Pina D'Agostino: If I could now call up everyone for panel two. I'm delighted to welcome Professor David Vaver, who's actually flown in just the other day for this. We're delighted that you could be with us. As many of you might know, Professor Vaver had a very big year last year. He was inducted into the Order of Canada and he also became a Fellow of the Royal Society of Canada. It's a distinct honor to have him with us today. So Professor Vaver.

David Vaver: Good morning. I won't spend a lot of time either introducing the panel or introducing the session. You'll see that there's one person not here physically. That's Maya Medeiros, who will be with us by telecommunications.

Maya Medeiros: This is Maya Medeiros. I'm just waiting. I guess I'll get a signal when I'm on. I don't know what order I'm actually on either.

David Vaver: Right. Terrific. Thanks. Oh, good. I can see you there. Terrific. Let me just say something quickly about the panel, Intellectual Property at a Crossroad, which should itself is a very suggestive title. Which is, I think, conjures up the idea of autonomous vehicle where approaching a T-junction or is it a T-junction or is it like a Piccadilly circus or a whole bunch of roads going off it and entering it a with a driver? Perhaps at what once was a wheel, but possibly no longer or maybe conscious, maybe having a snooze while the vehicle decides which road it's going to go on. Perhaps it does a little circuit of Piccadilly circus, a little bit and then goes off on one road. Can it come back from that road or is it one which cannot come back on?

I think the title is suggestive and we have an excellent group of people here to talk about it. The order we're going to do this in, is we'll have, Shlomit Yanisky-Ravid speak first, followed by Dave Greed, followed by Catherine Lacavera and Maya Medeiros will come in by telecommunication, visually and orally as well. Terrific. All right. Without any more ado, I'll just introduce you to Shlomit Yanisky-Ravid, who's a Professor at ONO Academic Law School, which I don't think has any connection with the late John Lennon's wife and also connected with the Fordham Law School as well. She's been one of the pioneers in academic writing in AI and we're very fortunate to have her. I will introduce her and ask her to give some comments. I'd like to keep everything down to 15 minutes if I may. I'm going to put on my little whatsit here, which will go for 14 minutes, so you will have a one minute wrap up.

Shlomit Y.: Okay. Thank you. I'm really happy to be here. Thank you [Pina] for arranging this perfect, organized and like you spoke about the [inaudible] and children, so I was really amazed to hear, that you have a triple antidote and you do all this work so wonderfully. I want to thank Michelle for working so hard and I know, because we were in touch for many months. Don't count this time please. I want to thank Dave as well. I know him from Israel as an outstanding student and I'm happy to see the progress he has made recently. The title is amazing. I think it's perfect title, Bracing the Impact. I think one of my inside state I'm trying to embrace, is the best solution I think should come from a combination between technology and legal scholars.

What I did when I saw it ... I'm focusing on AI and IP, intellectual property and

mainly copyright, but also patents. I'm asking the question of the ownership and accountability of dissertation where AI System, produce works of art or patents and inventions. I did it ... When I saw it inquiring that, so I was going after the programming itself, because I have a scientific background and created my own labs. I can truly understand how it works and start writing about that. I think combination. What happened today in this event is really something we should all follow, if we want to come to solution anyways. As I said, I'm focusing on AI and IP and I think AI and IP is really interesting for three reasons. There are more, but I want to just start with these three reasons.

One is that all AI systems that we can find in cars and medical devices and others, they are actually copyrightable. The second comment I would say about AI and IP, that I think AI and IP the combination of both of them, demonstrate the best example of how AI can be human-like and how creative the systems can be and we'll soon show some examples. The third and the last is one of the solutions that I propose, came out from the copyright regime. Alright, I'll better start.

Do I have something to move the slides?

David Vaver: Push very hard.

Shlomit Y.: All right. I want to start with this painting which was sold in an auction in New York few months ago for almost, it was $432000. It was made by AI as were already said. I think that is really important. They didn't use a lot of data, so I'm not sure the outcome is that. Available, but I think it's really important to see that AI systems are really here. I wrote a lot. You can [inaudible] and see my articles also about the data and transparency, but mainly I was focusing on AI and IP and I will speak here about three of my article; Generating Rembrandt and When AI Produce Patents and AI in Copyrightable Issues. All right. I'll start with the first point. My first point is trying to explain after having a lot of discussion with students and scholars that AI systems are already here and they're productive and they're creative and sometimes there are so creative that we cannot distinguish between a painting or works of art. It was made by AI and the other one.

I'm sure you all know Turing, from The Imitation Game but what we are going to do here is a short quiz. I know it's not the usual way of conducting or speaking in a conference, but I want you all to take part in this quiz. Alan Turing in 1950 tried to answer the question, "Are AI or machine different from human?" And they said, "Let's think, let's try. If we go by what we call Chinese room." Yeah, it was titled after a few years. But I mean, can you distinguish, can you differentiate between what's made by AI and a machine? If not so, the machine won in a way and actually passed the exam.

Okay. Let's see. Let's try to listen to, two back melodies just shortly. One was ... We'll have several questions like that and several examples. One was made by AI and the other one was made by human, and I'm asking you, which one was made by human? The other alternative that is possible in this quiz, in this Turing test is that both of the works were produced by AI. But I mean, there is no way that both

of them were produced by human. Okay. Let's start with A. Bach A. Thanks.

Okay. I think you got it. We don't have those time to enjoy Bach. Then Bach B please.

Who thinks? I want everyone to raise their hands that Bach A was made by human. All right. And Bach B? Human? All right. The answer is actually that Bach B is human and Bach A was made by Artificial Intelligence. Then we are going to see some painting. This is painting photos I took from an exhibition I teach at Fordham Law School about, [inaudible] the challenges of advanced technology, AI and Blockchain. We went to see an exhibition down in Chelsea where all the painting were produced by AI, and not just by AI, but also in one Ai system was producing the painting. The other one was titling them. Because when you produce works of art with an AI system, you get millions of results.

It takes a lot of time to go after that so in other AI just choose. Okay. Who thinks ... and I took the same topic. Who thinks that A1 is made by human? No one. All right. Who thinks A2 is made by human? All right. The answer is A1 was made by human and A2 was in that exhibition. Now we'll take comment. B1 made by human? Hesitating. Even the hesitating itself. And B2 human?

