TRANSCRIPT OF PROCEEDINGS

In the Matter of:)
)
COPYRIGHT ON ARTIFICIAL)
INTELLIGENCE AND VISUAL ARTS)
LISTENING SESSION)
)

Pages: 1 through 140

Place: Washington, D.C.

Date: May 2, 2023

HERITAGE REPORTING CORPORATION
Official Reporters
1220 L Street, N.W., Suite 206
Washington, D.C. 20005-4018
(202) 628-4888
contracts@hrccourtreporters.com

LIBRARY OF CONGRESS
UNITED STATES COPYRIGHT OFFICE

In the Matter of:)
)
COPYRIGHT ON ARTIFICIAL)
INTELLIGENCE AND VISUAL ARTS)
LISTENING SESSION)
)

Suite 206
Heritage Reporting
Corporation
1220 L Street, NW
Washington, D.C.

Tuesday,
May 2, 2023

The parties met remotely, pursuant to the notice, at 1:00 p.m.

ATTENDEES:

MARIA STRONG, Associate Register of Copyrights and Director of Policy and International Affairs
MARK GRAY, Assistant General Counsel
EMILY LANZA, Counsel
NICHOLAS BARTELT, Attorney-Advisor
DAVID WELKOWITZ, Attorney-Advisor
JORDANA RUBEL, Assistant General Counsel
JALYCE MANGUM, Attorney-Advisor
J. SCOTT EVANS, Adobe
BEN BROOKS, Stability AI
ALICIA CALZADA, National Press Photographers Association
SARAH CONLEY ODENKIRK, Cowan, DeBaets, Abrahams & Sheppard
KARLA ORTIZ, freelance concept artist
CURT LEVEY, Committee for Justice
REBECCA BLAKE, Graphic Artists Guild
JAMES GATTO, Sheppard Mullin
ALEX RINDELS, Jasper AI
PAUL REINITZ, Getty Images
LUC BOULET, Professional Photographers of America
HEATHER WHITNEY, Morrison & Foerster
DANIEL TAKASH, The Niskanen Center

ATTENDEES: (Cont'd.)

ZARA VARIN, Dual Wield Studio
DANIEL GERVAIS, Vanderbilt University Law School
JAMES SILVERBERG, American Society for Collective Rights Licensing
TOM LOCKLEY, Grey Owl Audio
MATTHEW CUNNINGHAM, Cunningham Concept Design
BRIAN FRYE, University of Kentucky College of Law
NETTRIC GASKINS, freelance artist
PHUC PHAM, Freelance Solidarity Project
ANKIT SAHN, Ajay Sahni Associates
JEFFREY SEDLIK, PLUS Coalition
PATRICIA SIGMON, artist/art director
DELANIE WEST, Be Super Creative

P R O C E E D I N G S

(1:00 p.m.)

MR. GRAY: Hello, everyone. Thank you very much for joining us today. We are going to start our session now.

Welcome to the United States Copyright Office's Listening Session on Artificial Intelligence and the Visual Arts. Today, we are going to be discussing a variety of issues in the visual arts space.

My name is Mark Gray, first off. I'm an Assistant General Counsel here in the Office of the General Counsel.

Before we start our first panel, I would like to introduce Maria Strong for opening remarks. Maria is an Associate Registrar of Copyrights, as well as the Director of Policy and International Affairs here at the U.S. Copyright Office.

Maria?

MS. STRONG: Thanks, Mark, and welcome, everybody, to the Copyright Office's Public Listening Session on Artificial Intelligence and Visual Arts. In copyright law, works of visual arts are broadly defined as pictorial, graphic, and sculptural works. Some examples include two-dimensional and three-

dimensional works of fine graphic and applied art, photographs, prints and art reproductions, maps, globes, charts, diagrams, models, and technical drawings, including architectural plans.

Because the visual arts include a wide variety of works, today, we will ask broad questions to facilitate discussion across each participant's area of expertise.

It's likely that almost everyone on this webinar has seen various images that deep learning text-to-image models can produce based on text prompts. We've heard concern from artists and photographers about what the training and deployment of these models might mean for their livelihoods and their industries both in terms of the input of their own images into these models, as well as the excitement and concerns related to the outputs.

And the purpose of our session today is to discuss these issues. We want to hear how the public is thinking about policy issues raised by these technologies.

To begin to address the copyrightability and registration issues raised by works generated using AI tools, the Office recently issued new registration guidance in mid-March. That guidance makes clear that

applicants have a duty to disclose the inclusion of AI-generated content in works submitted for registration. It outlines how to do so, how to update pending applications, and how to correct the public record on copyright claims that have already been registered without the required disclosure.

There was a lot of interest in today's event. Unfortunately, we were not able to accommodate all requests to speak. But this is not the last chance to share your views on AI with the Copyright Office. As we've said before and we'll say again, there are two more listening sessions happening later this month. And down the road, we will be requesting written input through a public notice of inquiry. Please visit our website, copyright.gov/AI, for more information and resources on our AI initiative.

Finally, we thank our panelists in advance for contributing to today's discussion and conversation. This is a complex topic and a deeply personal one for all our panelists, whether they are users or developers of AI technology, artists whose works help train that technology, or creators contemplating how AI will affect their careers. We are all looking forward to a thoughtful and respectful dialogue.

Let me turn the mic back to Mark Gray to outline the various logistics for today's session. Thank you.

MR. GRAY: Thank you very much, Maria.

So, as a quick reminder, before we get into specifics, today's listening session is the second of a series of listening sessions that we are doing here at the Copyright Office going through the end of May. Each of our sessions is going to look at different topics, different types of works, and, as a result, is going to have different panelists and may even use different formats.

So, after today, we have two more sessions scheduled. There is a session on May 17, Wednesday, which will be focusing on audiovisual works, which would include movies and video games. And our final session will be on May 31, which will focus on musical works and sound recordings.

The purpose behind these listening sessions is to inform the Office's overall AI initiative, so some of the questions our panelists raise may be ones that we seek to explore further in written comments later this year. So please keep in mind that while there are a handful of my colleagues here from the Copyright Office on today on video, the rest of the

Office is in the audience and is listening, and all of this is going to help inform our work.

The schedule for today, the session format is going to be two panels of different sets of speakers, followed by a third segment where a set of additional speakers will get the chance to share brief remarks.

We are making a video recording of this, both of this session as well as the other three. We are trying to get those online within three weeks of each session taking place, so please keep your eyes peeled for that if you have any friends or colleagues who don't have the opportunity to watch the session today.

Before we get started, a few Zoom housekeeping notes. If you are a panelist who is not speaking at the current session, please keep your camera and microphone off and on mute. And then, likewise, if you are a panelist, please keep your camera on and be ready to go off of mute when you're speaking.

We will be recording the session today. As I mentioned, the recording will try to go up in about three weeks from today. And we have enabled Zoom's transcription functionality for those of you who are

interested in following along with captions.

The way we're going to do the first panel is we're going to start with a brief introduction and short statement by each of the panel participants if they so desire. We'd like you to try to keep those to two minutes. We're going to be keeping an eye and the moderators may need to cut you off if it goes a little long just so we can keep everything on schedule.

After those introductions and brief remarks, we're going to do a moderated listening session. The panelists have received a set of broad questions in advance. Those are meant to prompt and guide a discussion, but panelists and participants are welcome to share any other relevant perspectives or experiences they think are important for the Copyright Office to hear.

If you are a panelist, please try to use Zoom's Raise Hand functionality, and we will try to call on you in the order that you raise your hands just to keep the conversation organized.

Please keep in mind this is a listening session and not a debate, so there will be other opportunities in the future for people to engage more directly with competing views. But the purpose today is really to help the Office air out a variety of

ideas and issues and perspectives for us to guide our own thinking.

As a final note, I see we have some questions in the Q&A from the audience. Unfortunately, this is a listening session for the participants. We are unable to accommodate audience questions. So thank you so much for your interest. Please keep your eyes out on our website for future public participation and comment opportunities, but we cannot take any comments today, unfortunately, from you.

With that, I'm going to hand it over to our moderators for the first session, Emily Lanza and Nick Bartelt. Emily is a Counsel in our Office of Policy & International Affairs. And Nick is an Attorney-Advisor in the Office of the General Counsel.

Emily, the mic is yours.

MS. LANZA: Thank you, Mark, and welcome, everyone. We will begin with introductions in the order as stated on the agenda.

So, first up, Scott with Adobe, would you like to go ahead?

MR. EVANS: Sure. Thank you for having me today. My name is J. Scott Evans, and I'm Senior Director and Associate General Counsel at Adobe.

For over four decades, Adobe's mission has been to empower our creative community with the tools that they need to express their imagination and earn their livelihoods in areas like photography, art, music, filmmaking, and design. AI and generative AI specifically have profound impact in these areas, so we really wanted to make sure that we, as we harness the power of this new technology, we're doing so in a way that empowers creators.

Last month, Adobe launched its generative AI technology, Adobe Firefly. Firefly's initial text-to-image model was designed to be commercially safe; that is, it was trained on images licensed from Adobe stock, openly licensed content, and content in the public domain. We want our tools to be good for enterprises and the creative community.

When it comes to copyright, we know that the issue of training is one where the creative community has concerns. For this reason, through a technology Adobe developed called Content Credentials, we're enabling artists to attach a do-not-train tag that will travel with their content wherever it goes. With industry adoption, it is our hope that this tag would prevent the training on content that has the do-not-train tag. We are working with generative AI

technology companies to respect these tags.

From an output standpoint, for much of our professional creative community, generative AI serves as the front door to the creative process. They're changing the image. They're adding colors, editing, adding elements. They're adding their own human expression to the work. So we need a way, a transparent way, to track this expression.

Here, Content Credentials can function much like an ingredients label. They'll show you where the image came from and what edits have been made to it. So, for generative AI, it gives the creator a way to show that they started with an AI generated image but, most importantly, to demonstrate the human creativity they brought to the work.

Finally, Content Credentials will bring a level of transparency that is much needed with the age of generative technology. Adobe is automatically attaching a Content Credential to images created with Firefly to indicate the image was generated by AI. We're working to drive transparency standards so that together we can deploy this technology responsibly in a way that respects creators and our communities at large.

I thank you for having me today, and I look

forward to engaging further on these issues.

MS. LANZA: Thank you, Scott.

Next up is Ben with Stability AI.

MR. BROOKS: Well, thank you to the Copyright Office for hosting us here today.

I lead public policy for Stability AI, a leading developer of open source AI models designed to unlock humanity's potential. These include, as many of you know, the latest versions of Stable Diffusion, which is a model that takes a text prompt from users and translates that prompt into a new image. Users can interact with these models either through a hosted service, like an app or an API, or they can freely use, integrate, and adapt the open source code subject to our ethical use license. Stability has also launched a number of other image models as well as a suite of language models.

Stable Diffusion is a type of latent diffusion model. So these models use content to learn the relationship between words and visual features, not unlike a student at a public gallery. Based on this acquired understanding and with creative direction from the user, these models can help to generate new works. In this way, AI should be understood as a tool to help artists express

themselves. It's not a substitute for the artist.

Instead, AI can help to simplify the creative process. It can help existing creators boost their productivity as part of a wider workflow. And it can also help to lower barriers to entry for people who simply don't have the resources or training to realize their creative potential today, including those with life-altering injuries or disabilities.

As with other assistive technologies, from paintbrushes to cameras to editing software, the user ultimately determines the content and use of any generated images.

I do want to acknowledge today the depth of feeling on these issues among creators and developers. AI is changing rapidly, and we understand that it can feel highly disruptive. We welcome a dialogue with all members of the creative community about the fair deployment of these technologies. And through the session today, I can share some details about how we're working towards that goal in practice, whether that be through new training techniques, authenticity standards, and best practices for things like opt-outs. So thank you very much.

MS. LANZA: Thank you, Ben.

Next up is Alicia.

MS. CALZADA: Hi. I'm Alicia Calzada. I'm the Deputy General Counsel for the National Press Photographers Association.

First, I really appreciate the invitation for NPPA to be a part of this event. This is very important to us and our members.

Before I was an attorney, I was a photojournalist for 20 years. And through NPPA, we serve the -- we are the nation's premier organization for visual journalists. We serve still photographers and videographers, and, frankly, most of our members do both.

Some of the things we do include working to support the First Amendment rights of visual journalists. We also advocate for their copyrights and for greater copyright protection and for a strong copyright system. We also have a code of ethics that is the industry standard among visual journalists, and that is, of course, a very important piece of what I hope we'll get into today.

NPPA has a few concerns related to AI. The first, of course, is copyright protection for photographers against unauthorized use of their images and unauthorized copying. So we do support legislation that accomplishes that.

For us, it's not just about money. As I mentioned, we care about ethics. And for visual journalists, their reputation is one of their most valuable assets. And so the right to control the use of their image and protect against misuse is very important. When their photos are used in an unethical manner, it impacts them. It impacts the entire industry, frankly.

And we also think that news consumers have a right to know the source and the authenticity of the content that they're consuming, the news that they're reading and watching.

Finally, a concern that we are monitoring is that journalists, like many photographers, do use technology in some ways that are, in fact, quite ethical, and so we're watching what the Copyright Office is doing as they frame the question of what is copyrightable. We understand that something entirely AI-created might not be copyrightable, but we want to make sure that in making policy we don't risk the copyrightability of photographs that for generations, frankly, have used special timers and triggers, such as the kind of things a sports photographer or a nature photographer might use.

So those are some of the things that are

sort of on our radar related to AI, and we're definitely looking forward to this session and to continued conversation on these issues.

MS. LANZA: Thank you, Alicia.

Next up is Sarah.

MS. ODENKIRK: Hi. Thank you very much for including me in today's conversation. My name is Sarah Conley Odenkirk, and I'm a partner with Cowan, DeBaets, Abrahams & Sheppard. I co-head the Los Angeles office and also the Art Law Practice Group.

My deep involvement in the implications of emerging technology and visual arts goes back almost 30 years with my dedication to representing artists and also working to establish public policy around visual art in public spaces. The combination of these elements in my practice has positioned me well to do a lot of advising around the impact and implications of blockchain technology and now AI, both from the standpoint of the impact on creators as well as on public policy.

It's crucial to maintain the focus on the impact the technologies have on artists and artists' abilities to continue to create and innovate. This becomes complex when we cannot easily determine when, where, and how potential copies and other copyright

infringements may be occurring.

In order to explore possible futures, we need to start by breaking down the processes used in AI into their component parts as the analysis will likely suggest different solutions at different points. Figuring out fair, enforceable, and economically sound solutions to questions raised at the point of training AIs will differ from determining how to treat the output artists coax from these platforms. We also must distinguish the generic generative process employed by commercial AI platforms from the more bespoke process of generative art as a medium employed by artists.

So I urge the Copyright Office to consider the impacts on artists in light of the new structures that are made possible with these technologies and to adhere or even strengthen principles underpinning the copyright law that support balancing the interests of artists' innovation and creativity with the market forces. It may be time to consider more than just guidance, more than just analyzing what is considered copying or protectable.

