot of romance books. You already find in there -- there are all kinds of scammers. And the scammers are not authors, they're people who are just really good at scamming the system.

So they take existing books, they put new covers on them. Amazon has pretty much put an end to that. They change the words around a little bit, maybe new characters. And in the last year or two, now, we see these called cut and paste books where they actually will take different passages from like 20 different books and put them together and create a new book. And using clip farms they manage to get these books way up in the listings. And once you're high up in the listings on Amazon, then it's -- then you're being promoted to readers and you - you sell even more books.

So I think that there's a real potential here for using AI by copying existing works to create new works. And something that we need to protect against. So I read a lot of the submissions to the PTO's notice of inquiry. And I think a lot of people in the industry said, well, you know the fair use -- fair use will sort out the okay uses from the uses -- the competing uses. Because if you're -- if you actually scan -- not scan but copy works, ingest a lot of works in order to create competing works, that's not going to be fair use.

And the argument is that one, it's not transformative -- which I agree with. And two that under the fourth factor, the courts will find that it does negatively impact the market for the works. The problem is --

MS. ALVAREZ: I don't want to interrupt

too much. We have about ten minutes left.

MS. RASENBERGER: Oh, okay.

MS. ALVAREZ: I want to make sure that Meredith has a chance to go.

MS. RASENBERGER: Okay, okay. So let me -- so -- the problem there is that courts are not necessarily going to agree because the damage -- fair use cases are not set up to look at mass uses like this. The courts will say, this particular work isn't necessary -- the market for it isn't going to be damaged. Really the market that's damaged is the market for works like this in general. It's the overall ecosystem. And I don't think fair use is the right -- the right context to be making policy decisions about these things.

So I am going to jump ahead and say that what I would recommend is that there be some kind of licensing -- collective licensing. And -- because, you know, if we don't proactively address these issues I think we are making a decision to give preference to AI over human writing. And I -- we need to be conscious about that. A collective licensing system is how copyright has approached situations in the past where the transaction costs of licensing them one-on-one are too high. I think an extended collective licensing will be most effective. And I am -- I am sure that blockchain could be used somehow to make it automated, but I have no idea how. So better minds need to figure that out.

Also, the -- just real, real quick -- that AI could infringe without being liable, direct liability, the Cablevision case. You could find that there's no -- nobody in the company that owns the AI that has volitional conduct. So we need to think about the volitional conduct and secondary liability. You could easily find no secondary liability also, particularly when we get to the point where AI is creating new AI machines. So we need to come up with rules about who is going to be liable in those cases. And then I think licensing -- collective licensing -- some kind of automated collective licensing might be the answer there, too.

MS. ROSE: So luckily Jason actually covered a good swath of the uses. I was going to talk about Janelle Shane, but I am happy that she has a lot of love in this situation. She just recently wrote a book called, "You Look like a Thing and I Love You." And this was from when she tried generate pickup lines through neural net. And one of the results was "you look like a thing, and I love you." So now that's the title of her book. Highly recommend going out and picking up -- she's just wonderful.

So a couple things I want to talk about, and I will keep it brief. One, as Jason mentioned, we have to sort of deal with the reality of the situation, where it is right now. There are tools available. There are neural nets available for free or extremely low cost on the internet that you can play around with. You can go home today. You can probably pull some of them up on your phone right now and start feeding it data. These are wildly democratized tools.

They are very popular, especially amongst the younger set. And you have situations where my nephew, for example -- 16 years old -- he plays around with neural nets in his spare time because he actually saw Janelle's website and saw some of the stuff that she was doing, thought it was hilarious and began to start generating his own neural nets.

And so you have this sort of pocket of communities online who are exposed to these things largely through humor and start to play with them, sort of as-is. As Mary pointed out, a lot of the times, as you get further down, the longer the output, the less and less likely it is to be considered passable for something generated by a human. The one I would like to read, actually -- just because it is probably one of my favorite things -- is Janelle Shane, again, tried to generate recipes through a neural net. Some of the results, one of them was just iced fridge water was the name of the recipe. And then it was -- but it involved shrimp. I mean, it was -- it got pretty wild.

So she had one that was called chocolate baked and serves cookies desserts. And you can tell the exact moment where the AI got bored if you look through the ingredients list, which was one cup butter, two cups peanut butter, one cup sugar, one teaspoon vanilla extract, three eggs, one teaspoon baking powder, one cup white cocoa, one cup milk, one cup horseradish.

(Laughter.)

MS. ROSE: She baked these, God bless her. She said someone actually ate a quarter of one of them at the party she was at. So clearly there's a market there. And if you read through, at the very end, the last step in the actual process is, add chicken broth. So --

(Laughter.)

MS. ROSE: Nobody is going to be mistaking most of these things for Julia Child any time soon. But we sort of need to deal with the reality of the fact that these are out there.

Communities are using them organically. And these communities -- I deal a lot with fan communities online, which tend to be interesting as they're sort of representative of self-organized examples of these communities, which are largely in touch with each other and sort of have shared practices.

We hit a situation both -- to put this sort of -- get to the more important point -- I am trying to sort of pare down here. The most important point is that a lot of these communities online that engage in this kind of work, and this kind of sort of unpaid labor really do this -- they have a very different set of expectations about what is being done with the work than I think you're going to get from talking to folks who are primarily coming from an industrial perspective.

These are folks who generate these things without expectation of paying or being paid. These are things that are circulated onto the internet for laughs in many cases. They are things that are circulated as kind of a here -- here's an interesting concept work. And you have very different standards of what I sort of think of as folk copyright, that happens online and in these self-organized communities. It tends to place very different values than what we necessarily have enshrined in the statutory copyright scheme. They tend to value attribution. They tend to value acknowledgment. They tend to value non-monetization in many cases.

And some parts of this overlap with copyright law as written, and some of them do not. But we're sort of hitting a point where that delta between the sort of folk copyright that takes place on the internet -- which is, whether you agree with it or not, whether you think it's a net benefit or not, is very much the reality of how people think copyright works -- is running up against issues of how it actually works and how we have constructed it to deal with very specific priorities about monetization and, you know, moral rights which sometimes map and sometimes don't.

But there's a very different set of priorities that have been ensconced in statutory law and how folks think copyright operates in sort of a common-sense way. And this is an interesting forum because text-based neural nets and AI generation are very low-processor power compared to a lot of other things. It's much easier to do than it is with video, for example, or sort of complex video game software where you have a lot of different sort of action conditions.

Feeding something Harry Potter chapters and generating a Harry Potter fan fiction through a neural net is a thing that has happened. And it happened several years ago. It's relatively easy to do because it is much easier to do. And if anyone is curious, the title of that fan fiction was "Harry Potter and the Portrait of a Thing That Looked Like a Pile of Ash." They actually printed it out in -- they made -- the mocked up a cover with the title on it and printed it out. It's like a six-page thing. For some reason the AI thought Chapter 13 was all that it needed to print, so it starts at Chapter 13.

You know -- and we laugh because these things are fundamentally very silly for now. I think there's certainly going to be a point where we might have a competition problem where people are generating things online that are in fact competitors for professionally-produced-by-humans products. But I think, you know, before we -- I urge caution in that as we start to make policy on this, we have to make it based on the reality of where people are because we are already, in many aspects of copyright policy broadly, facing real tensions between how people think copyright operates -- and it's not for lack of trying to educate, in many cases. It's just how people intuitively think what a common sense copyright system prioritizes and looks like -- and the complexities of a very, very intricate set of laws that we have basically duct-taped one onto another since the 18th century.

And as we are moving into this age where we have potentially conflicts at scale, we need to be very careful about how we structure that. And whether we structure that grounded in people's understandings, or whether we structure it grounded onto this several-centuries-old system which, in many ways, in no way reflects people's understandings.

MS. ALVAREZ: All right. Thank you, that was perfectly timed.

(Applause.)

MS. ALVAREZ: That's the end of our session. So thank you everyone. Yes. We will be on a lunch break now. I'm not sure what time -- we're running a little late. So I'm not sure when we should be coming back. If we want to -- what? 12:50 -- what is it? Okay, 2:50. Back here. 1:50? 1:50.

(Whereupon, the above-entitled matter went off the record at 12:50 p.m. and resumed at

1:54 p.m.)

MS. ROWLAND: Thank you guys. So thanks for coming back for our afternoon of copyright and AI. And we're going to kick off this afternoon with a panel on music, which is going to be led by our general counsel, Regan Smith. And so Regan, please take it away.

MS. SMITH: Thank you. Thanks for coming here. All right. So I'm really excited about this panel, which is about application of artificial intelligence for music creation.

And I'm going to start by introducing the panelists, and we're going to have this more as a discussion format I think than some of the other panels, which I think is going to be a lot of fun.

So starting on the end furthest from me is Alex Mitchell. He is the founder and CEO of Boomy, which is a company using AI to create instant music products.

And he previously founded Audiokite Research, a leading market research platform for independent music and has served as a consultant and advisor for companies, including Shady Records and Current Media Groups.

Next to Alex is David Hughes, the Chief Technology Officer at the Recording Industry Association of America. His responsibilities at the RIAA include developing and promoting technical standards, outreach to the broader technical communities and representing the industry's interests on technology issues.

Before that, he was a Vice President of Technology Strategies and Digital Policy at Sony Music. Next to David is E. Michael Harrington.

He is a musician and musicologist, who has served as an expert witness in hundreds of music copyright cases, including those including music by Pharrell Williams, Led Zeppelin and Taylor Swift and the Civil Rights anthem, We Shall Overcome.

Dr. Harrington is also on the board of

the Future of Music Coalition and has taught as a professor for multiple institutions, including the Berklee College of Music.

And then next to me is Joel Douek. Joel is an award winning composer who has scored over 80 documentaries, including those by Sir David Attenborough, hundreds of TV episodes, blockbuster animes and feature films.

He is the cofounder of EccoVR, a company that creates music and sound design for virtual reality experiences and is also on the board of the Society of Composers and Lyricists.

So thank you all for coming here. And I'm really looking forward to this panel also because it is the only one where none of the panelists are attorneys.

Oh no. Joel said he got a law degree during lunchtime, but I don't think he passed the bar, so I'm looking forward to getting into some real takes. And no one here is giving anyone legal advice.

So before we dive in -- at least that

anyone should rely on. So for better or for worse, the music ecosystem has often been at the forefront of technological changes.

And depending upon your perspective, it either positions music creators and their business concerns at the vanguard of progress, as the canary in the digital coal mine, or as both, and the AI technologies are going to bring even more change to the process of music creation, discovery and consumption.

And I think to start, David, why don't you lay out kind of a landscape view of some of the roles that you see AI technologies might be expected to play in the music industry now and in the upcoming years.

MR. HUGHES: Okay. Thank you, Regan. It's a pleasure to be here today and to talk about something that's not blockchain. I am very excited. You have no idea.

So I guess I'd like to start by saying, and a lot of people said this, that this AI and computer generated music is nothing new.

I have recollections of an album made by David Cope, who is -- he's a professor at UC Santa Cruz.

And he had an album called Classical Music Composed by Computers. That was 1997, and that was a minor hit released by Centaur Records. But it was very interesting at the time.

And what he did was he was suffering from writer's block, and he was a computer guy and a composer. And so he said, well what if I just write an algorithm and just feed a whole bunch of stuff in here.

And one of the things he did was just took all of Beethoven's sonatas, fed them into a computer, pumped them out, listened to endless variations I'm assuming that the computer made, picked one, recorded it, put it on an album.

And just so you know, if you look at the credits under composer, it actually says computer generated composition. That's how he credited the composer, not himself, which I thought was kind of gimmicky at the time but may have copyright implications in the future.

No, I didn't say that because I'm not a lawyer. But the reality is it goes way back. Computer generated music goes back to the '50s.

So 1951, Alan Turing, in an effort to try to have a computer do something that he thought only humans could do, turn his computers -- and you may remember Alan Turing from Breaking the Enigma Code or being in a movie with Keira Knightly. Yeah.

He's also going to be on the new 50 pound note in the UK. Did you know that? That's big news. Anyway, in 1951, he created some music on his computer. So it goes way back.

But now what we're seeing is really AI and how we define AI as a different part of our conversation maybe, but computer generated machine learning AI, I'm going to bundle them all together, impacting every aspect of the music industry, pretty much every aspect.

So I'm just going to go through a list quickly to give people some ideas because recently I was talking to an executive, not a major label executive but a music industry executive who said pretty soon, there's going to be music on Spotify composed by AI.

And I said, what year is it? There are tens and tens of thousands of production music, relaxation music, study music, yoga music already on all the streaming services, so this is a reality.

This isn't something that's going to happen, but let's take a step back. Some of the ways that AI can contribute to this whole process is we've seen AI.

I'm going to call them systems or tools, whatever you want to call them, that create ideas for lyrics. I call them lyric kernels, so it's gives you an idea to write your lyrics.

There's another one that does the same thing for melodies. It starts giving you some melodies to work on, so then if you had writers block, you have somewhere to start.

And then we see another company hypothetically creating -- generating beats. And Alex will talk more about his company. And Sony Computer Science Laboratories has a number of initiatives.

One was the composition of the song, Daddy's Car, which as trained on all the Beatles' catalogue and pumped out a song that sounds pretty much like if my junior high school band had tried to write a song that sounded like the Beatles.

It's pretty awful, but it's identifiable as a Beatles song-ish. And then we see more sophisticated companies, like Endo and others doing personalized audio tracks to boost mood or to be used for specific circumstances and so on.

But it keeps going. It's not just this sort of idea of kernels and beats. We've got companies and AI software out there that are doing real compositions.

So if you go on Spotify and you look

for AI playlists, one of the top songs is Lovesick by Taryn Southern.

And she trained her AI to write composition based on 19th century public domain pieces and then picked one, wrote the lyrics to it, went in the studio, recorded it and now hundreds of thousands of people, which sounds like a lot except on Spotify unless you're in the billions, it's really not -- but are streaming this track, okay.

And it goes even further though because now you can get an AI that starts writing lyrics. And then you do text to speech technology. Then you layer midi technology on top of the text to speech, so now it's singing according to the melody that the AI has written.

And it can create everything from scratch to finish for a sound recording. Might not be good, but the AI is there. And it's happening already.

And then I want to quickly talk about how it impacts everything else in the music industry. So we talk about generating music to match a certain issue.

So you feed a video in there, and it'll match the mood and compose something in real time to create background music for your video. Or if you needed to engineer, there are AI that will help you mix.

There's AI that'll help you master. If you want to have your song remastered, you can go to a remaster engine room, pay thousands of dollars or you can use one of these AI tools to remaster for much less.

These are available now. And my present concern is that some of my close friends are audio engineers, and I'm not sure how this is going to impact their livelihoods. But hopefully they're old enough that they'll retire before they're all out of jobs.

Then we see other technologies doing some interesting things like recording decomposition. So this is where you take an old recording, maybe a stereo composition that was just recorded in a couple tracks.

It goes in there, separates all the tracks out into what we would call recording stems and it pulls out the drums and the base and the vocals and everything.

You strip out the vocals. Now you have a cut okay track. And that's, you know, may or may not if it was licensed correctly, that's a great idea -- and could create all the stems from something that never even had stems.

It was recorded live with maybe just a stereo microphone. Then the next level is now we have a sound recording. We've got AI that are analyzing, predicting marketplaces for those human created songs.

There's a company working on AI that helps A&R people match songs to singers. There's everything. It's like blockchain four years ago for AI. If you can think of an idea, there's a company out there doing it.

There's AI that are generating these generic and thematic and personalized music lists, so you, based on your interests and the only music that you like, it'll create a list that's personalized for you.

We see, for example, in the news just a week or two ago, iHeartRadio laid off 850 people and when asked why, they basically said because now we have AI to create all our playlists. We don't need all those people.

And so it is impacting the music industry severely already. And it goes beyond making a sound recording or even marketing the sound recording in that sense.

Recommending concert locations, venues, days of week, who should open for the act, set lists in concerts, all of these things, people are starting to dabble with AI to do all of these things, anything to increase the ticket sales across the board.

There's another company called Muzio that listens to the sound recordings and automatically creates metadata. It pulls out all the, what we call objective metadata.

That's beats per minute, key signature, the language that it's being sung in, the absolute pitch for classical music, if you know. If it's an A, if it's 440 or not 440, it's technical.