All right. The answer is B2 was made up by Artificial Intelligence and B1 by human. Okay. C and D. Who thinks C was made by human? All right. And D? Is it human? The answer is both were made by AI systems at the same exhibition. Last but not least, that's jazz music we produce in our lab just to understand process before jumping in to the legal questions. If you can feel these play. Jazz A.

All right. That one is just less than a minute. All right. That's the last. Questions for quiz. Who thinks A was produced by human? It's no one. Who thinks B was produced by human? All right. Both of them were produced by AI system in our lab. [inaudible]. Yeah. Part of the things I'm doing is being the head of the AI IP project. All right. In one of my articles, I was inspired by this project of Generating Rembrandt, was it German? Yeah. A German university was producing this painting that you see with the frame after exposing their system for a lot of data from Rembrandt paintings that exists. You can see it in Amsterdam still. Anyway, so that's point number one, it exists. It's creative. It's very difficult to distinguish between a real AI system and human. My second question would be are AI human-like?

Because there was a wonderful explanation, just in the first panel about how it works. I speak that part learning by examples, learning by finding patterns and similarities. But it's very important for me to emphasize that AI system, neither copy or infringe any other music or photo or painting that we are exposing it to. It just finds patterns similarities and create new works of art.

In one of my articles, I try to point out like 10 features that make, under my judgment, AI System, human-like. There are those in a scale. But I think the more features you can find in AI systems, the more AI like them. I think AI system can be

creative, autonomous, unpredictable, goal-oriented, rational, evolving and that's something, that their programming is never ending process. Data collecting has free choice, can communicate with the Internet and get this data even without our permission. Even without the engineering knowing, actually that this happens.

All right. Let's go on to point number three. Point number three, I was asking myself, this introduction that we just heard is perfect for this point. I'm asking myself, "Who can be, when it comes to questions of honorbility and accountability, who can it be?" I don't know if you heard that the first case about AI ownership just came to court, but it wasn't decided yet. It's the Disney court about the Beauty and the Beast. An AI system produced the outcome of the mass that was used in one of the film of Disney and others film as well. It was the mask of the beast. The first film who really owns the software itself, claimed that the revenues and the income from the movie belongs to them.

The title provider was actually a Chinese film and the user is Disney. Who's going to win? Who should win? As scholars, we are more interested in this question, Who should win? Many scholars are giving this answer, that the AI systems are so creative and so human-like. They should be the owner of their own products. If you take it to other fields, it's really, I would say even bizarre or really something that cannot be right away taken and being implemented into the industry. The other more, I think, proper way of thinking, which is a bit wrong, but we'll try it without that, is looking for the man behind the machine. That's the chess player where people told hundreds years ago that it plays itself but actually there was someone hiding under the table and move the players, that looks like they moved themselves.

I think this way of thinking, we're looking after the man behind the machine, that's something which is very important. I'm taking like two more minutes. All right. Thank you. All right. Who can it be? Who are the candidates? I titled this model, the Multiplayer model, and these are all the players that take place in this process, when AI produce works of art or doing other things with autonomous cars and medical devices. Most of the countries, of the legislators are attracted to see the software programmer as the one to be entitled. I will think that the right solution, and I'll briefly explain why, it can be the trainer but the trainer has nothing to do with copyright ownership or patent right?

It can be the manufacturer, the user, the operator. There are many cards in this process and we have to consider, we have to understand that this process is never ending because in each and every piece of new data, all the formula of the AI system change. I'll just briefly mention why I don't think the software programmer like most of the legislators in the world, they might be wrong. That's why Bracing the Impact is really an important conference and important question. If you understand that AI system is actually human-like, going to the software programmer, I don't think that's solution, it can be a solution in part of the cases, but not in all.

The other thing is saying that the software programmer is the owner or the

accountable. It's like saying that the one who created the camera should be the owner of the fellow that was taken by this camera or the one who produced the computer, invented the computer, is the author and accountable for the songs or lyrics that were produced by using this computer. If you want to read more, you can read Generating Rembrandt. The other thing that was mentioned is the black box. I'm working with a lot of startups, and I see myself even in this jazz and music or other stories that we produce, we don't know how the system works. We have no idea and many of the startups programmers I speak with, they said, "We created the medical device but don't our prediction of heart disease or some other disease, but we really don't know how it works."

They have some other argument. When we go to this question, if it's not their software programmer, so who can it be from all these multiplayer model? I would like to give the answer that we have today and just ... [inaudible] sorry for taking the extra time with my proposal. The solution that was made in the US after this Monkey Selfie, a monkey that took a selfie of himself where David Slater, the photographer was taking photos in some resort and the US court said ... because first it was just an interesting story, but then an NGO for protection of animal rights rotate into court and US court says, and now it's adopted by the [inaudible] office as a regulation or guideline, that copyright exist only when human are involved.

If it's a monkey or a machine, there was no copyright at all. But if you take it to the Disney case and saying there is no copyright at all, I don't think that's the right solution. Some countries adapted this solution like US here, UK but many others. They said that the one who arranged the necessary arrangement for the creation should be the one. You see again, looking for the men behind the machine and maybe miss the old thing about the machine being creative. I don't think that should be the solution. Some propose the state and it might be a solution for some kind of system like autonomous weapons, but not at all. Many goes after insurance, but I think even insurance which covered damages, this solution doesn't answer all these question on who is accountable? Who should pay the price? How we would prevent from other risk?

With that really brief summary of four years of working and it will all appear in the book, which I'm working on one of the final draft, I would like to say it's in publish very soon, is what I'm proposing. My suggestion is taking the AI meant for higher doctrine and seeing the AI system as our agent. With that, implement the accountability and ownership on the user unless otherwise proven. I think this ... I'll just firstly say two or three comments about that on my argument. I think it better reflects the idea, the understanding that AI systems are really creative. The second thing, I think it works more efficiently with finding the man behind the machine.