I would love to see the Copyright Office take the lead in championing technical solutions that meaningfully address the way content is scraped,

sourced, and used, and explore realistic ways to track IP rights and compensate creators.

MS. LANZA: Thank you, Sarah.

Next up is Karla.

MS. ORTIZ: So back in April of last year, I saw a website called Weird and Wonderful AI Art. It had the names of many of my peers alongside work that looked like theirs but wasn't. I thought it's just a new experiment. Well, I asked my peers whose name I saw on that website, and no one knew what this was and no one had been asked to be a part of it.

So we tried to reach out to the folks who were running the website, folks who are also selling merchandise that looked like the studies they were doing. We asked them to please take down the work of the artists who didn't want to be there. Instead, we got ghosted. I thought this was small enough to ignore, but little did I know this would be my first encounter with generative AI.

Fast-forward to September-ish, and larger generative AI models like Midjourney and Stable Diffusion are now mainstream. So I research again, and I am horrified by what I found. Almost the entirety of my work and the work of almost every artist I knew was scraped and utilized to train these

for-profit models. I was mortified that this was done without anyone's consent, credit, or compensation, that once AI models are trained on our work, our work could not be forgotten, and that generative AI companies were even encouraging users to use our full names to generate imagery that can look like our work.

For example, Polish artist Greg Rutkowski, who in December, between Midjourney, Stability AI, and the very problematic un-Stable Diffusion, Greg's name had been used as a prompt for image generation about 400,000 times.

If there is one thing I want everyone to remember is that this hyped technology is entirely fueled by the ill-gotten data it is trained upon. It is unlike any tool that has come before as it is an innovation that uniquely consumes and exploits the innovation of others. No other artistic tool is like this, and I know. I've used most of them.

In my opinion, to reward tech that relies on the proceeds of theft by granting it copyright would just add insult to injury.

Oh, also, my name is Karla Ortiz. I am an award-winning artist who works in film, game, TV, galleries, you name it. I worked on *Magic: The Gathering*, *Guardians of the Galaxy Volume 3, Loki*, and

most notably known for my design of *Dr. Strange* for the film adaptation.

I am also a plaintiff in one of the first class-action lawsuits against generative AI companies, specifically, Midjourney, DeviantArt, and, yes, Stability AI. Hi.

MS. LANZA: Thank you, Karla.

Next up is Curt.

MR. LEVEY: Hi there. I'm Curt Levey, President of the Committee for Justice. We're a nonprofit that focuses on a variety of legal and policy issues, including intellectual property, AI, tech policy.

There certainly are a number of very interesting questions about AI and copyright. I'd like to focus on one of them, which is the intersection of AI and copyright infringement, which some of the other panelists have already alluded to. That issue is at the forefront given recent high-profile lawsuits claiming that generative AI, such as DALL-E 2 or Stable Diffusion, are infringing by training their AI models on a set of copyrighted images, such as those owned by Getty Images, one of the plaintiffs in these suits.

And I must admit there's some tension in

what I think about the issue at the heart of these lawsuits. I and the Committee for Justice favor strong protection for creatives because that's the best way to encourage creativity and innovation.

But, at the same time, I was an AI scientist long ago in the 1990s before I was an attorney, and I have a lot of experience in how AI, that is, the neural networks at the heart of AI, learn from very large numbers of examples, and at a deep level, it's analogous to how human creators learn from a lifetime of examples. And we don't call that infringement when a human does it, so it's hard for me to conclude that it's infringement when done by AI.

Now some might say, why should we analogize to humans? And I would say, for one, we should be intellectually consistent about how we analyze copyright. And number two, I think it's better to borrow from precedents we know that assumed human authorship than to invent the wheel over again for AI. And, look, neither human nor machine learning depends on retaining specific examples that they learn from.

So the lawsuits that I'm alluding to argue that infringement springs from temporary copies made during learning. And I think my number one takeaway would be, like it or not, a distinction between man

and machine based on temporary storage will ultimately fail maybe not now but in the near future. Not only are there relatively weak legal arguments in terms of temporary copies, the precedent on that, more importantly, temporary storage of training examples is the easiest way to train an AI model, but it's not fundamentally required and it's not fundamentally different from what humans do, and I'll get into that more later if time permits.

But I think the good news is that the protection for creators of the works that are used as training examples can and will come from elsewhere, where the generated output is too similar --

MS. LANZA: Thank you, Curt. I'm going to have to -- sorry, I'm going to have to cut you off there.

MR. LEVEY: Okay. Sure.

MS. LANZA: But we'll have time during the question to continue.

MR. LEVEY: Sure.

MS. LANZA: Rebecca, would you like to go ahead, please?

MS. BLAKE: Yes, I'm happy to go ahead. And I'm apologizing in advance for the construction that has just started up outside my window.

My name is Rebecca Blake. I'm the Advocacy Liaison for the Graphic Artists Guild. The Graphic Artists Guild is a trade association representing the interests of visual artists other than photographers, illustrators, designers of all stripes, production artists, cartoon and comic book artists, animators and others.

Our mission is to protect the economic interests of our members, and in that vein, we've long advocated for greater copyright protections for individual artists, fair labor and trade practices, and policy which supports small creative businesses. We welcome this opportunity to weigh in on AI generative technologies.

Our members include artists who have embraced generative AI in the creation of their own original works and artists who, for various reasons, have not adopted the use of generative AI or, in fact, see it as a threat to their livelihoods.

While we support the ethical, legally compliant development of AI as a tool for visual artists, we have serious concerns about the copyright and ethical questions raised by AI generative technologies. These include the inclusion of copyrighted material in the training data sets without

permission or notification, which we see as a copyright infringement not excused by fair use, protections for artists' works as inputs into AI generative platforms, the unfair competition in the marketplace resulting from the massive generation of images which may ape existing artists' styles or replicate artists' works, confusion with the registration of works containing AI-generated material, and existing barriers to the affordable registration of works created by visual artists other than photographers. And I hope we can go more into this in the subsequent questions.

MS. LANZA: Thank you, Rebecca.

And last but not least, James, would you like to conclude the introductions, please?

MR. GATTO: Yes. Thank you. Hi. My name is James Gatto. I'm honored to have the opportunity to share some views here today on the important copyright issues with AI. I'm a partner in the D.C. office of Sheppard Mullin, where I lead our AI practice. I've been an IP attorney for 35 years. I'm also a member of the ABA IP Section AI Machine Learning Task Force, but the views expressed today are solely my own.

I've been doing work with AI for about two

decades, but, like others, I'm seeing a significant increase in that work due to the meteoric rise of generative AI. Clients have a lot of questions.

I applaud the Copyright Office's initiative to issue preliminary guidance on the examination of applications involving AI-generated content. I know there's great debate in the community on these guidelines, on authorship issues with AI, the level of human involvement needed, and issues with joint authorship. I hope these listening sessions will result in the Copyright Office keeping an open mind on whether to tweak their guidance and provide further clarity on some of the procedural aspects of the guidance.

Some of the issues for which clarity would be helpful are the following:

When does the level of detailed input or prompts by a human provide sufficient basis for the output to be deemed original intellectual conceptions of the author and therefore protectable?

What is the relevance of predictability in the authorship analysis? This concept was part of the basis for the Kashtanova decision but does not appear in the guidance.

What level of detail is needed to comply

with the duty of disclosure regarding use of AI?

What is the copyrightability of a work where a human uses AI-generated content as inspiration art but does not copy it?

And what are the criteria for determining if AI-generated content is more than de minimis such that it should be explicitly excluded from the application?

AI is a powerful tool, and to promote the constitutional mandate, the Copyright Office should develop policy that promotes rather than deters its use. As a result of the guidance in the Kashtanova decision, at least many companies that rely on copyright protection for their content, including game companies, artists, and many others, are concerned about using generative AI and in some cases restrict or limit employees' use of it. That's not consistent with the goal of promoting the use of technology.

So we hope through these sessions we get to a happy medium where artists' rights can be respected and tools can be used to facilitate the creation of their expressive works.

MS. LANZA: Great. Thank you, James, and thank you all for those introductions, and welcome again.

So, to begin our discussion, I'll start off

with a question. How is the training of artificial intelligence models affecting your field or industry? What should the Copyright Office know about the technology's use of training materials when considering the copyright issues related to training?

And, also, please be specific in your answers in terms of kind of which part of the visual arts ecosystem you're talking about.

So, great, I already see hands. So, Karla, you're the first on my screen. Can you please go ahead.

MS. ORTIZ: Yeah, absolutely. So, basically, the training of artificial intelligence is already affecting my particular field of entertainment, specifically concept art, illustrators, anything that requires a painter. We're already seeing the effects of these tools, you know, in our industries.

Something to consider is the training of these tools is very important. When considering these tools, you can't just focus on the output. You have to see the entire process as a whole. And as a whole, these tools, you know, particularly, you know, some of the tools around here, like Stability AI and specifically LAION, under the pretext of research,

gathered 5.8 billion text and image data from across the Internet to train various AI/ML models for commercial purposes. Again, it was trained upon for research and then switched immediately for commercial purposes. Technologists like Andy Baio call this, you know, loophole data laundering.

But another thing that's important to note is that this was done without consent, credit, or compensation. The work of myself and almost all of my peers are in those data sets, again, and also our names are, you know, encouraged to be utilized as prompts so that users can get something that mimics or feels similar to our work.

I personally am of the belief that the work generated by these models is impressive only because it is based upon the works of artists. And, again, this was done without consent. And we're not even talking about all the issues when it comes to propaganda, identity theft, and so on.

One of the things that I will say as well that the copyright should consider, and I won't take much longer of anybody's time so people can have their say, as one of the few artists in this panel, you know, there's various others as well, but I'm a teacher, and I can tell you that anthropomorphizing

these tools to equate it as human-like is a fool's errand. I've spoken to countless machine learning experts, such as Dr. Timnit Gebru, such as Professor Ben Zhao, and they all agree that it's not what's happening. This is a machine. This is mathematic algorithms. You cannot equate it to a human.

And to further add and to give the perspective of an artist, an artist doesn't look at a bunch, like 100,000, images and is able to generate, like, hundreds of images within seconds. An artist cannot do that. Yes, I have my influences, but it's not the only thing that goes into my work. My life, my experiences, my perspective, my technique, all of that goes into the work.

Furthermore, something that I feel like a lot of people miss in these discussions is technical artistry, and one of the hardest things you can do ever in the arts is be able to successfully mimic another artist's style or another person's work. It's the hardest thing. I consider myself masterful. I can't even do it. In fact, it's so rare that they even have documentaries on Netflix showcasing the few artists that can successfully mimic, let's say, a Leonardo da Vinci.

And depending on what that artist does with

that successful mimicry, if they sell it or if they do anything commercial with it, you know, that could potentially be called forgery. So I don't know why --

MS. LANZA: Thanks, Karla. Sorry to interrupt.

MS. ORTIZ: Oh, yeah. No, no, no, it's totally cool.

MS. LANZA: But, yeah, we just have a couple --

MS. ORTIZ: No, no, no. Totally great. Just wanted to drop that in.

MS. LANZA: All right. Thank you, Karla.

James, you're next on my screen. Please go ahead.

MR. GATTO: Great. Thank you. I'll try to be brief.

So, I mean, obviously, one of the core issues with training AI model on copyright-protected content is whether it's infringement and/or if fair use applies, and, largely, that's going to be a fact-specific question depending on the details. I think that to the extent there, you know, are any policy considerations or guidance the Copyright Office, you know, can provide in that, that might be helpful, but there is a pretty significant existing body of law on

that, kind of the broad legal test.

I think some of the areas that should be considered, consistent with what Scott said from Adobe, there's a lot of tools out there that can be used that help mitigate the problem, and whether those tools should be mandated or, you know, some other role the Copyright Office can play with respect to them would be helpful.

Should AI tool providers be required to be more transparent on the content they use to train their models? I think that's an important issue.

Should there be greater use of tools that prevent AI from using copyrighted works to train AI? Similar to how robots.txt works to prevent search engines from indexing certain web content. The technology is there and some of the concerns can be abated if these tools become mandated or just widely used.

And the last point I'm going to make is maybe not directly relevant to visual arts, but just, you know, there's other content that using it is not a problem because it's licensed, whether it's open source software that's being used to train AI code generators or like images that are under a permissive license, like Creative Commons, as long as there's no

prohibition on commercial use, the use of it may be permissible, but the question is then are there licensed compliance obligations that need to be met and, you know, whether and how those should be dealt with in these contexts. Those are just a few of the issues I think would be helpful to consider.

MS. LANZA: Thanks, James.

Alicia, you're next on my screen. Please go ahead.

MS. CALZADA: Thanks. That's a really interesting point about things like Creative Commons that actually do have conditions to, you know, what seems like on the surface an unlimited license, but, actually, there are things you have to do in order to earn that license.

Back to the question about how it affects our industry, the primary concern, as I mentioned earlier, in our industry really is an ethical one, and journalists rely on copyright as a means of controlling how their work is used.

And it's one thing to say, isn't it neat what this computer can do while you're, you know, just goofing off with friends or doing research or that kind of thing, but, you know, when these works start being used to create deepfakes or images that are used

to promote civil unrest, there are a lot of ways that news images can be abused through this kind of a process in very, very negative ways.

And the journalism industry really is concerned about where that's going to go and how it impacts the industry as a whole. You know, we already have editors who have for decades, you know, paid very close attention, you know, to work that comes in to ensure the quality of the sourcing and that kind of thing. But, on some level, there's things out there in the world that we worry about people seeing and thinking is journalism when it really isn't.

MS. LANZA: Thanks, Alicia.

Next, Curt, you're next on my screen. Please go ahead.

MR. LEVEY: Sure. Let me first briefly finish what I was saying about the good news for protection for creators, despite the fact that I do think it's getting harder and harder to distinguish between what humans do and machines do. But, regardless of how they're trained, where the generated output is similar to one of the examples in the training data or really any preexisting work, it's a derivative work or an outright copy, and the licensing requirements for derivative works need to be as

strictly enforced as for non-AI works.

And then, second, and some of the others have alluded to this, since the source of the training data is typically unlicensed data, I should say publicly available data or web scraping, we need strict enforcement of the website or database terms of service. And Mr. Evans mentioned a do-not-train tag. That's a good example.

Also, when you said what should the Copyright Office be aware of, I wanted to say a little bit more about temporary storage and why that's not fundamentally required.

Generative AI learns from a very large number of examples, and so does a human artist or author. The artist or author is not born with that ability here. He or she learns from countless examples of art, photography, music, written works, et cetera, and, you know, more and more the human views those examples on a website. The human may purposely make copies of the examples he used. And even if he doesn't purposely do it, his computer makes a temporary copy as he views the image, reads the written work, et cetera. Yet we all dismiss that copying as fair use, you know, if we even acknowledge it at all.

So what AI training does is not very different. For convenience sake, the examples are put in a database, which a learning algorithm cycles through, and that is temporary copying. But humans, like I said, often copy for convenience sake as well. And once the AI cycles through the examples in training, the examples can be thrown away.