Subjective measures, like the genre and the tempo and the mood, is it happy? Is it sad and all that stuff. And there's AI doing that now. And then identifying the tracks of artists and the metadata, putting it all together, figuring out the metrics, trying to figure out how to maximize royalties.

You name it, AI is impacting the whole chain from beginning to end. And I've said my piece.

MS. SMITH: All right. Well, let me follow it up a little bit on your piece. So if you're comparing AI to the blockchain of four years ago, I'm not --

MR. HUGHES: That's only because we had a lot of panels about blockchain, which were mostly not productive.

MS. SMITH: Well, we're going to have a productive panel.

MR. HUGHES: Oh, this is great. This is awesome. This is exciting.

MS. SMITH: Right. So an earlier panel mentioned that TikTok has acquired Jukedeck. Large companies are acquiring AI technology companies.

And you just mentioned using machine learning for ID tracking at a time where we're seeing consolidation of collective management organizations globally.

We can see the U.S. effort with the Music Modernization Act to build this database. It's going to match sound recordings to their underlying musical works.

Are you seeing AI incorporated into viable businesses for royalty processing and creation issues?

MR. HUGHES: In the presence of the MLC I actually haven't seen anything, but that doesn't mean it's not there.

MS. SMITH: Right.

MR. HUGHES: Certainly a combination of fingerprinting technology plus the database plus training some systems on how to match, it'll happen if it's not happening already.

MS. SMITH: And then focusing on music creation, Alex, can you speak a little bit from the technical side --

MR. MITCHELL: Sure.

MS. SMITH: -- for what Boomy is doing?

MR. MITCHELL: Yeah, absolutely. I think, you know, I would start by just saying, you know, today there's been a lot of discussions of AI.

There's been discussions of GB2. And, you know, there's -- I want to stress that there isn't one, sort of one size fits all way to produce original works or recordings.

Certainly, our users are advanced as saying I want to make a hip hop beat that has cello at this BPM with these qualities. But it can also be as simple as pressing a button.

And to, you know, sort of put some numbers on it, our users have created in the last six months about 350,000 original works and recordings because it takes five seconds to create something original.

If we're looking at AI through this lens where there's like -- the only way to do it is to train data, then I think that's very problematic. But that's not the way this stuff works.

I mean algorithms are the past, present and future of music. I really believe that. I think that's kind of what you're hinting at there. If everybody in this room isn't aware, I'm not a lawyer.

I do have friends who are lawyers. And I ask them, you know, can we pull, you know, a statutory definition of artificial intelligence from anywhere and they sent me the National Defense Authorization Act, which is probably the most recent definition.

Of course, it only applies to that bill. All right, fine. I'm not a lawyer. But that second definition in there was broad enough, and I can read it.

But it was broad enough where it isn't a stretch to say that the way it was defined in artificial intelligence would apply to Jimi Hendrix or Brian Eno or any one of the musicians we traditionally consider to be creating music using algorithms.

Guitar pedals are algorithms, right. So when we produce music, we're using a variety of different algorithms, not that different from sort of chaining together, you know, different types of methodologies.

And we are working on lyrics. We are working on vocals. We're using a variety of approaches to do that. So it's a fundamentally creative process to build what we refer to as music automation systems that can do a lot of different things.

MS. SMITH: So can you -- is there

more you can share about how Boomy does use artificial intelligence to create songs? It's not a guitar pedal company, right?

MR. MITCHELL: Right.

MS. SMITH: Your tag is save the songs you like, reject songs you don't, and teach Boomy to make songs you love.

MR. MITCHELL: Sure.

MS. SMITH: So if it's not doing this through just using algorithms with inputs, how is -- what makes Boomy stand out?

MR. MITCHELL: Sure, I mean without giving up some of the really interesting things we're doing kind of in the background there, I can give you one example.

So we have focused -- there's been a lot of research on how to generate compositions from feeding past compositions. And generally speaking, this produces many data, which on its own sounds pretty bad because you need instruments and you need, you know, software instruments or real instruments to actually play those notes.

So we have taken an approach where we're looking as much at the production side of things, right. So we have systems, for example, that can create, you know, large volumes of sounds or different sounds for different purposes and then look at all those sounds and decide well, these are the sounds that are grouped together.

It's harder than you think because, you know, you could have a random, loud, silly guitar sound. You could also have, you know, a soft, interesting bass sound.

Those things might not go together, right. And so we're doing -- again, that's just one of many examples of an algorithm that we created that can create sounds, that can understand those sounds, how those sounds might work together.

Those use traditional training processes. Those are trained on samples that we purchased the rights to, also samples that we created.

And by analyzing them, you know, part of our, you know, IP and part of why our music sounds so much better than some of the past, you know, attempts of this stuff is because we're looking at that production.

So that's one of example of something where a copyright and, you know, it isn't really about the composition. It isn't really about feeding it, data that's been copyrighted.

It's just an algorithmic approach to solving a specific problem, again, which I feel like even though we're taking it to this interesting, you know, extreme maybe historically, it isn't fundamentally different from the way musicians and artists have thought about creating music up until this point.

MS. SMITH: I think that's a good time to switch to the composer, right. Joel, how do you think about creating music, and how do you use algorithms or a set a rules apply into your practices? And how do you see AI being useful or not to your craft?

MR. DOUEK: I mean I think the first important point I want to make is that it's not black and white.

Really, it's a spectrum from on the one hand, you know, the beginning of little tools that assist various parts of the process of the composition all the way to what we might imagine as, you know, full on composition without any involvement of, you know, the composer at all.

And I think right now we're kind of in this area here, and so it's -- there's technology that we're already bringing into, you know, our daily work habits as composers through a range of different things.

But fundamentally what we want as a composer or a songwriter is to try to deal with some of the more chore like tasks of composing. And AI's can actually help quite a lot with that kind of stuff.

We want to be able to liberate ourselves, accelerate those processes and liberate ourselves to really focus on the creative side of things.

And that can, you know, involve both structural things, and I personally don't use AI in, you know, this rule based, algorithmic way to help me write music.

But I use other aspects of it, so for example, in mixing and mastering there's very clever AI machine learning tools that can help kind of analyze and mix and improve it for you almost immediately.

So I look at it as at this point we're helping us solve some of the chores. But that's going to change. It's going to change over time. And I think the line is going to get a little bit thinner.

And I think it's exciting. I'm definitely not one of those composers who feels that unlike the art critic comment that was written before, that this is a horrible travesty in the world of art.

On the contrary, I think it's very

interesting. We were talking about earlier. I think it's going to kind of generate its own influence on the future politics of music.

One aspect that I think in terms of music production that keeps coming up for me is that composing is really the same thing as improvising. That's really what we're doing.

You just sit there and you, at the piano or the guitar you noodle around. You just -- until you find something. And what's interesting about that process and where that might also feedback on the way artificial intelligence is done is that, if anything, it's not about a rule based algorithmic way that our brains are working.

Quite the contrary, it's about relinquishing a good deal of conscious control and letting these things happen without the self-criticism, the self-editing.

So there's some ethical implications there. I guess what I'm saying is that if you really want an AI to write good music, you've got to free it completely of all boundaries.

So we'll see where that goes, but there is interesting work being done in that area of rethinking how AI is actually trying to write music and out trying to improvise.

By giving it a different kind of attitude, if you want, where it works in a much more probabilistic way, it might actually bear more fruit. Yeah, so the short answer is I think it's only days for us as composers.

It's working its way in, definitely it's a tool. It's super useful. It'll continue to be, and I think as long as keep that dialogue open between the developers of these tools and the composers so that we maximize how it can serve us, how it can help us to write music, this area here at the extreme is one I think we're going to get into in the rest of the discussion. So I'll kind of let anyone else chime in there.

MS. SMITH: Okay. So Dr. Harrington, maybe you can speak a little bit from your perspective as a musician, as a musicologist.

I think, you know, particularly for Western music and popular music, what many regard is pleasing is based on certain conventions of composition that's evolved over centuries that may be, you know, somewhat predictable algorithms.

Does that music especially favorable for machine learning? How does this change the type of music that we can expect to see created?

MR. HARRINGTON: I think that's a great question. I also think, going back to something Joel just said is one thing AI can do and computers in the old days, calculators, could do is even free up some of the chores.

Like if you read a book, if you read a 300-page book by Stephen King and say oh wow, how did you love page 87 or what are your favorite sentence or favorite clauses? You're not going to think that way.

So in a book, a lot of things have to move to get to the important points, the important characters. It's often the same with music. Not every note in Taylor Swift's last song matter. These parts do.

And to go back also to what's pleasing, everything -- I can sound really cruel and objective and heartless on this and so on. Everything is math based.

If you write with three chords, your rules are don't use a whole lot of other chords. If you arrange an AI -- you mentioned about artist and repertoire.

If you want to find music that works for Hootie and Blowfish, he likes to sing in a certain range. And some other people want to sing in a different range.

So AI can be used in all this, but there's math underlying everything, whether it's the three chords, whether it's cerebral jazz. And a lot of times this is just helping us create it more easily seeing.

One thing that interested me a lot was when I was growing up and learning how to play classical music after I had my Beatles phase, which never ended by the way, but of course.

Bach has this Prelude Number 1, C major. Every beginning pianist plays it, so these nice chords. They're three notes or four pitch classes, three or four.

And it occurred to me, this is perfect. And I thought well, how come Bach didn't use the others, you know, because he could. And then it occurred to me as a challenge.

Okay, use the others. So what I did, I found out later in terms of copyright, that's quite the derivative work. My work is completely derived from him.

It's the negative in terms of a photograph, and what could you do with that. And I want to apply -- well, I'm also interested in all the math of this, how we got to these points, what sounds good.

I used to study. I remember I found out Mozart was great one day. I always thought like you pulled one over on me. Mozart is awful.

I mean I love Bach and Beethoven and all that.

I thought Mozart was predictable. So I went to set up the theorem and all that. How could I prove my point? And well, I'd go to the library and open up Vivaldi and Beethoven, Scarlatti.

I'd look at the score I hadn't seen. I'd cover it and say, can I write what's coming. And I was really good at it. Telemi was simple. Scarlatti's easy.

I got to Mozart. I tried one. Ah, I'm wrong this time. So I kept doing this all

-- not all afternoon. After about 30 minutes, I realized Mozart does the most clever, mathematically complex things in weird places.

This is simple. This is simple. This is simple. Holy God, what's that doing? It's like it makes no sense, and you can't predict it.

And that's what interests me a lot in melody, harmony, rhythm, lyrics, when you get to some strange, you know, why do these words go in. And what's the math behind it all?

And to me, AI is, it's computer music. The stuff you mentioned, David Cope and Turing, I studied with Iannis Xenakis.

And we would look at all the math involved in something and not only just melody but density class, timber class and all these other things.

And I found all this does come together, and AI can just enhance it and make it easier for all of us to create. And if you want to do easy underscoring, you don't want to pay someone, yeah, it's going to take away jobs and playlists, which once playlists came about I thought like who's going to tell me what to listen to. I'm appalled at that, but.

MR. DOUEK: Can I?

MS. SMITH: Sure, yes.

MR. DOUEK: Yeah, so stylistically, I mean we always love to beat up on yoga music, but I guess, you know, I think we -- on the panel, we could all agree that, you know, certain kinds of music and may be easier for an AI to attempt and to succeed at.

And, you know, so music that doesn't get judged for being formulaic, but in fact it's one of its features. I think that can be a good thing.

I'm not sure whether jazz is then by definition the nirvana of, you know, what an AI could hope to achieve just in terms of its flexibility and inventiveness and improvisation quality. Just putting it out there.

MR. HUGHES: Yeah, so in terms of genre and how formulaic is it, I have to agree. So I think that country music is going to be biggest sort of paradox for AI because the underlying melodies and so on are very formulaic.

But the lyrics are based on human experience. I'd be very interested how long it would take an AI to come up with this song, Tammy Wynette song, D-I-V-O-R-C-E, unless that AI spent time in a kitchen fighting with their spouse in front of their 4-year-old AI, they probably wouldn't come up with it. I'm just going to guess.

So, but going back to what Michael said about math, I think a lot of this composition is formulaic. And we think that that's so bad.

But I always remember this interview that I saw with Chuck Berry where they asked him how did you write all these hits. And he said well, of course he used the three chords.

He knew it had to be three minutes. He knew the hook had to come in 45 seconds in, all the standard stuff. But they said well, what about the topics.

He says, I used to sit in the soda jerk shop and listen to the kids talk. And he said kids only talk about four things, school, cars, music and falling in love.

That's it. He said, why waste time writing about anything else. The kids don't give a shit. So that's what he did. It is very formulaic. He used algorithms.

And then he layered on top of that the

unpredictability. So the songwriter, Desmond Child, I don't know if you know who he is, but he's a pretty famous guy and one of my friends, best friends.

He wrote songs like You Give Love a Bad Name. And he just did an interview not long ago for Bob Lefsetz. And he said when he was ready to write songs, the best lesson he ever learned was every line should have a contradiction.

You Give Love a Bad Name, he goes, that's a hit song. Now we have to write it. And I don't know if that can be trained. Maybe you can have an AI spit out hit song titles and then you have to actually do all the hard work.

MR. HARRINGTON: If I could say something about the unpredictability, too, in the -- what you said about Chuck Berry -- AI does miss the big picture.

And I'll give you an example. Why would Garth Brooks, him -- I'm living in Nashville. I have to -- I should mention Garth, right. Why would he use an Egyptian drum in a hit song? And he did.

He used a doumbek. The AI couldn't have come up with that. Or how about Jimi Hendrix writing a lovely Viennese waltz? By that you're thinking, do I know a Viennese Waltz. Yeah, Manic Depression.

Yeah, it's beautiful. You dance a certain way to it. Why did he use -- why did his drummer use brushes on the opening song in the second album?

And you could go on and on with these moments that AI can't do that because it's against what it learned. Back to, I think it's learning from us and improving on us sometimes.

MS. SMITH: What do you think, Alex?

MR. MITCHELL: Sure. So what's not hard to do, what's not that hard to do anymore is create sort of a best guess copy of a style that came before.

I remember seeing some research some years ago before Boomy about an AI system that was trained to create folk music. Do you remember this? I can't exactly remember what it was.

But it was trained on American folk songs, and it could faithfully reproduce American folk songs. And someone who studied fiddle and used to play American folk songs, in the back of my head I was like why, right?

No, seriously. It's a serious question. There's people that make American folk songs like all the time, and they're amazing at it.

So I think the utility question, right, of is there really so few people who can make American folk that we need to automate this? And what we're hitting now and what you're talking about is the real challenge, which is that it's very hard to do something new.

It's hard to allow an automated system to make those kinds of mistakes and to screw up in such a way that is still creative, right.

So this balance between creativity,

which necessarily means making mistakes and accuracy, which means I'm kind of just doing something that the people do -- people can do already, I think is crucial for the crop of music companies to think about because it comes down to, you know, this is a nascent market.

This is still -- there's a lot of different approaches being tried. And I question sort of some of the utility of just making a best guess copy of a jazz song or an American folk song because there's plenty of people who do that already, right. And you end up with the lowest common denominator.

MR. HUGHES: Sure.

MS. SMITH: Yeah, so I mean when we start to get copyright in this -- I think we heard earlier copyright protects, you know, the bad art along with the good, but the whole, you know, goal is to encourage the production of more of the good.

In the case of AI technology, does it matter? Is it likely that certain genres are more right for wholesale substitution versus what you started saying of taking a kernel or building off a piece, a component of a song?

It might be in the case of country. Three chords can be predicted, but we're not quite sure if AI can find out the truth. But in the case of hip hop, is it -- that seems more layered and perhaps more difficult for an AI created song to be fully pleasing without human editing, interference and contributions.

MR. DOUEK: I mean, I personally think it is just a matter of time. Utility or not, I think it's going to get done. We're going to do it because that's what science does, if you like.

And this is really under the umbrella of science. Why do we do it? Because we can, and we need to find out. So it's going to happen, and then we'll have to kind of retroactively look back on it and say, you know, was this meaningful or was it not.

So I think all those stars of music are going to eventually going to come. I think, you know, drilling a little bit more into the question of what makes music human, we've talked about this a little bit, is one of the ones I came up with is imperfection.

You know, a lot of time when we record a live instrument, it's not to introduce what you'd expect, which is the level of virtuosity. I mean that's a given.

Hopefully, they play in tune and everything. What it is, it's the nuance and it's the mistakes and it's that imperfection that comes into it.

And somehow for us, that translates into maybe a sense of heart and soul. I don't know. What makes music human, it's hard to, you know, get past that area without getting into the metaphysical.