I think it also unveil the people who are really using the machine to make them more aware. It doesn't mean that's the software programmer being thrown out of the game, but if there are so many, the multiplayer model would bring us to no solution. The same thing that happened with the Tesla accident. All right. With that, I will end and if I'm trying to think about the best solutions, sometimes I wish I could have an AI system for finding this solution or for accountability and

ownership to ask them, what would be the best solution?

Thank you very much. I'm really happy to be here.

David Vaver: Thank you. That's all right. That's great. I know most professors can't stop talking with shorter than a 50 minute period. Second up is, Dave Green, who's the Assistant General Counsel in Canada, of a small startup based in Seattle called Microsoft. Very much look forward to hearing what you have to say Dave.

Dave Green: Thank you [inaudible].

You'll put up slides. While they're putting up the slides, let me just say, this is my first visit to Toronto. When I go visit a city for the first time, I don't take taxi cabs. I typically will walk the city. I'll take public transportation just to really feel the pulse or the energy of the city. Toronto is an amazing city. It has just the amount of energy and authenticity and it's just a treat to walk around. I'm only sorry that I didn't add a couple of days to explore but I'm definitely coming back.

The prior speaker I think, did a wonderful job of talking about the, what if? Right? The capabilities of artificial intelligence. I'll give a couple of similar examples. It's fun when you see someone showing some slides and you've got similar slides and sort of like going into a ball and you're wearing the same outfit. It's a little embarrassing. I'll take a different tack. The tack I'm going to take now is not so much talk about the what if, which I think was an excellent presentation about what the law could look like, but really just take a summary of what's the law today? How do intellectual property laws address and how well equipped are they to address the current capabilities of AI and the current promise of AI.

There was an excellent slide in the first presentation that talked about AI. From Microsoft's perspective, this is how we think about AI. You've got a set of technologies here on the right hand side, that have existed or that are being developed. They're applied through machine learning and through a set of deep neural network learning and other types of intelligence learning to produce or augment or simulate those things that we typically associate with human capability; vision, speech, language, knowledge.

How should we think about AI and IP? How well equipped are contemporary laws to deal with the challenges that I think were well addressed? First is, I think we have to start from a policy perspective and really fundamentally from a copyright perspective. The first question I think we have to ask ourselves is, what behaviors are we trying to incentivize with the legal structures that we have in place? Are those behaviors adequately protected and adequately addressed by current legal structure? Or do we need regulation or changes or alternatives to continue to incentivize those behaviors? Assuming we deem that those are the right behaviors that we want to incentivize.

Obviously from a copyright perspective is, I think everyone knows. We're trying to incentivize the creation of works of authorship. The traditional way of doing so is,

to preserve to the copyright owner that exclusive right of copying, subject to important exceptions. The one exception, I'm not going to address in this panel, I'll address it in the next panel, is the use of copyrighted works to train and do machine learning. I'll address that separately because that's its own topic. I really want to focus on authorship.

I think the examples that we've seen really demonstrate that AI is here and AI, some state, is capable of producing useful works and certainly expressive works that we can identify as being equivalent to those produced by humans. I think the Harry Potter example is a really interesting example. This is an example of a laboratory that utilized from basic machine learning and text recognition and then trained that model on a specific subset of Harry Potter novels. The task was to see if the AI could then write its own chapter, in a vein and with the capability that one would associate with, as being a Harry Potter work. If you've read the actual chapter, it's interesting. I think the title was Harry Potter and the Portrait of what looked to be a large pile of ash.

Okay. Expressive, non expressive. The Harry Potter example is interesting to me because it really goes into the examples that a prior speaker had spoken about, which is, at the very core it's about machine learning. Right? It's about training a machine to, number one, be able to recognize words and place those words and a structure that resembles a sentence. Then from a deep neural network perspective, that was about taking that capability, that developed model and capability and educating that model to see if it could produce its own expressive work without being directed either by data or directed by a human as to what output it was going to produce. Again, I'll talk about the machine learning model in the next panel, but I want to talk here about the expressive work and to see how well adapted copyright is to protecting that work and whether it should.

We've already seen the example of the work that was produced, the Rembrandt style work, that was produced by artificial intelligence. I think the important thing here is normally people care about works that are produced in the application of copyright law because they generate value. What I would tell you is I think that's the wrong attribute. I think the fact that this painting, sold for 45 times what its initial estimate was and produced a pretty significant valuation is somewhat divorced from the policy initiatives behind copywriting and behind incentivizing the behaviors that copyright incentivizes. Whether or not they produce value as an outcome, not necessarily a policy objective. At least that's from our perspective.

The challenge with existing copyright law obviously now is that, it's really developed around this concept of personhood and the United States Copyright Act. In the constitution it's very clear that there's a requirement that a work be produced by humans. I guess the question that I'd ask is, "Does it matter and why does it matter?"

To answer that question, I think that the prior presentation was a wonderful job of, if it didn't matter, how could the lobby adapted and what are the difficulties and complexities of addressing non-human authorship? But I think that's a fundamental

question that we have from a policy perspective is, do we want to insist upon the requirement of humans and of person's hoods? If we don't, if we capitulate and ask that the law protect particular works that are not developed by humans, what does that say about downstream issues that's typically associated with copyright; issues of intent, issues of infringement, of reproduction, and certainly liability associated with all of those acts.

Well, I'm going to use an example. The example I'm going to use here ... By trade, prior to Microsoft, I've been at Microsoft for over 10 years. But I was with a small Bill Gates startup. There actually is such a thing or was such a thing called Corbis, which was a digital media company. It was a collection of one of the world's best collections of photographs. To use this example, I want to use it because this example to point out the issue of, if we are going to apply copyright, whether we apply copyright to the human attributions, the human contributions to artificial intelligence and works created by artificial intelligence, or whether we apply copyright and copyright protection to those same aspects as contributed by an algorithm. When does copyright attach? It's a critical point, right? Because there are a lot of steps in the chain of producing a creative work. Not all of those steps are necessarily protected by copyright. When they are, typically the law's required that there be a human being associated with those creative contributions.

