



Creative Economy and Generative AI - October 4, 2023

Madeleine Varner:

Hello and welcome to the FTC's roundtable discussion on the creative economy and generative AI. My name is Madeleine Varner and I'm a Senior Technology Advisor in the Office of Technology with a background in investigative research. Today we'll be discussing artificial intelligence tools that can output content on command, also known as generative AI. We've gathered professionals from a broad range of creative fields to discuss how these tools are reshaping their respective lines of work and how they're responding to these changes. Participants today include screenwriters, actors, programmers, editors, musicians, models, and more. Thank you all for being here and voicing your unique perspectives.

Before we begin, please note that the FTC is recording this event, which may be maintained, used, and disclosed to the extent authorized or required by applicable law, regulation, or order, and it may be made available in whole or in part, the public record, in accordance with the commission's rules. In practice, this means that we'll be sharing a recording of this as well as a transcript on the event webpage after this concludes.

Now, we're going to be hearing some opening remarks from Chair Khan. Chair Khan?

Chair Lina Khan:

Great, thanks so much. Hi everybody. So great to be here with you all today. The FTC is hosting this roundtable to hear directly from creators about how generative AI is affecting your work and livelihood, and I'm just so glad that we have this opportunity to be engaging in a very timely and important discussion. I wanted just to provide some backdrop for what the FTC's role is here. The FTC was created 109 years ago against the backdrop of an industrial revolution that had delivered enormous technological progress, but had also concentrated power and control in the hands of a few. For example, the advent and expansion of railroads meant that farmers could now move their wares across the country dramatically expanding the number of markets that

they could reach, but consolidated control over the railroads also meant that farmers were often at the mercy of a single company that had the power to arbitrarily hike rates and set discriminatory terms.

And the anti-monopoly movement and antitrust laws that followed were fundamentally about securing terms of fair dealing across markets to ensure that Americans be it farmers, small businesses, workers, or consumers could engage in commerce on fair terms no matter what the technological advances of the day were, and so Congress created the FTC to enforce these rules of fair competition, and a key part of our mandate is ensuring that these core principles continue to apply even as technologies and business practices evolve. Lawmakers actually explicitly gave the FTC the authority to conduct detailed market inquiries so that we could track and keep pace with new market developments.

Today as we see growing use of automated systems, including those sometimes marketed as artificial intelligence, we again want to make sure that we're keeping pace, that we're fully understanding how these new tools could be affecting people on the ground in positive ways, but also potentially in harmful ways and potentially unlawful ways.

One particular area, of course, is generative AI, and I think we're all waiting to see what the full impact of generative AI will be, but there's little doubt that this technology could really transform how we live and work and communicate. And at the FTC, we're in particular looking closely at how some of these tools could turbocharge fraud, entrench the dominance of firms that control the necessary raw inputs like cloud services and computing power, and potentially lock in business models that incentivize the endless surveillance of our personal data. As these technologies involved, we're fully committed to using all of our tools and authorities to maintain open, fair, and competitive markets without tolerating exploitative or deceptive business practices, and we think it's extraordinarily important to not assume that there's some inevitable endpoint about how these technologies will be used, but instead recognize that the laws and policies against which these new tools are introduced will significantly shape who benefits and who is harmed. That's really what we see as the FTC's rule.

We've been very clear that we want to make sure the market understands that there's no AI exemption to the laws on the books, so all of the laws that already prohibit unfair methods of competition or collusion or discrimination or deception, all of those laws still entirely apply. I know that generative AI in particular poses a unique set of opportunities and challenges to creative industries. We've already heard significant concern about how these technologies could virtually overnight significantly disempower creators and artists who may watch their life's creation be appropriated to models over which they have no control. Really, really eager to be hearing from all of you.

Over the summer, I had a chance to meet with some of the writers who were part of the WGA and picketing and really understand from them what their concerns were. I was really thrilled to see that they were able to reach an agreement that includes certain protections for writers from how AI could be deployed.

Very much recognize that this is a fast-moving dynamic situation, but it's clear that for enforcers and regulators to be keeping pace and understanding what's happening on the ground is going to be absolutely essential, and we really couldn't do that without you all, so thank you so much for taking the time to share your experiences and views with us today. Our Office of Technology, which we launched earlier this year, is really a critical part of this effort. We wanted to make sure we have the skill sets on board to help us understand how are these tools really working and know what's really going on. And they, in close partnership with the Bureau of Consumer Protection, Bureau of Competition and Office of Policy Planning, were essential to putting today's event together, so my deep gratitude to everybody who was involved with that as well. Really looking forward to hearing from you all. I know Maddy and others have teed up some key questions and then we can have a discussion. With that, I will pass it back to Maddy to kick it off.

Madeleine Varner:

Thank you, Chair Khan. We're going to turn it over to Commissioner Slaughter for her remarks. Commissioner.