But I can tell you that on a daily basis as composers when, for example, we're dealing with a small budget and you need to get a big orchestral sound, we will habitually just record a few instruments live and then layer it on top of sampled instruments.

And with the sampled instruments, we'll use our AI effectively to humanize. I know the irony is not lost on me, that we use the computer to humanize the playing performance style so that it kind of matches and blends, you know, as something that is, you know, something that feels human.

So what makes music human? I don't know. It's an open question, and it's definitely going to end up as a metaphysical question because we'd like to believe that there is something kind of above and beyond, something intangible that we as composers bring to the table.

Is that true? Is that not? I think we're going to find out all of this. I'm not going to proclaim anything. I think we're going to find out all of this in some like giant singularity moment.

MS. SMITH: Can't wait.

MR. MITCHELL: He was staring at me

when you say what is -- I mean what music isn't human I think is another angle to look at that through, right.

MR. MITCHELL: According to Queen, synthesizers.

MR. HUGHES: Well, sure, right. I actually think that's a great point. I mean Queen used to tag all of their albums with no synths. And there was a conversation about synthesizers and how it was going to affect instrumentalists.

A more recent example would be auto tune. I think there's a lot of sort of protests against auto tune. And now find me something on Billboard that doesn't, you know, that doesn't use auto tune, Melodyne.

I think this is something that is, you know, when we look at these automation technologies, particularly as it pertains to the production side and the composition side where like now I have something, if I'm reacting to it, like I remember the first time I was moved, right, by something that, an algorithm we created.

It was crazy. It was like this absolutely sounds -- it had sort of the imperfection you're talking about. We had modeled that as a piano in a room, sort of had a tape hiss.

And it was like this is -- I mean we're way past, you know, that's obviously robotic or that's obviously a computer. I think ultimately the arbiter is going to be the market.

Ultimately, it's going to come up to do people want to listen to it. Do they want to use it? Where are they using it? And I think that's what's interesting.

But I don't know that I've ever heard a piece of music that I would not consider human. I have a personal belief that there's no such thing as AI generated music.

There's people at every step of the way when you create these systems. Algorithms don't spontaneously just, you know, happen. I

believe there's intention and --
MS. SMITH: Well, certainly there's --
MR. MITCHELL: -- there's musicality still necessary to create an automated system that produces something that frankly is worth listening to at all.
MS. SMITH: But break that down a little bit because I think earlier we heard about some of the Copyright Office's practices in registering and requiring some human active creation.
And, you know, going back there is a certain case I know. I'm talking about layers, about taking a photograph. If you are a developer of a system of algorithms that doesn't know what the work is going to be, is that the same as creating the output of the work overall. I mean, is that what you're saying?
MR. MITCHELL: Again, it's -- personally, that's what I believe. It's hard to draw a line, right. It's hard to draw a line between somebody who, I use the guitar as an

example because it's a common example like pre-computer use of algorithms in creating music.

It's strange to me to say well, this pedal, well that's okay. But this other pedal if you use some type of algorithm or training process, we're going to kind of treat that differently.

And certainly people who designed those systems, right, they're not going to do it in a vacuum. They're going to do it because they have something to express.

You know, we work not only with our, you know, internal resources, but we're also sort of looking out at the world and seeing, you know, researchers and composers, people who have unique ways of creating music with algorithms, many of which don't actually use any of the systems we've talked about today and which are creating things that are fundamentally new.

They sound different. I can't really put a genre on some of these algorithmic works. I don't understand why you wouldn't look at those people as artists.

It's this kind of new type of artist that is maybe a little more common than we think it is of people who are adding their own theories, their technical backgrounds into creating these things.

And I would say if somebody never learns anything about music theory but uses code and creates a program that creates music, I don't understand why we would call that person anything other than a musician.

MR. DOUEK: I was just going to say --

MR. MITCHELL: You know, that's an opinion.

MR. DOUEK: Well, what keeps coming up for me is the idea of autonomy because the guitar pedal, it's still very much doing the bidding of the musician. And I think where we're getting to this part of the spectrum is where --

MR. MITCHELL: Yeah.

MR. DOUEK: -- things start to become autonomous. That's where the fears lie at many different levels. And I think that's kind of a different question really at that point.

MR. HARRINGTON: I was just going to say about I agree that AI is human. I mean it came from a human. There's no such thing as non-human music. And I use the example of George Gershwin, who went through periods of study with Joseph Schillinger, who is doing math and music and plotting parametric equations to come up with melodies.

Summertime was supposedly involved in that way, and I can show you a lot of beautiful melodies that are oh, it's kind of like that curve, and it's this curve.

And so what. It's an image, just like if you saw mountains coming and going. You can plot the melody. That's been done for centuries. And what's different from a parametric equation doing that? That's human.

MS. SMITH: Well, I want to ask you a follow up question as to the role that AI is going to have and what we're seeing in music, which has been, you know, whether actual or not I would say definitely perceived uptick in infringement cases in the last couple of years, which I know you have been involved in many of these cases.

So there's been recent headlines involving Katy Perry, Taylor Swift, Drake, Migos, Stairway to Heaven, and if you followed the 9th Circuit's Blurred Lines decisions where Judge Nguyen in the dissent criticized the decision as allowing the Gays to accomplish what no one has ever before with copyright and musical style.

So what is the increased use or discussion on AI to create wholesale or assist in the creation of music? What is that going to do to the landscape and these cases for infringement do you think?

MR. HARRINGTON: Well, I have a huge dark cloud over my head on this because I got dragged into the Blurred Lines appeal. 212 musicians used my stuff in several places.

And suddenly I'm getting all these

congratulations for being involved and didn't know that I was, and then Ken Freilich said do you want to coauthor a brief and so got involved.

Got involved in the Led Zeppelin appeal, the Katy Perry appeal. What I hated about especially the Blurred Lines is the first time in history no one copied a melody.

No melody was copied. Sing it, right. No. No lyrics, no chord changes, no rhythmic features, no sampling. It's the -- it reminds me of standard, which is awful.

And this is the part that leads me to thinking about AI. If you say like Sony did when they had this Daddy's Car song you mentioned earlier, it's built on 45 songs fed into a computer.

And out came some stuff and I reverse engineered it as I think I can, you know, as best I could. And I found that a lot of it is not really, it's -- a lot of it's human.

But the fact that you said it was Sony doing Beatles and the reason was because they own the catalogue. They're not going to sue themselves.

However, this a big problem if you say who your influences were because Pharrell and Robin said oh yeah, Marvin Gaye and especially that song. So you're influenced by that song, that now means you're ripe for lawsuits.

Publishers find this stuff out, who your influence is, especially if you say I want to get under the hood and see what was programmed or what did AI learn.

If it learned George Clinton, then Bridgeport is going to start suing, which they like to sue almost as much as the Gays like to sue. So I think there's some problems with that in that it shows influence, awareness, maybe the use of.

Well, it should be -- Paul McCartney said one time the Beatles are the biggest plagiarists there are. He said we copy from everyone. We're "Knickers Extraordinaire", that British saying, whatever that means.

So to me, I'm worried about that. And I can already see where Daddy's Car could be sued. I do a lecture on this, so I've got all the examples. I could do the defense, too.

And here's what I'm looking forward to. When does AI start to sue? And I'll give you one reason. Go check Daddy's Car, you know, on YouTube and then check the cover of it.

There's a woman covering, and she's not doing it that well. Have you heard it? Oh my God. It's hysterical. And she's earnest. It's nice. She means it, you know, but that could -- there's a section of the copyright law that says you shall not change the basic melody or fundamental character.

And that's sort of us like doing European moral rights, but if you don't like the version then uh oh, AI's going to sue for that or AI is going to create something, you guys. I don't know.

MS. SMITH: I think to get to where you're trying to go, you need to I guess recognize a copyright in AI. And I think that part of what we're circling around is whether a lot of these compositions or productions are created by an AI or whether they're created by a person using AI technology and whatever point in the process.

Or maybe it affects authorship, or if you do get to the point where it's an entirely work created without human involvement whether copyright should attach.

So I think that is a good time to discuss Boomy's approach to ownership of the music because Boomy asserts a copyright in each of the works, right. And then your model is you would sell it for $5 to $20 to --

MR. MITCHELL: Yeah, to the professionals --

MS. SMITH: -- the professional.

MR. MITCHELL: -- who need those rights or need that ownership, sure.

MS. SMITH: Okay. And then do you -- who is listed as the author of the works?

MR. MITCHELL: Yeah, so we're sort of maintaining just the best practices of now which is that it's basically the developers and sort of the team that is assigned the song writing, you know, credits so to speak.

But in the case of the professionals who want to integrate, for examples, works that they create with Boomy, you know, in their -- they need specific licenses, they need synch licenses, we're happy to sell, you know, that composition and that recording to people.

We do it for $20, actually as low as $5. And we feel that's fair. And if they don't want to purchase it, there's still a lot they can do with it. But we just sort of continuing to again assert that ownership and assert that copyright.

And we're sort of covering that under the contract that our users agreed to when they use the free, you know, application that we produce.

MS. SMITH: So a user, you know, agrees to the terms which is that whatever they say, I want a song sort of like the Beatles but, you know, mixed up with something else.

MR. MITCHELL: You can't exactly do that.

MS. SMITH: Okay, can't exactly do that. That's not what Boomy does, guys.

MR. MITCHELL: That's a little -- yeah.

MS. SMITH: Okay. Give me an example of what I could tell Boomy then. I don't want to get you in trouble.

MR. MITCHELL: You can say give me a rap beat at 90 bpm that's got cello and, you know, synth.

MS. SMITH: All right. And then can I say I don't like that, make it faster?

MR. MITCHELL: Sure. Yeah.

MS. SMITH: Okay.

MR. MITCHELL: You can go in. You can manipulate, and you can change it. You can also edit it. You can delete sections. You can rearrange sections, you know.

So when it comes to some of this infringement, right, which is a really interesting question when it comes to the way the music is being produced.

You know, I'll give the lawyer answer which is it depends. It depends. If somebody takes something that Boomy creates, and they turn it into a song that they're heard and they download it and produce it, then there's probably situations under which you can consider it infringement.

But I think it really, again, it just depends on that person and their use of the technology. The notion that we would create something that's infringing sort of just the ether, it's unlikely.

It's mathematically very unlikely in which case we'll produce the DMCA and all the, you know, relevant requests. But what we're getting into is just to pertain to best practices in some of this gray area.

MS. SMITH: Well, I want to drill down into that because I think what you've described and, you know, coupling that with what Dr. Harrington just described is a system where it seems like it's really charms you have with the user which are in part perhaps driving the authorship to vest in Boomy, right?

MR. MITCHELL: Yes.

MS. SMITH: We're not really probing whether it has to be. It is because there is an agreement that you're describing a situation where there does seem to be a human directing towards a predictable outcome, which may not be the whole universe of AI music.

But what you're saying, I think that when you're saying best practices, who is determining the best practices? Are you looking at other competitors in the marketplace, or yeah?

MR. MITCHELL: Remember, there're attorneys that we pay money to, to write these contracts for us.

MS. SMITH: Right.

MR. MITCHELL: And some of these attorneys, they represent multiple, you know, AI music clients. You might be friends with some of these people.

And there's starting to be kind of, again, in lieu of sort of guidance or, you know, exact law and where specific instances might lie within this gray, you know, contracts and sort of contract law is what we're going to rely on.

And we're going to do, you know, basically what the attorneys, you know, tell us is the best way to do it. We don't have a choice, right? That's pretty much all we can do.

But I will say that, you know, there are aspects of copyright law and there are aspects of this that make it kind of difficult to create systems that we might want to do.

For example, if we wanted to grant song writing ownership to every single user for every single song, if that was something we wanted to do, it would be -- I said that to our attorney.

He just laughed. He's like, the amount of liability you'd be opening yourself up to, all the issues.

MS. SMITH: Because of infringement concerns or --

MR. MITCHELL: Infringement concerns or just the pure management of the data and managing all these different publishing contracts, it can get messy because nobody's ever really done it sort of this fast or this frequent before.

And so it's not as if I'm blaming, you know, the law, but we have to, you know, work within the framework that we, you know, that we have.

MS. SMITH: Yes. Okay. I want to switch to David for a second though. And can you, from the recorded music perspective, talk about how labels might be thinking about ownership in the AI context.

I know Warren Music just took a stake

or signed a deal in a company called Endel which -- and has a practice of listing software engineers as music authors.

How do you see sort of the -- do you agree with Alex's characterization of best practices or what other things we should be thinking about how the music community is looking at these issues of ownership?

MR. HUGHES: I think right now the major labels in particular are following it very closely and that's it. First, I'm not a lawyer. Second of all, I'm not going to speak for them. And third of all, I think there's a lot of balls in the air.

MS. SMITH: Joel?

MR. DOUEK: You know, again coming as a composer and also as a Board member of the Society of Composers and Lyricists, which is about as close as we've got to a union, I guess my salient hope is that, you know, we can find ways during all these changes that protects the ability of composers and songwriters to generate revenue.

I think there's often the mistake because we make comparisons of like oh, let's suggest the Beatles and Beyoncé. And somewhere in that narrative we're looking at being incredibly powerfully wealthy.

And what gets discarded is that the vast majority of composers and songwriters are not in that bucket and just need to feed their families, pay their mortgages.

And that will go away. And so the question kind of makes me want to ask is does anyone care if real music written by real humans goes away or not.

Because if they don't really care and everyone is satisfied with whatever an AI might ultimately quite brilliantly create in the future, then it's a different discussion.

But I think where we are now is if we're talking about let's say the ethical use of AI, the idea of protecting artists, composers, songwriters as this moves forward and as we try to shape these laws, it has to be mindful of that. It really does.

MS. SMITH: Well, you know what. I'll jump in there and I see maybe Alex wants to say something because you know what, like I care, right. I mean I bought a fair trade coffee this morning.

Are we going to enter a world where we see an attribution on a playlist whether music is human created, primarily human created. Is that going to be a trend that we're going to start to think about in our music consumption and how does that -- I know we've got no lawyers on the panel, but I think it does relate to some of the attribution and moral rights issues that were raised on earlier panels just in terms of giving credit where credit is due and recognizing what is underlying the creation of this song or a recording that we're listening to.

MR. DOUEK: Yeah, but at the same time, you know, I know firsthand that a large popular streaming platform is actively using AI music because it owns it and therefore doesn't have to pay out any kinds of royalties in that it can therefore help its bottom line.

And so there is this, you know, it's a race to the bottom musically. And, you know, we've seen parallel type things happening in terms of how we accept the fidelity of music where the vast majority of younger generations are quite content with, you know, low bit depth mp3s and nobody really cares anymore about hi-fidelity and pristine sound and that kind of quality. So there are parallels in precedent.

MR. MITCHELL: For sure, and honestly I couldn't agree more. I mean I couldn't agree more. I think we also -- we didn't stop hiring wedding photographers because we all got iPhone cameras.

And maybe more pertinently, we didn't stop appreciating great singers or listening to great singers just because auto tune came out, right?

So I think this question on ethics is

crucial, and it's probably important to talk about, you know, why we're even doing what we're doing. What auto tune did was it allowed people to express themselves musically who could not sing.

And why couldn't those people sing? Well, maybe they couldn't afford lessons. Maybe they didn't have all the access to the time and resources necessary and the training, right, that we have been lucky enough to have received in our lifetimes to do.

The reality of music is that it leaves out the vast majority of people who don't have access to those resources or just simply don't have time to create music.

And I think that is related to the issue that you're talking about. If people don't understand music, if people aren't -- symphonies, right. I mean I have great friends who play in symphonies.

And every year there's less and less people going to a symphony simply because I think there's less and less people who are just equipped to appreciate it or really even understand, you know, what's going on in the music that, you know, symphonies perform.

And so I think, you know, I'm not saying AI is going to like fix everything. What I am saying is there's a reason why we're releasing our technology for free.

And it's because we want to enable people to be able to create music. And when we build our technology, we test it on an iPhone, and we test it on a $10 Tracfone that you can buy in different -- there are countries where, you know, the smartphones, they get the Internet and they get smartphones.

And these phones cost $5, $10. Our technology works on those phones because we are thinking very globally. We are thinking broadly about this. My hope is that people graduate honestly from what they're doing on Boomy and we get this -- we get it all the time.

We get can I do this? Can I do that?