The trained model, consisting of millions or billions of weights, analogous to the synaptic connections in the human brain, retains no copies of the training examples. Human memory, on the other hand, does remember at least some specific examples. So, in some sense, there's less of an infringement danger with AI than humans. But, to be fair, neither humans nor AI depend on retaining the specific examples they learn.

So, again, the problem with relying on the temporary copy argument is that it's not really necessary. You could train the AI model by having it scroll through the very same images or written works that the human learns from. In fact, the AI model could learn from, you know, data being relayed by a mobile robot that, you know, visits art galleries throughout the nation. Someday, you know, that may be how it's done. Think Google maps. Either way, my

point is that hanging one’s hat on temporary copying is skating on very thin ice.

MS. LANZA: Thank you, Curt.

Next up is Rebecca. Please go ahead.

MS. BLAKE: Yeah, gosh, there's just so much to unpack from that previous answer.

Very quickly --

MS. LANZA: Oh, Rebecca, you muted yourself. Can you unmute?

MS. BLAKE: I'm so sorry.

MS. LANZA: That’s all right.

MS. BLAKE: Very quickly, some of our members completely eschew using AI image generators. They're concerned about the ethical concerns with the way the image data sets were built. They're worried about copyrightability. And they're worried about exposing their clients to infringement.

Other members of the Graphic Artists Guild, in fact, use AI image generators. For the most part, we're hearing that they use it for ideation but not for the creation of completed works, or they use it to generate elements of a much larger work, for example, to create background graphics.

We do have one member who, in fact, has a career in AI generative for an AI image generator as

part of that new generation that has achieved a career. However, we've been trying to gauge job loss, job creation, job loss, and we're in very, very, very early days to be able to do that. It’s something we need to start tracking now that these generators have been out almost a year. However, we do hear a lot of anecdotal evidence of job loss. It's in particular sectors. That is hampered by the fact that many of the artists working in these areas, in fact, sign NDAs or are reluctant to go on the record discussing projects that they’ve lost because they're afraid of retaliation. They work in a very small industry.

Of our members who do use generative AI, one member stated that he was able to take on larger projects with a smaller workforce. So that does indicate that generative AI permits a streamlining and less hiring of artists. And another member stated that because she uses AI generative technology, she was able to cease contracting to a certain number of designers. So, again, that indicates a benefit to one member but at the loss of others. So that's speaking to the job market.

But I want to address two other things. First of all was this idea, this equivalency of the way machine learning works to the way human learning

works. This is a false equivalency for a very, very major reason. When a human learns to draw, they will ape, they will copy the styles or the works, existing works of illustrators they admire. This is very common in the learning practice.

But there are ethical considerations, copyright considerations, and best professional practices that professional illustrators follow that take them away from the wholesale copying of either a style or, in fact, of an image itself. This does not occur with machine learning. The machine is not driven by a creative process, a desire to develop one's own style, one's own mark, one's own creativity. It's simply reiterates a style that it has been learned on. So there is no equivalency in the outputs.

The second thing I just wanted to touch on very briefly was this idea that there can be tags or codes or metadata which is embedded in images which, in fact, permits one to track whether or not an image can be used for inclusion in a data set, whether it can be ingested into a platform, et cetera.

There's a huge issue with that, which is that Section 1202 of the Copyright Act permits the removal of copyright management information, including

metadata, if that removal is done knowingly or with reasonable grounds to know it will induce, et cetera, et cetera, infringement.

We believe that section of the Act needs to be modified so the removal of CMI, including metadata, without permission of the copyright holder is prohibited, regardless of whether or not it's done knowingly to permit infringement. We see metadata and CMI as key to being able to protect artists' works in an AI environment, but that failure in Section 1202 needs to be addressed.

MS. LANZA: Thanks, Rebecca.

Next up is J. Scott. Would you like to go ahead?

MR. EVANS: Sure. You know, at Adobe, we believe that if AI is done right, if this is done right, it benefits both creators and consumers of content because it does nothing but amplify human creativity and intelligence. It doesn't replace it.

And so what we see as a major issue here is that creators now have limited resources to attribute their work, especially when generative AI comes into play. One of the important things we need to do as a collaboration with artists and technology is to put creators at the forefront of this technology.

Creators want control over their work that is used in generative AI training, and we need to give them the tools in order to make those decisions. We know many creators that are very excited about this technology and want their creativity to be used in training these models. They are very excited about them. But we do understand that there's a segment of the community that is not excited and wants the ability to prevent the use of their art in training, and they should have an ability to do so.

And that's the reason Adobe has developed the Content Credentials. We worked very hard with setting up an open-source industry standard with the Coalition of Content Provenance and Authenticity, the C2PA. It's an open standard that platforms and hardware manufacturers can put into their products that will allow you to put these Content Credentials that will surface them to users and developers of AI technology so that those cues can be followed, and that's something that we think that's very important.

We also think there may be technology where artists could harness this technology by training models based solely on their own style or brand and then commercializing that and having that technology and understanding that there are different ways that

this technology can be used is very important.

One of the ways the Copyright Office, I think, can help in this is to encourage industry to adopt these open standards that will give artists the ability and tools in order to identify whether they want to participate or don't want to participate and encouraging that kind of proactivity among the companies that are developing this technology to give artists a tool to control their creative work.

MS. LANZA: Thank you.

And, Sarah, please go ahead.

MS. ODENKIRK: Thank you. I think I'm going to be reiterating a number of things that have already been said, but, first, I'd like to say that, you know, I think that there's a lot of reasons to be concerned about AI in general. There are big issues, big global ethical issues that definitely need to be addressed. Unfortunately, I think we do need to somewhat separate those questions from these questions that we're talking about with regard to copyright issues in order to parse through things. Otherwise, we're going to very quickly get sidetracked with, you know, scary potential future possibilities, which I don't think we should ignore, but we need to separate that out of the copyright conversation for now.

There's clearly a lot of potential in addressing some of the training issues through metadata and through some of the tools that Mr. Evans was speaking about, as well as some other tools that have been developed and people are looking to in order to protect the content.

And I'd like to underscore what Rebecca said with regard to Section 1202 and, you know, really needing to be concerned about the way in which the metadata can be taken off of content and thereby allowing it to be misused and really keeping creators from being able to track that data.

So I think that the final point that I want to make is with regard to paying attention to the purpose of the use that the images are being scraped and collected for. And if what we're talking about is using those images for the purpose of creating a commercial venture, a commercial product that's to be used to earn money, that's a very different use than artists looking at images and using tools in order to generate art, and while they're obviously connected, I think we need to look at them very separately in terms of figuring out what policies and laws and approaches we can take to protect creators in the front end of that process.

MS. LANZA: Thank you, Sarah.

So, before we move on to Question 2, I just want to make sure everyone had a chance to speak.

Ben, would you like to add anything to Question 1 before we move on?

MR. BROOKS: Yeah, thank you, Emily. I think just on this question of impact, I think these go to a broader set of issues around style and authenticity, and I do have remarks on training specifically for later. But I want to reiterate what I said at the beginning, which is that we see AI as a tool to help artists express themselves, but it's not a substitute for the artist.

That said, we obviously support efforts to improve creator control over their public content, and we're focusing those efforts in three areas in particular. So one is around access to content. Today, already, data sets like LAION-5B respect protocols like robots.txt that indicate whether a website consents to automated data aggregation. But we're also developing new ways to help creators qualify the use of that public content for AI training.

So one of the things we've done is we've committed to honoring opt-out requests from creators

in the next wave of Stable Diffusion releases. And going forward, I think this was a point alluded to by J. Scott, we're also exploring new standards for opt-outs so that the opt-out metadata will travel with the content wherever it goes subject to some of the problems that have been flagged just a little while ago.

The second area we're focusing on is authenticity of content. So we're working to implement content authenticity standards, like C2PA, with the Content Authenticity Initiative so that users and platforms can better identify AI-assisted content. By distinguishing AI-assisted content, these standards can help to ensure that users apply an appropriate standard of scrutiny in their interactions with that content. It can help to limit the spread of disinformation through social media platforms. And, ultimately, it can also help to protect human artists from unfair mimicry or passing off.

And the third and final point I just want to make is the work that we're doing to improve the quality of data sets. So, for example, by improving the diversity and reducing the duplication in training data, we can help to mitigate the risk of things like overfitting, which is where the system erroneously

overrepresents certain elements of a particular image from the data set. So, for example, if you've only ever seen sunsets, you might think that the sky is always orange.

In addition, by improving diversity in our data sets, we can be more representative of diverse cultures, language, demographics, and values, all of which can help to mitigate the risk of bias in those outputs.

So I think the final point on this question I'll add is, you know, we believe the community will continue to value human-generated content, right? We carry a complex digital camera in our pockets everywhere we go, yet we continue to value painting. Likewise, Photoshop didn't destroy photography. We have machines that can run faster than athletes, but we continue to place a premium on sport. And the same will be true of visual arts in the post-AI creative economy, particularly when we have some of these content authenticity standards in place.

MS. LANZA: Thank you, Ben.

I'll turn it over to my co-moderator, Nick, for the next question. Thank you.

MR. BARTELT: Thanks, Emily, and thanks, everyone.

I think the focus of the first question that Emily had asked was a little bit more on the input and the training. So I think we'll shift the discussion a bit to ask, what should the Copyright Office be aware of regarding how these AI systems, and some of you have already touched on this, how these AI systems generate works of visual art?

And then, as sort of a subquestion there, I'll ask, because I know we have limited time, is that, you know, are there any copyright considerations that vary based on the type of visual works that are at issue there?

So I see James's hand first.

And I know, Karla, we had lost you a minute ago, but I see you're on there too, so we'll go through.

Go ahead, James.

MR. GATTO: Great. Yeah, so, obviously, the operation of the AI tools vary, and each case is fact-specific. We recognize it's a challenge for the Copyright Office to give guidance for all scenarios, but there are a number of fact patterns that are common. And I think what would be helpful, one thing that would be helpful, is kind of like the Patent Office did with patent eligibility guidelines, if the

Copyright Office could provide examples of situations that are used in generative AI that they would deem to be copyright-protectable, that would be helpful.

The other thing is, just to take one use case, so I do a lot of work with NFTs as well, and there's a lot of concern around the use of generative AI for some of these NFT projects. If I create NFTs that represent images and, for example, I specify two images of dogs, each having a different collar that I designed and different colors, and I use AI just to generate the permutations of those artistic elements that I created under my control, saying, produce those permutations, AI should just be deemed a tool, even though it's output from a generative AI.

The question is, as you scale that up and maybe some of the parameters are a little bit looser, where does the line get crossed between it being my creative expression and it being too much input from generative AI? That's one practical use case that we've seen in a number of these NFT projects.

So I think that there's many other -- I'll be brief so others can talk -- but there's many other, I'd say, common use cases that we're seeing, and I think that any input or guidance or examples that the Copyright Office could provide would be very helpful

to, you know, assist those in trying to figure out where the line is and recognizing that, you know, there are fact-specific differences. Thank you.

MR. BARTELT: Okay. Thanks, James.

Actually, because, you know, Karla had her hand up before we had switched the question, we'll go Karla and then Ben and then Curt.

So go ahead, Karla.

MS. ORTIZ: Wholeheartedly appreciate it.

So something that I think the Copyright Office should be aware of regarding how AI systems generate work of visual art, there's been some talk at the idea of, like, whether these models copy, remember, memorize, whatever the word, overfit, whatever the word really is.

Something I'd like the Office to know is that studies are being done concerning these issues. For example, there's research from the University of Maryland and the University of New York. They did a study that found diffusion models generated high fidelity reproductions, which is basically plagiarism, at an estimated 1.88 percent, and it is estimated by these researchers to be higher.

Cursory numbers, but let's take a look at like Lensa AI, which uses Stable Diffusion, has about,

you know, this was around December, has about 25 million downloads, and gives users about 50 trials each. At 1.88 percent, that's potentially 23,500,000 generated images that could be similar to training data. And we see this consistently with, like, for example, Steve McCurry's famous Afghan girl can be perfectly plagiarized by these tools, as it happened in Midjourney.

Another thing that I'd like, you know, folks to consider as well is, like, the music, the, you know, Stable Diffusion -- no offense, Ben -- but Stable Diffusion already has made the case for us. I mean AI companies have already made the case for us.

For example, Dance Diffusion was a music program developed by Harmonai, which has links to Stable Diffusion, and as they trained their model, they stayed clear from copyrighted data and only did public domain. And one of the things that they quoted on is, because diffusion models are prone to memorization and overfitting, releasing a model trained on copyrighted data could potentially result in legal issues. Why was this done for the music industry but not the visual industries?

And this also goes to the point to why opt-out is inefficient, regardless of what it should --

the standard should be opt-in, because opt-out places an undue burden on people who may not know the language, who may not be online, who may not even, like, know what's going on. Companies cannot just arbitrarily grab our copyrighted works, our data, and just say, this is ours, and then later on we have to remove ourselves, which is why opt-out is important.

The other thing that's really important, again, is transparency. And I know that Adobe, you know, is mentioning this, but, for example, we need to really know what, for example, open license works mean. We really need to know and have an open data set to see exactly what it means so that licensors, you know, can actually, like, fulfill their licenses. And --

MR. BARTELT: All right. Thanks. Thanks, Karla. So I want to make --

MS. ORTIZ: Oh, I forgot the guidance on you, but you guys go on ahead. If we have time later, let's do it. I'm sorry.

MR. BARTELT: Okay. Yeah. I just want to make sure we get through everybody in our remaining 10 or so minutes.

MS. ORTIZ: Yeah.

MR. BARTELT: So I'll go back to you now,

Ben. I think you mentioned, what should the Office be aware of of how these AI systems generate works?

MR. BROOKS: Yeah, look, I think it's important that we properly characterize the training process, right? These models are not, as is sometimes being described, you know, a collage machine or a search index for images. These models review pairs of text, captions, and images to learn the relationships, again, between words and visual features, right? So that could be fur on a dog or ripples on water or moods like bleak or styles like cyberpunk. And with that acquired understanding and with creative direction from the user, those models can then help the user to generate new works. So, in this sense, training is, we believe, an acceptable and transformative use of that content.

But there are some good instinctive examples as well. Stable Diffusion notoriously struggled to generate hands, right? So it produced three-finger hands or 12-finger hands because it doesn't know that a hand typically has five fingers. And it isn't searching a database of the many images with hands, right? Instead, it has learned that a hand is a kind of flesh-colored artifact typically accompanied by a number of sausage-like appendages, right?

And that all has real implications for how we should think about AI training and generation. In other words, these models are using knowledge learned from reviewing those text-image pairs to help the user generate a new work. They're not using the original images themselves. And those images are nowhere in the AI model.

MR. BARTELT: Great. Thanks, Ben.

And, Curt, we'll go to you next.

And just so everyone knows, we have about two minutes for the remaining five people with their hands up before this panel ends. So, if you have any concluding remarks, you know, kind of work them into whatever you have to say here. Thank you.

MR. LEVEY: Well, let me answer the question, but also, in a sense, these are concluding remarks.

