

Emily C. & John E. Hansen Intellectual Property Institute

**TWENTY-SEVENTH ANNUAL CONFERENCE
INTERNATIONAL INTELLECTUAL PROPERTY
LAW & POLICY**

*Fordham University School of Law
Skadden Conference Center, Costantino C
150 West 62nd Street
New York, New York
Thursday, April 25, 2019 – 2:45PM*

SESSION 4: THREE CONCURRENT SESSIONS

4A. Artificial Intelligence (AI)

Moderator:

Laura Sheridan
Google, New York

Speakers:

Massimo Sterpi

Gianni, Origoni, Grippo, Cappelli & Partners, Rome
Neurography: The Rise of the Neural Artist

Carlo Scollo Lavizzari

Lenz Caemmerer, Basel

***Textbooks for AI and Clean Data to Train Machines? – How Machines Learn
and What It Means for Authors, Publishers, and Media Businesses***

Katharine Stephens

Bird & Bird LLP, London

***Patentability of Artificial Intelligence and Machine Learning:
The EPO Publishes Guidelines for Examination***

Panelists:

Shlomit Yanisky-Ravid

ONO Academic College, Israel

Mihály Ficsor

Hungarian Copyright Council, Budapest

Celine Castets-Renard

University of Toulouse Capitole, Toulouse

MS. SHERIDAN: Welcome, everyone. My name is Laura Sheridan. I'm from Google and I'm based in our New York office. I'm also a Fordham grad, and I've come to this conference a whole bunch, and honestly it's one of my favorites. The reason that I think it's so wonderful is that it brings a truly international perspective to very complex and very cutting-edge IP issues, and this panel is no exception to that.

We are going to be digging into IP issues around artificial intelligence (AI). We'll be looking at topics like copyright and patent protection for works that have been created with AI or AI-related technologies. We are also going to explore what goes into AI technologies and IP issues that arise around the data.

Our speakers and our panelists are experts in their fields from across the globe, and we are very lucky to have them all here. I'd like to have them briefly introduce themselves by name and affiliation.

MR. STERPI: Good afternoon. I'm Massimo Sterpi from Gianni, Origoni, Grippo, Cappelli & Partners in Rome, Italy.

MR. LAVIZZARI: Good afternoon. Very nice to be here at Fordham. My name is Carlo Scollo Lavizzari. I'm a copyright lawyer practicing in Switzerland. I'm also admitted in England and Wales and in South Africa.

MS. STEPHENS: Hello, I'm Katharine Stephens. I'm from Bird & Bird.

PROF. YANISKY-RAVID: Hello and good afternoon. I'm Professor Shlomit Yanisky-Ravid, originally from Israel. I have been a Visiting Professor at Fordham University School of Law since 2012, where I teach the course IP and the challenges of advanced technology, AI and blockchain. I'm also the head of the CLIP AI IP & Blockchain Project here at Fordham Law School. In addition, I am a Fellow at the Information Society Project at Yale Law School. I write intensively about IP and advanced technology in general, and more specifically about AI and IP. My book about AI and IP (AI systems that produce works of art) will hopefully be published soon.

MR. FICSOR: Mihály Ficsor. I am a member of the Hungarian Copyright Council, but here I think that I will speak more about my experience as former WIPO Assistant Director-General in charge of copyright, the way we dealt with this issue — artificial intelligence, computer-generated works, etc.

PROF. CASTETS-RENARD: Hi, everyone. Thank you for inviting me. I'm Celine Castets-Renard. I'm a professor at Toulouse Law School in France and I am a Visiting Professor here at Fordham Law School where I teach a class on EU information law and policies.

MS. SHERIDAN: Okay, thank you.

With that, let's get started. I'll turn to Massimo for our first presentation on the topic of neurography.

MR. STERPI: What is neurography? Neurography is actually one of the latest embodiments of computer-generated works. The word was created by Mario Klingemann, who is one of the most famous digital artists who started experimenting with computers in the 1990s. I have a couple of quotes from him:

"That's what a digital ultimately is: a huge number." It's a combination of huge numbers, of numbers coming from different sources. The exact definition that he gave of neurography is "neural photography" or "photography created by neural networks."

Aaron Hertzmann said neurography is "a photography of imaginary things," things that do not exist in reality but were created by the mind of a computer.

But is it really the mind of a computer without any human intervention that creates these images, or actually is the situation a bit more complex?

The timeline of images created by algorithm and neural networks goes back to the 1990s. So it's not the first time that we have this, but over the last five years, I would say beginning in 2014, there has been a revolution.

Why? Because rather than using one single computer program to mix up images or other graphic elements with some random addition of a computer program, in 2014 a computer science expert, Ian Goodfellow, created the first Generative Adversarial Network (GAN), which is a combination of two neural networks playing one against the other, one creating certain things and the other making a selection according to certain criteria. Why was that made? In order to improve the quality of the output, as the random output was of uneven quality, whereas selecting the output made the quality of the resulting works much higher.

Then, working on this concept of the adversarial network, we have an even more recent addition, the so-called Creative Adversarial Network (CAN), created by Ahmed Elgammal, another computer science expert who then became an artist.

Let's move into what neural networks are and how they work.

A neural network is not an algorithm. It is a combination of algorithms. It is a structure combining various electronic elements and algorithmic elements in order to process certain kinds of information. The result may be images, texts, or music.

An adversarial network is a combination of two different elements: (1) the generator; and (2) the discriminator, the element that selects and/or diversifies the final result. But let's concentrate on one point which is fundamental for our discussion today. There are many humans involved in this process.

- The first one creates the program or puts together the structure of the neural network.
- The neural network by itself cannot do anything; it must have content to process, and content must be fed by another human (or research according to criteria indicated by a human).
- When content is put in, there are two adversarial networks that fight, and actually they must be set appropriately, so there must be training of the two networks, and a human does that as well.
- Finally, when the images are created, they must be selected, and once the images are selected they must be printed out or titled, so there is another creative process done by humans.

All these processes, at least for the time being, are normally carried out by humans.