Can I add this? And what we have to say is no, Boomy doesn't do that. Go download and learn how to, you know, make a DAW. There's free and cheap tools where you can go ahead and learn how to make music.

You can download what Boomy does and, you know, put it in these other systems. So our hope, and I think again when it comes to the ethics of it, is to just increase the overall opportunity for people to create music in the same way that auto tune did but maybe in a more advanced and maybe more extreme way. So that's the why, right. And I'm -- yeah, that's --

MS. SMITH: So I just want to say one thing because I know we don't have a platform representative on the panel. The idea that someone is using AI created music to save money on royalties may or may not be true.

But I don't know. From a copyright perspective, it may not be fundamentally different from, example, an ad agency that is using in house musicians instead of commissioning something, right.

So I wanted to say that, and I also wanted to give this panel, which I think has been really interesting, opportunity to give some last thoughts.

I think the next panel is about AI ethics that we have sort of teed up very well, but this is, you know, anything else you'd like to share.

MR. HARRINGTON: I was going to say this is -- technology threatens business as usual until it becomes business as usual. And it threatens music until musicians start to incorporate it.

The drummers had a terrible thing happen to them in the early '80s, drum machines. Well, thank God, they had drums when they hadn't got the drum machine because at the gig you might want some of those sounds.

And I think we're always going to adjust to it. What you said about AI not being there for the turmoil in the kitchen, where you can do the chords but you can't do the words, I think AI -- I looked at some AI recipes to see how language is used.

Let me just read a recipe. Use two large bones, sliced chicken or salmon. Think about that, two cups chicken stock or mayonnaise. 1,100 versus 3200 calories. Cut the snow peas into quarter inch cubes. We're going to need humans for a long time.

MR. DOUEK: I get requests like that all the time.

MS. SMITH: All right, guys.

MR. DOUEK: Yeah, okay. I guess following on that thread of protecting, you know, composers' earning capacity I don't know that I have formed a really strong opinion about this, but I'd like to see -- we can never reverse engineer ultimately how the AI has created its music, what music it ingested and what choices it's made based on which of those, you know, large amounts of training paradigms that it's been exposed to.

So I'd like to see some kind of a blanket license type of model that is basically saying well, at some point in its history, the AI has gained the knowledge and the musical acumen from the body of music that is out there.

And therefore, we will then pass a fractional amount onto registered songwriters and composers. I would like to generate a pool effectively of money that can help during these transitional moments until we figure it out.

MS. SMITH: Anyone else?

MR. HUGHES: I guess going back to one idea we had before, in the same way that we have a hi-res music logo to promote high resolution music, maybe we need a free trade organic music logo if it's actually composed or played by human beings.

Then we at least know a real person is getting paid. I don't know. Then all of our music is gluten free.

MS. SMITH: All right.

MR. MITCHELL: All of our music gluten

free by the way. We built that in. I want to add two points. The first is that I feel like a crazy person sometimes that what's happening on Boomy is it's just people making music.

I actually don't think -- it's not as different as it is new both in terms of the

-- maybe the user experience is new, but the underlying concepts, we only get to do it because of the digitization of music over the last 20, 30, 40 years that's been happening naturally.

So this was always kind of a natural endpoint of where creation was going to go. And I think we have an obligation to see where else it goes. You can't turn the clock, you know, back on some of this stuff.

And the second point would be that we don't have a technology environment today in which tech companies, right, Silicon Valley companies like us can get away with stuff, right.

There's so much scrutiny. People will read our EULA, and we're constantly in communication with our users, with the industry, with all of the relevant stakeholders to make sure that this doesn't happen in an ethical way and that we are not just, you know, flooding the market with a bunch of crap.

That, in my view, is not the right way to build this business. Transparency and sort of being up front about these issues and being, you know, having a perspective of it isn't finished.

We're building this, and we're going to build this along with these perspectives. I think it's going to be frankly probably the difference between success and failure for us in the long run.

MS. SMITH: Okay. So I think that is a great point to wrap up on before we move to bias and ethics in AI. And on behalf of the Copyright Office and WIPO, thank you so much to the panelists for your participation.

MS. ALVAREZ: It's not a break just so you know. We're just switching panels here. I think that people are coming up here. So next, we're going to be talking about bias in AI although apparently this music panel has really sparked some discussion.

MS. ROWLAND: Okay. Thank you so much. Now we're going to turn to our discussion about bias an AI. And Whitney Levandusky at our office is going to moderate this, so I'll let you take it over.

MS. LEVANDUSKY: All right. Thank you, Catie. Hi. Good afternoon. Thank you to the music panel for starting the conversation about bias and artificial intelligence and ethics and artificial intelligence.

This is a topic that touches on all matters involving AI. It is an issue if you are a consumer of AI, if you are a creator of AI, if you are a creator who has found their works incorporated into artificial intelligence.

So I'm really excited that we get to spend some time on this topic. And one of the things that's kind of repeated itself throughout the day is the tension between computer and human.

So in copyright we think about it as the computer generated versus the human generated. And when it comes to bias and AI, when we find ourselves surprised that the outputs of an AI system don't reflect the best of us and sometimes reflects the worst of us.

We often talk about it in terms of objects. We talk about bad data. We talk about the algorithm. But within all of these situations are humans, humans making decisions.

They're selecting, coordinating, arranging the data that goes into the system. They're building the algorithm. They're making decisions based off of the raw output that the AI is making.

And so, for today what I'd like to discuss, we're going to be talking about those objects that are involved and that implicate -- that kind of come into play when we talk about bias, but we're also looking to center humans.

If you are creating an AI system, what

might you be thinking of. If you are a consumer, what are your questions? And Dr. Ulrike Till in the international session said that the important thing with intellectual property and artificial intelligence is to ask the questions, that we may not have answers all the time, but it's important to keep asking the questions.

So that's my job as moderator. And as moderator, I get to ask questions of two great experts. And I'd like to spend just a moment introducing both of them.

So to my left, I have Amanda Levendowski. She is the Associate Professor of Law, and she's the founder of the Intellectual Property and Information Policy Clinic at Georgetown Law.

And then I also have Miriam Vogel, who is the Executive Director of EqualAI, which is an organization that is committed to addressing the question of bias in artificial intelligence.

So without further ado, I would like to turn to Miriam and ask her to set the scene for us. What is bias in artificial intelligence?

MS. VOGEL: Well, thank you for that great question. Bias in artificial intelligence is inevitable from my perspective. I should take a step back and tell you what EqualAI is because most of you are probably not familiar with our new organization.

We were created a year and a half ago by Arianna Huffington, Rob LoCascio of Live Person, Jimmy Wales of Wikipedia and others who started seeing how AI is being used in ways big and small throughout our day and lives and that it's actually influenced by bias.

It turns out that is I would say inevitable because if you look at what machine learning is in particular but AI in general, it's pattern recognition.

It's getting to the answer quicker. It's solving for the answer by recognizing patterns. In the same way, if you look at what is implicit bias, it's recognizing patterns and making decisions accordingly.

So there's no way you're going to, from my perspective, remove bias from AI because it is necessarily biased. What we can, should and need to do is look at the ways that it is unconsciously biased.

We need to set the guard rails of what are the biases that we will not accept, and how do we solve for them.

MS. LEVANDUSKY: Thank you. And so we talk about unconscious bias, and maybe one of things that we can think about is transparency, right. So if bias is inevitable and we are looking towards moving our unconscious bias to conscious bias, how do we inform people about that, like the bias that is reflected in an AI system?

MS. LEVENDOWSKI: Happy to take it.

MS. VOGEL: I'd very much to hear what the professor has to say. So I think that's actually the most important piece of what we can do here together.

All of us have a responsibility once

we recognize this problem. We all have this responsibility I would argue, as a consumer, as a person, as a parent, as a classmate.

Whatever your role is, you have a responsibility to do something about it once you know that it exists. We have to hold companies accountable and we have to get our lawmakers to act on this because there's going to be this inevitability of companies answering to shareholders.

They have their bottom lines. They are set up for profits. They are not necessarily going to be looking for all the potential harms or downstream users of their products.

That's incumbent upon us to demand of them and to make that a part of the expectation for the products that they are creating.

MS. LEVANDUSKY: And so, are there best practices that companies and researchers can implement so maybe it does look less angry at the shareholder meeting?

MS. VOGEL: Absolutely. And for

anyone who wants more detail, I do a half day workshop. So I'll give you a brief snapshot today of what I share and what we talk about to that experience.

But basically, it's a two-part situation. First of all, it's recognizing all the different touch points where our bias can seep into AI. And that's essentially looking at all the human touch points.

It's looking at the design. It's looking at the data that's being used to feed the AI program. It's looking at who's testing this AI to make sure that it's not harming and hopefully is helping more people.

And so we have a variety of solutions that we ask companies to take on when creating AI, but we also have a five pillar framework that we recommend that all companies follow if they're serious about combating an AI.

We actually take a step back before the AI creation and before the hiring of employees who will be impacting this AI. And we say that if you're going to be serious about combating bias in AI, you have to invest in the pipeline.

You have to make sure that we have more people of color, more women, more voices from more regions. It's not going to be solved for us by looking at one or two categories.

You need to be thinking about the universe of who would be impacted by the product and making sure that they are participating in the AI creation.

The second pillar is looking at your HR systems and calling it hiring, promoting with your values. So we know that there's implicit bias in hiring patterns and also in promoting.

For instance, if you look at women in the tech space, there're estimates that they are between 20 and 30 percent of the workforce. If you look at the executive level, it's half that.

And if you're talking about people of color, it's also half that. So we need to make sure that in addition to the implicit bias of the humans who are in the hiring and promotion that we're not doubling down on discrimination through the AI products that are being used more and more with companies large and small across the world in hiring and promotion.

Third, we say that you have to evaluate the data set. You have to see, based on the data that you're using, where are the gaps. If you're talking about an AI for healthcare product, you're going to have limited data sets.

If it's insurance data, you're not going to have people of certain economic brackets who can't afford insurance. If it's genomic data, it's on average 80 person or more people from European descent.

So where are you going to find gaps that you need to change -- there are ways to change it. It's getting different data, or it's solving for over representing, under representing based on the population you want to be serving and that's not available to you in that data that you're using.

But you're not going to be able to see everything in the data because you're talking about terabytes. So then you have to dissect your data. That's the next pillar.

There are tools available, but there's also testing, just interrogating the data, solving for if it's a hiring program. Who are great candidates that are being left out if you test this system.

And we can use data, and I know Professor Levendowski will talk more about accessing different data sets. But it's imperative that you test out the AI products that are already in existence.

And this is true for both those who are developing the AI but also the majority of us who are the users. We are equally responsible for ensuring that the AI that we're using and consuming doesn't discriminate or produce outcomes that are different than what we want.

And then finally, we say the final problem is reevaluating who your team is because a lot of people will tell you there's nothing we can do about the current biases if it's a reflection of us.

Our coding department is mostly white males. That's who's available to be coding these days, and our leadership team is what it is. But there's a really simple solution for that.

And one of the companies that sees things I mean Live Person has this really interesting work around where they create chatbots. So rather than just relying on the chatbot creator or coder or programmers, they actually went to the call centers and had their input in the chat box they are creating.

And it turns out not only was it skyrocketing in terms of the consumers satisfaction and the employee satisfaction, the products sold off the shelf because they had the input of the various perspectives, not just those who had been creating the AI initially so.

It's a much longer answer than you were probably looking for, but I'll stop myself because I could go on.

MS. LEVANDUSKY: No, so it seems like we've got some good tenets to think about from a corporate standpoint.

So your human touchpoints, I really like that where you're taking a look at every time that human hands are involved in the AI system, and that's from hiring, that's building your team, looking at the skill set, making sure that there is sort of some effort put in on in terms of representation and making sure that there's a variety of perspectives and then taking a look at your data.

And it seems like it's both focusing on maybe the specific issues with your data but then also the general trends that come with certain sectors of data. And so I wonder, if Professor, you can drill down that data issue.

MS. LEVENDOWSKI: It would be my delight. Question. How many of the people in the room actually know how AI works, like they could explain it to a friend over dinner?

Oh, a pretty good number. So this is perfect. What I'm going to do is I'm going to use an easy example using my favorite topics of conversation, my cat.

And I'm going to explain to you how bias works its way into these AI systems. And I'm going to make it as simple as humanly possible though I'm sure a computer could do better.

So if I'm looking to build an AI system that recognizes cats because I want to walk into the world and whenever there's a cat nearby, I want to know about it. I think you all relate.

When you're walking around and you're thinking about that, I would first have to start with training data. I would need to train the artificial intelligence algorithm, the most common application currently that's commercially viable is called machine learning, something you've probably read about in the news.

So if I'm creating a machine learning

algorithm to recognize cats, I have to start with training data. And unfortunately, I've looked up how much it costs to license one image of a particular type of cat.

And it's $175. And you need thousands and thousands of images to effectively train an AI system to be accurate. So I'm probably not going to want to license that.

Instead, I'm going to turn my attention to an easily available, legally low risk source of data. I want to know that I own the intellectual property rights so that I'm not dealing in licenses or even thinking about fair use, though I would love to think about it later.

And I want to make sure that it's easily available, not something that's going to take a lot of computing power for me to actually integrate into the algorithm.

And I have exactly the perfect source right in my pocket. It's my camera's photo roll. I have an incredible number of pictures of my cat just sitting and ready to become data for this machine learning algorithm.

So I take the data that I selected. I feed it to the machine learning algorithm, and I realize that I have two different but equally important problems.

Miriam touched on the different potential human touch points at this stage in the process, and we've now touched on two. One of them is the selection of the data, and the second one is the creation of the algorithm.

And depending on how I create this algorithm, I'm going to get biased results based on the data I've input. But if I train the algorithm to focus on the colors of what a cat is, it's only going to pick up on the colors of my specific cat because that's 100 percent of the training data.

So the machine learning algorithm is going to think that a mélange of gray and orange and cream and black is what makes a cat. But if you looked around any urban environment, that actually describes an enormous number of buildings, patches of dirt, brindle dogs, none of which are cats.

So I've essentially created a biased algorithm that specializes in false positives. Now we can go in the other direction and I say okay, pretend that the colors don't matter. And we're not looking for tortoise shell cats only.

We want to look for something different. It's going to look at a potential picture of a cat and say I want the features of a cat. I want big pointy ears. I want big fluffy fur. I want a long tail.

Well, now we're going to get some false negatives because if you know anything about cat taxonomy and based on some of your faces, you're like of course I don't. That's a very weird thing to know about.

If you do know anything about cat taxonomy, you have Scottish folds, which have folded down ears. You have Minx cats which have bobbed tails. You have cats like mine, a Devin Rex, that have short fur.

So you're going to get some false negatives, some things that are cats that the algorithm is not going to recognize as such. Now this might seem very like dumb and whimsical when we're talking about recognizing cats.

But if you're talking about facial recognition cats, but if you're talking about facial recognition algorithms that are used to make decisions about whether people are engaged with our criminal justice system, false negatives and false positives are not just engineering goofs but are life changing due process problems that we are still grappling with.

So that's the basics of how this final piece of how the human touchpoint gets integrated. It's the outputs in the data sets. So now we've got this full stack at how bias can get introduced at every step of the process.

And if we zoom out, there's even a bigger meta question about the human touchpoint and bias, which is should we building this system in the first place.

MS. LEVANDUSKY: And so there are some systems now, and you talk about facial recognition, that there are some companies who want to bring like facial recognition to the consumer sphere.

And so it's now governments and the private individual who have these tools at their use. How should an individual think about these things, and what makes a conscious consumer?

So this is both an ethical issue in terms of engaging with a product that will have a bias of some sort and then also an accountability issue.

Miriam, you talked about how we as individuals can hold businesses to account. So as an individual, how do we navigate this bias issue?

MS. LEVENDOWSKI: I think facial recognition is such a good example because it's something that people can experience in a way that commercially makes their lives easier.

If you're one of those folks who uses

facial recognition to log into your phone, you might find that to be a really easy, frictionless way to not have to remember yet another password.

But we can also think about it being used in its dual use, right, by governments. And I think this is one of those places we're saying decreasing the bias in the algorithm is maybe not always going to be the satisfying right answer because if you think about the way these systems are going to be deployed by government perhaps, they're going to be used for surveillance.

We statistically know that that surveillance is largely going to be targeted at communities of color, the way that it often is, and that's where a lot of the policing is going to occur.

And the problem with that is that having perfectly accurate facial recognition doesn't solve the underlying bias of how that technology is used.