A couple of the panelists feel strongly that machine learning is not like what humans do, so let me say more about why I believe it is very similar.

The human brain consists of neurons connected by synapses of various strength. So, when a human sees an example, those synaptic strengths are slightly modified. Modification takes place slowly. But, you know, given a lot of examples, there's a lot of

modification and learning. That is how we learn.

Neural networks consist of artificial neural networks connected by artificial synapses. When the AI is shown an example, the synaptic strengths or weights are slightly modified, and, again, over time, the modifications add up, and we call that learning. And I realize this similarity is difficult to see because, for one thing, humans are very invested in seeing themselves as being very different than machines, you know, myself included.

So my advice to the Copyright Office is to look past that difficulty and use the similarity to your advantage to guide your policy development rather than reinventing the wheel.

And I would also suggest that your position that only product of human creativity can be copyrighted is something that, you know, you should think about. I don't take a position either way, whether copyright should be granted to AI-generated works. But the Copyright Office should keep in mind that it likely will have to recognize AI authorship as AI becomes more sophisticated and the philosophical distinction between human and machine creativity becomes harder to sustain.

You know, already generative AI can pass a

Turing test of sorts in that it's hard to tell the difference between human creation and machine creation, and that'll be all the more so once artificial general intelligence becomes a reality.

And I'll stop there.

MR. BARTELT: Okay. Thank you, Curt.

Sarah?

MS. ODENKIRK: Thank you. And apologies for noise. I'm in the corner of a conference room in the midst of a big conference trying to find a quiet space here. So just a couple of things.

I think that, you know, even if we assume that an AI platform has preemptively licensed all of the content that it's using to train the AI and that everything underlying is fine and licensed, we really need to look at the way in which the artist is using that content and how they're interacting with the AI platform in order to come up with their eventual artwork. And that's something that is going to have to be explored in a little bit more detail rather than just having really broad blanket rules about what can come out of that AI interaction and collaboration where artists are using it as a tool.

To go back to something that James brought up with regard to generative artists, I think that

point really warrants a little bit more exploration as well because we do have a whole group of artists and they've, of course, received quite a lot more attention lately, especially with the NFT marketplaces opening up. And that has to do with artists who are generating their own algorithms using their own data sets to put the artwork in and generating work. It may be that it's quite a bit generated by the AI platform and not by the artists themselves except that the underlying content is, in fact, created by the artist, including the algorithm. So these are really different ways of looking at tools that kind of emanate from the same place but end up being used in different ways and creating different results.

One just last point is that I think that oftentimes we're going to have to still look at and rely on traditional means for figuring out whether an output is substantially similar to either somebody else's work or that, you know, there are copyright infringements in that output. This is where we're going to have to consider things like the fact that many artists don't even register their copyrights. And so whether they have access to the court system because of that, it's really going to impair their ability to seek some sort of resolution of those

issues.

And this is where maybe we can take a look at the Copyright Office's Small Claims Board as a place to perhaps think about resolving some of these issues because we are going to have a tremendous number of artist-related issues that maybe don't rise to the level of needing to go to federal court but definitely are going to need to be resolved in a meaningful way for those artists.

MR. BARTELT: Thank you, Sarah.

J. Scott?

MR. EVANS: Thank you. I mean, at Adobe, we look at the laws that exist today. We look at the guidance that you provided. And after many hours of thinking through this, we view text prompts as conceptually an idea, an idea that is put into a machine that will give you many different expressions of that idea. So the expression seems to be being completely generated by a computer, and under current law, that is not copyrightable.

But, as we said in our opening statement, most of the creatives that we have talked to don't use generative AI as the end product. What they do is they use it as inspiration or a jumping-off point, and then they take that product and they do other things

to it, just like they would if they took a piece of art that was in the public domain and they created a new work that had elements of creativity and human creativity to that. So that's what we think is important.

I think what would be helpful is some additional guidance that you heard from our first speaker from the Copyright Office that gave examples of what you're looking for. And I hate to say this, but it probably would be more helpful if we had a new form that was specifically designed for AI-created works that drilled down on these particular issues so that some of the concerns that we hear from filmmakers and from photographers who use some sort of AI in the generation of the work, but not to the extent or in the way that we're talking about AI that generates the work itself, is used.

And so I just think that, you know, we need specific guidance. We need malleability from the Copyright Office. As the technology changes, as you just heard, we may need to change that guidance. We need the ability to do that and to follow these issues very closely.

But, at Adobe, you know, we think that we're headed in the right direction based on the current

laws that exist today, the guidance that you all have provided.

MR. BARTELT: Great. That's good to hear. Thank you, Scott.

And, Alicia, you are next.

MS. CALZADA: Thanks, and I'll be brief.

I just, I hear efforts to minimize the value of the work that are input into these systems, and I just want to make sure that the Copyright Office always keeps in perspective that without the input, which involves copying works in their entirety, there's no output. And so we can't minimize the value of the works that are used to generate AI. That would be inappropriate. Thanks.

MR. BARTELT: All right. Thank you, Alicia.

And just to wrap up, I think we have about two minutes left, maybe a minute each. We'll go to both Karla and Ben, and then we'll wrap up this session. So go ahead, Karla.

MS. ORTIZ: Thank you. So, again, I just want to reiterate that current practicing leading experts in the AI/ML machine-learning industry warn of equating machines to humans. We should heed their warning.

Furthermore, no matter how extensive the

prompt will be, it still relies on the training. And machine-learning companies want to automate prompts, as said various times by Stability AI's CO themselves. So I don't know how much of a standard that will be if it will eventually be automated.

And, lastly, honestly, for me, I'm concerned that AI-generated material will gain copyright, that the proceeds of theft will be rewarded. And, potentially, as someone who is likely to be affected in a whole industry that I see from the inside being affected by these tools, I'm very concerned. Thank you.

MR. BARTELT: Thank you, Karla.

And, Ben, with our remaining minute or so left, go ahead, please, and give us your closing thoughts.

MR. BROOKS: Yeah, just a last point on the guidance.

MR. BARTELT: Sure.

MR. BROOKS: Look, we acknowledge and accept that there is a threshold of authorship below which a work with negligible human input may not qualify for registration. I think our concern is that in the guidance and in the caption of the decision, the Office does not necessarily account for all of the

ways in which human input might rise above that threshold, and some of them have been discussed previously on this session.

A creator, to get a desirable input, may fine-tune that model using their own existing content to evoke their own existing style. They might provide detailed prompts that narrowly define the range of possible outcomes. And they also may, as J. Scott mentioned, use the image as a jumping-off point. And they may refine that initial image many times using traditional editing software or further prompting. Any one of these factors may qualify that work for registration. And so a user who has clear expressive intent and takes steps to steer these tools in a particular direction should be able to register their work.

I will just finally add that, again, we want to make it clear that training is not about stitching images together, it's about learning hidden relationships. But we do understand that there are other ways to improve creator control of their content, and, hopefully, some of the things we flagged today can be considered by the Office in your future deliberations.

MR. BARTELT: All right. Well, thank you,

Ben. And thank you to all of our participants in the first panel. I think we learned a lot, and we appreciate you all participating.

So, with that, I am going to hand the mic back over to Mark Gray.

MR. GRAY: Thank you, everyone. Thank you to all the panelists. This is the end of the first session. We are now going to take a 10-minute break.

(Whereupon, a brief recess was taken.)

MR. GRAY: Hello everyone, welcome back. For those of you who've only recently joined us, we're going to do a quick reminder on Zoom housekeeping.

For panelists who are speaking but not at this specific session, please keep your cameras turned off, please keep your microphones on mute.

We will be recording the session today. That recording will be made available on our website in hopefully three weeks. And we have also activated Zoom's transcription function today for those of you who would like to follow along with captions.

We're going to start our panel with brief introductions for each speaker, and, optionally, everyone is allowed to do a short statement if they desire. Please limit any statements to two minutes. We will have to watch the clock and keep things moving

if you go over two minutes.

And after those introductions, we're going to have a moderated listening session. The panelists have received the questions in advance, but keep in mind those are just intended as prompts and guidance for discussion. We, of course, welcome you to share any relevant perspectives or experience that you have that you think is important for the Office to hear.

For those of you who are on the current panel, please use Zoom's Raise Hand function. We will try to call you in the order that you raise your hand. That will help us keep the conversation nice and orderly.

And then, as a final reminder for those in the audience, we are, unfortunately, not accepting audience questions today, so, please, no need to submit questions or raise your hand. We will have opportunities for public participation in the future, including through a written comment period. But, for today, we are trying to focus on the panelists and the information they have to provide. So thank you very much for your understanding.

With that, I'm going to introduce our moderators for the second panel, David Welkowitz and Jordana Rubel. David is an Attorney-Advisor in the

Office of the General Counsel with me. And Jordana Rubel is also my colleague, who is an Assistant General Counsel.

David, the mic is yours.

MR. WELKOWITZ: Thank you, Mark. We'll start with the brief introductions.

Jasper AI, Alex, would you like to start, please?

MR. RINDELS: Yeah, thank you. My name is Alex Rindels. I'm Corporate Counsel at Jasper. I want to thank the Copyright Office for holding these really useful listening sessions. Jasper is a generative AI tech startup that develops and deploys software tools to assist businesses large and small and individuals in their content creative processes.

As it relates to the subject matter of this listening session, Jasper provides a software tool called Jasper Art that receives text inputs from our human users and filters them through generative AI foundation models from various providers and ourselves and spits out output in the form of art, whether that looks like photography or any other computer-generated art.

Our users are typically within the professional marketing and professional creative

spaces, and they use these outputted images in furtherance of things like marketing campaigns. And, oftentimes, it's not just the end images that they use, but they use those images for ideation to further their campaign.

And because many of our users and customers are creative professionals themselves, we receive two things, a lot of positive feedback that these tools have really freed up their creative processes so they can think more about the things they want to create and direct their efforts to the human side of the creative process rather than the mechanical side of it, and they're really grateful that we're able to aid them in that process.

But then, two, we also have an ear to the ground on the concerns that they have in terms of, you know, their work or others' works in the creative space being misappropriated and used for wrong reasons or not having the right rights to use those, and we really take those seriously, and we want to work with our customers and everyone in the community and in the Copyright Office to make sure those are handled in a responsible way. Thank you.

MR. WELKOWITZ: Thank you.

Getty Images, Paul, would you like to begin?

MR. REINITZ: Hi. Yeah. Thank you for the opportunity to speak at today's session. My name is Paul Reinitz, and I am Legal Advocacy Counsel at Getty Images. Getty Images is an established and respected member of the global media. Our growing content library includes over 520 million visual assets representing the work of more than 516,000 creative contributors.

Getty Images believes AI and generative models hold the potential to provide significant benefits. However, we see significant risks if the current development and deployment of these technologies are left unchecked.

First, bad actors can use generative AI tools to easily create deepfake content and distribute it widely via Internet platforms without check. This undermines public fluency in facts, trust in institutions, and democracy in general.

Second, generative models are trained on the existing work of creators. It is important that longstanding IP rights that protect these creators and sustain ongoing creation are not ignored. Transparency and respect for third-party rights are key components of mitigating these risks. Similar to the position expressed by your Office in its recent

registration guidance, we believe that the use of AI tools in the creative process should be disclosed.

Further, we believe that, one, AI generative content should be identified as synthetic, and two, to give rights holders visibility into the use of their work, records must be kept detailing how generative models were trained. The latest draft of the EU AI Act codifies transparency regulations of this nature, and we encourage the USCO to collaborate with the EU with the goal of harmonizing standards. We believe that supporting similar policy in the U.S. is crucial to the Office's mission of promoting creativity and free expression for the benefit of all.

Thank you for your leadership in this area. I know that the creative and media industries more broadly welcome the opportunity to provide further input.

MR. WELKOWITZ: Thank you.

From Professional Photographers of America, Luc.

MR. BOULET: Thank you. Greetings. My name is Luc Boulet, and I serve as the Government Affairs Manager for Professional Photographers of America. We're the world's largest and oldest photographic association with a membership of 35,000 creative

professionals. Our members include small businesses dotting every city and town across the nation, typically operating one- to two-person shops and earning an average net income of $38,000 per year.

PPA acknowledges the profound impact of artificial intelligence on our world, and we firmly believe in the responsible development of this revolutionary technology that will uphold the principles and objectives of the copyright system.

We also recognize the potential harm that AI may cause to the intellectual property sector, our leading net export valued at an astonishing $1.6 trillion. Our greatest concern with AI is the unauthorized scraping of a photographer's life work. This work is made public and available online by a photographer to promote their business and is then used to create new works by AI engines.

AI tools are being designed to directly emulate an artist's style at the click of a drop-down menu. While style is generally not a copyrightable attribute, the act of copying a photographer's image is a violation of copyright. Transparency is possible and achievable. Best practices from corporations, research institutions, governments, and other organizations that encourage transparency around AI

training already exist.

With this discussion, we hold two main objectives. They are, number one, it is essential the rights of copyright holders are respected as AI develops and AI laws and policies are formulated. And number two, new determinations on AI policy should be based on the foundation of preserving the rights of copyright holders and new rules and policies should be carefully considered to achieve this goal. Thank you.

MR. WELKOWITZ: Thank you.

Next, from Morrison & Foerster, Heather.

MS. WHITNEY: Thank you for inviting me to participate. My name is Heather Whitney, and I am an attorney at Morrison & Foerster and a member of the firm's AI Steering Committee. Previously, I was a Bigelow Fellow in Lecture and Law at the University of Chicago Law School and a fellow and faculty affiliate at the Berkman Klein Center for Internet Society.

Today, I'm speaking on behalf of my client, Kristina Kashtanova, author of *Zarya of the Dawn* and *Rose Enigma*. The Office recently refused to register images Kashtanova created using Midjourney, essentially on the grounds that Kashtanova could not be the author because they could not predict ahead of time what the output would be. In March, we submitted

a copyright registration for *Rose Enigma* and are waiting to hear back on that application.

My point today is a simple one. The Office is not writing on a blank slate when it comes to the copyrightability of outputs created with the assistance of generative AI tools. Images created with these tools are visual works, and the Office should treat all visual works the same.

Today, however, the Office's treatment of AI images diverges substantially from its treatment of photographs, with the bar for copyrightability much higher for AI images. This inconsistent treatment threatens to destabilize the registration process for visual works. This inconsistency shows up in several areas, but briefly I'll mention predictability and what it means to be the mastermind of a work.

On predictability, photographers receive copyrights in photos without knowing what their photos will look like ahead of time. Garry Winogrand, one of the most influential photographers of the 20th Century, said that he photographs to see what something will look like photographed. Wildlife photographers register photos taken by cameras on motion sensors, photos where the photographer was not present and had no idea what the image was going to be

until they looked later. Photographers do not have to predict how their works will look to be the authors of them, and AI artists should not either.

On the mastermind, to be the author of a work, the Office has stated that one must be the mastermind of it. We are all the authors. We could all register the endless, endless photos we take with our phones. If we are the masterminds of those photos, where our creative contributions are so minimal, it is hard to understand why AI artists, like Kashtanova, are not the masterminds of images generated after setting far more parameters and making far more creative choices.