If you're a photographer and you're creating an image, you obviously compose that image. You make a set of selective creative decisions about which lens to use, how to set that lens. What is your composition going to be? You might even dress up the composition. Even when you take that particular photograph historically, that wasn't the end of it. Obviously there was a lot of work in the processing and in the old days you used to take a photograph and subject it to chemicals and light and make a number of decisions about how to render that particular negative, for example. Ansel Adams, quite a very famous photographer used to compare the negative to the composition. He used to say that the print, the process in the dark room, was the performance of a particular work, both of which could potentially be protected by copyright.

Again, the law as it's equipped today, will recognize some or all of those creative contributions to the extent that they combine to produce an expressive work. Right? The law currently looks at those human contributions. What decisions that are made by the human being to produce that particular work? What happens when some of those contributions are initiated by a nonhuman, as was the case in the Naruto case? Now that case wasn't necessarily fundamentally about copyright. It was about many other things, but certainly copyright underpinned that particular litigation.

In this instance, I think none of us would argue that, that's a highly expressive work. Right? The composition of that work, all of the selection and arrangement, if the shredder had been pressed by a human being, we wouldn't be arguing about copyrightability. But the fact that a nonhuman did that, we suddenly are thrown in a state of chaos as to whether or not, that work should be protected by copyright, whether it generates value. Again, go back to the policy incentives behind

copyright. What behaviors are we trying to incentivize? If we do award copyright protection, at what point does the human contribution cease to be protected and the work contributed by AI becomes unprotectable?

A lot of this is old wine in new bottles. We've dealt with this issue before when, back in the 1860s, when photography was first versioning, the supreme court had to deal with this exact issue. The claim at the time was that, photography is nothing more than mimeography, right? It's just simply capturing facts that explains nothing. It's blind to the spirit of the realm, I think was the claim. This is a very famous case involving a photographer who had done a portrait of Oscar Wilde and sold that portrait as a series of lithographs and obviously copyright wanted to incentivize the creation of very highly expressive works by giving to that author the right to control that expression contribution. It may simply be that the law is actually very adequately equipped. If we break down the steps that are associated with creating an art of work of artificial intelligence, we'll get into that in the next panel. We'll have an understanding that human beings actually do contribute a fair amount of expression.

It's a question about whether the law is equipped to recognize that expression. If the law is not equipped to recognize that expression, if those contributions are deemed sweat of the brow or otherwise non-expressive, then there really is a fundamental question about whether or not, copyright should step in and should be altered to protect the output of the expressive work. What about patents? Right? How am I doing on time, by the way?

David Vaver: About five minutes.

Dave Green: Good. Okay. I'll zip through on the patent side. But I think it's an important question because obviously, when you're dealing with artificial intelligence, it's not just simply about the creative contributions and to the extent that copyright laws are incapable of contributing, of protection for those expressive elements. There's always patent law that potentially can step in.

Well, when you look at artificial intelligence and you apply it to the steps in a patent law, I think what you conclude is, "Boy, there's something there perhaps, right?" It's highly technical and we have a wealth of patents that protect technical processes. Certainly AI is innovative, right? There's certainly novel and useful aspects of artificial intelligence. It's not just the underlying technology, but it's certainly combinations of that technology and applications of that technology to resolving challenges in the world in novel and useful ways. Perhaps from the face of it, patent law might be able to fill in the gaps where copyright laws and has decided that, because of human authorship, that it's not going to step in. Well, what's the challenge of course?

The first challenge I think is that a lot of AI inventions really stem from algorithmic processes and patenting algorithms obviously from the US perspective is not allowed. It may simply be that algorithms trained on particular data models and then retrained or applied to solve new issues may just simply not be obvious ...

They may be obvious. They may lack the novelty in order to obtain patent protection. Certainly the challenge from an ownership perspective is a lot of these processes take place and are increasingly taking place in cloud environments. There's a huge amount of difficulty in detecting that infringement and being able to then understand and apply and do the analysis necessary to determine whether or not, there's been patent infringement.

We've talked a lot about data and we certainly talk a lot about data in the prior panel. The interesting part, the innovative part of artificial intelligence does sit in the training data and I want to differentiate the training data from the raw data. Raw data itself is just simply that. It's raw data. It hasn't been formed, it hasn't been modeled, it hasn't been labeled. It hasn't been assembled in a particular way to produce a result. There's a fair amount of engineering and work that needs to take place on raw data in order to make that data useful for AI.

Certainly the AI engines and tools are capable in some degrees of patent protection. There's a fundamental question about whether the combination of those tools to produce a unique algorithm based upon trained data can rise to the level of novelty and originality necessary for either copyright protection or for patent protection. Obviously, different jurisdictions around the world will have different perspectives in applying patent law. Just like in copyright law, there's a set of fundamental questions about whether the law today is fully equipped to protect not just the aspects but the output that AI delivers.

But the funny thing is, that's not necessarily slow the industry down. I think if you look at these charts from YPO, you'll see that the amount of patenting activity, particularly in a number of different sectors, transportation sectors, a variety of other sectors is exploded over the last couple of years. So certainly, doesn't to look like from the volume of patents that are filed on this particular area that there's a particular challenge that's being understood. Obviously as those, claims go through patenting process and review and get rejected or get litigated, we'll have to see what challenges emerged or what issues emerged from that.

But I think from today's perspective, I think we end in a couple of different thoughts, which is clearly AI is just exploding. It's exploding not just because of the availability of data and the democratization of tools, it's just exploding more generally. It does remain to be seen whether and how courts can grapple with the concept of authorship and what elements ... There's a number of cases out there currently existing in copyright law that go back 15 years that look at factual components and recognize copyright ability because of the selection and arrangement and the judgment that was applied in determining that particular output and have given it a minimum level of copyright ability. Certainly from a patent protection, while there are certain limitations, it doesn't appear to have slowed down the patent activity.

I think what I'll leave you with is, we can imagine and I think prior speakers and future speakers will imagine the challenges and the concepts and how to shape, what would be if we were to expand the concepts and the scope of intellectual

property to protect AI derived, AI produced output. I think from our perspective, intellectual property is doing quite well. It's furthering the policy initiatives that it set out to do. It's certainly not slowing down. No one's slowed down. I think the amount of activity and the growth of activity suggests that at least for the time being, we're in good hands with our current intellectual property statutes and provisions. And with that I'll let the other speaker come in.