And that's why I think it's so important for consumers who are engaged in this question to think critically about whether these systems are something we even want to be introduced into our society.

So for example, facial recognition has been shown by research done by Georgetown's Privacy Center called perpetual lineup if you're curious in looking at the information, statistically shown to be biased amongst certain groups, against people of color, young people and women.

And so while you might be able to say introduce new data to compensate for those biases, even having perfect facial recognition doesn't solve the underlying problem about whether it's going to be used for over-surveillance against particular communities.

And you can think about this in other types of applications of machine learning as well. We still need to ask that interrogation ethical question as consumers of do we want this.

And the answer across multiple jurisdictions so far has been a resounding no.

Washington currently has a bill introduced. There's other bills that have been introduced in the suburbs of Boston to ban face recognition use by law enforcement in part because of these concerns that it's almost a sort of unsafe at any speed argument, right.

Bias is certainly a huge part of the problem, but it doesn't provide us all of the answers, as you were saying earlier.

MS. LEVANDUSKY: And Miriam, from a consumer's perspective, from an individual perspective, how should you think about or encounter these issues?

MS. VOGEL: I like the approach that some have taken when a consultant to engineer happened to be affiliated with the company tested the company and saw that it wasn't able to identify his face properly, he sounded loud on social media.

I think making sure that you're testing out. You should have access to any public product, and if it's not recognizing you or serving your properly, then you know, we can't wait until the unmanned vehicle can't detect a female figure and has an accident before we correct for it.

So I think making sure that people are raising their hand and saying I want to make sure this is safe for me. I want proof this is accessible to me. Again, there's a variety of harms.

Nikon tried to make an AI program to enable people to be better photographers. And they apparently had no one Asian on their team because if you had Asian eyes, it would label them as closed.

So, you know, people should be demanding that this AI works for them and whether it's a fatal harm or a harm along that line, it's the company's responsibility.

And at the end of the day, the argument I make to companies and that I think consumers would certainly support and should is that it's their product's reputation on the line.

It is their trust and integrity. And Google is able to overcome that lapse in their AI identification, but could a smaller company that we're less reliant on?

And then you'd have to think -- I also talked to companies about the employee morale. Think about how at a company like that when they missed and when they've harmed a consumer whether it's by not being inclusive in their product or an actual physical harm, think about how that impacts employee morale where they're trying to create these AI products that are going to save the world, make it more efficient, make it more fun, make it safer.

There's a real impact. Right now we know there's a war to get the best employees possible. So if you're in the AI space, you can't afford to be harming employee morale in that way.

And then there's legal liability, which I found with most of the companies I talk to that's often a point where I get their attention. So I think that we've seen a bit of case law so far trying to decipher where the liability resides, who's responsible in these oversights in the AI space. I think we can expect to see a lot more.

MS. LEVANDUSKY: So you talk about liability as sort of something that can hook people, can hook businesses. And so it seems like there's often a business decision and a legal analysis that happens.

And I wonder, Professor, if you could just speak very briefly about how individuals make choices to engage with or avoid trademark and IP issues in the selection of their data.

MS. LEVENDOWSKI: Sure. Well, because so much of the data that we were talking about earlier is protected by copyright law, there's a real question about how to pick data that's going to be legally low risk and easily accessible to use for machine learning.

And unfortunately, because of those hurdles related to potentially licensing all of those images, it's not necessarily that people are just using them anyway, although that's often the case.

But one of the other things that happens is people turn to biased, low friction data, easy available, but definitely biased. Examples can include exclusively public domain data to train your natural language processing algorithm, which may seem like an incredible rich source of text.

But it's only going to be as -- basically be as woke as the year 1925, which I think we all agree was not that woke. And that's a potential problem for the training data.

Another possibility would be turning to Creative Commons license data, which has an incredible number of benefits as well. It's a more contemporary corpus of work.

However, if you're look at who contributes to the largest resource of Creative Commons licensed data that's freely available to use, that's Wikipedia.

And that's largely dominated by white male editors from Western countries, so the information is biased. The information I love, or the example I love to use about Creative Commons license data is Rob Gronkowski, tight end for the New England Patriots, not playing this year.

But the article about Rob Gronkowski is about 3,000 words long. The article about the first woman admitted to the New York State Bar doesn't exist.

So it's not even about what the articles say that you're feeding back into the algorithm. It's whether or not the articles are even available to train these systems.

So in both of those cases, you can see how being concerned about intellectual property liability might channel creatives and AI to choose biased data because it seems easy.

And in the paper that I wrote about this issue, the largest example is the Enron email data set, which if you don't know, most of the machine learning that you've interacted with in the last five years has been trained on 1.6 million real emails sent between the executives of a Houston oil and gas company that collapsed under investigation for fraud at a federal level.

And you can think about whether there might be some biases embedded in the Enron emails. And I think that there's actually some research that shows that the answer is yes.

It's not representative socioeconomically, geographically, not in terms of race and gender certainly. And all of those biases are then fed into a system that can potentially amplify them.

So all of those are in response to trying to avoid liability for intellectual property infringement. And instead, we get biased junky data that's informing the AI systems we interact with every day.

MS. LEVANDUSKY: So AI, we call it machine learning. We call it machine process, but there's actually -- all of this bias implication is because of the human involvement, whether it's in the creation of the data, the selection of the data, the production of the AI or the deployment of the AI.

MS. VOGEL: Even in the design.

MS. LEVANDUSKY: Even in the design of the AI.

MS. VOGEL: It's --

MS. LEVANDUSKY: Yeah.

MS. VOGEL: You know, if you're making a hiring -- sorry.

MS. LEVANDUSKY: No, go.

MS. VOGEL: If you're talking about a hiring algorithm -- most companies I work with now are using AI somewhere in their hiring process whether it's deciphering through resumes, video analysis, et cetera.

The person who is designing that AI has a bias in mind of who the ideal candidate is. So aside from the fact that the data from the company or the data in the AI program their using is very likely to be biased, you can have it starting out asking the question because an algorithm is often described as an opinion.

Who are you solving for? Who are you looking for? And I think we can only expect people to be creative, as creative as their mindset allows. And that's why this age old issue begs for an age old answer.

I wish I had something, a sexy AI answer I could give, but you have to have diversity of --

MS. LEVENDOWSKI: There's no sexy AI.

MS. VOGEL: Well, I'm going to leave that. You won't be limited by our bias. You have your own biases to whether or not that's the case.

MS. LEVENDOWSKI: That is my bias.

MS. VOGEL: There it is. And we can all share our biases as part of this. We can raise our hands and share. But at the end of the day, a diversity of thought and perspectives I think is the best tool.

And that's why the consumer awareness

is the best tool to make sure that we reduce this problem. But I'm sorry. You were

--

MS. LEVANDUSKY: No, it's great. I just want to spend just one moment to see if there are any questions in the audience, anyone that might -- yes. We've got a question in the back. Yeah.

MS. TANEN: Hi. I'm Becca Tanen. I'm a librarian here in the Copyright Office. I have a couple of questions about equal AI. So you talked about the five pillars of your recommendations for companies.

So how are those enforced? What are some of the, you know, follow up processes that you might have for companies that have committed to those recommendations and just in terms of incentive for signing on to those recommendations.

You mentioned legal liability, but I think we've seen and it's been discussed in the panels today that a lot of the legal regulations haven't caught up to the technology that's being made.

And unfortunately, large tech companies are often the exception when it comes to those legal rulings, so would love to hear a little bit more about how that is being enforced and followed up on.

MS. VOGEL: Well, thank you, Becca for that great question. And I'd like to enlist your help going forward so we can put some of that to good use.

So enforcement of our pillars. I wish I could tell you that I am an enforcement agency or aligned with one that could enforce this in some way.

I do have the opportunity to brief lawmakers and regulators. Next week, I'll have another opportunity to do that at the FCC. And I briefed at the Fed last month.

But at this point, as a nonprofit organization, we offer best practices. We're working with companies to figure out what the standards are that they could sign on for.

And an important part of upholding these standards is making sure that they are continually repeated. So it's not a one and done if you're looking to reduce bias in AI.

It is a rinse, wash, repeat, repeat, repeat, repeat because as we know, AI is iterative and these biases will continue to pop up in different or similar ways over time.

So at this point, it is absolutely voluntary, but would love your thoughts on how we can make that more mandatory.

MS. LEVANDUSKY: So we'll have to end it here. And so thank you so much. I want to thank Professor Levendowski and Miriam Vogel for their presence today on this panel. I think --

MS. LEVENDOWSKI: Well, thanks for having us, Whitney.

MS. LEVANDUSKY: Oh, anytime. You guys --

MS. LEVENDOWSKI: This has been fun.

MS. LEVANDUSKY: We'll put the stage

up for you anytime. So thank you. It seems like AI is much like Groundhog Day. It is something that we repeat and learn and learn and learn until we can improve and move on to February 3rd. Thank you.

Oh yes, and we'll have a 10-minute break here. Thank you.

(Whereupon, the above entitled matter went off the record at 3:25 p.m. and resumed at 3:36 p.m.)

MS. ALVAREZ: Hello. We're going to start again after the break so we'll let people come and take their seats. This is going to be a panel that is different than the others because we're talking about the marketplace and things, the products and services that are out there that focus on AI.

Before we get started, though, I have a housekeeping note. There is, if there is a Kathleen Burke in the room, let us know. We may have some of your things, so please come and see us at the desk out front.

But with that, I will turn it over to Mark Gray. Mark, take it away.

MR. GRAY: Great. Thank you, Catie. Hi, everyone. Our next panel discussion is on consumer AI and AI in the consumer marketplace. Joining us today are three esteemed speakers. We have, starting with my left, Julie Babayan, from Adobe.

Julie is a senior manager of public policy at Adobe, and she focuses on global technology policy and issues such as artificial intelligence and intellectual property. To her left is Vanessa Bailey. Vanessa is the global director of IP policy for Intel. She previously worked for the law firm, Jones Day and spent 12 years at Nokia doing licensing, FRAND, patent litigation and a bunch of other cool stuff.

And then, last but not least, is Melody Hansen. Melody is an IP partner at O'Melveny & Meyers, and she is the chair of the firm's Automated and Connected Vehicles Group.

And so just -- I guess on a logistics

note, we're going to start kind of with a quick, about a five minute presentation with them talking about some interesting, AI consumer applications they're doing and then we're going to have a panel discussion afterwards.

So with that, Julie?

MS. BABAYAN: Great. Thank you. My name is Julie Babayan. Hi, there. I'm based in Washington, D.C. with the Global Government Relations and Public Policy Team at Adobe. You may know Adobe from some of our most well known products, such as Adobe Photoshop and Adobe Illustrator and Adobe Acrobat. But let me just give you a little bit more background on Adobe and how it fits into today's discussion.

Adobe's business is comprised of three cloud based solutions and all of them use artificial intelligence. We have the Adobe Creative Cloud, the Adobe Document Cloud and the Adobe Experience Cloud.

And what we're talking about here is really specialized AI that's designed for specific purposes. So in the creative space, when at Adobe think about AI, we think about it in terms of how do we use AI to help creative professionals do their jobs.

And it's worth noting that we're unique in the copyright space because at Adobe we develop AI powered tools for creative professionals who then use those tools to create copyrightable works.

And we've really built an industry on helping creative professionals express themselves. And we vigorously support the ability of creative professionals to protect their work through copyright and to realize economic value from their creative works.

And AI is all part of this because it helps us deliver tools for creative professionals, and they appreciate the advantages that AI can bring to them. So, for instance, graphic designers can use AI assisted search to search for stock images on Adobe stock.

And film makers can review footage and

have it suggested to them and other creative professionals can have their tools, like Photoshop, customized to their interests and their areas of focus and their skill level.

So in all of these examples, we're talking about AI helping you do your job better. And I thought it would be helpful to take a closer look at a concrete example of how Adobe is using AI right now and the research that goes in to making these innovations possible.

So let me first just say that our research organization, Adobe Research, is really amazing and we ultimately incorporate many of the insights and discoveries from Adobe Research into our products.

So let me just walk you through some recent research. In this case, researchers wanted to colorize black and white photos, and there were all sorts of reasons why you might want to do this. You can gain new insights from historic photos and look at them in a new way.

And this is something where, you know,

colorizing black and white photos is very tedious and time consuming to do it, to do by hand. So actually, it's a problem that's very well suited to AI.

Researchers were able to start with colored photos and then convert them into black and white photos. And that is your training data. And then basically you were able to have the AI system predict the color version, given the black and white photos.

So just as you and I can look at a collection of photos and gain insights and knowledge from the content, here we have the AI system that is gleaning information about the objects that are depicted here, the various components, the variations in light and shading and shadows and how they behave.

Once researchers trained these AI systems, they noticed something that was interesting and that was that the AI on its own did a really bad job of colorizing black and white photos.

And you can see from here the automatic results are mostly gray and brown, so not the greatest photos to look at. And there's some reasons for this. Some of the colors in the training data are really well defined, such as skin tones, for instance. But other colors are not well defined at all, like clothing.

And ambiguity, as it turns out, is a very, very hard problem for an AI system to solve. So this gets us to Adobe's approach, which is: let the artist decide. The AI gets you most of the way there, but it's ultimately the artist that's in control and the AI is merely a tool to help the artist.

So with this, you get much better photos. And this is just one example. And maybe you wanted slightly different color tones here in the clothing or maybe here's another example as well.

And I bring this up just because I think it illustrates Adobe's approach to AI in general, which is it's ultimately about giving people useful tools and letting them decide how they're going to use them. And this is a good segue into Adobe's approach to public policy on copyright and AI.

Recommendation number one is promote policies that expand access to data, avoid outcomes that limit access to data. AI opens up creative possibilities, but truly realizing these possibilities requires reasonably unrestrained access to data and often including copyrighted works to train AI systems.

So what would be a bad outcome is if copyright law were interpreted to limit access to data, because as we heard in our last panel, outdated and insufficient data also contributes to bias. And outdated and insufficient data also would put the U.S. at a disadvantage as compared to other countries that have clarified their law explicitly to permit such uses.

Which leads me to recommendation number two, which is harmonize international copyright laws to promote AI and, specifically, continue the international trend of text and data mining exceptions.

So a number of countries have grasped the importance of AI and the role that data has to play here. And they've adopted copy rules to facilitate the development of AI. These are often referred to as text and data mining exceptions or TDM exceptions.

And Japan, for instance, recently amended the Copyright Act to add exemptions. Other countries such as Singapore, Australia, China, Thailand are also looking to update their copyright laws to further facilitate machine learning. And the EU recently adopted limited TDM exceptions as well and could be exploring further refinements in this area. So currently the U.S.is a world leader in AI, but the U.S. could quickly fall behind if our policy is to limit access to data.

And that finally leaves you -- I'll leave you with this last one, which is provide guidance in the United States to enable AI to flourish. Here in the U.S., we do have the fair use doctrine and related case law that supports the legality of processing copyrighted material for the purposes of training AI models.

But we would still recommend specific guidance to promote certainty to enable AI to flourish. And providing guidance, establishing a clear right to use copyrighted materials to train AI systems is consistent with the goals of copyright law to promote the progress of science and useful arts and ultimately to help us encourage innovation. Thanks.

MR. GRAY: Thank you. Vanessa?

MS. BAILEY: Okay, I'm going to wait for my slides to come up there. But good afternoon. My name's Vanessa Bailey. I am the head of IP policy for Intel Corporation, based here in Washington, D.C.

So basically we have a lot of really cool AI inventions that I will show you a couple of them just to let you guys know what's going on in the tech sector in AI. But how did we get here? What do we really need for AI to flourish?

And one of the big things is data, right. So we're looking at a big data explosion. And we know that nowadays, especially among young people, there's like a lot of data going through on your phones, connected vehicles.

A lot of data, I have numbers up there. There's just a lot of data going through. With that, obviously you need compute to make the data make sense, to make -- be able to process the data. And of course, the goal here is connected devices.

Looking at the different solutions, you know, I speak on different AI panels and usually Microsoft's up there with me. And I give them so much kudos for having Common. I don't know if you guys know who Common is, the actor bringing AI to the masses, I say, with his, you know, agriculture commercial because it really does -- lets you know the usefulness of AI.

A lot of different areas that we see

AI, education, government, health. There's a lot of different health applications. You know, better research, quicker results.

People talk about this friendly robot that you actually know, like when they give you medicine, the robot has a special temperament because it's programed to have to deal with different personalities. You know, my mother would hate that, but, you know, different people have different things.