In short, whatever the test for authorship, the test should be applied consistently across the visual arts. Thank you for having me.

MR. WELKOWITZ: Thank you.

From The Niskanen Center, Daniel.

MR. TAKASH: Hi. Thank you. My name is Daniel Takash. I'm a Regulatory Policy Fellow at The Niskanen Center. We're a 501(c)(3) public policy think tank. We work on a wide range of public policy issues, from employment and poverty welfare to climate and intellectual property.

Thank you so much to the Copyright Office

for putting this together. Thank you so much to everyone for attending.

And so I'd like to make three global points, I guess, to frame my comments.

First, it makes absolute perfect sense that the Copyright Office would be one of the first agencies to host formal discussions around the nature of artificial intelligence. Just the way that this field is developing and things are shaking out, they have implications for copyright law, perhaps before more than any other field.

And I would also like to commend the way the Copyright Office has dealt with the practical implications of artificial intelligence, especially as it relates to the registration of AI-generated works, both as a matter of policy and a matter of law as it exists today. I think they've gotten it right.

But the second point I'd like to make is that as important as copyright is, I don't think it should be the final word or even necessarily the most consequential word on developments in AI. This is a technology that perhaps will stop developing today, and it will just remain a novelty, in which case IP will remain the most important framework under which to regulate it.

But the sky is the limit, and the potential of artificial intelligence across virtually all domains of life are extraordinarily consequential, and for that reason, I think it's important to keep it in perspective. If we're talking about general safety, job dislocation, or other issues that are separate from, even though they may be related to, intellectual property, I think it's important that they take priority in any discussion.

And then, finally, I would encourage everyone to remain somewhat forward-looking. The role of copyright is to promote the progress of science, and it would be problematic for us to use that policy regime in order to limit it.

We can all look back and laugh at John Philip Sousa, you know, bemoaning the threat of phonograph-dooming music. And I want to be abundantly clear that the threats posed, that fear of creative destruction, is real, sincere, and precedented, and we must be prepared to say what other policy domains separate from copyright law should be best used to address the problems to the extent they exist.

Thank you.

MR. WELKOWITZ: Thank you very much.

From Dual Wield Studios, Zara.

MS. VARIN: Hi. Thank you so much for the opportunity to speak today. My name is Zara Varin. I am the Art Director and Senior Product Designer at Dual Wield Studios. It's a company based and founded on making things we love, inspired by the things we love.

My personal background is a little all over the place. After getting out of the Marine Corps, I established a career as a game developer and got a foot into the licensing industry as well. I've worked on everything from video games and comics to action figures, costumes, TV, movies. The point is I've gotten to work on a lot of cool stuff.

Before all that, though, I was a fan artist, and I still am. Uniquely, in my day job, we strive to champion and partner with fan artists and indie creators to elevate them in official collaborations with IP holders in ways that center those fan works while compensating and crediting them fairly, all while safeguarding the IP's representation. We bridge the gap between licensors and fandom to ensure creativity isn't stymied but celebrated.

Within the licensing industry, we're an outlier in our prioritization of fair wages, ethical production, and credit to the artists that we work

with. For many artists, that credit underscores their portfolio of work in acting as a resume.

So, in order to be effective at what I do, I have to keep up with constantly changing technology and tools, and I worry about the current state of generative AI. It grossly undermines credit to the vast swath of works informing their data sets.

For artists whose name has been used to populate guidance for these imitative generators, their actual work is becoming mired in a sea of manufactured imposters. It purports to be built for us, but the methodology informing its data sets and lack of clear ethical foundation indicates otherwise. When data set training is reliant upon things like underpaid workers combing through horrifying content to identify and tag it, it demonstrates a worrisome set of priorities.

Lastly, the most disheartening aspect of this is seeing how it has scared and dissuaded so many folks, from those just starting to explore creative expression to those who have spent countless hours honing their skills and often sharing to encourage others. Many people whose work was used nonconsensually to train these tools are being jeopardized and devalued in their own industry by

them.

I think generative AI genuinely has the potential to become a robust tool for creators, but it requires far greater ethical application and scope before it's something I'm comfortable considering incorporating into my workflows.

I'm grateful to discuss things further with all of you. Thank you so much.

MR. WELKOWITZ: Thank you very much.

From Vanderbilt University Law School, Daniel.

MR. GERVAIS: Yes. Thanks, David. Thanks for having me as a panelist today.

My name is Daniel Gervais. I'm a Professor at Vanderbilt Law School where I teach AI and IP law. I've written extensively on AI and IP and on the regulation of AI, and also on the legal and functional differences between human and machine learning, intelligence, and thinking. It's all available on SSRN, so if anyone wants to have a look. And, in fact, I just posted a summary checklist of issues on AI and IP to both Twitter and LinkedIn that you may find useful.

As to today's topic, there are five legal questions I'm interested in, and I don't think we'll

have time to get into all five, but here they are.

The first is, obviously, is the scraping or text and data mining legal? And, here, I'm interested not just in U.S. law. I'm also looking at foreign international law.

Second question, does the machine infringe when it produces a new work? And I think the analysis here needs to be separate for different rights in the copyright bundle. I could come back to that.

The third question we've been already discussing, I mean, in the previous panel, and I'm sure we'll get back to it in this panel, is, can the machine be an author? Here, basically, I strongly support the guidance from the Copyright Office. I'd be happy to explain why.

The last two questions I don't know if we'll get to would be, is the AI algorithm itself protected by IP law?

And the fifth is, is the data set, if you want to call it that to simplify, copyrightable or protected?

I'm also interested, and I'll end with this, at a deeper level on the alignment, by which I mean, is it possible to align the future development of AI in this space with the needs for humans to grow and

develop and to use art to communicate, both to send and receive new ideas and messages?

So I'll end here and look forward to our conversation. Thank you.

MR. WELKOWITZ: Thank you.

And, finally, from ASCRL, James.

MR. SILVERBERG: Hi. I'm James Silverberg. I'm the CEO of the American Society for Collective Rights Licensing. I'm also a former law professor, and for more than 40 years, I worked litigating copyright cases across the United States.

ASCRL is a not-for-profit corporation that represents tens of thousands of illustrators and photographers. We collect and distribute collective licensing fees. These come from collecting societies in foreign countries which have implemented proven and successful collective licensing systems to compensate artists for non-author, non-title-specific content use, that is, compensations distributed to authors in spaces where the use of content is certain, but the identification of the specific content or authorship is not always possible or is impractical. These spaces are like the AI space.

ASCRL is interested in exploring legislation to provide for collective licensing solutions for the

use of illustration and photography in AI platforms and believes that subscription and advertising or other fees should be collected to compensate authors on a nonspecific basis for the ingestion of their material into AI platforms.

We need to be aware that the current copyright paradigm is not well suited to the promotion of art and authorship in the context of how AI generates visual artwork. The constitutional authorization for legislation to reserve to authors the rights to their creations first found form in the copyright laws. The existing laws are focused on prohibitions against copying expression and permissible exceptions, like the Fair Use Doctrine, with the objective of preserving the economic benefit of artwork for authors with the intention of expanding American culture.

The constitutional purpose of reserving authors' rights remains outstanding, but the current copyright model does not work well when dealing with AI-generated works where the technical lines of expression copying are not always crossed or where existing fair use factors become a part of the equation.

For example, one of the main challenges with

AI-generated works is that the existing copyright laws are ill-equipped to preserve the economic benefit of authors when their material is learned or ingested or when uncopyrightable styles are appropriated. Even when infringement does occur, it can be difficult to detect and identify and prohibitive costs can be associated with pursuing legal action.

Much of our discussion today is already focused on the failures, challenges, and uncertainty of applying the existing law and debating its application in the context of injection, and this in and of itself may be proof that copyright law is at best problematic and uncertain as a solution to the problem of author protection.

For these reasons, ASCRL believes that we need a new way of thinking about how we should implement the constitutional premise that we reserve to authors their rights and their ability to receive compensation because the current copyright system is not achieving and cannot really achieve the constitutional goal in the AI context.

So, to address this challenge, ASCRL recommends that we do not entirely focus on the niceties of infringement, issues of interim copying, fair use factors, and that we move towards

legislatively implementing collective licensing systems like those that are currently used very successfully in many foreign countries. These systems serve our constitutional objectives and facilitate licensing and the use of AI and create a more balanced system that recognizes the needs of the AI community as well as the authors whose works or work attributes are ingested into these systems.

We are hoping to level the playing field by requiring non-title-specific, non-author-specific compensation where works cannot be specifically identified in order to compensate for uses where specifics of use are not available.

MR. WELKOWITZ: Thank you, James.

MR. SILVERBERG: I appreciate your inviting me on the panel, and I look forward to our discussion.

MR. WELKOWITZ: Thank you very much, James. And thank you all for introducing yourselves, and welcome again.

And to begin the discussion, let's begin with this first question. How is the training of artificial intelligence models affecting your field or industry? And what should the Copyright Office know about the technology's use of training materials when it is considering copyright issues relating to

training? Please be specific in terms of which part of the visual arts ecosystem you're talking about.

Let's see. And please use the Raise Hand feature. Ah, good. Zara, you're first on my screen.

MS. VARIN: Hi. Thank you. It's Zara.

MR. WELKOWITZ: Oh, I apologize, Zara.

MS. VARIN: It's okay.

MR. WELKOWITZ: I apologize. Sorry.

MS. VARIN: Thank you.

So I occupy several different unique spaces. I have a foot in a lot of different worlds, both as a game developer and within the licensing industry, and something that we've been seeing a lot across the board is that generative AI has kind of become a digital gold rush. And the training for informing that digital gold rush has come from a lot of my peers. I believe Karla touched on this during the first part of this panel.

But it's very disheartening when you Google search somebody's name and what comes up now is a proliferation of images that were created utilizing their art style but were not created by them themselves. And that's creating uncertain authorship and a great deal of uncertainty for people who are actually in these fields who are thinking maybe I

don't want my online portfolio available at all. Maybe I should safeguard or gatekeep my work in a way where I have to have password protection so that I can control the access. And it's essentially denigrating a lot of the work that people have been sharing online since the Internet's inception.

And it's incredibly disheartening as well within a production environment that these training methodologies are being done nonconsensually as well. Even people who have attempted to reach out utilizing the opt-out options that some of these models have kind of after the fact offered to artists, they're not having successes with opting out, or they're still seeing their work featured in those models.

And it's, I think, very telling that if you ask one of these models to create work based on an artist who wasn't part of that data set, it cannot. It'll try. It'll do something. But you're not going to end up with something that is aping the style and imitating something.

They're very sophisticated calculators by taking your set of inputs, running it through the algorithms to arrive at an output, but there's no creativity inherent in the process. And the training data utilized to do that is where the creativity

exists. The prompt might be a couple of cool words that somebody's put together. And I don't want to punch down on folks who are starting to learn how to do prompting. But the training data sets themselves are pulled together from a vast quantity of many works, folks who are here representing some of them, and that's causing many creators to second guess whether they want to be creators at all.

MR. WELKOWITZ: Thank you.

Luc, you're next on my screen.

MR. BOULET: Thank you.

This question implies that the visual arts industry understands the extent to which AI companies are using and profiting off of the works of others. And this is quite the opposite. Without a transparent and open AI system, it's impossible to determine which AI-generated works are incorporating the copyrighted material of others. And this leaves both the public and copyright holders completely unaware of which photographic works are being used by AI engines and the extent of their use. This is why it's imperative for AI companies to disclose their sources and methods used for when creating their final product.

There's growing anecdotal evidence of AI work displacing works created by human authors. A

publishing company looking to support their article with photographic evidence may turn to AI to produce a bespoke image, and others are inputting their selfies into an AI server to generate their latest LinkedIn headshot.

It's no surprise that individuals with the least bargaining power are most vulnerable to the negative effects of AI systems. Photographers who rely on their work to make a living often display their images in digital galleries and are promised with the option to opt out of having their works scraped in the future. However, these promises are empty as the timeline for when they may choose to do so has yet to be announced or is unclear at all.

Copyright owners may sometimes choose not to license their work, and that wish should be respected.

MR. WELKOWITZ: Thank you.

Paul?

MR. REINITZ: Thanks.

So, yes, I think that the short answer is that these technologies are having a big impact, and we expect that the impact is going to continue. As I said in my introductory statement, we believe that there's a lot of potential for these systems to help creativity, but there is also a lot of potential for

harm.

Now, to talk about exactly how it's impacting our industry, you know, I would like to point out that, and this kind of goes into the second question as well, that, you know, high-quality content along with captions that explain that content is really valuable in the machine learning process, and because of that, you know, there's a lot of demand for our content. And we're doing licenses out there. There's a big demand in our industry for people to come to us knowing that we have collected rights over nearly three decades of our existence and that we can license safely for these usage.

We also see that our customers are using this technology already. We recently did a scientific survey, and over half, I think it was 56 percent, of our customers that were polled are already using AI in their current workflow. Now, you know, we think that there are lots of great ways to use it and it can help promote creativity. But, again, you know, at the end of the day, these customers are more excited about licensing content from us for commercial use because they know that there's legal certainty in a license. And I think that hopefully we can get to a point where there's confidence in these tools that somebody can

use the output for their creative work.

I'll stop there.

MR. WELKOWITZ: Thank you.

Daniel, you're next.

MR. GERVAIS: Thanks. Three quick points.

The first is I think the Office in its capacity as advisor to Congress and the courts should bear in mind the international obligations of the United States, in particular, the TRIPS agreement, three-step tests, for example, in providing advice on what can be done without permission or payment.

Second, the question of scraping, of course, can happen in many different countries. Many other countries have already adopted exceptions, the EU, Japan, others, Singapore. So one way or the other, you know, text and data mining will happen, but there will be limits, and whatever needs to happen beyond those limits will need to be licensed. So perhaps the Office can play a role in facilitating licensing. And by licensing, I don't mean just that if you're an artist, you get paid once because your work gets scraped and you get paid because there's a licensing regime. A license can also impose contractual limits on the use and reuse of material.

Now anything that's scraped can lead to an

output that will potentially create a commercially competitive product, but at least there is some conversation, some exchange of consideration.

So I think those are the three roles I would see the Office potentially playing. Thank you.

MR. WELKOWITZ: Thank you.

James?

MR. SILVERBERG: So I wonder if we'll really succeed in disentangling the many problems of AI. Is the problem really whether there's a copyright infringement? Is the problem really identification of works? Is it opt-out? Is it content access restrictions? I think all of these things are important when we're looking at the ingestion problem. But is that really where we should be focusing our inquiry? Isn't the problem that we need to have a clear system, call it copyright or something else, that makes sure that the authors are compensated?

MR. WELKOWITZ: Thank you.

Heather?

MS. WHITNEY: Thanks. I just want to quickly mention something about the contribution of artists who are using these tools in terms of how it has outputs. So I think there's a misconception that all of these people are just sort of typing in some

generic words and then there's an output. I strongly recommend people just Google control net models, Stable Diffusion, and you can see the different kinds of models that are now used on top of kind of like an extension of Stable Diffusion in order to have a much, much greater control over the images that are being used. It's not just the model that's doing the work.