A Generative Adversarial Network (GAN) is a combination of a generator and a discriminator. The generator generates images by randomly mixing the input data. The discriminator selects the images that are most compliant with the pattern detected in the initial set of data; it tries to select those images that are more "average" with respect to the pattern detected, and thus more recognizable by the human eye. This process creates specific aesthetics which is recurrent. This has been called GANism.

This is how it works. The random noise is actually the material you put into the neural networks.

The generator mixes it up and creates an image.

There is a training set of images.

The discriminator then will select between what it calls "real" (acceptable) images and "fake" (not acceptable) images that the humans will not recognize as artistic imagery.

The Creative Adversarial Network (CAN) instead operates according to a different process. There is as well the generator, but then there is the deviator. This is a very interesting concept. Actually, it doesn't select in and out the images created by the generator but rather twists the images, trying to reproduce what happens with artistic styles; for example, where someone started creating a cubist image and then other artists started painting in a cubist manner by twisting the style with their own creativity. The CAN therefore simulates the evolution of an artistic style.

These examples are not only theoretical. This is not something that is only happening in laboratories, but actually these works are on the market and are in art galleries.

The painting "Portrait of Edmond de Belamy" by the French collective Obvious was sold at an auction last October in New York, and from an initial valuation of \$7,000–10,000 it fetched over \$400,000. By the way, I was bidding on that piece on behalf of a collector, but his maximum budget for this piece was lower and we did not get it.

The Creative Adversarial Network creations were exhibited in a Chelsea gallery in New York a few months ago. What is interesting is that they were presented as a collaboration between Mr. Ahmed Elgammal, who is the software programmer, and AICAN, which is an algorithm created by the same programmer.

There are, of course, very important copyright issues. If the works were created really without any human intervention, then there would be a difference in treatment between an Anglo-Saxon country like the United Kingdom (and others like India or Hong Kong), which have rules on computer-generated works that (fictionally) attribute the copyright to "the man behind the machine," the man who made the arrangement, and the countries that do not have specific legislation (such as the United States or the EU Member States other than the United Kingdom). In fact, for the time being at least, both in the United States — in the *Naruto v. Slater*¹ decision and in the opinion of most scholars — and in most of Europe, the lack of human intervention would prevent any copyright protection. All the litigations or disputes so far have not been about the presence or absence of copyright in artworks generated by or with neural networks but, rather, they were disputes among the various humans involved in the creation process who were claiming to be recognized as co-authors and/or be entitled to share the profits of the exploitation of the artworks so created. You will find these cases on the Conference website.

Therefore, for the time being, this intervention of the computer is not exclusive of the human intervention but, rather, it is a cooperation between humans and machines.

MS. SHERIDAN: Massimo, what do you propose with respect to a solution where there is heavy human intervention, it's not just the machine running autonomously? What are you suggesting in terms of copyright protection?

MR. STERPI: When there is an intervention of algorithms and humans are still involved in the process, it is very difficult to use the usual concepts that are typical of the copyright law, especially that of authorship. Authorship intended as the expression of the personality of the author is a concept that is not very appropriate to describe what is happening.

There is a contribution of multiple factors, some of which are technical. You can think of an algorithm as a digital brush or as a new kind of technique that you can use to create artworks or as an independent contributor to the final result of the creative activity.

As many people – physical and digital — are involved in a multiplayer model, a multiplayer solution is probably more appropriate for these issues in which whoever makes

¹ 888 F.3d 418 (9th Cir. 2018).

a contribution should also be somehow involved in the sharing of the profits. You will find more on that in my presentation on the Conference website.

There is a new website called Ganbreeder that permits everyone, any of us, to go and take two images and to mix them through Generative Adversarial Networks. The process is tracked, so every person who added or twisted an image is credited. The images on that website have a list of credits including any person who contributed to the creation and development of each image. The matter being still very fluid, for the time being this website doesn't have any legal notice or any legal indication, but still, in the long run, probably everyone involved in the process could be considered as a co-author or a co-owner of the copyrights.

Another model that has been proposed and discussed is that of collective works. When you create, for instance, a collection of essays by various authors and publish them in a book, normally the author of the book is the so-called "curator" of the book, the one who put everyone else together, taking the curatorial responsibility. Probably we are evolving from a romantic model of individual creation and authorship into a curatorial model where some humans will coordinate the work of other humans and machines and they will take the curatorial responsibility.

The third hypothesis we can work on is to create a sort of sui generis right to protect these kinds of creations, where we reward more the effort of the creation rather than the personality of the author's expression. We have a model that can be an interesting point to start with, which is the database protection that we have in the European Union for non-original databases in which what is rewarded is the effort to put together all the data that comprise the database.

MS. SHERIDAN: Thank you.

Are there reactions from the panelists to any of that?

MR. FICSOR: A lot. I'd like to mention that last year in December there was a European Union-China academic forum in Shanghai, and the same options emerged there.

But what is interesting is that the same options emerged already in 1991 at a WIPO worldwide symposium we organized together with Ralph Oman who was the Register of Copyright at that time. We were very, very close to get even Al Gore to participate, but finally he could not come. As I have said, the same kinds of options emerged there already.

Of course, we spoke about the hired work, the employed-author model and the producer model, all the contractual arrangements and we discussed, yes, the collective works model too, which existed and still exist mainly in the European Union.

We also mentioned the possibility of a sui generis protection system where it's not possible to prove that there is a human contribution. You should remember the credits at the end of the film. In the case of AI, the credit list may be even longer, but there is some human contribution.

What was interesting, what we discussed in Shanghai last December — from Europe, Joseph Straus was there for industrial property and I for copyright — was that there is no interest of the operator of an AI system to say that, no, there is no human contribution. There is no interest because it would result in a lower-level and shorter protection. Therefore, it would be difficult to go from copyright to a sui generis system, but it may still be tried.

MS. SHERIDAN: All right. Thank you.

Let's turn now to Carlo's presentation.

MR. LAVIZZARI: Thank you, Massimo, for a very interesting exposition on art and artificial intelligence. This demonstrates how broad the subject is.