Looking at media, you know, the Olympics is coming up. We have things called 3DAT, which is the 3 D athletic analysis, and it's like the overlay on the athletes. And it's actually quite amazing. I didn't actually realize what it was doing until somebody says, no, you know. You can't tell because it looks real time, but it's actually predictive and it's actually estimating what the athlete's going to do next.

And that's what it's -- the overlay is like a one millisecond behind them. But it knows where he's going, and I'm like, wow. Who knew? I thought that was, you know, somebody sitting there trying to move, trying to figure out where they're going.

But with this AI, we actually have all this data in there and we're predicting where the athletes go and what's going to happen next? And we have the overlay on top of it. And apparently at the Tokyo Olympics is going to be some quite amazing things for the audience to see with that.

Looking at some of the practical applications of some of our AI, we have something called TrailGuard AI, which is a part of our AI For Good program. And in that we have partnered with, I think it's called the Leonardo DiCaprio Foundation and National Geographic to deal with the poaching problem in Africa of elephants.

And so the problem statement was that there's just too much land. There's not enough bandwidth for the rangers to deal with the poachers. And apparently there's some stat that every 15 minute an elephant was being poached.

And it seems quite amazing to me that

that was the stat, but, you know, that's what they said it was. And we looked at, well, how -- what can we do, what's the problem here? And the problem was that you have to manually survey your thing to see if a sensor was tripped or a camera was tripped. And you had to manually go over there and then you make the judgment call of whether that's a person or not.

And then you had to, you know, tell the ranger to hurry up and get over there and deal with the problem. And they were missing a lot, obviously, as you can imagine.

So with that, we had this thing called TrailGuard AI, which deals with the on camera inferences. And it basically will detect whether it's -- you can't read that; the fonts are off there. But it basically detects, is it a human, an object or animal. And it makes a decision based on that and it will instantly send out a ranger; no manual checking needed. And it has a very, very, very high rate of accuracy. And so that's one of our AI for Good programs.

Another program that we have which I've actually witnessed and it's quite amazing, quite frankly, it's called Wheelie. And what it is, it uses the Intel Core chip, the RealSense camera and the OpenVINO software, which is a Computer Vision, our Computer Vision AI product. And it helps people who are mobility impaired such that they don't have any functionality of their limbs, and they can use facial expressions to move the move their wheelchair.

I think my next one, if we can do the video of it, just to give you a more of the sense of what it actually does. Yes.

(Video played.)

MS. BAILEY: And so that application is actually used, I think, at the University of Houston Hospital. And it's being used in a lot of the different Texas hospitals. One of the people helping with the training data for that is a Brazil -- sorry. That's not -- yeah, there we go. Other different aspects of some of the applications are the health applications. Now people are called healthineers instead of engineers when they use AI, using it in a real time MRI analysis of a cardiovascular issues.

A lot of -- it's actually used quite uniquely. I was actually, you know, from the legal side, involved with the BrainIAK project, which is really cool. It's Intel's first venture into opensource for using AI in the neuroscience area. And it is quite an amazing project that we're still working on.

One of the applications, I guess, a lot of people can see is our autonomous driving department, which is a very big department for us. Obviously, we acquired Mobileye and we are heavily into autonomous driving. You know, there's a lot of different aspects there, but, you know, from our sense of what we're good at in that area and why we're excelling is that the fact that Mobileye can use all the sensor data and every day to create a very detailed map very quickly.

And we're using that for -- one of the

examples that shows up there, I don't know if you can read that, that we did in Las Vegas, we went through the roads in over 248 miles, 16,000 drives, and we created a map in 24 hours, which, if you think about it, is quite amazing.

And of course, that's using a lot of the data that uploads -- here, let me -- I was going to -- oh. Yeah, if you want to run that, I'm just going to show you just generally why I talk, what -- no. There we go. Yeah -- what autonomous driving looked like because a lot of people have different thoughts about what it is. At CES actually, recently, Intel announced that we're actually at the Autonomous Driving Level 5, which is quite amazing, actually, because people are thinking that we're really not there.

Level 5 -- I don't know if you have a different the levels. So zero is where there is no autonomy, where you have full control over everything in the car. Level 5 is completely autonomous. There is no driver in the driver's seat.

At CES, now Intel announced we're actually at Level 4, which is a great advancement because it means that then, you know, there's a driver in the seat and he can take control over it, but all the functionality of it is there. So there's a lot of different things going on. Obviously, there's a lot of copyright things going on in that space. I tell people, you know, generally I'm like a patent person. I was a patent litigator for 24 years.

And I tell them now it's really cool to be a copyright person because it's so fun. It's all this technology. And there's really a lot of very hard and interesting questions being asked that need to be answered, especially in the different, various tech areas.

But so, you know, there's a lot going on in AI. And there's a lot further we can go in AI. But we are really getting to a good place, I think, rather quickly.

And but I think, though, that the -- as far as copyright is concerned, I still think that copyright laws are still adequate. I think we're doing fine with what we have and that smart lawyers are figuring things out.

And so maybe there's not the biggest conundrum in the world that maybe everyone thinks. But maybe I'm just not optimistic about everything in life, so.

MS. HANSEN: So, as Mark mentioned, I'm Melody Drummond Hansen and I have the perspective on this panel of being an outside counsel.

So rather than being responsible for advising primarily one company, although there are many hats for both of you within your companies and different types of products with different, potentially different, competing interests, you know, I have the pleasure of advising companies kind of across the ecosystem for AVs and connected vehicles and hearing from them, what keeps them up at night, what are their questions; what are the, you know, what are the different technology models they're using, what are the different data models that they're using.

And one is one of the things that I like to remind folks of -- I mean, how many people who are very familiar with autonomous vehicles, follow it closely? We got some super fans. So maybe you'll tell you if I -- if you want to correct anything. But, you know, the DARPA challenges in the United States were back in the early 2000s. You know, in 2004, you have challenges for vehicles to drive -- I mean here, off road, in this picture, with being able to do their own navigation, their own detection of obstacles and navigation of obstacles for some of the challenges.

You know, some of them were in desert locations was, you know, where rock is one of the biggest obstacles to the technology. And then in urban environments, which we all know create different challenges for autonomous driving.

But you have back, you know, 15 years ago, cars perfectly capable of driving themselves, navigating novel environments, sometimes without a lot of input or a lot of advanced knowledge. And that was part of the design challenge, is to create courses that would be unknown or have had details only known at the last minute.

And then, you know, today I live in -- I work in Silicon Valley; I live in San Francisco. And I see a lot of AVs on the road every day from different companies. And on the roads -- I mean, to kind of echo Vanessa's point -- the technology really is capable here on the roads.

We may all have an experience with the lower levels of autonomy, with our assisted driving techniques. I mean, for me, when I drive with my cruise control on, just regular old cruise control, I see how stupid everyone's driving is because you start to notice when people are braking for no reason or braking because they're going under a tunnel or braking because they're going around a curve or just making a dumb decision that slows everyone down.

So I think there's a lot of promise and people acknowledge the promise. And echoing Julie's point, you know, the latest DOT guidance on AV, it's called a AV 4.0, they really put it in terms of we need to enable this technology so that we can remain competitive in the world and so that we can remain a world leader.

And so I think there's that recognition these days. There are a lot of diverse market players in this space. So in addition to, you know, the Waymos of the world or the Zoox of the world or the Cruise of the world, we're actually kind of designing around cars or designing their own cars.

There are other players who may supply, you know, only LiDAR or only software or only, for example, chips or data storage that are enabling the space. And all of them are potentially affected by the rules that we're making.

Like Vanessa, too, I come from the patent background, so it was more fun to be in the copyright space a bit. But here is a picture of a device that's an aftermarket kit that you can add onto an existing car and make it self driving, and there are a number of players in this space who are doing that as well.

So I thought it would be helpful to have some examples of the types of data that are being collected by AVs. And this is from The New York Times and citing Google and depicting a modified Lexus. And you know, at number one, on the top of the vehicle, is the LiDAR, which is sort of like a point on, too, of light and radar. It can map, in the loose sense, an environment by shooting out light and receiving beams back.

Underneath that is a camera at number 2. And then on, what's labeled as 4 on here is radar, to be able to detect information around you; 5 is another smaller LiDAR. And in the trunk, at 3, is a big fat computer to try to analyze and deal with all this data.

I mean, terabytes and terabytes of data being generated in a single day. So what does this look like? You know, here's an example of LiDAR are from Luminar. This is the website that it came from. But this is driving down the Embarcadero in San Francisco. And you can see it's we're picking out different objects. And in the picture that Vanessa showed and we'll see another example of this, whether it's with cameras or with LiDAR, there is often metadata attached to that, often real time to kind of draw boxes around different types of objects.

So we need to be, as the other panelists were talking about, we need to be able to distinguish between a pedestrian and another vehicle in addition to being able to judge distance, stop signs and the like.

And so, you know, one of the things I joke about you know, folks want to ask about with AVs, how are we going to train the algorithm as to who it should kill when faced with two options?

And my first response to that is always, well, you've driven a lot. How many times have you had to decide who you should kill during your daily commute?

But in reality, some of the more practical, immediate problems have to do with being able to accurately represent, you know, detect traffic cones in different positions and navigate these types of everyday obstacles, map those and make accurate depictions.

And so you can imagine that these have a lot of questions around them in terms of you've got a lot of video or photographs. You've got datasets that are potentially annotated with this type of metadata, either in real time or using its own AI.

And folks really don't know what to do about this. They really don't have a sense of how this will be treated. And I think different market players have different perspectives on it.

But to echo the point on the previous panel, the notion of having to get the rights to the hundreds of thousands of photographs of traffic cones or something like that is a kind of daunting task. And yet everyone acknowledges that the industry benefits from sharing data and being able to share data about the routes that you're traveling, about the ways that you've identified obstacles about existing routes.

And there are a lot of efforts to do that these days. Waymo has released certain data sets, Lyft has released certain data sets with this type of thing. But there are a lot of questions around the rights on them. And to the extent folks are familiar with them, the data sets have different types of information on them. Some have video, some have 2D photographs. Often they have this attached metadata.

And they have terms. I mean, they're called open data sets, but they have terms attached that, you know, one of the data sets was sort of criticized for specifying that you can't use it to train an autonomous car.

So with that in mind, I'll finish up this part and another pick -- these are the practical realities of the car, not what -- who they're going to kill. Can they drive in the snow? But that's the type of data we're talking about.

MR. GRAY: Great. Thank you, Melody. So I guess to start off and maybe to stay on the topic for a minute, you know, there's a lot of talk sort of in popular culture about, you know, self driving cars and what cities will look like in 10 or 20 years. How close practically are we based more where the technology is right now?

MS. HANSEN: I mean, I think we both have a lot of confidence and optimism.

MS. BAILEY: Right. That's right. Because I think, you know, as we've seen, that that was actually driving in Jerusalem when you saw it. But it was actually, you know, it wasn't staged. That's real driving.

So I think, you know, I think we're pretty close to being there. I mean, the whole issue, I find, is not where the technology is, but whether people will be accepting of the technology, right.

So when you drive in your Tesla, you know, it's doing pretty well until you know, you want to get all the way over an exit at the last minute, right, and that's when everything falls apart.

So, but I think that with the technology, we're getting better, but the whole key is data. And I would echo what Melody said. I think, you know, from an in house perspective, when you're looking at this and everybody's asking, you know, what are we going to do about this because, you know, I would even say more than that. But, you know, we need hundreds of thousands of data when, you know, say, for OpenVINO which is Computer Vision. You know, you've got to train dataset. And where are you going to do the data?

And when you get it, how can you not know if there's something in that's copyrightable? And I think that we just rely on the current case law and the current cases for fair use and just rely on the fact that, look, we're not using this for an expressive value.

We've got to be able to turn over the data because safety is number one in autonomous driving and you only get safety by knowing all the variables, dealing with all the variables and having access to what every variable can look like in different countries, right? Because when you have autonomous driving, I mean, we do have the car, right? It has a sensor data. It's coming in. You might be in, say, I don't know, Jerusalem.

But then you're sending it to the cloud, and the cloud is in, you know, San Jose, California or something like that. So you have a lot of very, you know, things to consider, especially from a copyright perspective.

But it all comes down to the fact that you need the data, but that we are -- you know, we're closer than I think a lot of people realize on autonomous driving.

Go ahead.

MR. GRAY: Oh. And then I guess sort

of as a follow up question then, what are the specific, I guess, copyright law issues that come up that may be unique or particularly specialized to the autonomous driving space?

MS. BAILEY: So for me, I think that -- I don't know if it's unique, but there are different aspects as like ownership, right. I think, in autonomous driving, a lot of the contracts that you see are between the OEMs and the, you know, the person who owns the AI model, the trained data set.

From my perspective, like you have different contracts for different leads. So, well, for us, you know, we have the sensors on the car and we have the road work map and we have all of the -- we have the AI itself.

So for us, you know, when the -- the things, that sensor is sensing things in. You know, a lot of copyright issues comes with that, like what if it saw a billboard and there was copyright on the billboard. And now we took it in and we sent it to the cloud and now its fixed, and what does that mean in the law in Africa because that's where it's sitting on the server.

So you've got a lot of different international aspects to consider that make it unique and fun. We also have that aspect of then, well, you know, when it comes back, you know, you've got the trained AI model. It's going to push information back to the car to give you the green boxes and tell you where things are.

But then, you know, in some countries where the driver's trying to say that, well, we're the driver and we tell the car where we want to go, from home to work. And so we want copyright or whatever it takes in from the cameras. And we're like, wow, really? Because, you know, there's a lot of different variables going on.

And, you know, and you read at least U.S. case law, it doesn't seem that that might be the right answer. And then there's just a lot of different unique variables, because you've got the information coming in. You've got the information going out. You've got the information churning in the cloud and you've got the information coming back.

And there's a myriad of different copyright aspects involved in all of that, from video to imaging to everything.

MS. HANSEN: Yeah, I guess to echo that point, and I would say the questions are really -- I don't think they're unique to AVs. I think they have a lot to do with training AI machine learning.

But there are a lot of questions around data sets and there are myriad ways that these, that data sets are made. There are different players who are offering them. You may not have insight even to where the initial data came from if you're using other people's data.

And particularly with some of the data sets that are available, I think there are a lot of questions about how will those different types of data be treated, you know, whether it's visual data, location data or other information, are there different ways it'll be treated?

And I'll say, you know, there's a lot on these open data sets because they're being made available. They have terms. The terms are variable. And I should mention, we have in the room one of the foremost open source experts, Heather Meeker, who works at my firm. I have the pleasure of working with her on a lot of issues that involve the kind of convergence of questions and copyright with specific questions related to open source as well.

And so I think there are a lot of questions around that. You can ask questions. We kind of talked about this in our prep. You can ask questions at every stage of it, right - author, infringement, and fair use.

But I think a lot of the questions right now are around the default rules on fair use, because there are many in the industry, not everyone, but many in the industry who feel, just in their gut, you've got to be able to use this data.

But when you start to try to apply it, it's not -- it may be not so clear what the reason would be.

MR. GRAY: Right. And then I guess to sort of stay on the data set training side, Julie, what kind of data sets is Adobe using for their products? You know, are they looking at a lot of open data sets? Are they generating their own? What's the thought process there?

MS. BABAYAN: Yeah. So I guess some examples of some data sets that I know of that are -- and, you know, I won't speak for Adobe specifically because, you know, that are out there are, you know, the types of open data, the types of data sets that we see that are image collections that universities and other research institutions put out there for research purposes.

And, you know, for instance, some of these image collections provide thumbnails or URLs. And for many researchers, these can be just tremendous resources for doing the type of research that leads to new insights and innovations.

MR. GRAY: Cool. And then another question I had is, you know, when we when we spoke before we, when we prepared for the panel, Vanessa, you had a really interesting point about the different ways different product groups look at their sharing or sort of maybe not wanting to share the training data that goes into a product where the product you're selling, maybe the fact that it is so accurate because of the data set.

MS. BAILEY: Right.

MR. GRAY: Could you talk a little bit about kind of different perspectives on either sharing or just maybe keeping a little bit more closed, the data set you use to train?

MS. BAILEY: Right. So one of the things we're looking at, because we have different -- different business units have different needs. You know, I would say one of our products might need to be trained. And so we need access to data. But then on the flip side, depending on the business model where you are in the ecosystem and the value chain, you might be looking at the fact that, well, you've collected a lot of data that someone else needs.