And if you look at what they can do with these things, with these kinds of different kinds of models, setting up the composition, choosing sort of the angles, the lighting, all those kinds of things, and you compare that to the things that are said in the Compendium about what is required for someone to have a copyright in a photograph, this is just absolutely without question equal to or greater than the kinds of things that people are doing in that context.

So I just think it's helpful that people understand the technology and how it's evolved, and I think the control net models are really a way for people to start to understand where that's going.

MR. WELKOWITZ: Thank you.

Alex?

MR. RINDELS: Yes. Thanks. I'd just like to briefly respond about how the output is positively

affecting the industry that we're in.

So, in two regards, the professionals who use our service as creatives, they're able to unleash their creative ability that might have otherwise been tied up in time-consuming creative processes. So many of them use Jasper Art to create ideas that they then build upon or to create end products that they use in marketing campaigns or in other business uses or commercial uses, and this greatly frees up their time to use their creative potential, their intelligence, their efforts in other productive ways.

And secondly and probably more importantly are the output in tools like Jasper and others, Stability, it allows people in the creative space who otherwise could not have created output like this to create output, and I'll explain briefly.

So we also have a text-generating tool. And we've had numbers of customers who have dyslexic disabilities or otherwise who would have otherwise been unable to create output, and they regularly respond to our tools and tell us how grateful they are that they're now able to take what was in their head and put it to paper. And, previously, it was just basically a mechanical limitation that prohibited them from doing so. And now, with tools like this, it's

untapped, it's unleashed their potential that otherwise was untapped in the industry.

So, in the first part, it's freeing up space for people who already have the potential to create art or other output and it gives them more time. But, secondly and probably more importantly, it also allows those who otherwise could not have created artistic output to do so.

MR. WELKOWITZ: Thank you.

Daniel?

MR. TAKASH: Thank you. So, with respect to the images that are training, Niskanen's policy with respect to all the works we put out, images included, is Creative Commons provided, there's attribution, because we do our best to contribute what we like to think at least is good-quality work so that images can be trained. We are also fond of using AI-generated images just as a substitution for stock photography that we would otherwise license, or we can have some bespoke images that particularly complement some work that we're generating. So we have an interest in making sure that material is out there so that the models can be the best that they can be.

That being said, we recognize that there ought to be tools available and the law should

accommodate the use of those tools to allow the authors or the rights holders of works, in this case, visual artists or visual arts, although this principle can be extended to other contexts, that, you know, that creates some ability to remove or make it so that it's much harder for their works to be learned upon.

That being said, we should recognize that copyright law does not necessarily provide an absolute veto to the rights holder, and there are plenty of examples, you know, say what you will about the quality or the desirability of this outcome overall, where, you know, use can be even in a way that the original rights holder may not approve of.

That being said, in order to ensure quality and respect, incentives to produce arts, I think the best way to square the circle is to focus on a regime that deals with remuneration and financing, which deals with much larger policies, as I alluded to earlier in my comments.

MR. WELKOWITZ: Thank you.

Paul?

MR. REINITZ: Hi. Thanks. Yeah, I just wanted to make a quick comment about what Daniel just said. I mean, you just said that you, you know, as a consumer of stock photography, sometimes use generated

content as a substitution. I think that that's fine and I think that the market is probably going there. But I would like to point out that if you are going to be using a substitute and that model that creates the generative content was trained on unauthorized content, you are substituting a product without compensating the artists that were needed to make that model.

And, you know, I just, sorry, I just needed to call that out because it just, it’s so real, as you say that, as a user of stock photography.

MR. WELKOWITZ: Thank you.

James?

MR. SILVERBERG: Just to follow up or elaborate on Paul's point, which I very much appreciate, the problem is also particularly paramount for small businesses and medium-size businesses that involve authors who do not have vast amounts of content to aggregate and to license on a large-scale basis. And so, while there are content aggregators that are able to do that on a much larger scale because they occupy a larger market force, hundreds of thousands of individual creators and authors don't have the capacity to engage in a licensing transaction with an AI platform, nor do they have the capacity to

enforce copyrights in copyright litigation or possibly even to identify the use of their work.

And so that is why ASCRL is advocating for a remuneration system that will prevent a complete market displacement of these constituents and which will make sure that they get compensated.

MR. WELKOWITZ: Thank you.

Daniel?

MR. TAKASH: Thank you for circling back.

Yes. So just to be clear, with respect to Paul's point, I am fully cognizant of the, I guess you could call it, recursive nature between the body of works available for training and the output of it, so I'm fully aware of that.

I think the problems -- or not problems, the challenges artificial intelligence creates with respect to copyright is a difference of both degree and of kind, which is why I'd like to underscore again the support for something that moves away from more traditional models of infringement, alluding to -- I can't say I'm super familiar with the model James brought up originally, but something that focuses more on remuneration moving into a very different technological paradigm in order to, as I mentioned before, square that circle.

MR. WELKOWITZ: Thank you.

Okay. I think I'm going to turn the mic over to my colleague, Jordana Rubel, for the next question.

MS. RUBEL: Thanks, David.

I guess I'll start off by maybe just asking a question that relates to the point we were just talking about, which is more solution-oriented thinking about if we get to the place where we are just talking about remedies here. Maybe we can start with James to give a little bit more detail about what ASCRL's proposal is, and if other folks want to speak to that, you're welcome to raise your hands as well.

I'll turn it to you, James.

MR. SILVERBERG: Thank you. So I appreciate in the United States there's a general lack of familiarity with collective licensing systems, particularly in the visual space, because we don't have any. In foreign countries, there are collective systems in a number of different areas. I'll just mention two, for example, to illustrate how they work.

One would be in the area of library lending. As we all know, in the United States, there's a first sale doctrine. If a library buys a book, they can lend it forever and never pay for it again. But, in

foreign countries, a system's been introduced as a secondary or adjunct system to the copyright law involving what's referred to as secondary rights.

For library lending, people don't know what books are being lent, people don't know which authors are involved in the lending, people don't know which photographs are in what books or what illustrations are in what books, but fees are paid into a collecting society in order to compensate the authors for the lending of this material. It's done on a non-author, non-title-specific basis. We don't know exactly what the fees are for other than they're for the lending itself, and then algorithms are established in order to make a fair allocation of the revenue to the individual authors whose works might be embodied.

It's a form of rough justice. It's not specific. It's not a perfect system, but it's a way of creating compensation. It is analogous to a system, sort of like a SoundExchange, where fees are being paid for the use of recorded music that's digitally streamed, and that goes into SoundExchange, one of our directors on our Board of Directors, a former executive director of SoundExchange, and those funds are distributed, but they're distributed on a more specific basis where it's known what material is

used.

We would be proposing a system where AI platforms would be distributing compensation to collecting societies in order to distribute those according to a fair algorithm which would compensate illustrators and photographers. And this was also espoused by Authors Guild in the other listening session for works in the writing space.

MS. RUBEL: Thanks.

Daniel, did you want to respond to that, or do you have any other thoughts on this topic?

MR. GERVAIS: Just a footnote just to be clear, James is right about public lending, but the way it works is not random. I just want to make that very clear. So, in countries where they have it, I don't know, Germany, Canada, and others, the authors must register their works, and then the collective actually will only pay -- so they will survey certain public and private libraries, usually mostly public libraries, and will only pay if the book is actually in those libraries. And the reason they don't pay per, you know, the number of times that the book was taken out by a user at a library is partly privacy.

And so the second regime, though, that is a little different is this thing called extended

licensing that I know the Copyright Office has published a number of reports about where, basically, a collective is given the -- basically, it becomes an opt-out, essentially, so the collective basically has rights to represent a class of right holders, except those who opt out. This is very, well, I can't say very common, it's common in Europe, and other countries are looking at it. Some countries have functional equivalents of that.

That is a different thing because what it does is it gives a collective a right to license, but the basis on which then authors will get paid or publishers or whoever else has rights is not predetermined. It has to be decided case by case in the appropriate way. It's not a black box. I want to make sure people understand this isn't just a black box of money that gets, you know, paid somehow. You can use data to actually apportion the funds correctly. Thank you.

MS. RUBEL: Paul?

MR. REINITZ: Yeah. Thanks. And thanks, Daniel. I was actually going to bring up a similar issue.

I would say that, you know, I think the idea of collective licensing is a good one but only if it's

really necessary. Really, you know, it's a complex system to set up. There's a lot of administration, and, you know, there can be a lot of inefficiency in it.

From where I currently sit, as I described in my opening statement, we're seeing a lot of interest, and we're doing a lot of licenses directly with people that want to or organizations that want to license content. Now I understand that that doesn't work for everyone, and it's much harder for an individual, you know, photographer to say do that on their own, and maybe we need a solution like what James is suggesting.

That said, you know, along with what Daniel described, I think the idea of an opt-out and following these models of extended collective licensing is essential because, you know, if an organization wants to opt out, they should be able to do so.

MS. RUBEL: Thank you.

Zara?

MS. VARIN: Thank you.

So, to your question, it's been really interesting to see this conversation contextualized and framed by different industries. So, for instance,

film and music and the music industry as a whole, they have a far more rigorous enforcement of their copyright on their works. Artists, especially within, like, the games industry or visual effects industry or working as freelancers, don't have a single unified front to advocate for us.

So there's efforts to unionize, but we don't really have any sort of large collective that is coming to attend events like this, that's speaking on our behalf and helping to advocate for the rights of artists and creators in those industries.

I think folks on these sorts of creative teams doing freelance work and within, I think, more of the cutting-edge tech industries where there's a lot of overlap between utilizing technology in our day to day and rapidly adapting to what that new technology is could definitely use some sort of remunerative residual system or something along the lines of what I think James was getting at and what Paul has also touched on. But that has to work and advocate alongside credit to the artists who have helped build that system. That system wouldn't work without all of the pieces that have built that learning model.

So I think step one is figuring out, how do

we backtrack and ensure that all of the existing learning models are clear, transparent, ethical, have defined what their scope is, and also define what efforts they are making? Because the onus is on them for building these tools to do so in a manner that is not going to undercut people that are already working in those creative industries, because, again, I do think there's a lot of great potential in these tools, but they are tools. They are not creative.

Artists necessarily are called skilled. It's not talent. Talent is kind of a very worrisome eugenicist concept. Skill is a thing that you have to work at. There is no skill inherent in the AI generative process. The only skill present is coming from the works that are scraped to put everything together. So whenever we're contextualizing all of this, I want folks to keep that in mind, where is the human labor and where is the skill coming from?

MS. RUBEL: Daniel?

MR. TAKASH: Oh, sorry, two Daniels. It gets confusing.

MS. RUBEL: Sorry.

MR. TAKASH: Yes. So, with respect to the model for remuneration and licensing based on the work, I think that cognizant of the transaction costs,

which are not insurmountable but are real and there, and additionally, the fact that artificial intelligence is nascent and may not yet, underline yet, have a clear path towards commercialization and ability to generate a source of revenue that can easily be extracted.

And balancing, I think, those shortcomings at present with the potential for artificial intelligence, again, across the economy, you know, in a wider way and getting it off the ground, I think it may be worth considering, and I value everyone else's input, some type of alternative source of revenue in the initial phases of any type of licensing that may exist. A popular idea that comes around every now and then, particularly floated by Paul Romer, would be a tax on online advertising revenue, I think that's an attractive opportunity for a source of revenue, but something that I would like to encourage folks keeping in mind at least in the early phases.

MS. RUBEL: Okay. I'm going to go to Luc next because he hasn't spoken on this point, and then I'll come back around to James and Paul.

Luc?

MR. BOULET: Thank you.

Our position is that an AI engine cannot

capture the beauty of a photograph without initially copying the image. And ingestion of copyrighted works by AI developers without proper authority constitutes copyright infringement on a massive scale and is of great concern.

The reason why developers are seeking out copyrighted works is because professional photographers create high-quality photos that are exceptionally well-suited for AI ingestion, and in this case, they're considered so valuable because quality of input determines their quality of output.

And I just want to also add that the priorities of individuals using copyrighted materials for AI ingestion, you know, must not and should not take precedence over the rights and interests of creators. The AI systems should not be built on their backs without their consent. And we must not compromise longstanding laws and policies that protect the rights of copyright holders in the pursuit of developing AI technology.

MS. RUBEL: Thanks.

James? Oh, I think you're still muted.

MR. SILVERBERG: I really appreciate what Luc said, and I want him to be right about everything that he said, but I'm a little bit concerned about the

copyright infringement assumption that is part of the discussion about AI ingestion because there are fair use issues and other issues there which make that final conclusion problematic.

Additionally, even if there is a determination in these court cases that there's an infringement, I'm not really sure what the utility of that is going to be for individual authors and small businesses who won't be able to monetize or possibly even identify the use of their material on AI platforms, particularly not through expensive litigation.

So I think it leaves us in a situation where we're still looking for answers even if Luc turns out to be right about everything that he said, and, again, I hope he is and he may be, but we're still going to have a problem.

And to follow up on Zara, I hope I said your name correctly, to follow up on Zara's comment, unfortunately, the onus may not be on the AI platforms to account for what they're doing because, unless what they're doing becomes clearly illegal, they have loopholes through the Fair Use Doctrine or other means, interim copying doctrines and other legal niceties, to evade liability. And they can sit here

today and present the argument that they're compliant with the law and may very well be.

But our problem isn't whether they're legally compliant and whether they're violating the copyright law. Our problem is, what do we do in order to make sure that the Zaras and the Karla Ortizs of the world get paid for the use of their style, content, appropriation of their efforts, when the current copyright system sort of is not really well fitted to doing that the way I would like it to be?

MS. RUBEL: Paul?

MR. REINITZ: Thanks.

Yeah, so I also wanted to respond to Zara and also the problem that James is pointing out. And I think, you know, again, as I said in my opening remarks, I think, you know, one of the best solutions we can do to mitigate these risks is require transparency. We need to have obligations on the developers of these models to keep records of the content that they trained on so that Zara or any other creator can tell if their content has been used.

And, you know, we are very early on, but this is something that we cannot go back and redo. We need to have these transparency obligations now so that there isn't just the ability to basically sweep

this stuff under the rug, and as the legal issues are decided, we need to preserve, basically, the evidence of what's going on.

MS. RUBEL: Alex, I'm going to give you a chance to jump in here.

MR. RINDELS: Thanks. Yeah, I'd say two points.

If society decides that we should pursue some type of remunerative system for this, I think being able to attach copyright protection to the end works themselves would actually be a pretty efficient way to collect fees for that remuneration. So a lot of our end users are constantly asking us whether the images they generate can be protected by copyright, and if they had to file an application like everyone else and that's part of that, meaning they have a duty to disclose that AI tools assisted in their generation of the content, maybe there's some sort of a fee that attaches to that.

My second point is I think it becomes very, very difficult for a remunerative system, although there are people much smarter than I who could probably determine where and into what portion of the fees, who they go to.