Thank you.

David Vaver: Thank you so much David. Our next speaker is Catherine Lacavera who is Vice President of Intellectual Property Litigation and Employment of another little startup firm called Google. Sounds she's in charge of a rather large department, but it sounds like a pretty easy job because when she has any sort of problem, all she has to do is Google it. But in event, I look very much forward to hearing her remarks on this.

Catherine L.: Thank you and thank you to the panelists. I'm not going to rehash the very helpful analysis that Dave did on patents and copyright. I was planning to actually go off the reservation a little bit and say that, while we're ... I think we are at a crossroads with AI, I don't really see a crossroads in IP. And in fact, I think as Dave alluded to, our existing patent and copyright systems are plenty robust to deal with the changes that we're seeing in the AI space. I think the much more interesting challenges we're seeing are on the regulatory and the social impact front in AI. Just to bring that home a little bit, I'll just touch on the patent space. Maybe we don't fully agree because, I think to the extent you think that the patent system is working well for software patents, I think there we would disagree.

I think at least from the view from where I sit, I've managed north of a thousand patent lawsuits almost always as the defendant. It's been a lot of costs. It's been a lot of burden. It's been a lot of upset for our engineers. It's not been a lot of upside, frankly, for the folks suing us, because I think our success rate is something like 99.9%. I query why they still do, but in any event, I would say when you hear from the innovators that the patent system is not working well for software, that's when you should listen. I think in the pharmaceutical space, the advocates of the patent system are the innovators, whereas and in fact, we in Microsoft are largely aligned on this. In fact, there was a time when we were suing each other over patents and not withstanding that, agreeing on the policy front and working together to evolve the patent systems of Europe and elsewhere, to be more effective for software patents.

I think everything that applies to software patents equally applies to AI. So to the extent there are things that need to be worked through on tightening up the patent system for software patents. I equally think that applies to AI space and so I don't think we could have a long debate and I would bend your ear on some of the challenges there, but I don't think they're unique to the AI space. On the copyright front, I think having seen a ... I've managed some of the largest copyright cases in the world along with the patent cases. If there was this real shift, paradigm shift in the application of copyright to the AI space, given all the AI work we're doing, I feel

like I would've seen it by now.

I do think there are some really interesting edge cases like the ones we saw, with the monkey, the Rembrandt. But I still think all the traditional principles of copyright of, "Is it a collage or is it a derivative work?" When you take only Rembrandts and combine all of them together into something that looks a heck of a lot like a Rembrandt, then I think you're tending towards the derivative work space. Whereas, if you take an amalgamation of wide range of different sets of data and produce a really transformative and unique output, then I think maybe you are looking more towards the space of novelty and copyrightability and you have something. But again, I still think, I'm not telling you anything new about copyright law that applies to AI.

Like I said, I think at least from where I sit, the more interesting questions are on the regulatory and social impact space. And if I could, I'd like to tell you I think, things that are exciting to me about ways that we're applying AI and then some of the challenges that we're facing and some of the ways that we're dealing with those. First and foremost, I think, obviously as alluded to and has been all morning, AI is having ... it's not new technology. And I would agree with Jonathan that certainly the AD and the availability of large data sets has made it have more impact. We're now able to take on some of these global challenges that we were not able to tackle before.

I would also argue that obviously compute power has made it possible to do some of the things that we weren't able to do before, even though the underlying algorithms existed. One big thing that we're tackling ... Of course we host the world's data. We host YouTube, for example, we host the world's videos and the Internet is this wonderful and sometimes challenging mirror of everything that's going on in the entire world. All the good things and on also the bad things. One of the biggest challenges, we're grappling with day in and day out on platforms like YouTube is, all the bad content that people are uploading, terrorist content, child abuse, all of these kinds of things. Artificial Intelligence is being deployed to great effect in content moderation and it has the dual advantage of doing human review of these kinds of things. First of all, it's not possible given the sheer volume of it, but AI is well suited to tailoring it. Also that the impact on the reviewers is quite serious having to review this kinds of content.

We're getting better and better at deploying AI algorithms to get ahead of that content and make sure that the Internet is a safer place. Translation is another area that I'm excited about that AI has had a huge, huge impact. We saw orders of magnitude improvement in Google translation by deploying AI in a way that we have this neural system that in effect looks at entire sentences, instead of just single word translation and we've gotten close to near human capability in translation and now we just launched this live transcribed, with a few helpful from my dad. If you know anyone hearing impaired, you can sit there with a live translation of a transcript or if you're watching TV and want the sound a little better, you can have a live translation and it's all based on Artificial Intelligence and translation capabilities.

Self driving cars I think was mentioned. Obviously there's AI is critical to that advancement. Generalized learning algorithms. You've probably heard about DeepMind beating the game of Go, but we've gone on to be tested in other games and with self training systems, machines are now training themselves to do the things rather than having human training. I could go through a huge long list of things. There's everything from diagnosing diabetes to earthquake aftershock monitoring to breast cancer detection. In the medical field, in the environmental impact field, like across the board and we're investing in these things, but we're also open sourcing our TensorFlow technology and seeing amazing applications. Third parties are taking up. We had some high school students develop technology for early detection of California fire outbreaks.

Just amazing things being done with technology. Having said that, as with all technology, when you deployed out into the world, again, there can be good uses of it. There can be abusive uses of it and we're grappling with this question of to what extent do we make these powerful tools available and powerful data sets. And we heard folks talk about the privacy challenges and you need the data sets in order to do the innovation, but to what extent can we make it available to third parties in order to operate on it. There's this grappling struggle between enabling the innovation versus all the privacy challenges and also the potential abuses of the technology. Along those lines, and I think in the regulatory mindset that being, I think where the bigger challenges are, we've released these AI principles.

I think Microsoft may have also done that. Some governments have done it. And I think that's where the real game is right now because, you don't want to overregulate in a space that is still so nascent and there are so many, really interesting applications still out there. But you also want to set up a framework where these abuses won't gain speed. I want to talk about, just a little bit about that, some of those principles that we're relying on, we've set up as council to review some of the applications. First and foremost, you're balancing to be socially beneficial in the applications that we will allow. We don't want to be overly restrictive in how people use open sourced AI technology. On the other hand, we have committed to not allowing its use in certain areas like weapons development or surveillance, these kinds of things you can see where, really powerful AI could become concerning.