So we need to protect it and so there's kind of this internal, I will call it, rivalry between, oh, well, we put all this effort into getting this data. We, you know, this is something we want to license out on very reasonable terms, but we still want to license it out versus the other core AI team.

It's like, no, everything's open source and so open data. You know, this is this is -- we're not using it for its expressive value. We're using it for, you know, to create metadata so that you know that that's a baby and not a dog. Surely, you know, everybody should have access to the data.

So there is this kind of you've got to figure out where you are in the business model on what you advise your clients, on what you base your business units on. You know, what they should be billing and looking out for on the legal front.

MR. GRAY: Okay.

MS. HANSEN: I was just, on the way here, I was at the airport and I overheard a conversation. It's not uncommon in the Bay Area that it's about tech. And I overheard someone saying, oh, yeah, well, we had this data set and it's for facial recognition.

It's really great because copyright doesn't apply to it at all. I mean, it just is like completely not what you're doing. And I thought, well, that's interesting because I'm going into a conference where a whole lot of smart people are talking all day long to try to figure out that answer.

So I think, you know, that's another area like the, what you're doing with the data set. I mean, there are so many different uses of this.

MS. BAILEY: And also, I guess it depends on closed versus open, right. So I think Adobe -- and I could be completely wrong. I'm not speaking for Adobe. But I would say that they have -- they use a lot of sometimes closed data sets. So it's their own and maybe it's more reliable.

And the data in AI is only good as -- well, I guess the AI is only as good as the data, right. So a lot of times we, I guess, in tech, can't get our own closed data set.

So, but if you've a closed data set, you could be more reliant upon the fact that maybe you own the copyright, hopefully -- that's what they're talking about -- versus need to get a copyright.

And, but then again, I find that, I guess, when you've got these big autonomous driving, you've got these kind of mobile applications, it's hard to have your own data. It's almost impossible to have your own data. It's also impossible to get it from just one source.

So, again, to have reliable, less biased data you have to look from different sources. And so with that comes, you know the issue of what are you going to do if there's one in the data set that is copyrighted and you don't know anything about it.

You never know, when you think about it, quite honestly. I mean, how do you know?

MR. GRAY: Great. Well, I know we're already running a little late on time, and so I've been told to wrap up. So I'd really like to thank you very much, Julie, Vanessa and all of you. Thank you so much for this conversation. This was very interesting.

MS. ROWLAND: Thank you so much. We are actually going to do our final panel of the day now. It's about digital avatars and AI, and it is going to be quite interesting.

So I think we're going to be switching out the panel here. So we are, we saved the best for last, I think. This is going to be a great panel that talks about all the interesting information that you can learn about using a digital avatar. God forbid that you yourself are, you know, in some other movie right now that you don't even know about, but we're going to learn about it now.

So we have two really great experts here to talk about it. We have Sarah Howes who is with SAG AFTRA, who's going to be talking about that perspective. And we have Ian Slotin from NBCUniversal, and he's is going to be talking about the movie studio side of this.

And I want to start with Ian, who's going to kind of set the stage for us today.

MR. SLOTIN: Great. Thanks, Catie. And so what we're going to talk about today is, as Catie mentioned, is not content in general, but professionally produced film and television content.

And what I'm going to go through is some copyright considerations. I'll go through that very quickly because it's a lot of repeating what's already been said today and then move on to sort of the right of publicity issues, which folks talk about a lot.

So in thinking through copyright issues on this, I thought I would take an example of something that's not professionally produced and not something you would think of as a popular film or television show. But it's just kind of an interesting use of the Internet of deep -- oh. Okay, there it is. So this is a meme that's on the Internet where users are using -- oh, and now it's gone. There it is.

MS. ROWLAND: Oh, no. There it is.

MR. SLOTIN: Okay, so anyway -- no? Oh, well, I'll just say what it is. Essentially, it is users taking Nicolas Cage's face and putting them into various films where he didn't actually appear: Lord of the Rings, Superman, Titanic and other things like that.

And so if you want to think about copyright issues related to that -- let's see if it moves to the next slide -- one before. There we go. So in thinking about the copyright considerations here, I've grouped them into two categories. So it's on standard clearance issues, which is essentially the same issues that apply for content that doesn't include an avatar like, you know, what elements are eligible for copyright, what parts are new, what parts are not.

And secondly, what other elements are appearing in the footage and do those need to be cleared. And then the novel issues are the authorship issue that's been discussed a lot today. And secondly, the use of training data to create the Nicolas Cage avatar.

So speaking quickly about the standard clearance issues, very quickly, because this is obvious to everyone in the room, the question of eligibility. You know what, you know, the likeness itself is obviously not something that can be copyrighted, neither can contextual footage if you know, if it's from somewhere else.

But perhaps the avatar's performance and poses, movements and dialogue, and maybe if there's a new plot, you know, those could be eligible for copyright. One note is that, you know, in the example of Nicolas Cage, you know, that's pretty much a derivative work. So that's not going to work.

Moving to the novel considerations on authorship, I think for -- while it is true, of course, that not -- that works created purely by machines are not eligible for copyright in the United States. I don't think we have to worry too much about that for professionally produced film and television content, because we're a long way off from a situation where a studio like mine will put out a film or a television show that's entirely created by a machine. There's just too much, too many variables.

We're just, as I think it's been shown here, that there's just too many moving parts there right now. So the question is more around, you know what we're worried about are disputes that may arise as to who the author is, who the human is, who did the lion's share of the creativity.

It could be the designer of the AI

algorithm, as has been discussed earlier today; who -- the person who selected the training data or the operator, the one who iterates on the results.

And from our perspective as a movie studio, our primary concern is to make sure that there's some clarity around that so that we know as a producer who's right, you know, who we need to clear the rights from -- you know, who we need to obtain those rights from.

Moving next to the training data issue, but in a situation where, like in the Nicolas Cage example, presumably a lot of footage of Nicolas Cage from various points would need to have been used to train the AI system to do what was just done.

And so the question would be, you know, using preexisting footage like that, you know, is it a fair use or not? And without judging that particular example, you know, our main points on this is that we don't need a special new fair use rule in United States to deal with this question.

The existing fair use factors were supposed to be technology neutral when they were enacted. And indeed, when courts have applied them to mass digitization applications, they have come up with nuanced results.

So, for example, in the Google Books case, that was determined to be a fair use, whereas in the TVEyes case, which was a similar situation, except that the application was different, that was determined not to be a fair use.

So moving now to rights of publicity, just as a quick reminder, the rights of publicity are state rights. They're not federal. And it has to do with using a name, likeness or identity for a commercial purpose. There's a mix as to which states recognized postmortem rights. It's not across the board.

And finally, when we're talking about film and television and expressive works, you know, there needs to be a First Amendment accommodation there because they are expressive works.

So what are the key tests that have come up in sort of accommodating the First Amendment? One is strict scrutiny, right, which will say that not really all users are exempt because a right of publicity statute is a content based regulation of speech.

And so under the Supreme Court's rubric, that requires a compelling government interest and a narrowly tailored solution to that interest. And in many respects, you know, simply, you know, wanting to have a say over what is said about you in an expressive work that is not defamatory, you know, probably generally would not rise to the level of being, you know, being a compelling interest.

Transformative use, this is a test that was developed in California, essentially saying how -- what was the transformative nature of the work and the use. And finally, the Supreme Court had one case that's a key note that

I'll mention a little bit later.

So strict scrutiny, this was a case -- this is a case involving The Hurt Locker. The Hurt Locker is about, the main character is a bomb disposal technician in Iraq. It was based on an interview that was given to a reporter by Jeffrey Sarver.

And Sarver sued, saying you based the movie on me. That's not okay. The 9th Circuit held that, you know, applied strict scrutiny, applied the Supreme Court case Reed v. Town of Gilbert, which had come out the year before, to say that essentially the California right of publicity statute is content based and is subject to strict scrutiny.

And by the way, for that reason, Sarver lost his case. And by the way, one of the points that they made was that Sarver wasn't even in the business of monetizing his persona. So there wasn't even a loss of income or anything like that in that situation.

Transformative use, just very quickly,

established in two key cases in the California Supreme Court, one relating to lithographs of the Three Stooges, the other relating to a very fanciful comic book.

And essentially, you know, one thing to note about the transformative use test is it still needs to sit within the rubric of the First Amendment. So it may very well be that a test that can be applied and makes sense in certain cases, but it must -- that must be because it's consistent with the First Amendment.

One way of thinking about that is the transformative use tests might actually be a way of distinguishing between expressive uses and merchandising uses, right. So the example of the lithograph below was a picture, but it was on, I mean, it was on T shirts and other things.

So the question is, you know, we're at sort of a line between merchandising and an expressive and since merchandising has a lower level of scrutiny, commercial speech does, you know, perhaps, you know, that this is a way of differentiating those things.

And finally, there is the only Supreme Court case about right of publicity and it's a human cannonball case. And, essentially, what happened is the local news showed a 10 minute -- 10 second story, which was essentially the entire act, right.

It showed him shooting out of the cannon. And Zacchini said, hold on a second. No one's going to the fair and see me perform my human cannonball act because they've already seen it on the news. And the Supreme Court said, yeah, you know, you got a good point there.

By casting the entire act, this is a substantial threat to the economic value of the performance. So a couple of caveats about this case, it's a very old case. It pre dates more recent Supreme Court pronouncements on strict scrutiny, and strict scrutiny wasn't applied in this case.

Another point, you know, well, one of the key interpretations of this case, possible interpretations of this case, is that while, you know, is that taking a person's performance that they have worked to establish that, you know, their livelihood, you know, there can be a right to publicity claim for that.

So now moving to, you know, avatars and depictions of performers, there's -- this leads to sort of a distinction between depictions of performers as themselves and depictions of performers performing overall.

So in the first example of showing up as themselves, I gave a couple of examples, you know, where avatars weren't used, but actors were used, right. So the top one is Once Upon a Time in Hollywood. The actor, Mike Moh, played Bruce Lee as himself. You know, Bruce Lee shows up on the film set and fights with Brad Pitt.

And then the second one is the FX series, Feud: Bette and Joan, in which Susan Sarandon and Jessica Lange, you know, played Bette Davis and Joan Crawford in a series about, you know, power dynamics in Hollywood and how women were treated in Hollywood.

So, at least from our perspective, it seems that these portrayals are all perfectly okay and that I think most people would agree that, you know, that the subjects of these works should not have the right to approve or not approve their depiction in these sorts of works that are not defamatory.

And the corollary to that is should -- does the technology make a difference? So it's Mike Moh, looks a lot like Bruce Lee. You know, he worked really, really hard to have the same mannerisms, the same voice, everything else.

And in fact, you know, if you look at some other portrayals, you know, people remark about how well an actor, you know, inhabited a role. You know, what's the difference between that and using a, you know, an avatar of the person?

From a First Amendment perspective, it seems that there probably isn't a difference. And that should probably be a guiding principle in terms of thinking through, you know, the permissibility of these things.

The second category, right, is performing a role, right. So this is a scenario where essentially, say, Universal says we're going to make Jurassic World, the next movie. We don't want to hire Chris Pratt to play the lead role. So we're going to use the -- starring the avatar of Chris Pratt, and we're not going to get his permission to do that.

That seems to fall under this Zacchini precedent, potentially saying, well, you'd be sort of essentially taking away this actor's livelihood, you know, that is something, you know, to be considered. The examples in here are actually both of deceased celebrities.

The first one is Paul Walker from Fast and Furious when he passed away during filming. The second one is Carrie Fisher in the recent Star Wars film. Our view would be that there's a difference between living and deceased individuals in the sense that, you know, once you're deceased, you no longer have the ability to earn a living.

You know, that is no longer taken, and under a strict scrutiny analysis, you know, the -- an interest of an heir to protect the reputation of their -- of the person when the actual defamation doesn't even apply after a person has passed away, you know, or the right to collect the money for it is not sufficient for that.

So a couple of other points on this. If there's going to be a consent required to use an avatar to perform a role, that doesn't sound like a right of publicity right. Because right of publicity relates to uses in commercial speech. And these are not commercial speech.

So if there's going to be a right like this, it's probably a, say, sui generis right. The second point is that any right like this needs to be appropriately tailored to exempt the kinds of performances that we think are okay when we use living actors.

So, for example, a biopic of an actor is probably going to have scenes in which that actor is performing a role as part of the biopic of the actor. Clearly, that's okay too. So. So just it's a very nuanced area.

And then the third point is that it probably makes sense to consider the concept of deception and fraud and passing off in this context, because thinking through the difference between having, hiring an actor to look exactly like someone and using an avatar, the difference may be that the second case is so realistic that people could actually think that someone performed when they didn't, that someone said something when they actually didn't say it, those sorts of things.

And if someone is taking active steps to profit from that, to deceive the public into thinking that someone, you know, endorsed something when they actually did not or were involved in the project when they were not, you know, that may be a framework to think through some of these issues.

And finally, just a note, you know, the distinction I mentioned about performers as themselves versus performing a role, it doesn't really work for other kinds of public figures. So, you know, musicians, athletes, political figures, you know, it can't be the case that, you know, musicians are known for singing or playing instruments.

But in a biopic of Freddie Mercury, you can clearly show Freddie Mercury performing music. In, you know, or showing Jackie Robinson, you know, playing baseball or showing Dick Cheney, you know, doing his usual Dick Cheney stuff. So that's pretty much everything I have.

MS. ROWLAND: Great. Thank you, Ian. And now we're going to turn to Sarah, who has maybe a little bit of a different perspective on the issue, from SAG/AFTRA.

MS. HOWES: Yeah, definitely. I mean, I would say that, like, it's important to note that, you know, Ian and I have more in common than we disagree. I do respect his work on the First Amendment, to a large degree. We do have some significant differences that always get nice headlines. But, anyway, so I'm going to start off.

And I'm going to talk a little bit more about this from a production standpoint, to be completely honest, just because part of this is, a lot of you are lawyers. And part of this is an opportunity to think about this from a perspective of sort of guidance on what the union kind of expects.

I'll spend a little bit of time kind of responding to some of the things Ian said about the law and things like First Amendment. But I really do want to focus on what is the technology, what are the advancements in the technology, you know, what are the union rules -- a little bit about how we see the laws.

And then I'm going to go into a pretty serious topic, which is how I spend about a third of my time now, which is, you know, Me Too related issues in the industry.

Okay, great. So this is taking about just generally some of the concerns we have about the uses of digital images and sort of how it runs into some of the things that Ian was talking about.

We have just the traditional, I'm sure everyone's familiar with the sports cases of using an athlete and a video game. I'll talk a little bit about using a musician in a video game.

We have the realistic -- what Ian was talking about in terms of, you know, bringing an actor to play a role with digital technology used. And you're talking about that in films. You're talking about that and TV shows, video games to give a realistic acting performance.

The next one is holographic live performances of musicians or actors. And this is happening right now. In fact, unfortunately, the leading hologram creator is also the person who created FilmOn. And also, he doesn't like copyrights very much and then also just got the largest sexual harassment judgment against him. So I'm very happy that he is the leading person on holographic concert technologies.

The last piece is just voice cloning. That's something that has been interesting from the work that I've been doing for the last four years is just how advanced this technology is getting and all the different ways that you can use it and how that impacts people like voice performers.

Great. So we're going to talk a little bit about just -- I want to have a little bit about sort of the old school way of doing some type of false depiction. So there's -- there was a film, Nymphomaniac. It was made overseas or it wasn't under one of our contracts. And the way that they revealed to depict Shia LeBeouf's character as engaging in simulated sex was they actually hired porn stars to come in and have real sex, right?

And then what they did was they didn't

actually use any type of CGI. They just actually edited it in such a creative way that it looked like it was him being depicted in these acts.

This is an important note. As I said right now, all of you are representing film companies, and just so you know, it doesn't matter if someone's a porn star or not. There is no actual sex in any union covered work, a body double as a principal performer.

And it is obviously for lots of reasons, very, very risky to have any performer performing actual sex. I'll talk a little bit about that later.

Okay, so then the next kind of, you know, more advanced form of doubling is where you want to actually see the person's face on to the body.

And so Natalie Portman -- this is a little bit scandalous because they didn't disclose that it was a body double at first and then it came out after she won an Oscar, but they hired -- so they hired a dancer to be the body double in this film.