I will delve into a completely different area, also linked to artificial intelligence, and adopt the observational standpoint of publishers and owners of copyright-protected

content created by authors, owners of databases, and data sets that are used in science, and many of them have factual works. I will lay out a little bit of the landscape in the time available.

The algorithms used to produce artificial intelligence are not so new as one might think. Many of them go back to the 1950s and 1960s, and I think it was mentioned this morning that many of them are available in libraries online and in open source.

What has changed is the computing power and the big data environment, and these algorithms — many of them prior art in a patent sense — together with the computing power and the big data produce this powerful mix.

Here you see, based on the WIPO report and some U.S. patent surveys, a sort of snapshot of where things were earlier this year, in January or February. Things change rapidly.

You can see who's in the race. Even though I come mostly from a copyright background, it's interesting to see that the patent filings are catching up with the published scientific literature about artificial intelligence or about algorithms that constitute one of the ingredients of artificial intelligence or machine learning.

There are many tools online that you can explore for yourself. This isn't something that is completely beyond us or exclusively for someone who is super-intelligent or has a science degree. You can build your own machine fairly easily, but you will find out that, whatever you do, you will have to train the machine, and for that you will need suitable data.

"Data First" is what the CEO of Novartis said when asked why there wasn't more progress on AI-developed drugs in his company. The challenge of getting clean data that is not biased as to the lab from which the data is coming or focuses on the wrong parameters is absolutely gigantic.

The algorithms were available even during the Cold War, as a famous example from the Cold War may illustrate. There were images of tanks from the Soviet Union available, and automated machines were already in the 1950s and 1960s developed to recognize Eastern Bloc Soviet Union tanks. This worked wonderfully in the laboratory environment, but in practice it didn't work at all because the machine learned to focus on the type of camera film used in Eastern Bloc countries instead of learning to recognize tanks or their muzzles or whatever.

To pick more positive applications of artificial intelligence, publishers already use AI features to do peer review, to scan for plagiarized manuscripts, and they are becoming better at it. You will find that in any business and legal practice AI applications will become commonplace fairly quickly.

What can be protected through patents or copyrights? I've put this snapshot together simply from reviewing the literature and the wonderful WIPO report that was released earlier in the year. You will find this is almost Software 4.0. It's not clear what this is. Is it really software? Is it a trade secret? People that I've interviewed said: "Well, actually, it's about the architecture of putting layers together and, again, a lot of value is in the data."

Where are the patents being applied for, whether they are valid or not? We heard a lot about Section 101² and the shenanigans around that. You will find that data and data sets play a very big role. From the perspective of publishers, I would say there is a gigantic opportunity in the future to not only publish for humans but to actually publish for machines bespoke.

² 35 U.S.C. §101 (2012).

If training machines becomes such a big market, then that should become part of the normal exploitation of a copyrighted work. Many more works will be read by machines than by scientists perhaps. People will use machines to review diabetes papers — there are about 16,000 — that describe the current medical knowledge, and the doctors will learn and read the most relevant ones. So the market will effectively be the machines.

Here I set out how the data used in machine learning might be protected.

Massimo already mentioned the database protection in the European Union, which also exists by the way in Mexico and South Korea. You can see good candidates for protection are copyright, trade secrets, database, contract, and then there are privacy issues associated with data that people might want to use in the public sphere.

Here is a list of what you can do to legally acquire data. Basically, there is no straightforward way of acquiring data to train machines, a standard procedure or an International Organization for Standardization (ISO) standard that would tell you how to go about this. In the long term it might be useful if industry by industry there could be best practices of how to train machines and to ensure that they're unbiased and properly set, especially where AI services would be socially beneficial.

Many of the proponents of machine learning devices or AI devices say, "Copyright inputs should be ignored because we're using just mere data," or "the use we make is non-consumptive." There is a lot of non-consumptive use of fact-based works. If you think of cookbooks — well, I suppose you consume what you cook, so in that sense it's consumptive, but not in a copyright sense. But also travel guides. People use fact-based works. People don't generally like fake news but factual news. Maybe those should still be protected to incentivize the use of qualitatively high-grade data, the integrity of which is assured.

It was interesting that Hugh Hansen mentioned charts and maps being protected by the U.S. Constitution, to say that maps are in fact also protected by copyright. What these machines use copyrighted works for is effectively the same as a map.

It's maybe not as clear-cut as saying "no protection whatsoever" or "full protection," but somewhere in the middle. If we look at the Google Map conditions, what we perhaps need in this fast-evolving environment is a best practice ecosystem that might be evolving and not immediately go to the questions of exceptions versus rights.

Thank you very much.

MS. SHERIDAN: I will turn to the panelists to see if you have any reactions to that.

PROF. YANISKY-RAVID: Regarding the data, if I could react to both of the great presentations, I would add few words about the "multiplayer model" — a term that I coined in a few of my articles, including "Generating Rembrandt: Artificial Intelligence, Copyright, and Accountability in the 3A Era — The Human-Like Authors Are Already Here — A New Model"³ and "When AI Systems Produce Inventions."⁴

When questions of ownership and accountability arise, policymakers tend to focus on the "man behind the machine," the software programmer. Policymakers tend to usually

³ Shlomit Yanisky-Ravid, *Generating Rembrandt: Artificial Intelligence, Copyright, and Accountability in the 3A Era – The Human-Like Authors Are Already Here – A New Model*, 2017 MICH. ST. L. REV. 659 (2017).

⁴ Shlomit Yanisky-Ravid & Xiaoqiong (Jackie) Liu, *When Artificial Intelligence Produce Inventions: The 3A Era and an Alternative Model for Patent Law*, 39 CARDOZO L. REV. 2215 (2018).

focus on the software programmer, where we are looking for the man behind the machine in questions of ownership and accountability. One should be aware of that fact.

But one should also be aware of the fact that it is the same (or similar) AI software for each and every AI system, regardless of the purpose for which it is being used. Sometimes the AI systems produce copyright works, other times similar systems are used for legal or medical goals, and so on and so forth. So I think it is neither the software programmer nor the software that directly creates the end product and other role takers in the process may be much more important when discussing ownership and accountability.