One, some AI models themselves aren't

storing the images so that when somebody puts in a text prompt, it's not going and retrieving an image and outputting it or some variation of it. The models themselves temporarily notice patterns in the images and in those patterns create an algorithm, and that algorithm overall is what the model is. And when you put an input for it to generate something, it simply provides a reasonable approximation of what the output should look like based on your text input.

So, for that, in that case, it would be very difficult to determine whose input was used to produce your output, so I think that would be difficult for attributing the fees as well.

MS. RUBEL: Yes, and we certainly heard that from Stability AI earlier this afternoon as well.

Luc, you're going to be the last voice on this question, and then we're going to move to talk a little more about the Copyright Office's registration guidance and related issues.

MR. BOULET: Thank you.

And I would just say to James's earlier point, that really leads me to the logical conclusion that there has to be an open AI system, that there has to be a transparency element added for the public market. And I would just add that just because it's

difficult to have transparency doesn't make it impossible. And, frankly, just because it is difficult, that is the road you must travel in order to use copyrighted content.

MS. RUBEL: Okay. Thank you.

Last question I want to pose to everyone, we have heard, the Office has heard in response to the guidance we released recently that, and I think Heather started making some points earlier in her responses along these lines, that there's possibly some things about the technology or how the users are interacting with the AI technology that the Office doesn't fully understand or appreciate.

So I'm interested in any thoughts you have about what the Copyright Office should know about how the AI systems generate content, what the participation of users might be in different AI models, and any other feedback you have about the guidance that the Copyright Office released.

Daniel, why don't we start with you. Daniel Gervais.

MR. GERVAIS: Thank you.

So, first, as I said in my introductory comments, I support the guidance. I think what's going on -- so, you know, machine learning is almost a

synonym of AI these days, and so, you know, machines learn when they produce output. There's a report online I wrote for the European Commission a couple years ago. I looked at every case where people said, look, the machine created this. And at least as of two years ago, there wasn't a single case where there wasn't substantial human edits.

But it's called machine learning for a reason, is that the machine will learn those edits, right? And as time passes, there will be more and more productions that have what I call no human cause really. And I think it's perfectly the right decision to say that doesn't have copyright. It would be weird to say you produce something, which is code that produces something, so you have copyright on the code, no question, right, if it’s human written. But whatever the code produces, you also have copyright on that. We don't have that. We've never had that, right? And this is, I think, a line that should not be crossed.

So what I think the comment might be would be there are ways in which humans collaborate. The doctrine’s very clear. When you take a public domain picture that's not protected anymore and you do something to it, we know, we can test for whether --

and you can test for whether there was enough done to it to be registrable as a new work, right? So the test is well known. It's just the technology allows humans and the machine to collaborate, as Alex and others mentioned.

So I would absolutely urge you to keep the guidance. Just in terms of its application, yes, there may be some ways to dig deeper into technology, but both the transparency and the exclusion of non-human works, I think, are very important.

MS. RUBEL: Heather?

MS. WHITNEY: Thanks. So just on the guidance, a couple of just points very quickly.

So one is, as I mentioned, and I think a few people mentioned on the earlier panel, it's not clear when you're reading the guidance what the relationship is between the analysis there and the analysis that the Office presented in Kashtanova letters because, as I mentioned before, that letter is completely dominated by a discussion of predictability, and the inability to predict what the output would look like was basically fatal.

But then, in the guidance, there's literally no mention of predictability at all. And while it's possible that the Office has decided that

predictability is no longer part of it, it still links to the Kashtanova letter as guidance to artists on its website, and I have heard from many artists through Kashtanova that they're just confused about what that really is supposed to mean.

The second thing is that the guidance itself has -- if you look at what the tests are for authorship within the guidance itself, they lead to different outcomes depending on which one you choose to look at. So, for instance, there's a lot of mention about Burrow-Giles with the idea that you are the author if you are the thing, basically, it owes its origin to you. And that's, like, a pretty easy thing to satisfy, and that's what's used a lot, I think, in photography. You click the button or you set up the thing and click a button and now you get the copyright. It would be clear that you would have a lot of copyrightable works if that were the test.

And then later you also say that basically you won't register works that are produced by machines or mechanical processes that operate randomly or automatically without any creative input or intervention by a human author. Again, that's like a pretty low standard, any creative input or intervention. But then, at other points, you talk

about this idea that you have to have ultimate creative control, that it's you are the one who is contributing all of the traditional elements of authorship.

So I think that that is a difficult thing to understand how these things are supposed to work together. And so, in the future guidance, it would just be helpful to sort of make clear what the relationship between those is.

MS. RUBEL: Thank you.

James?

MR. SILVERBERG: Yeah. I really have to echo Heather's sentiments on this point. You know, I think the Copyright Office has the guidance completely correct on the fact that non-human authorship is not copyrightable and is completely correct on the point that human authorship is copyrightable.

I think where the guidance, in my opinion, is a little bit off is something that Heather also pointed to, I think, which is that I think the guidance suggests that a substantial contribution needs to be made in order for the human authorship to be copyrightable. But, under the law, I think the standard for copyrightability of a work has a somewhat lower threshold than substantial contribution.

And so I think the devil really is in details. And I think greater clarity could be given. I think maybe Daniel suggested that examples be given in order to help clarify what you say in order to render the subject matter of your application copyrightable or not.

And just one quick separate point. I think it would be beneficial for the Copyright Office to have a policy where, in circumstances where there's a bulk registration of multiple works, the author can make a singular statement about what the copyrightable or non-copyrightable components are in all of the works so that a photographer registering 750 works doesn't have to do this 750 times.

MS. RUBEL: Thank you.

Luc?

MR. BOULET: Thank you.

I just want to state that the U.S. Copyright Office should not, from our position, and does not have the capacity to be engaged in investigations into the boundaries of what is disclaimed as AI-generated and whether or not there was sufficient human involvement in each case.

The current registration process for works with AI contribution creates confusion around what AI

material should be disclosed in a registration application, and then the guidance applies obligations to disclose AI-generated works without drawing clear lines around what those are. And we don't want the registration process to become more burdensome because the Office is launching investigations into canceled registrations.

And the cost of registration is already high with confusion already taking place for many artists, and that standard application severely limits the capabilities of visual artists to register their works in group registrations, leading to greater costs and higher barriers for individual artists.

MS. RUBEL: Thank you.

Zara?

MS. VARIN: Thank you.

So, before all of this, in my stress prep, I had the opportunity to watch some of the videos that were featuring other folks from the U.S. Copyright Office, and there were two issues that were cited by Shira Perlmutter, I believe. One is authorship, and the other was ingestion of copyrighted works from machine learning. I've already touched on the latter, so I want to go back to the authorship point, which a couple of other folks have kind of alluded to and

mentioned here.

I think an important distinction is, at what point does human involvement cross that threshold into authorship, right? And, right now, we don't really have a clear set of guidelines for what constitutes that, especially as it pertains to AI-generated pieces.

I really appreciate the way that Creative Commons phrases that copyright law's fundamental purpose is to foster human creativity. Copyright helps protect folks' creative works while ensuring there are clear distinctions and guidelines for what constitutes a copyrightable work. I think, with the growth in this area, there's a great deal of potential to revisit and redefine aspects of what those distinctions even entail.

And as other folks have mentioned, we're not necessarily going to be solving all of this, and that's not entirely on the Copyright Office, but you're setting precedents, and I think it's important to consider that as it's going to ripple through every other aspect of every single creative field and all parts in between.

I think that there's also, whenever registering stuff, things that I would want to

consider seeing. The use of any and all mediums, including generative AI, must be disclosed, which I think that that's already part and parcel of the process, but it needs to be, I think, more distinctively and clearly communicated. It's digital work. This part kind of sucks for some artists, but you can record and maintain a full record of the process that goes into creating something.

I regularly screen share or share a process. I use a program called Procreate that can start to finish share what the process looks like on pieces. That is an option that digital artists have and I think is something that, if not necessarily that but something similar, could and should be considered when evaluating the degree of human authorship for any sort of AI-generated works.

I think there's another factor here that we've skirted around, and that's fair use. The fourth factor in particular is pertaining to the effect on the potential market, and that's of extreme concern. If somebody's art style becomes a popular prompt fodder and the Internet is inundated with countless generated images in that artist's style, how does that impact the artist? The short version is it sucks. And I've seen many people that have been directly

harmed and have lost work because of the lack of really distinct enforcement and ethical development around these AI tools.

So I could keep going on the list, but I know that we're running out of time, and I want to give other folks a chance to speak.

MS. RUBEL: Okay. Thank you.

I think we'll hear last from Paul.

MR. REINITZ: Great, and I'll try to be brief. You know, I just want to say this has been a really good discussion, and I think that really good points brought up, especially on this issue.

I also would like to commend the Office on putting this guidance out there. I know it's really hard to get your hands around it. And, you know, one of the things could have just been ignore it. And, you know, I think that you've done a really great job in trying to get it out there. Is it perfect? No, but, you know, it will get there.

I think one thing in terms that could be helpful for you is to understand that it's not necessarily all or nothing when using generative AI. It's not like I have created an image and it's completely generative AI. There's also applications of it where you could be considered using it to, like,

modify an existing image. And I think that, you know, that needs to be taken into account.

Yes, the disclosure still needs to exist. And, you know, to simplify things, I think that there should be disclosure anytime there's generative AI used. But, you know, the line might be different. For instance, if somebody, you know, in Photoshop, some of these tools are already using AI, right? And if someone is using AI to, say, correct sharpness or color, and that's something that has traditionally been okay in part of the creative process, you know, using other tools, you know, that needs to be considered, that that's different than creating a generative image just from a prompt.

MS. RUBEL: Thank you. And thanks to everybody. This has been a really productive session. We appreciate all of your comments. And I like that we've styled this as a listening session. Really, the purpose was for all of us to be able to listen to one another, and thanks to everyone for sticking with us and sharing your perspectives with us as well.

I'll pass it over to Mark briefly.

MR. GRAY: Great. Thank you, everyone. So this concludes our second panel. Thank you very much for everyone on the panel for this session.

(Whereupon, a brief recess was taken.)

MS. MANGUM: Welcome back, everyone, and good afternoon. My name is Jalyce Mangum. I'm an Attorney-Advisor here in the Office of the General Counsel.

We're going to begin the final session in just a few. But, first, for those of you who are just joining us, I've got a few Zoom housekeeping points.

If you're joining this session, but you're not a speaker for this particular session, please keep your camera off and your mic on mute.

We are recording this session today, and the recording will be available on our website in a few weeks. The transcription function is also activated for anyone who wants to follow along that way.

In this session, we're going to ask each of our speakers to give brief remarks on the subject of artificial intelligence and visual art. Each person will be limited to two minutes, and I will be watching the time to keep us moving along.

We'll call on the speakers in the order listed on the agenda, and we're going to start first with Tom Lockley at Grey Owl Audio.

So, if, Tom, you're on, you can get started.

MR. LOCKLEY: Awesome.

So, before I start speaking, I would just like to thank USCO as well as the panelists for coming here today and speak on this issue. AI and its role in art is a deeply complex issue, and it's good to hear from all stakeholders involved.

My name is Tom. I'm a YouTuber, writer, educator, and investor. Over the past year, I've had the chance to participate in a number of AI art communities, including Midjourney and Stable Diffusion. I've also created an essay called *The Defense of AI in the Artistic Fields* in which I explore technologies like diffusion models and CLIP, along with the surrounding regulatory contexts involving them.

In my time here today, I'd like to share an adjacent though nonetheless important thought. AI is a powerful tool for equity and expression in the arts. Creators who lack access to expensive tools in studios or who are unable to produce art in a traditional manner due to illness or disability can rely on AI to provide alternate and often more manageable forms of expression.

I would know. Something I didn't mention in my introduction is that I have ulcerative colitis, a disease in which a person's immune system attacks

their digestive tract. This can put me out of action anywhere from weeks to months. It also manifests in other ways, from joint pain to a sometimes insurmountable fatigue. I can get around half the issue with tools like scheduled posts and videos, but AI fills the other half. I use it for thumbnails, book covers, character concept art, and branding. Yet the amount of human direction that goes into my work shouldn't be discounted. The process does not end with one prompt, and it can take anywhere from hours to days for me to make a piece that matches or evolves from the initial vision I had for it.

I'm not alone in this. In conducting research for my essay, I had the pleasure of speaking with a number of creators who work through their disabilities via AI, some of them for love of the field, while others use their art to generate income. This regulation of this nascent community can lead to damaging outcomes for those who participate in it. A lack of protection for AI artists unfairly penalizes them by opening the door to infringement and art theft by larger players who have both financial means and business motivations to lock up the industry in their favor.

Regardless of the outcome of today's

conversation, lessons we take from it and the precedents that we eventually set will come to define the next several decades. Let's get it right and make a system that works for all types of creators.

Thank you.

MS. MANGUM: Thank you, Tom. We really appreciate your comments and your perspective. That was really interesting.

Next, we've got Matthew Cunningham from Cunningham Concept Design.

Matthew, you're on.

MR. CUNNINGHAM: Hi there. Thanks, everyone, for having me and thanks for the great presentation so far. It's been really enlightening listening to everybody's angle on the whole thing.

My background, I'm a concept designer in the feature film and television industry. You may know my work from shows like *Star Trek: Picard*, Season 3, Amazon's *Citadel* that just recently was released, as well as a couple of *Godzilla* movies. I also work as a car designer for companies like BMW, Toyota, Hyundai, you name it. And I'm a former labor union leader with the Art Directors Guild. And I also teach design at the Art Center College of Design in Pasadena, which is the top design school in the world arguably.

My reason for being here is I'm part of a larger coalition based in Los Angeles, and what our concerns are fundamentally have to do with the copyright violations that are being employed by certain companies who are scraping the data and identity of artists. My concern is multi-pronged but I would say primarily is for the artists who are immediately being impacted economically, but also for successive generations of artists and also for the generations of past artists, where I think quite a lot has been discussed regarding the ability to imitate the style of artists.

This has very immediate economic impact, but I think also, through the historical lens, going back and, you know, if you visit a museum, you'll see some works of artists that are quite prominent, but these might become muddy waters in due time given enough of an evolution with the technology.

So I'm just simply here to speak for artists in both organized labor, unofficially, but mostly for myself and my colleagues in the film industry.

MS. MANGUM: Well, thank you, Matthew. We appreciate you being here and for you sharing your thoughts.

Next, we have Brian Frye with the University

of Kentucky College of Law.

Brian?

MR. FRYE: Thank you for inviting me to speak. I'm Brian L. Frye, Spears-Gilbert Professor of Law at the University of Kentucky College of Law.

I think we're asking the wrong questions about AI and copyright. Everyone is asking whether a copyright protects AI-generated works and whether training an AI algorithm infringes copyright. The obvious answer is no and no. Copyright only protects works created by people. AI doesn't even create works. It generates content, which we consumers interpret as works.