Somebody talked this morning a little bit, I think it was Jonathan, about the unfair bias and how that is being built into these algorithms. I think what we're seeing is, frankly, this concept of data poverty where certain populations, they're just not online. And the result of that is that their data is not being collected, not being built into the algorithms and then some of these amazing innovations like in the medical field simply don't work for certain populations. Obtaining diverse datas, it's a challenge. First you have to have participation by those populations. One of the things we did, we actually released an open source technology called "what if tool" that allows people to detect or at least tried to detect some of these inherent biases in their models. But there's another question about how much should we be doing to proactively control the datasets.

I use this example because I think it's powerful for me of like image search. Well if you look for doctors, it depends on your dataset. If you're looking at 100 years, you're probably going to get only white men. If you're looking at 10 years it's going to be a very diverse population. Should we be at least be transparent about the data set that we're relying on when you conduct searches and also maybe putting into the user's hands more control around what data set you're relying on. Because frankly, we don't want to be arbitrarily choosing a dataset. Either answer would be wrong. It depends what your question is. Are you looking at doctors for the last 100 years or the last 10? There's a lot to be done there. Also obviously built for safety, critically important, accountable to people. I think somebody mentioned that.

The ability to explain what's going on behind the algorithm is a hugely challenging problem. It's not going to be a good enough answer for me to say, "Well, nobody knows why this content, it was left up, go ask the machine." You're right. The regulators are going to want a better answer than that. Getting really good about being able to triangulate how these things are working and transparent about it. Incorporating privacy design principles of into AI and obviously, we talked about notice and consent and transparency but also portability.

You don't want to have the lock in problem, of all your data is locked up in one organization and you can't port it and high standards of scientific excellence, getting input from the larger community on what we're doing and making sure that we're accountable to that more broadly. And finally, of course, being available, I.e. opensource, accessible for use in accordance with those principals and with the limitations I talked about, not allowing uses for weapons or surveillance, et cetera. That's the way we're thinking about AI space right now, is it's again broadly deployed. It's in all of our products, been there for years, nothing new there, but also nothing new about this incorporation of privacy thinking, accessibility, open sourcing and bringing the world along if you will, to get to a better place. Happy to take questions.

David Vaver: Thank you very much Catherine. That was terrific and within time. Could we have Maya put up onto the screen now? Maya Medeiros. Hello Maya. Good to see you. Maya is a partner at a firm called Norton Rose Fulbright's, which doesn't just deal with plant breeder rights and flower arrangements. I see from her CV that, she and I have something in common. She was on the International IP program at Oxford some years back. I'm so happy to see you again and I look forward to your remarks.

Maya Medeiros: Thank you David and I hope everybody can hear me okay. I'm sure somebody will yell if they cant and thank you to Pina and Michelle and Aviv and Ian for organizing and ... I do remember having a nice beer in one of those amazing pubs in Oxford with you as well, many years ago. And thank you to the other panelists for the great job they've done so far. I think I probably don't have to speak right now, but to avoid duplication, I'm going to take a slightly different approach to looking at Artificial Intelligence and intellectual property and I'm going to discuss the importance of collaboration for AI development and deployment and then the impact and role of IT in that.

Collaborative development of intellectual property assets can help maintain leadership positions for AI innovation and protect freedom to operate this transformative technology and avoid starting from scratch and helps to connect silos of technology, data and knowledge. I just read a recent article, that indicates the top countries are the US, Japan and Canada and they're continuing to invest ... sorry, and China and they're continuing to invest heavily in AI. I also say, look at Canada as well as South Korea and Singapore, Taiwan, the UK, there's lots of competition for this development and deployment of AI.

I think to remain in this group and remain competitive collaboration and coming together and bringing together assets is very important. Want to look at a few examples of different collaborative efforts. Of course we've discussed health already and of course we need patient information, patients need to be onboard, care providers, you need some treatment data, historical data. We also need to get the technology players onboard as well as researchers and all grew up together to look at different news cases under the big health umbrella and they were mentioned earlier as well.

Another interesting collaboration that I flagged is a block collaboration. Recently there was an announcement that Microsoft and Amazon are collaborating using Microsoft business and productivity focus, Cortana to engage with Amazon nor consumer E-commerce focus Alexa. I think this is very interesting to see how these bots will integrate and collaborate effectively to provide a better solution for consumers. Jeff Bezos was quoted saying, "There's going to be multiple successful intelligent agents and they're going to have different dataset and different specialized skill areas." Bringing them together enables them compliment and enables a richer experience for our customers so that I can talk to one bot and it can go off and engage with others and save me from some time and make me more productive. I think that integration of bots with natural language processing and understanding is going to be fantastic.

As well as looking at map systems, we're going to need data from other sources and just more generally, there's going to be different participants in the AI ecosystem from industry research, technology providers and traditional industry players. We need to work together to facilitate this development and this collaboration will enable the use of data from different and multiple sources and this may provide better and more diverse data set to mitigate against the bias that Catherine just discussed and I think Jonathan mentioned earlier as well. Very important to help develop these diverse data sets by bringing people together and companies together.

When I use algorithms from technology experts and use data, maybe from traditional or non technical industries and to enable bringing artificial intelligence solutions to those industries as well. We also need subject matter experts for training and refinements. There might be a general purpose classifier for example, that needs to be refined or tuned to a specific use case. We're going to need to engage those subject matter experts in that collaborative effort. And we're also

seeing already a ton of open projects. Catherine mentioned open sourcing, TensorFlow for example, for a new application development. But of course with some restriction around use case. So definitely seeing a lot of benefits from collaboration specific to this AI development and AI in different use cases.

How can IP facilitate these multiparty collaboration to protect the AI innovation as Dave indicated? How can the law incentivize collaborative behavior? We want to look at protecting or maintaining freedom to operate for example, for transformative technologies. I think the importance of IP rights and freedom to operate come to come together in order to provide these more defensive assets to block or prevent others from claiming the exclusive right to something that the company's working on or the collaboration project it's focused on and that enables control and even permission to use, but could further other collaborations that there's control of important IP assets by a group of people or group of players that want to then enable collaborations on those to collaboratively developed assets. It'll be a cycle that we want to encourage.