The most important role taker is the data provider. The other speakers who mentioned the major role of the data have made a very good and valid point. I think data is the most important input in this process for the function and the end results of AI and machine learning systems. Therefore, the data has become the main focus of the discourse nowadays, especially in the context of regulation.

I agree that data can be biased, but I think that when addressing data biases it is tremendously important to mention that there are four issues / problems / threats in regard to Big Data that AI systems “swallow”:

- (1) Data can be biased;
- (2) Data can be incomplete, fake, illegal or unreal;
- (3) Data can violate privacy rights. Just because AI systems use data via the Internet it does not cure the violation of our rights. AI systems “take” and use data which can be defined as our own property without our consent;
- (4) Data can be copyrightable and the use of this data may infringe third-party copyrights.

The four Big Data threats raise questions like: should someone whose data was used by an AI system have any claim, any ownership, and any influence on the outcome that AI systems produce?

My point of view, as I explained in my articles “Generating Rembrandt” and “When AI-Produce Patents,” is I don’t think copyright stakeholders have legitimate claims over the end product (in regards to “their” data being used by AI systems to produce new works). I see the process of AI systems that produce works of art in a similar way to an artist or an inventor who was exposed to many other paintings, galleries, and ideas before they combined them all together to create a new work of art.

So I do think data is the main issue, and I think the solution should be focused on transparency and regulation of the data in the same way other entities like the FDA or the securities agencies function. Technically speaking, data should be kept within the AI system and, legally speaking, data should be regulated because data is the main issue when addressing AI systems.

MS. SHERIDAN: Thank you.

MR. STERPI: I have a comment. First, there is a necessary inter-relationship between AI and data. We tend to visualize AI like software; but software without data is not AI, it’s just software; it’s a computer program.

Second, the interaction between the mechanical output and the human(s) is very complex. One of the best definitions of the “Portrait of Edmond de Bellamy” was given by Pablo Klingemann, that the author, the French collective called Obvious, missed the opportunity to declare this as the first digital readymade.⁵ That’s an interesting comment,

⁵ James Vincent, *How Three French Students Used Borrowed Code to Put the First AI Portrait in Christie’s*, THE VERGE (June 26, 2019), <https://www.theverge.com/2018/10/23/18013190/ai-art-portrait-auction-christies-belamy-obvious-robbie-barrat-gans>.

meaning that what comes out of the machine is inert, doesn't have a meaning, a significance, but actually the artist can make it a work of art using that very physical object, even if it is of digital origin, and giving it meaning by a title, by a commentary with a signature, or something else. So the connection between the two is very fluid.

PROF. YANISKY-RAVID: But you should also take under consideration the "Black Box Problem."⁶ Many of the software programmers and the users of AI systems do not really know what is going on and how the systems actually function and produce the new works. If you combine the "Black Box Problem" with the importance of the data, the issue becomes more complex and complicated. You don't know what the software is doing, but you know what data the algorithm was exposed to.

MR. FICSOR: I'd like to show this. [He shows the front page of the WIPO publication containing the material of the 1991 Stanford symposium.*] This was a creation by a neural network, but authorship was still claimed by a certain Harold Cohen from the University of California San Diego. I obtained the permission to use the picture for him.

Then, it was also remarkable that he characterized the system as a computer program. In fact, it was a neural network, and it produced a lot of different such drawings. The machine learned on the basis of graphical representation of human movements and produced hundreds and hundreds of different versions, and Harold Cohen chose one which was the most beautiful. So there is also the choice of the operator which may be relevant and might be considered a touch of creativity.

Then it was colorized by him; so it was a kind of co-authorship between AARON, which was the name of the program, and Harold Cohen.

MS. SHERIDAN: Thank you.

Katharine?

MS. STEPHENS: Thank you.

On November 1, 2018 the EPO updated its Guidelines and included a short, one-page section on the patentability of artificial intelligence and machine learning.⁷ This was immediately followed by some headlines proclaiming that the EPO had given the "green light" to patenting of AI. Is this correct, and has anything changed? The answers are "no" and "possibly yes."

The EPO identified three types of AI patents in the Guidelines.

- The first was core AI, and by this I'm referring to computational models and algorithms for classification, neural networks, and genetic algorithms. These are, per se, of an abstract mathematical nature and therefore fall within the exception in Article 52 of the European Patent Convention and so aren't patentable irrespective of whether they can be trained or not.⁸

- The second type of AI patent is AI as a tool in an applied field. The Guidelines gave an example or two of this. They suggested that the use of a neural network in a heart-

⁶ The regulatory and implementation barriers caused by the un-explainability of sophisticated AI.

*Available on the basis of a Creative Commons license in WIPO's on-line bookshop at <ftp://ftp.wipo.int/pub/library/ebooks/wipopublications/wipo-pub-698e.pdf>.

⁷ European Patent Office [EPO], *Guidelines for Examination in the European Patent Office*, at Part G – Chapter II-5 (Nov. 1, 2018), [http://documents.epo.org/projects/babylon/eponet.nsf/0/2A358516CE34385CC125833700498332/\\$File/guidelines_for_examination_2018_hyperlinked_en.pdf](http://documents.epo.org/projects/babylon/eponet.nsf/0/2A358516CE34385CC125833700498332/$File/guidelines_for_examination_2018_hyperlinked_en.pdf).

⁸ European Patent Office [EPO], *European Patent Convention*, at 108 (Jul. 1, 2016), [http://documents.epo.org/projects/babylon/eponet.nsf/0/A3DF61084E7706E7C12584A400521D6F/\\$File/EPC_16th_edition_2016_en.pdf](http://documents.epo.org/projects/babylon/eponet.nsf/0/A3DF61084E7706E7C12584A400521D6F/$File/EPC_16th_edition_2016_en.pdf).

monitoring apparatus for the purposes of identifying arrhythmia was sufficient to overcome the exclusion on mathematical methods because it involves the use of technical means, in other words, the heart-monitoring apparatus. Well, nothing to report here because this is really repeating the current status of the EPO case law.

- But at the bottom of this one-page memo was a paragraph that was all to do with trained models and training of AI models. This paragraph states: “Where a classification method serves a technical purpose” — so far nothing new — “the steps of generating the training set and training the classifier [i.e. the algorithm] may also contribute to the technical character of the invention if they support achieving that technical purpose.”