Many years ago, Roland Barthes predicted the death of the author, and AI has written the author's obituary. Likewise, training and AI algorithm doesn't and shouldn't infringe copyright. AI algorithms don't copy works, they merely catalog rhetorical conventions and then deploy them to create conventional content.

We should be asking what AI can tell us about what copyright should protect and why. Copyright can only protect creative works, but courts and the Copyright Office have struggled to define creativity. Maybe AI can help.

An AI algorithm is essentially a nonsense

generator designed to produce banalities. In other words, AI is uncreative by design. An AI algorithm is a machine for regurgitating conventional wisdom. Indeed, we are amused when an AI hallucinates and fails to satisfy our pedestrian expectations. But we can be just as boring as any AI, and there's no point in copyright protecting banalities.

Maybe AI can help us limit copyright to works that are actually creative. It's easy, just ask AI to evaluate the creativity of works created by people to determine whether they deserve copyright. No one knows a fake like a faker, and AI is designed to identify banality. That's what makes it a killer app.

We don't know how to identify creativity, but AI can tell us what isn't creative, and maybe that's good enough to tell us what is creative, if anything.

Thanks again. I'm on Twitter @brianlfrye.

MS. MANGUM: Thank you, Brian. Really interesting comments, and thanks for the Twitter tagline. We appreciate that.

Next, we've got Nettrice Gaskins, who is a freelance artist.

MS. GASKINS: Thank you, and thanks for

having me.

In 2017, the Andy Warhol Foundation launched a preemptive lawsuit against photographer Lynn Goldsmith, who captured photos of the late musician Prince in 1981 for *Newsweek*. Warhol was later commissioned by *Vanity Fair* in 1984 to produce a pop art recreation of one of the images after licensing it for $400. However, Warhol continued to use the image for his portfolio, taking his own spin on the original photo.

In response to being sued for her own copyrighted photograph, Goldsmith filed a countersuit, and she lost. New York State District Judge John G. Koeltl ruled in favor of the Warhol Foundation. Koeltl argued that though Warhol used Goldsmith's photograph as a reference image, he removed nearly all of the photograph's protectable elements. Thus, Warhol did not violate the photograph's copyright.

Like Warhol, my clients license images from -- for me to use as references for the creation of AI art. Mimicking these references is not a goal. Reimagining, recasting, remixing, and restyling them are.

My first commission was in 2019, and I've been using a variety of AI tools to make art since

2016, long before prompt-based tools. I use AI to produce variations on reference images, and I use text prompts with keywords that the AI recognizes in order to generate unique visuals. I use image editing software to revise visual elements as well as layer images I've done, and I've done this since the early 1990s. Just one of my AI artworks merges multiple art styles and techniques, posing subjects in ways that go far beyond the original image references.

For me, as someone with a traditional fine arts background, AI broadens what is possible for artistic production, and these new tools have enabled me to become an art director, maker, and curator of my work. This puts me in a pipeline that previously excluded me. Thank you.

MS. MANGUM: Thank you, Nettrice. We appreciate you sharing your experience. Really interesting remarks.

Greg Hopwood was unable to join us, so we're going to skip right on to Phuc Pham from the Freelance Solidarity Project.

MR. PHAM: Hi, everyone. My name is Phuc, and I am here today to lay out my perspective as a photo editor and a photographer, as well as to represent conversations I've been having with

freelancers I organize alongside with in the Freelance Solidarity Project, which is the Digital Media Division of the National Writers Union.

Editors of born-digital works that are primarily distributed online, freelance digital content creators, not just those working in the visual arts space, like photographers, animators, illustrators, and graphic designers, but also in every other corner of the media industry, are among those the most likely to be impacted by generative AI technologies. And what I'd like to focus on during my time right now is just a simple example of, like, the sort of innumerable images that train these AI systems to generate their output.

You know, these companies that develop AI engines largely obscure details of how their technologies work. As such, it's nearly impossible for an individual artist to seek recompense for their copyrighted work showing up in these training data.

Those who wish to remove their images from these data sets have found tools such as "Have I Been Trained?" to even determine exactly which images were used to train these systems. You know, following that discovery, an artist would have to register their works in order to even defend a copyright claim. You

know, for example, a photographer whose digitally published works were scraped to train these systems, those works could number in the hundreds of thousands. And at $55 to register 10 images, that quickly can become an exponential amount of money.

This process is not only economically draining, could be, but would also demand time and focus away from an artist who definitely won't have an entire corporate department to sort of support these copyright claims like you're seeing with, like, Getty Images or a lot of these other bigger outfits that are, you know, bringing forth lawsuits for this type of usage of their work.

And beyond the impracticality of defending individual copyright claims, as someone who works with photographers to commission original works and who organizes to improve industrywide working conditions, I'm simply disturbed, you know, by the implications of these systems and how quickly they are sort of can quickly replace the human spirit and the novelty that working artists bring to these types of works. It seems to me like an untenable future for visual artists who are born-digital and otherwise.

MS. MANGUM: Well, thank you. We appreciate your remarks. Really interesting.

Next, we've got Ankit Sahn from Ajay Sahni Associates.

Ankit?

MR. SAHN: Thank you. Good afternoon, esteemed speakers and members of the United States Copyright Office.

By way of an introduction, I'm Ankit Sahn. I'm an IP lawyer based in India. And I'm the owner of the RAGHAV AI painting tool. I filed what was perhaps the first application at the USCO and the CIPO in the Indian copyright office where a human and AI were identified as co-authors back in 2021. I'm grateful for the opportunity to present my views today.

As we continue to rely on AI tools to produce works of art, music, literature, and other creative outputs, creators must be assured that their works will be protected under copyright law. Denying copyright protection to AI-generated outputs could result in a chilling effect on creativity and innovation.

Just as when cameras were invented, humans moved up the value chain by becoming photographers, using AI-based tools to create demands human beings to move up the value chain once again. We must recognize the increasing role that AI is playing in the creative

processes across all industries while also acknowledging the value of human creativity and talent and thus maintaining that balance.

The definition of human creativity in the context of copyrightability, therefore, has to be reconsidered. Works created with the assistance of an AI-based tool could be considered as a separate category of copyright, perhaps with reduced duration and scope of protection to incentivize human effort and promote innovation.

As Mahatma Gandhi said, “The future depends on what we do in the present.” Protecting AI-generated outputs under copyright law is a crucial step in this direction, and by providing necessary protection and recognition, we can continue to foster innovation, creativity, and ultimately benefit the creative industries, as well as creators of these AI-based tools.

If AI-assisted works are not protected, on the contrary, it could lead to creators suppressing the fact that they use the assistance of an AI tool to create a work, which would in turn be unfair to creators who have utilized AI tools to enhance their creative outputs and in any case would not be reflective of the correct factual position on the

copyright register on record.

Training on proprietary data, as many speakers pointed out today, is one of the key issues. Collective or possibly compulsory licensing of data could be explored as a possible solution.

In conclusion, copyright law must evolve with the times, recognizing the role that AI is playing in the creative industries. Protecting AI-generated outputs under copyright law, I believe, is a crucial step in this direction.

I am grateful once again for this opportunity. Thank you for your time and attention.

MS. MANGUM: Thank you, Ankit. We appreciate your remarks.

Next, we're going to move on to Jeffrey Sedlik from the PLUS Coalition.

MR. SEDLIK: Thank you.

MS. MANGUM: Jeffrey?

MR. SEDLIK: Thanks. Yes, I'm Jeff Sedlik. I'm President of the nonprofit PLUS Coalition at PLUS.org. We're currently collaborating with the IPTC on establishing metadata fields to communicate information about works that may incorporate AI and to provide for the expression of permissions and constraints on the use of AI or use of visual works

for AI training and for generative AI.

I'm also the former president of the American Photographic Artists and a photographer, filmmaker, and professor at the Art Center College of Design, where I teach copyright law and licensing.

Like other visual artists, copyright is at the core of my business. To sustain it and to support my family, I rely on revenue from licensing my works throughout the life of my copyrights. I offer licenses of my work to all manner of clients for all manner of media in all manner of purposes, including, importantly, artist reference use, the use of my works by other artists to adapt my work in new derivative works. The exclusive right to adapt a work is often overlooked but is, in fact, one of the core rights of copyright, no less important, no less critical than the exclusive right to reproduce, distribute, display, or perform a work.

The practice of offering, seeking, and granting artist reference licenses has been in broad use for more than a century, and the use of visual works as AI image prompts falls squarely within the definition of artist reference use. In fact, many of us have offered up our works for paid licensing for AI training and artist reference use in AI generative

works for years.

This fact, combined with the fact that many stock photo agencies are now routinely licensing their works for AI training, establishes that a marketplace for licensed use of visual works for AI training exists and is therefore usurped when works are used for AI training without authorization. Meanwhile, thousands of copies of my work are included in the LAION database and other databases of, collectively, billions of images used to train AI systems, many without attribution. The more I license my works, the more copies are included in the training sets, and the closer the appearance of generated AI works to my original creations. But, as a professional, I must license my work in order to sustain my business. It's a Hobson's choice.

Some suggest that contrary to the core provisions of copyright law, anything on the web is fair game and that if artists don't want their works used for AI training or image prompts, we should remove our works from the web. But the web is the primary medium for licensed use of our creations, and we necessarily depend on the web to monetize our works.

It's also very important to recognize that

many, perhaps the majority, of copies of our visual works displayed on the web are infringing copies made without our knowledge or permission. These infringing copies are then blindly scraped for inclusion in databases like LAION, which is arguably a database built on infringements.

MS. MANGUM: Jeff --

MR. SEDLIK: In addition, we must not forget that the widespread aggregation of visual works --

MS. MANGUM: I'm sorry. We're going to have to -- I'm going to have to interrupt. It's been a little over two minutes.

MR. SEDLIK: Okay.

MS. MANGUM: But feel free definitely to submit your remarks when there are opportunities later to do so. Thank you so much for coming.

MR. SEDLIK: No problem.

MS. MANGUM: Next, we're going to move on to Patricia Sigmon, who is an artist and art director.

MS. SIGMON: Thank you so much. My name is Patricia Sigmon, and I am a professional artist and art director who has primarily worked in 2D and 3D animation. I will be speaking on my personal experiences with AI-generated images and my concerns about how they have compromised the ability of artists

to retain work, as well as how those problems may multiply in the future without more regulation.

At a previous position as an art director, my team of artists realized that character design references we had been given were AI-generated. Company leaders did not consult me about the use of AI beforehand. The team came to me with several concerns, including fears about their jobs being replaced, ongoing legislation against AI, ethical concerns, and the poor quality of the images they were given. While AI-generated images seem like they would be an easy shortcut for gathering references, they often include nonsensical details that do not translate well to designed outfits in 3D models.

A major concern was also that work produced by artists for the company could be used to train AI without their consent. There has been ample evidence from previous panelists that the goal of many people who use this technology is to imitate the styles of existing artists. It is not a stretch to imagine that without legal guarantees against this, a company could hire artists, train AI with their work, and eventually use it as a replacement for them as the output quality improves.

I compiled all these concerns and statements

from the artists and brought them to the heads of the company. I was fired two days after that meeting.

I am part of the first wave of artists affected by major companies adopting AI image generators. My goal is not to totally disavow the use of this technology or prevent individuals from using it. I understand that the march of progress cannot be stopped and that AI image generators have uses outside a professional setting. The problem is that companies will always be more concerned with their bottom line than the people they employ. It is imperative to protect the ability of artists to earn a living from a craft that they've perfected through a lifetime of practice without that work being cannibalized.

Thank you for your time and for inviting me to this discussion. I think it's really important to have, and I've appreciated hearing everyone's perspectives.

MS. MANGUM: Thank you so much for sharing your experience with us, Patricia. We really appreciate it.

Last but not least, we've got Delanie West from Be Super Creative.

Delanie?

MS. WEST: Thank you, Patricia. Thank you

so much for that perspective. I appreciate you too.

Thanks for the invitation to share my perspective. I'm Delanie West, Founding Creative Director of Be Super Creative. I have 30 years experience in creative marketing, business, and brand development, and I've led creative and product development for U.S., European, and Japanese brands. I serve in a leadership capacity at Women in Toys, Licensing & Entertainment, as well as the Graphic Artists Guild, and other creative industry organizations, but the views I express today are my own. I'm here today to share a perspective of a small business owner doing the work of developing creative for product development.

As a value creator who celebrates this new technology, I'm also concerned about the impact of AI on consent, credit, and compensation. AI can revolutionize the creative process, but we must ensure ethical methods to generate outputs that respect intellectual property and current law.

I've long worked aside legal IP counsel to protect and defend copyright and patents, and the same process is necessary for ethical AI implementation. I ask that we embrace AI while being mindful of ethical considerations and adhere to the current and future

intellectual property laws.

The use of AI in product development has tremendous potential for boosting creativity and productivity, but we must approach it with care and responsibility to ensure practice for consent, credit, and compensation.

In summary, AI has been a game changer for product development timelines for me, but we creators need the guardrails that enable users to respect the rights of all parties involved. Thank you.

MS. MANGUM: Thank you so much. We really appreciate everyone who shared their perspective and their experience.

I'm going to turn it over to Mark to close us out.

MR. GRAY: Thank you very much, Jalyce.

And to echo Jalyce’s remarks, we truly do appreciate, both I, as well as the rest of my colleagues here in the U.S. Copyright Office, we do appreciate you taking the time to talk to us today and to share all of this information. Of course, we're going to continue to think about all of these issues and all of the things that you've told us as we work on our initiative examining issues of copyright law and policy and the intersection with artificial

intelligence technology.

So looking forward, we have two more listening sessions on the calendar. Our next session is on Wednesday, May 17, which will be focused on audiovisual works, including movies and video games. Our final session is going to be on May 31, focused on musical works and sound recordings. The audiovisual session, unfortunately, is no longer accepting signups. We are wrapping up the selections for those panels now. The music session remains open until May 10. So, if you are interested in speaking about musical works or sound recordings, you can find more information on our website at copyright.gov/ai.

As a reminder, we will be making video recordings of this and of the other sessions available to the public on our website. We're aiming for about a three-week turnaround on those.

And, as we've said before and as we will continue to remind people, this is not the final word. This is not the final chance to speak to us. There are many more chances coming down in the future. We look forward to hearing from you and thank you very much for spending your time with us today. Thank you.

(Whereupon, at 4:00 p.m., the listening session in the above-entitled matter adjourned.)

CERTIFICATE

CASE TITLE: Copyright and Artificial Intelligence
Visual Arts Listening Session

HEARING DATE: May 2, 2023

LOCATION: Washington, D.C.

I hereby certify that the proceedings and evidence are contained fully and accurately on the tapes and notes reported by me at the hearing in the above case before the Library of Congress, U.S. Copyright Office.

Date: May 2, 2023

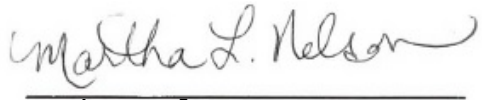

Martha Nelson
Official Reporter
Heritage Reporting Corporation
Suite 206
1220 L Street, N.W.
Washington, D.C. 20005-4018