As well, we want to look at Defensive Licensing and those are more freedom to operate type licensing, where you have the royalty free perpetual abilities to use the technology. So you might not own it, but you might have this free use of the technology from that collaboration, which might be just as important to the goal you want anyways for being involved. In order to facilitate, with collaboration IP rights, we need to really define clearly the scope of these IP assets, to facilitate the sharing. Companies or individuals or groups, they'll feel like they're giving away things by coming together. It's a sharing in a very controlled manner to enable and facilitate that behavior. We clearly want to define these IP rights.

While there are challenges to protecting AI with the traditional IP framework, I think Dave and Catherine both did a great job saying, "I think there's still adequate protection for AI." At the current framework, they'll work and so let's look at that. And Dave did a really good job, I think reviewing the challenges of IP protection. I will touch on those briefly, but I think he covered off copyright and patent quite nicely. And we'll see these challenges maybe evolve and become more complex or maybe have some clarity, as there's an increase of litigation or perhaps legislative changes in this space. It's still a booming space and while it is old, we're not seeing a ton of litigation. Some of these issues that we'll maybe see some clarity from different jurisdictions on needs. When you're looking at AI innovation and looking at IP rights, I think it's important to break it down into the different components.

I think you don't want to just do a broad brush and say, "Oh, this is the AI tool and off we go." You really want to dive into the technical details of a collaborative project, looking at the hardware components, software components and that can include algorithms as well as data sets, as well as the interfaces for algorithms and data sets and also for end users. I think people assume that's a really simple solution or something simple to implement, but actually that can be a technical issue there. Just this interoperability or the broad interface. And then also looking at the classifiers of the rules, it felt that the algorithms are focused on, but they might vary for specific applications and use cases as well as network topology and

training data set. These are all the very important components of AI and different players, may be bringing different pieces together and we want to make sure we layer the intellectual property rights to cover these various components because there is no "one size fits all" solution.

As Dave mentioned, there are patents that play people are getting patents for AI innovation, covering technical processes and machines and improvement. But there is some challenge looking at patent eligibility when yes, as you mentioned, when you scrape it all the way, what's actually left and in working in this space. I think there's a lot less, particularly when you're looking at new use cases and really diving in deep into the technical effects of applying something more general, which may be a known process or a known algorithm at a high level to a specific use cases where the technical effects or small challenges that are overcome, that actually might be the huge advancement for that specific use case. Although it may seem narrow, it's still quite a broad form of protection.

I'm looking at tangible results, will help you distill the patent eligibility of components of AI. You also want to understand the patent landscape, who owns what? Figure out, maybe you're collaborating with parties that maybe don't own a lot of ... maybe don't seem the important players for a specific space. Having a pulse on the patent landscape might enable you to invite other parties to the table, that you might collaborate with based on this, at least based on their public patent filings, who knows what their actions building. But that could give a pretty strong indication of what they're working on. And of course keeping an eye on freedom to operate, to avoid walking into an area that's highly litigious and avoiding, that might actually cease the project from moving forward all together. Of course there are challenges with detectability as Dave mentioned as well.

Thank you Dave for doing a great job and particularly looking at cloud implementations. I'm thinking about how to claim the invention. A software code, of course is going to trigger copyright as well as the data and the application programming, the interfaces and different new user interface elements and graphics. Of course we're going to have the challenge with computer generated data and work. We already talked about examples really into music. I liked that demonstration earlier on as well as paintings. When we also have the issue of dynamically changing code, so is the programmer really the author, if the code is updating altogether and so these are issues that will be resolved hopefully or if not at least people will have some comfort in knowing that, that might not be resolved but go ahead anyways for these important cases. With copyright, at least from the Canadian lens, we're always looking at skill and judgment for the tests of originality.

When you're looking at data you need to consider, how much skill and judgment is used for this date of generation, particularly when machines are involved in generating that data. There are case law in Canada, at least around machine generated data and skilled individuals were involved in setting parameters for the data collection. There could be likely a copyright attaching to the date that was generated. Just keeping in mind and walking through, "Hey, is there human skill and judgment involved in this?" With that provide adequate protection for what

we're bringing to the table, to the collaborations that we make sure we're sharing without giving away.

Then of course stepping back and looking at that idea versus expression. Tension that we see in copyright law and copyright as it intersects with patent law as well to make sure that you're covering the process as well as the actual implementation and the code, and not enabling somebody to learn something and then go off and independently create without copying something and still maybe not short of a substantial taking. Looking at clearance issues because data mining, at least in Canada, it's unclear whether that is a fair dealing under copyright.

How was the system components actually generated in the first place? From the training process to make sure the other parties participating in the collaboration have clearance or comfort. That they can use this data set or these algorithms without worrying about painting the overall project IP with maybe some clearance issues early on. As well, we're going to have trade secret and confidential information as we talked about. The difficulty of keeping it a secret as well as balancing the need for transparency and other ethical issues that come up. I'll just touch on that briefly because I know that I'm a little bit short on time and our panel is a little bit short on time. Finally we want to look at designs. I think with industrial designs at least in Canada or design patents in the United States. We're seeing them more and more important.

I'm looking at those for hardware form factors, but more particularly in the user interface space. Seeing how those are being applied, particularly user facing or business facing interfaces and how those elements can be protected. There might be interesting visualizations that are generated by an AI tool, for example, enables to give more insight into raw data by an increased and improved visualization. So definitely some important innovation there as well. As well as minor trademark issues around youth and on the products when you're looking at software might not be as [inaudible].

Finally you always want to consider contract provision in an intersection between property and contract rights. While somethings might have comfort in the contract arresting it, I'd always want to lay our proprietary rights there to ensure I'm not just dealing with that one party that I'm contracting with and they can try protection more broadly than that. Contract provisions typically are important around licenses to use data and algorithms.