Question: Does this mean that the EPO is prepared to be a little bit more lenient over what is patentable?

To answer that, let’s first look at what the EPO is already doing.

There’s a Board of Appeals decision in relation to an Intellectual Ventures patent entitled, “Method of using image recomposition to improve scene classification.”⁹ This involves training a semantic classifier to generate an expanded set of images for improving both the training of a classifier and the classification of an image. This succeeded because it was all to do with image processing; since this is a specific technical application, the patent was allowed.

But these Guidelines could be read as indicating that a European patent would in principle be granted to an algorithm if it is possible to credibly link the training method or the generating of the data to a reliable and repeatable technical effect. It would still have to be shown that there are specific features that enable the training method to achieve the technical effect or there are specific features of the training data set. If this is what is intended, it will be a significant step forward in what is theoretically patentable, although as an aside I wonder whether it would be possible in practice to identify the necessary specific features.

But, if it is correct, then two points follow. First, it appears to be introducing an aspect of patent law that is usually only encountered in the biotech and pharma fields, and that is of plausibility because of the specific features that have to be shown to credibly enable the training method or plausibly enable the training method. Second, although nothing in the Guidelines changes the inventive step analysis of these types of claims, they indicate that the technical effect may be found in the method of training or generating the training data.

There are a couple of things that the Guidelines do not deal with, and I will touch on these very briefly.

The first is the test for inventive step and sufficiency, which refers to natural persons. In Europe a skilled person is considered to have at his or her disposal the means and capacity for routine work and experimentation. In future, will a skilled person encompass someone who has access to AI? If this is the case, then what will happen to inventive step? Will it be raised so high that actually nothing will be considered inventive in the eyes of the law, even if it was surprising and inventive to the human?

Also, the Guidelines don’t deal with the very difficult issue of inventive AI systems. Inventing is at present seen as a human endeavor, and it is not possible for AI to be considered the inventor. But in that case, who is? The current frontrunners, in my view, are the users and the devisers of the algorithm.

These issues will need to be considered in the future, particularly as computer-implemented inventions, including super-software based on AI, are tipped to be one of the

⁹ Case 1286/09, In Re Intellectual Ventures Fund, 83 LLC, 2015.

main drivers of future innovation, according to the EPO's 2017 study entitled *Patents and the Fourth Industrial Revolution*.¹⁰

Patent numbers in this area are on the rise. I would commend to you *WIPO Technology Trends 2019 – Artificial Intelligence*, a study of technology trends in artificial intelligence, which was published earlier this year and which Carlo also mentioned. There's a wealth of data in it. WIPO looked at all AI patents since the 1960s and found that there were nearly 340,000 of them. Although this only represents 0.6 percent of all granted patents, 50 percent were published in the last five years, mostly in the United States and China.¹¹

Governments are keen to promote innovation in this exciting and fast-moving area, and the patent system is often seen as encouraging innovation. However, this has to be balanced against the concerns that some types of patents, including computer programs, those for business methods — and I would include AI — could be seen as doing more harm than good.

I would call the EPO's new Guidelines on AI a first step at systematically setting up the rules to achieve that balancing exercise. But, because the case law and the patent examination process at the EPO is able to adapt within the confines of the European Patent Convention as the technology evolves, these Guidelines are almost certainly going to evolve as well.

So, as others have said before me at this conference, watch this space.

Thank you very much.

MS. SHERIDAN: Carlo?

MR. LAVIZZARI: Katharine, I have a follow-up question on the intersection, let's say, of copyright and patent law, and that is sufficient disclosure. I think one of the reasons why software became popular in the field of copyright is that sometimes software producers wanted to avoid disclosing the software and thought that by patenting they would have to disclose. Should there not be a system where the training data, like with the microorganisms that are deposited, is somehow deposited?

MS. STEPHENS: No, I don't believe that. I know where you're going with that, but I'm afraid I'm on the side of being open with your invention. It's part of the *quid pro quo* for a patent that you publish the invention such that those who follow you can, once the patent has lapsed, follow you and reproduce the invention. So, I'm afraid no.

PROF. YANISKY-RAVID: Your argument is very interesting. I would like to further focus on comparing copyright and patents.

On the one hand, I think that putting together an IP regime of copyright and patent laws and AI is an important point in discussing AI systems in general. This is because an AI system is software, and it might be patentable and copyrightable in a way. Additionally, the discourse about IP and AI can reflect how creative the AI system can be. And, finally, we can adopt solutions from the copyright regime, such as "AI work made for hire" by seeing the AI system as an agent.

Having said that, my question is: Don't you think all patent laws and even copyright laws have become outdated, inapplicable and/or irrelevant when addressing

¹⁰ Yann Ménière et al., *Patents and the Fourth Industrial Revolution: The Inventions Behind Digital Transformation*, EUROPEAN PATENT OFFICE [EPO] (DEC. 2017), at 84 [http://documents.epo.org/projects/babylon/eponet.nsf/0/17FDB5538E87B4B9C12581EF0045762F/\\$File/fourth_industrial_revolution_2017_en.pdf](http://documents.epo.org/projects/babylon/eponet.nsf/0/17FDB5538E87B4B9C12581EF0045762F/$File/fourth_industrial_revolution_2017_en.pdf).

¹¹ World Intellectual Prop. Org. [WIPO], *WIPO Technology Trends 2019 – Artificial Intelligence*, at 23, 32-33 (Jan. 31, 2019), https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf.

software that operates by itself in an autonomous way, when the software programmers cannot predict the outcome, when the software keeps on evolving and/or changing, and when the programming process has become never-ending? That's my question. Do we need to start thinking about new AI regulations or should we stick to patents and copyright and modify just some terms here and there?

MR. FICSOR: May I say something? Isn't it familiar to you the situation we are faced with? We have been there already once – not the ladies, because at that time they were still in kindergarten – but in the 1970s and 1980s there was a similar discussion whether software should be protected by a *sui generis* right or patent and copyright.