And she performed the scenes and then Natalie Portman performed the scenes. And, of course, the professional dancer did a little bit of a better job. So what they did was they put this mask on her and then they used CGI and, you know, some basic advanced tracking technologies. And this can be done with the software, After Effects, which is a pretty affordable consumer product. And then they just did CGI blending and adjustments.

Peter Cushing, this is where he gets a little bit interesting because this is really one of the biggest examples of what I call just like digital human technology, where you were able to create, even though you could, you know, if you looked close enough, you could tell that it was digitally done.

But I know people that literally were like, wow, he looks great for his age, right. They went to the film. They were duped. And this process was really using a lot of different types of advanced motion picture sciences and special effects.

It was starting with, as you guys have seen through the last decade, right, performance capture Guy Henry, who was a very well respected British actor, decided to take on this role, which is good for him. Also, he got paid a lot because he wasn't really depicted as it. And he came in and they were able to do the basically mapping and sort of CGI with his face. They were able to capture all of his movements. And then the way that they made it look like Peter Cushing is that they couldn't just use it from images and old footage of Peter Cushing.

They were sort of lacking the ability to really get his bone structure and his actual face mapping. So they were able to track down, because Peter Cushing was known for being in a bunch of like British horror movies, and so they were able to track down this old face mold that they had.

And so they were taking -- because

when you're taking somebody from an older era before the digital era, you don't have things like which I'm about to show you, like on set scanning, right.

So, nowadays, particularly if you're doing any type of large budget or action adventure movie, part of a condition of your employment from being a background performer to being the leading star is to do various types of 3 D, 360 degree scanning so that people have data basically of your entire body.

There are a number of my members who do not like this. They find it very intrusive. The public people who have talked to the press and media about it are Donald Glover was very upset about this. He did not like engaging in this. He had some concerns about the cultural and misappropriation, potential abuses of his body being scanned. And then also Jessica Chastain did not like being scanned in this way. They found it intrusive.

But this is becoming commonplace.

It's a condition of employment, including for background actors.

Okay, so then this is where we get to artificial intelligence, right. So all of those processes that I was talking to you about were very expensive and they were very time consuming. The Peter Cushing example took 18 months and millions of dollars, right.

And then in January 2017, an article dropped about artificial intelligence, deep fakes. Okay, so I'm going to show two examples. I'm not going to show the whole thing, just enough for you guys to see it.

So the first one is a fantastic professional impersonator. And then the ability to sort of take his impersonations -- yep, that's right, and then use deepfakes to have him depicted as the person.

(Video played.)

MS. HOWES: Okay.

(Video played.)

MS. HOWES: Okay, So for purposes of

time, we'll move on to the next one. It's great video, though. You can watch it. So the next one is -- was really scary, to be honest, for actors, when they saw this.

So this is putting Harrison Ford into the Solar movie.

(Video played.)

MS. HOWES: All right, so we can have that example. Oh, it's almost done anyway, just going to stop it.

So when that video came out, let's just say I got a lot of responses from the members and even directors started seeing it. And, you know, because to be honest, that was pretty much production ready, right. That's what was sort of scary about it.

And let's see here in the next slide. Okay, so then this is the part where I wanted to talk a little bit about voice cloning, because I actually think voice cloning has a lot more potential harms. And I honestly think it's going to impact your average middle class performer a lot more, because if someone's going to take an image, right, there has to be a pretty valuable amount of money to that image for it to be worth taking.

Voices -- how many people in here, you know, can name the actor who played Little Mermaid or the actor that you hear at Walmart, right?

Those people make their living as voice performers, podcasters, videogame performers. And they're not the people you see on the red carpet. So Adobe Voco, this is one form of technology.

So what's interesting about the voice crowning and the new voice technologies is actually, you know, synthesizing and re editing that's been around for a long time. What they're saying is new is the ability to insert a new word. So have somebody say a word that they've never said and you think it's that person saying it.

So there's different ways to go about

doing this. So Adobe Voco, basically what it does is it takes a recording of someone's voice and it breaks it out into these tiny little bits of 80 different types of sounds that are common in the English language and it quickly rearranges them to have you say a word.

So Jordan Peele did a great presentation of this where he came in and he said something like, I love to kiss my wife when I come in the door and they he changed it to I love to kiss my dog when I come in the door, right, so they were able to insert the word dog.

So then the second one, with a link, this is a company that kind of started, which I'll have you click on, a couple years ago. They're using artificial intelligence to do voice cloning.

So I'll have you go to the bottom. And I'm going to explain a little bit about how this is done. So I did this, and I am not a professional actor, and it sounded a little bit robotic, right?

And then we got, you know, people like Richard Masser and Harry Shearer, people who are top voice actors to do it and it sounded like it was them. So what they do is it's literally a minute of capture where they say a bunch of random phrases. And then within a minute you have a voice clone, and then you can just literally type in sentences.

So they're -- an example, if you click on the little blue on the left, that's a person's real voice.

Oh, I'm sorry. Go to the Mozilla Firefox at the bottom, like the other app. I have it uploaded on there. There we go. Doesn't work on --

(Video played.)

MS. HOWES: So then go to the next one.

(Video played.)

MS. HOWES: So that was done with one minute of audio. Yeah, okay, so then I want to go back to the PowerPoint. Okay.

So there are union rules around some of this stuff. So I just want to spend a little bit of time talking about, because especially since a lot of people here have been talking about what your source material is, what your underlying licensing material is, is that if you're a production and you license out or you yourself use any type of existing footage that was made under our contracts, you actually have to go back.

And that new time of use that's important that you're going to use this footage again, you have to get permission. And if you don't get permission, there's going to be contractual damages from the performer.

It's really important to say that it's the time of use. You can't just have a contract that waives all re use for the future. You literally have to come back and say, oh, I want to use this scene from the latest Disney movie. I want to use it in the next sequel. You have to get re use permission.

And then there's the other aspect to this that has to do with copyright enforcement. So, for instance, if somebody uses your clip and they didn't get permission, the producer has to show us, they usually write me a letter that says, you know, here's the fair use reasons that the other side is going have.

This is really important under like the Lenz decision where you have to sort of show that you thought about fair use before you start enforcing things. And then if they don't have a good fair use argument for the person who did it without authorization, that the producer has to enforce their rights as a copyright owner.

In this last year, in 2019, we actually entered into our first ever Netflix agreement. So not only are they bound by our larger TV theatrical agreement, they also came to the table and we negotiated a special agreement. And inside of that -- people don't really know this, but background actors are covered by SAG AFTRA, but they're only covered a certain amount.

And they have to, like let's say you're in L.A., you have to have, I think it's like 35 background performers who are covered union performers.

And the rule basically says that you can't use like digital doubling of background actors, which is pretty common in like fight scenes and stuff like that, as part of that count.

Oh, there's -- sorry, there's one more rule that is actually just the rights of the depicted person, which is, as you can imagine, where I see this voice cloning really being used in the future is going to be dubbing.

All over the world, movies are dubbed for those markets. As you could imagine, it's going to be a huge market incentive for like, let's say, you know, what's the last movie you think, one of the last movies you made?

MR. SLOTIN: Us?

MS. HOWES: Yes.

MR. SLOTIN: 1917.

MS. HOWES: 1917, right, so it's in

English, I would imagine. Yes. So you probably want to have it be in Chinese or whatever it is. And this technology is going to allow you to have -- what's one of the actors in 1917?

MR. SLOTIN: That I don't know.

MS. HOWES: Oh. To be fair, I should have known that. Okay, so you want to have him, with his voice and his resonance and his kind of performance, have him doing it in Chinese or Japanese because obviously foreign markets, I believe they're important to you, right? Yes.

Okay.

And so, you know, this is going to be really kind of harmful to the dubbing community, right, because they make a lot of money doing this. But then on top of that, we'd have in our contract, if somebody -- say that you're Antonio Banderas and you speak Spanish, you have a first right to do the dubbing yourself. And that's in the contract.

So the reason they -- you're thinking, oh, there are some union rules, you know. So can't you handle all of these digital replica issues through union contracts, it's really important to talk about sort of why that's a problematic way of thinking about it.

Sometimes people are like, well, can't this just be handled by union contracts? So the way that -- and I'll try to be quick here -- I apologize -- the way that this works is that we are the exclusive bargaining partner for actors performing in audio visual works.

That does not mean that every audio visual work is union covered, right. The reason that we're able to hook NBC and all these other companies to doing union projects and coming to the table is I imagine you want to hire our stars and our members, right.

So in order to hire, they are bound, the stars and the actors are all bound to Global Rule One, which means they cannot do any type of cover jurisdiction or they're violating their role as a member.

And then on top of that, like I said,

we are exclusive bargaining agent for performing and labor. We're a labor union, right. So if someone is doing licensing of athletes and video games, that's an image licensing deal, right? That's not a labor issue.

And so that's the reason that it's really important to have external rules. Okay, great. So right of publicity here, we went into that and I'm running out of time. I didn't make a couple of notes just to respond to a couple of differences that we have.

You know, the way that read the Sarver decision about California is that I, a hundred percent, agree that biopics should be exempted in statute and I, a hundred percent, believe that if you were doing a biopic and having somebody depicted in their real life, that that should be bound by strict scrutiny.

I read the Sarver decision, and I'm not -- we're going to go back and forth on this. I just encourage you all to read it, which would be a great thing. I read it is that California law still has the transformative test for that, for the right of publicity, generally.

And in the case itself, it actually goes through this long laundry list of like we're not talking about greeting cards or advertisements or merchandise or performance theft, is one of the examples that they use that in the traditional right of publicity framework.

And then it says we're talking about a biopic and right of publicity doesn't apply to biopics. And if it did apply to biopics, they would have to survive the strict scrutiny, which I think is a very big distinction. Yes. But there are states that have strict scrutiny as the standard, notably Florida and Nevada.

Okay, so, Gwen Stefani, this was a case out of California Court of Appeals. Oh, okay, and so this is a case that talked about basically like inside of the right of publicity, while citing Zacchini, said that basically a person, performance theft, if you are taking someone, a literal re creation of somebody doing the activity for which they're known, that that is a right of publicity violation unless you have some sort of -- you've transformed it under the transformative test to the point that it's not just like stealing the economic value of that digital recreation.

Let's see here. So this is an example of something that was disputed but it never resulted in a lawsuit, and they ended up just changing it, but an example of how this can happen.

So, Ellen Page is one of our members and she was doing a video game and promoting it where she did her performance capture and all these different work. It was a covered contract for the video game called Beyond Two Souls. And right before Beyond Two Souls was supposed to release, The Last of Us came out. And a whole bunch of people were like, wow. Ellen Page is in The Last of Us. She was not. So the actor that was hired is Ashley Johnson, who I depict on the left. As you can tell, they look nothing alike, as is being depicted in that little bucket that says The Last of Us original.

So Ellen Page, on a Reddit thread, someone said, oh, man, isn't it an honor to be depicted in this video game? And she did not think that was an honor. She thought she was being ripped off and she didn't appreciate it because she had another game coming out and she -- that was a lot of money to have -- she was basically competing with the digital version of herself.

Settlements were made or whatever. All of a sudden, it was adjusted a little bit. But if you were on the video game review, YouTube channels or whatever, people playing video games, the joke is that it's Juno.

Okay, so now I'm going to talk a little bit about a more, a very serious issue that's come up. And so everyone's probably heard of deepfake pornography.   We were just talking about in the space of creating like, you know, Nicolas Cage memes and sort of putting, you know, people into Hans Solo movies.

This is much more serious, right? It's not a laughing matter. This is a form of image based sexual abuse. There's been multiple studies being done. It's actually almost impossible to find one of these deepfake porn videos, which is using still images to put people into very graphic pornography, as performing those pornographic works.

99 percent of this has been of women, so it's very gendered. And, you know, from our members' perspective, this is both a form of basically revenge porn and sexual abuse. It's very traumatizing. They are not happy about it. There are entire websites dedicated to fake porn that have a thousand profiles of our members of different members. It's being monetized on the Internet.

So my members view it as both a form of sexual abuse and a gigantic form of commercial exploitation. To give you an idea of how much money the porn industry makes, the worldwide revenue of porn is $96 billion dollars a year. To put that in perspective, and the people from ESA can correct me, but the worldwide video game market is $81 billion dollars a year.

And as you could imagine, in the same way that people want to go out and see Scarlett Johansson's movies because they're a fan, people are going out to watch Scarlett Johansson porn because they're a fan. And so this is massive exploitation. It is very much predicted that these sections of these websites are going to be the most popular for these companies. And these are major websites that are hosting deepfake porn. And some of them are advertising it. But they're careful not to actually use the people's names or their images so that they don't have any liability under Section 230 of the Communications Decency Act.

Okay, finally, I'm going to end on an even less wonderful topic, which is the reason that SAG AFTA is also so committed to having laws about the digital manipulation for sexual content is because we are in the process of really deep diving into what is a systemic, horrible practice in the industry of nudity violations in this industry.

So these are allegations. Like everything in Me Too, unfortunately, as everyone knows in this room, there aren't a lot of prosecutions and there aren't a lot of instances of liability or judgments or people bringing cases because of things like blacklisting and silencing and just fear of people coming forward with their stories.

So these two women came forward with their stories. These are allegations. We have -- I cannot say it happened, but if it happened, it's pretty horrific. Emilia Clarke was on HBO, HBO's Game of Thrones, and she talked about how she was basically pressured her entire time to perform in the nude and to engage in sex acts. And Ruth Wilson recently came out that the reason she left The Affair was because of how much pressure she got to do those scenes and how some of them would be like, oh, you're going to do like a basic sex scene and it would turn into like basically a form of rape.

So this is why, you know, doubling falls into this and why deepfakes is very scary to us. So just to give you an idea of how sort of doubling violations happen, you're sitting there and you're trying to pressure the actor to do this nude scene or the sex scene, and they just say no.

So we got this one call from this woman. It was actually a pretty large production. And she was like, can you guys just like, tell them to stop bugging me? I don't want to do a nude scene. And then a few months goes by and all of a sudden, oh, my God, they're using a body double. Please help me stop them, right?

The response is actually pretty similar to when you ask people who are making deepfake porn. Hey, why did you do this? Why did you use a body double to depict this person naked? And they say, well, it's not her real body. I don't get what the big deal is.

These are very, very harmful. If you a performer who decides to do a nude scene or a sex scene, it's going to be harder for you to do commercials. It's going to be harder for you to get roles in kids' movies. It's going to live with you for the rest of the life, your life. It's going to show up on Mr. Skin, porn websites, things like that.

So when it comes to deepfakes, as you could imagine, we're not very happy about this technology, being able to give independent film producers work, which is a lot of where this abuse happens, a really cheap tool to just exploit a bunch of performers. And so that's why we worked very, very hard. California now has a new law, 1708.86, which gives victims of any type of either deepfake porn or manipulated content of manipulated performances, a civil cause of action.

When it comes to copyright, because we're here at a copyright forum, I do want to point out that I consider this to be probably the best example of a moral rights violation of a performance. If you were to manipulate it into being sexual, the right of integrity under the Berne Convention is making something that causes massive reputational harm. And I feel that this law falls into that compliance with the U.S.

And then that's it.

MS. ROWLAND: Thank you, Sarah. We are running quite late. No, no, it was great. It was very interesting. I think that instead of me asking questions, because we are running late, I would like to give the opportunity for anyone out there to ask a question.

So before we wrap up, does -- oh, look, Ben. Hello, Ben.

PARTICIPANT: I just wanted to add, like movies, video games are expressive works and will be protected by the First Amendment. If you want to know more about our views on the right of publicity and digital avatars, feel free to come to my video game law class at Georgetown next

MS. ROWLAND: No, no. Anybody else? Okay, well, with that, thank you guys so much. That was really educational and great.

And I think we are supposed to -- we are just going to wrap this up, so we're not going to give any closing remarks except to say thank you all so much for being here.

Thank you to WIPO for partnering with us on this. And we will continue to work with them and keep you up to date about what's going on. If you don't subscribe to our NewsNet, you should because then you'll learn about all of these events in the future. And with that, I hope you have a great evening. Thank you for coming.

(Whereupon, the above entitled matter went off the record at 5:01 p.m.)

A

E

J

C E R T I F I C A T E

This is to certify that the foregoing transcript

In the matter of: Copyright in the Age of AI

Before: US LOC

Date: 02-05-20

Place: Washington, DC

was duly recorded and accurately transcribed under my direction; further, that said transcript is a true and accurate record of the proceedings.

Neal R Gross

Court Reporter