If one party's bringing a component to the table, that they've collaborated with other parties for, you want to make sure the licenses have extent to that as well and just checking that. Use different intellectual property rights, to fine and protect components of the collaboration. That would include the background intellectual property as well as any project generated intellectual property to facilitate the collaboration so everybody knows they can keep what they're bringing to the table and figure out some ownership around what they're generating as part of that collaboration. YOu want to make sure the treatment of background intellectual property and project generated IP, motivates participation by diverse stakeholders.

There's a proplayer coming to the table with a large amount of background IP.

The parties will need to give comfort to that player that we're not sucking that in and building and it's all together now. It's actually going to be, maybe some limited use rights to that background IP for the collaboration, but all improvements and what have you, will be owned by that initial party. As well as need to think about the project generated IP. Are they're going to be ongoing access rights by other parties to this. They want to just work together for a limited project and going to walk away. They're going to need ongoing access rights to IP. Looking at those issues from the contracting and the terms that the collaboration will be important to motivate participation.

You also want to look at the IP ownership models for AI innovation. We're already seeing some interesting things come up looking at artificial intelligence. These case example developed from collaborations or is there going to be joint ownership model around the IP for the AI innovation or is there going to be a single owner than granting licenses to other participants for ongoing access or improvement? Or is there going to be an IP holding company that owns the project IP? These are different considerations that will have to be worked through when looking at this artificial intelligence collaboration to enable success and hopefully an example or a project going forward to look at it and copy or at least take parts of that to encourage this collaborative development.

Multi-collaborations and final thoughts. These multi-party collaborations are important for the AI development and deployment across broad use cases and working with diverse stakeholders. I think IP rights can facilitate sharing as opposed to prohibiting sharing and still provide protection for the parties and their AI innovations and anything resulting from that. I'd like to see a positive light tall there as opposed to IP hindering these collaborations. I think I'm out of time anyways so I hope that added something to the discussion.

David Vaver: Thank you Maya. Did it on 15 minutes. Very good. Well, we have some time for questions. Questions from the floor? Well, people are getting themselves together. I just wanted to just perhaps say something about the international aspects of all of this, because a number of you sidled up to them, but we could develop them a little bit more. I was struck by your comment, Catherine, too about access. Very easy to forget that [inaudible] is wired for Wi-Fi and I was in conversation with a copyright scholar in another jurisdiction, who was doing a lot of work on copyright and educational publishing.

I said a lot of these stuff must be available online, and she said, "It may be in our capital city, in the countryside, there is no Wi-Fi, there are hardly libraries or hardly books." I think we have to keep that side of this question in mind as well. On the international side, question of human authorship. Yes, that's very important. Also the question of human inventorship on the patent side as well, and whether or not, the international conventions allow you to start providing for non-human ownership or protection, when the international conventions haven't been updated to cater for this bidding. You may have your local protection but if it's

students I think crosses the border. You're in a different situation, your rights may not be recognized. I wonder whether you have some comments on those matters.

Shlomit Y.: Yeah. I think international convention should intervene in this topic. Unfortunately, I'm visiting WIPO, two times a year and I'm working with them. I don't know if you know, Mitchel Wood, who is the head of the copyright division at WIPO (World Intellectual Property Organization). Until recently, they didn't want to move forward because that should come from the country. But they are ... I mean everyone is wireless, were as much as them that all of these AI issues and Internet are match beyond specific nations. But recently they establish a new division that works on AI. But only on for now on a research project. It might be in the future being implemented.

David Vaver: Yeah. Dave?

Dave Green: Maybe I've gone a little far overboard when I said most laws around the country insist upon a human component to authorship. That's not necessarily true in all contexts. If you think about corporations for example, corporations are a perfect example of a non-human entity that can be liable, can own copyright, can potentially under legal regimes, create and certainly own expressive works. Estonia has actually taken a really interesting perspective. They've actually thought of AI as a corporate entity and created a set of principles, rights and responsibilities that would attach in the same way that they attached it to traditional corporations.

But I think that's an interesting and novel approach. It carries with it, all of the complexity that we've heard on this panel. I do agree with you. I think if we are going to go the route of altering intellectual property laws, not only do we have to worry about international treaties, we have to think about more than just the concept of authorship and ownership. We have to think of all the downstream implications that are associated with them.

Catherine L.: Yeah, and I actually liked Dave's analysis of, again, what are we trying to incentivize with our IP system? What would be a reason for creating a right of ownership in the machine? At some point, somewhere down the stream, some human is involved in the creation of it, whether or not their contribution was creative enough to create a work of authorship or something that warrants copyright or patent or other protection. I think that's still the same question we ask in all contexts. I think going the additional step of saying, "Oh, we're going to create a new kind of IP rights for the machine." I think number one, I'd be asking, why are we doing that? Why are we detaching it from the programmers that generated the AI or the input to the owners of the input, if it's a derivative work or the owners of a unique output, if it's so creative that it warrants it. I would just ask that question.

I actually weirdly have a survey of the laws of the world on the ownership, at least in the copyright context. Most of them tie it to a human at some form or another. I would query why we would change that.

David Vaver: Yes. I was going to ask Maya, the issue of collaboration that may be distributed

internationally, how you handle that?

Maya Medeiros: Yeah. I think international collaborations we're already seeing that. Already irrelevant are the intersections of different laws and if there's parties from different countries, how you come together and at least in specific to the Canadian context, we do have this notion of maker, a reference in our act in addition to an author. We already see other types of creators, in participants come to play, at least in Canada. I think another specific example we're seeing being relevant today, at least when they're reviewing the Canadian copyright act, is this notion on data mining. I've been looking at international approaches to data mining and whether there should be a fair dealing exception or a fair use, to use the US term for data mining. Particular when you look at the issue of bias.

Maybe providing some exception or some limited use for data mining might enable more diverse dataset because if the company's acting being very conservative, they won't touch data unless they have a clear license to use that data. Some traditional open datasets have bias in them. I think it's Wikipedia that is authored by 97% males for example. So thinking about what the implication of that is, could the training system and the training process.

David Vaver: Thank you. How are we doing?

We're pretty ready to wrap up. Well, thank you. That for me, was extremely informative session. And I hope, and I'm sure that it was for you too. Thank you very much panelists.