I contributed to the mistake - that the international community opted for copyright - as the head of the copyright administration of Hungary; it was at my proposal that we recognized copyright protection in statutory law as the third country in the world in 1982. I was one of the main allies of Michael Keplinger to promote copyright when at February 1984 we produced a breakthrough at a WIPO meeting – after which the same year, in June-July, the following countries opted also for copyright: France, Germany, Japan, the United Kingdom. I have to recognize retrospectively that a *sui generis* system probably would have been better.

Actually, what was applied was not truly traditional copyright. For example, in the Software Directive, in fact, *sui generis* provisions had been included about interoperability and so on.¹²

One of the questions is for the case that we decided to maintain copyright protection for such products: for seventy years after the death of whom the term of protection would be calculated? For the picture I have shown, Harold Cohen identified himself as the author, but on the picture he indicated the name of the program, AARON, the neural network.

70 year-protection after the death, of let's say, the operator who has produced these pictures may just be an argument in favor of a *sui generis* system. I am not a chauvinist copyright lawyer, so we should think about the available options twice or several times, and, of course, we should be cautious because when the AI or the robot says "I" it may not be so dangerous, but when they say jointly "We" it might become really frightening.

PROF. YANISKY-RAVID: Of course, if the U.S. copyright decision about the monkey's selfies is valid, that there is no copyright for works of art that were produced without any human intervention, what does this conclusion mean when a non-human (i.e. AI systems) produces works of art? Does it mean that we have no copyright over these works? Who owns these works? Everyone involved? Or maybe it's the state, if not one actor? Can the machine be the owner? Who is entitled to the \$432,000 that the seller received in an auction for a painting that an AI system produced?

That brings me to another question that I would love to hear some answers to, possibly from the audience. What happens when there is no regulation? Who becomes the regulator when you don't have any regulation? For example, when deciding that there is no copyright over works produced by AI systems — "Okay, there is no copyright over works produced by AI systems" — Who decides? It might be no one. But it might be the private entities, so should we invest private entities in a new originality concept?

MR. FICSOR: What about your new originality concept, Shlomit, because, as I have understood, it was one of your ideas that we should invent a new originality concept for AI.

¹² See Directive 2009/24, of the European Parliament and of the Council of 23 April 2009 on the Legal Protection of Computer Programs, 2009 O.J. (L 111) 16, 19.

PROF. YANISKY-RAVID: I wrote around ten articles suggesting ten different solutions, but I'm here to hear and learn more from you.

One of the solutions that you pointed out suggests amending and modifying the copyright law by changing the term "originality" to include works produced by AI systems, rather than a human, as original.

Another solution suggests treating the AI system as an agent and creating a new "AI work made for hire" regulation, not strictly and directly based on copyright laws, because we have different attitudes in Europe than in the United States, but regulating a new "AI work made for hire" rule that transfers the accountability and ownership to the user.

When we have to choose among different alternative solutions, we have to address the theoretical justification and ask ourselves: What do we want to gain? On whom do we want to impose the responsibility? It's not enough to say, "Let's vote for 'strict liability' as a solution" because the question regarding who is going to pay the price remains unsolved.

MR. FICSOR: I think that the perceived need for a new originality concept may be a good argument for a *sui generis* system. If it were not the same as for copyright, it might serve as a basis for a copyright-type *sui generis* system as was the idea of Dr. Arpad Bogsch, the late legendary Director General of WIPO in the 1970s for software.

MS. SHERIDAN: Let's turn for a moment to the fact that we have activity all over the world with respect to these issues. There's activity happening in the United States and in Europe and in Asia, and it feels like there's a need, or at least a drive, for some kind of global consistency on this.

Celine, could you speak generally on is that a need, is that a likelihood, and maybe give some background on the complexities of trying to get not complete consistency but at least agreement on some of these core issues?

PROF. CASTETS-RENARD: One of the main issues is to think about something universal, not only in an American perspective, because there are more flexibilities in the U.S. copyright regime than in the European copyright regime; for example, the United States has the concepts of fair use and work made for hire. We don't have these concepts in the European law and I think it's less flexible. For instance, one needs to work to train the system to create some artificial intelligence art. Using previous works might be an infringement in Europe because we have no exception to consider this situation. In the United States probably these kinds of issues could be covered by fair use.

In Europe, even with the new copyright reform that enacts new exceptions for text and data mining, that may not cover these kinds of issues because the new text and data mining exception is only for the purpose of scientific research and not for creation.¹³ However, the text finally adopted contains new provisions to cover the text and data mining in a general perspective, not only for research purposes, to integrate the new AI activities.

This is just an example to show in consideration of the different conceptions of copyright or author right, it's more or less difficult to create another regime or an evolution of copyright to consider the AI issues.

MS. SHERIDAN: Reactions to that?

MR. STERPI: I would like to comment on the actual need for copyright protection for a computer-generated work and the possibility of having other models.

¹³ Directive 2019/790, of the European Parliament and of the Council of 17 April 2019 on Copyright and Related Rights in the Digital Single Market and Amending Directives 96/9/EC and 2001/29/EC, 2019 O.J. (L 130) 92, 113.

There was a very interesting article in *The New York Times* about the exhibition of works created through the Creative Adversarial Network of Mr. Ahmed Elgammal where this programmer-turned-artist said that considering the infinite generative capability of the CAN, you can just continue pushing the button and you would have an endless number of images.¹⁴

The exhibition in Chelsea was the first experiment to print these works on canvas and exhibit them in a gallery. But actually they were thinking about offering a service of image production, and there was the idea that, like in fashion, you could have a new collection every six months, as you could take all the imagery of the last six months out of fashion catalogs and magazine and generate images that are perfectly up-to-date, “à la mode” so to say, and actually offer them for large spaces like a hotel or a corporate headquarters. They can change the images every six months and then throw the old ones away because it costs nothing to reproduce them. Actually, in that model the price would not be justified by the uniqueness of the piece, but by the speed and visual appeal, which are typical characteristics of fashion by the way.

Whereas copyright protection was granted both in France and in the United States to permit artists to make a living from their artistic endeavors, when the machine is the creator are we sure that we really need copyright protection? The machine doesn’t have to be incentivized — it creates because it creates.

Of course, there are already examples of machines that are practically autonomous. I’m thinking about the so-called Distributed Autonomous Organizations (DAO), which theoretically work, exist, and produce what they need through their activity. An example of a DAO can be a Bitcoin exchange that earns a commission for every transaction, and therefore has the money to buy services it may need and be completely autonomous from any human. At that point does a DAO need copyright protection? Not at all.

MR. FICSOR: Massimo, you mentioned France. It reminds me that France in the 1970s, early 1980s had a solution for software, a specific one. They said: let us keep it within copyright, but let us apply the model of the protection of works of applied art. After all software is the product of kind of applied science. Let us have a lower-level and shorter protection, which is allowed in that case, twenty-five years from the making. It also could be a logical, reasonable solution here, too. It would be nice if WIPO, after 1991, once again put AI on its agenda.

MR. LAVIZZARI: If you look at all the consultations taking place all over the world, they are about defining what is socially beneficial artificial intelligence. It might need to be a bit slow. We could create a man-versus-machine or woman-versus-machine race to say if it’s created by artificial intelligence, then no protection or only twenty-five years of protection; if it’s a human, then you have full copyright protection.

It reminds me of the traditional knowledge discussion, where previously people who didn’t want to pay for copyright said: “Ah, that’s a traditional knowledge work; I’m not going to pay.” Once you have traditional knowledge protection, it becomes a gambit. The infringer has to choose: are you going to go left or right? So this may become the same if we create a second protection scheme for artificial intelligence creations.

It is important also to make sure that the public trusts in whatever is created by artificial machines, and that is where I think the transparency is needed and why the data and the prominence of the data is key.

MS. SHERIDAN: Celine?

¹⁴ Gabe Cohn, *Up for Bid, AI Art Signed ‘Algorithm’*, N.Y. TIMES (Jan. 22, 2018), <https://www.nytimes.com/2018/10/22/arts/design/christies-art-artificial-intelligence-obvious.html>.

PROF. CASTETS-RENARD: Just to follow up this discussion, I don't know if everybody here is aware that two weeks ago we had some new Ethical Principles for AI that were enacted by the European Commission based on a report from AI experts.¹⁵ But these Guidelines are not very helpful for our discussion here because they only concern the point you mentioned concerning the ethical use of AI for automated decision-making especially and not the IP issues. For instance, some principles on fairness, transparency, accountability, and human autonomy were adopted. They are very important principles but not very related to this issue here in my opinion.

I don't think this kind of discussion can really answer the question of protection of art generated by AI, patentability of AI, and we really need to think about this question within the copyright system or the patent system.

Just an aside, if we consider *sui generis* rights in this field of property rights, that is a field of ethics because the questions are not the same at all. Property and ethics are two different questions and two different regimes.

MS. SHERIDAN: With the patent and copyright issues being different but also having so many similarities and overlap between issues like creators and inventors, how do you see the interplay between the two? How should this play out?

Katharine, do you have any thoughts on that to start?

MS. STEPHENS: When it comes to patents, what is being protected is effectively the engine behind the invention, whether you regard that as a software program, and therefore protected by copyright, which in this respect I don't think is going to be terribly useful, or whether that is patentable. That's one question.

When it comes to the inventive aspect, here I really do struggle as to whether we should be granting rights to AI and what the economic benefit is of doing so.

MR. STERPI: The fact that we do not grant patent protection to AI inventions does not mean that they cannot enjoy another form of protection. In fact, there is a very famous alternative route for protection, which is secrecy. Imagine that I know that I have no patent protection on something generated by an AI. What would I do? I would keep my algorithms secret, and at that point I enjoy the protection of trade secrets law. Not granting protection doesn't mean that everything becomes public domain; simply that I have to use an alternative route.

MR. FICSOR: Of course, the grandfather of all the IP rights, protection against unfair competition may also be there.

PROF. YANISKY-RAVID: Unfair competition is a very important point, but it doesn't solve the problem of ownership and accountability.

While I am sitting here and listening to this very interesting discussion, I think of my former students, Dinara Kamalova, who brought the *Disney* case,¹⁶ which is very interesting and relevant to the AI-IP discourse, and, more importantly, these issues are now being considered in other U.S. courts. *Company A* has seven (or more) patents and a bunch of copyrights over an AI-like system that produces masks. Disney was using this AI system to produce the mask, for example, for the film *Beauty and the Beast*. Now *Company A* is suing Disney for all its revenues derived from the technology used.

So here is the question: if one programmed an AI system, which is autonomous and keeps on evolving and creating new and unpredictable works of art, such as the mask (depending on the data the AI system keeps on "swallowing"), and the AI system creates the new mask that Disney used in the movie *Beauty and the Beast*, who is the owner of the end product? Can the AI system be the owner, does the mask belong to the system; and, if

¹⁵ *Commission Ethics Guidelines for Trustworthy AI* (Apr. 8, 2019).

¹⁶ *Rearden LLC v. Walt Disney Company*, 293 F.Supp.3d 963 (N.D. Cal. 2018).

so, who is entitled to the income? Would the owner be *Company A* or the user which is in this case is Disney? This is a question the courts have to decide.

How should the court decide? Are they going to use patent law? Are they going to use copyright law? What are the theoretical justifications for giving the rights over the mask to Disney or another entity?

MR. STERPI: Probably the response is that as long as humans will be involved, they will have a natural tendency to try to appropriate the result and get control. When the machine would be in control, maybe the situation will be entirely different.

PROF. YANISKY-RAVID: But even the U.K. solution that speaks about multiplayer is not sufficient. One has to consider “The Tragedy of the Anticommons: Property in the Transition from Marx to Markets,” a famous article by Michael Heller, that says if we have too many players that want ownership, nothing would be created.¹⁷ It is the same phenomenon that happened with Tesla’s accidents: the software programmers were blaming the camera, and the camera manufacturer explained that “the Shadow is to blame.” At the end of the day the court could not decide, as everyone blamed each other. However, we have to come to a decision.

MR. LAVIZZARI: Yes, the “success has many fathers and failure has none” syndrome. But I think in physics articles it’s not uncommon to have more than one hundred people credited for contributing to an article. So that could be done.

I am a bit concerned that authors may become like the person who still fills the steam engine with coal even though the train is now running on electricity. As a practicing lawyer, I would recommend people always add a human person to any artificially created work just to be safe in this transition period.

But I think the societal question is: what do we want for a beneficial use of these new technologies and how do we ensure that humans can also continue to create?

MR. FICSOR: What bothers me is we haven’t dealt with the 225 questions and comments that the members of the public certainly wanted to ask here.

MR. STERPI: We have a question.

QUESTION [Benjamin Rubin, Nokia]: If the programmer is the enabler of the expression by programming the software, the neural network, and but for those algorithms the software which creates the AI, where you start having expressions that will continue on and there could be some learning that occurs at the time, it would seem to me that for that data that’s being created, that expression — I guess this goes back to the example of the art that was provided — wouldn’t the enabler be that human? Wouldn’t that be the case?

My other question is with regard to the patents. Why could you not have patents? It seems that you’re going to have learning models that will be getting more and more efficient with AI and, as a result, there will be significant inventiveness and efficiencies that will spur business and growth in the economy, all of those things which are supposedly the benefits of patents and inventions. I would think that also would be patentable.

This covers the two different issues that I’ve been hearing about for the last half-hour.

MS. STEPHENS: Shall I deal with the patent one first?

QUESTIONER [Mr. Rubin]: Your choice.

MS. STEPHENS: I do agree with you, actually. Perhaps I was a bit flippant earlier. If the human is using AI to help devise a patentable invention, I do think that certainly the user could be named as the inventor. It might also be that the joint inventor would be the

¹⁷ Michael A. Heller, *The Tragedy of the Anticommons: Property in the Transition from Marks to Markets*, 111 HARV. L. REV. 624 (1998).

devisor of the algorithm. I think these questions have to be thought about when they arise. But certainly I think the user is in pole position.

When I was being a bit flippant I was thinking of the AlphaGo program that had come up with entirely new moves for the Go game.¹⁸ At that point, if AI is absolutely working completely autonomously to generate those new moves, I feel there must be a cutoff at some point — in contrast to where, as you were positing, there was a human involved in the process — at one stage the AI is working completely autonomously.

QUESTIONER [Mr. Rubin]: But those would then be expressions.

MS. STEPHENS: Yes.

QUESTIONER [Mr. Rubin]: It's not about the algorithm at that point. Those are expressions that are starting to get removed from the creator of the AI itself. It didn't seem to be algorithm-based.

MS. STEPHENS: Which segues to the copyright question.

PROF. YANISKY-RAVID: I want to point out the threat. We run a lab where we produce music and stories. You have to understand that once you have this software, I have no clue what it actually does. You press the button and you get tons (millions) of outcomes.

Two years ago, I co-authored an article with an engineer, Xiaoqiong (Jackie) Liu, that was published in *Cardozo Law Review* about AI systems that produce patents.¹⁹ We wrote about Koza's invention of a system that produced patentable electronic cycles and started obtaining patents that covered the inventions that the system produced. My point is that with AI systems you can produce millions of inventions (patents) in a few seconds.

Not only can AI systems produce patents, they can also submit the patent applications (as a standard process that we just saw used in traffic appeals in Stanford). I think it's a bit risky to say that inventions produced by AI systems can be patentable.

MR. LAVIZZARI: That may be the near future. But for now, for instance, if you told AlphaGo "lose" or "be happy you've won," there would be no response. I think for now there are always humans that in some way are involved.

QUESTION [Richard Vary, Bird & Bird]: Do you think when we're looking at copyrightability of works of art or patentability of inventions created with AI that we are now looking at what the intent is of the human being behind it, whether the human intended to use this tool to create something, or whether he created a tool which then autonomously did something? Do we now have to look at the intent of the person?

MR. STERPI: Even if your question wasn't expected, I have the answer in writing!

QUESTIONER [Mr. Vary]: You can pay me later.

MR. STERPI: I have a quote from Mario Klingemann with respect to his art created with a neural network and algorithms: "I tried to create algorithms or trade models in a way that gives them many degrees of freedom. But ultimately whatever output they generate is still within the space of possibilities I defined. Nevertheless, one of my long-term goals is to reach a point where the computer is able to continuously evolve its own artistic personality."²⁰ I think this is a perfect description of where we are and where we are going.

¹⁸ An artificial intelligence program created by Google. In 2017, the program defeated the Ke Jie, the number 1 go player in the world. See Paul Mozur, *Google's A.I. Program Rattles Chinese Go Master as It Wins Match*, N.Y. TIMES (May 25, 2017), <https://www.nytimes.com/2017/05/25/business/google-alphago-defeats-go-ke-jie-again.html>.

¹⁹ *Supra* note 5.

²⁰ Interview with Mario Klingemann, Resident, Google Arts and Culture, EyeEm Festival (Sept.16, 2017).

QUESTION [Robin Jacob, University College London]: Assume that's so. When is the work created and how long does the copyright last for? Life of author plus — who's the author?

MR. FICSOR: And what about the protection of the underlining computer programs themselves as "collective works". And who would represent neural networks if we granted rights to them? When would they "die", if at all, from when we could count the seventy years? There are a lot of questions.

MR. STERPI: The response is that currently in continental Europe or in the United States the work goes directly into the public domain because there is no copyright protection; therefore, there is no duration to calculate. The U.K. law, as you probably remember, talks about fifty years from the creation, but here of course there is no human in the picture.

What is also interesting, to comment a bit further, is that the two neural-networks-related art disputes that are known so far were not legal disputes about the economic rights (because in each case the use of the algorithm was under a free license, so there was no copyright infringement because copyright use was authorized), but actually the disputes were about the moral rights to be recognized as the author of the final work. This is interesting because at this point, exactly as Carlo said, it is more similar to what happens in the scientific world, where you do not want necessarily to get money out of it but rather to be recognized as part of the research group.

MS. SHERIDAN: Thanks, everyone.