Rebecca Slaughter:

Thank you, Maddy, and thanks to Chair Khan and all the FTC staff members who've worked to make today's event possible, and thank you in particular, to all of our roundtable participants, artists from such a wide range of creative media. Today's discussion involves two very different concepts, art and technology that are connected by an essential input, humans. Art is fundamentally human. Humans may use technology to assist in creating art, but something cannot be art without human input. Technology is, by definition, not human. Yet technology including generative artificial intelligence requires human intelligence. While humans may endeavor to make generative AI that is ever more intelligent, it cannot and will not replace human creativity. The value of creative arts to society is so fundamental that it is enshrined in the Constitution. You all are the humans who have mastered a craft and you share it with all of us for the benefit of the public and society as a whole.

In return, such works of art may be granted copyright protection. Copyright provides your livelihood and the ability to continue to create and promote further creation and learning. Generative AI poses important questions and concerns about how copyright law and policy must be applied or adapted to continue to both protect creators and promote the useful arts, but copyright is not and cannot be the only tool to address the deeply personal concerns creators hold about how their works are used. The Writer's Guild of America has demonstrated the power of collective bargaining to secure important rights on how they will interact with, use, and be subjected to generative AI. Many but not all states have laws that provide rights to publicity, which may provide avenues for legal protection and compensation, and as the chair noted, the FTC's prohibitions against unfair and deceptive practices and unfair methods of competition apply to applications of AI just as much as they have to every other new technology that's been introduced in the market over the last hundred years.

These are powerful tools we can use on behalf of creators, workers, and consumers, and there may also be gaps in the law that need to be filled. As artificial intelligence processes, uses, and applications evolve, we cannot lose sight of the fundamental truth that technology is a tool to be used by humans. Humans are not and should not be used by technology. One of my favorite things about the work of the FTC is how it matters to real people in their real everyday lives, how it matters for real humans. I return to where I started, that art is fundamentally human. As the FTC considers how its tools can help fight fraud, deception, and unfair methods of competition related degenerative AI, it's critical to hear from stakeholders and that's why I'm so grateful to all of our speakers for their time and insight today.

Thanks again and I'll pass it back to you, Maddy.

Madeleine Varner:

Thank you, Commissioner Slaughter, and just to note, my remarks today don't necessarily reflect the views of the commission or any individual commissioner.

The Office of Technology is fortunate to be able to organize today's roundtable alongside staff from across the agency—I want to give an immense thank you to those in the Bureau of Consumer Protection, Bureau of Competition, and Office of Policy Planning for your efforts leading to today.

We look forward to hearing the lived perspectives of those from across different creative industries, ones that are actively being transformed by the development of new generative AI systems.

Our Office works to keep pace with emerging developments in digital markets, including in the rapidly involving generative AI space. Our software engineers, researchers, and practitioners use their deep expertise to examine the different layers of these technologies—including training data and infrastructure used to develop AI models and the models themselves in order to better understand how these technologies are built and the ways it might impact market participants and pose harms to consumers, workers, and small businesses.

In addition to understanding the technical underpinnings of generative AI, we root our internal expertise in the day-to-day lived experience of those who are most impacted. Shining a light on how emerging technologies impact people and communities is an important way for us to orient our law enforcement and policy work.

We know that generative AI relies on a critical input to function, large, diverse data sets of human content. One method to build such data sets is web scraping, often performed unseen, and without the knowledge of creators whose work is being collected. We have heard from individuals, some of which are participants today, that this dynamic deeply impacts those who share work online, particularly creative professionals whose livelihoods can depend on having a public portfolio or presence to attract customers. Today, we'll have the opportunity to hear directly from these people.

Today's roundtable is an important opportunity to ensure that we are assessing the full range of the impacts of generative AI on creative communities and help ensure that we are using our full set of authorities to tackle unfair or deceptive acts and practices and unfair methods of competition in these fields.

We're grateful to the creative professionals who have lent their time today to share their experiences, and we're looking forward to an engaging and insightful roundtable.

And with that, we're going to kick things off with Duncan Crabtree-Ireland, the National Executive Director and Chief Negotiator for SAG-AFTRA.

Duncan Crabtree-Ireland:

Thanks so much, Madeleine. Can you all hear me okay? I'm hoping so. Hi everybody. I'm Duncan Crabtree-Ireland. I'm the National Executive Director and Chief Negotiator of SAG-AFTRA, and we are the union that represents over 160,000 members who are the faces and

voices that entertain and inform the world, and I just want to say thank you for the opportunity to speak on a topic that's so important, not just for creative talent, but for workers across all industries.

As AI technology, and generative AI in particular, affects anyone who's concerned with consent and with protecting their own intellectual property rights. Generative AI, generative artificial intelligence, poses a threat to the livelihoods of many. But to be clear, we at SAG-AFTRA are not opposed to new technologies and we're not opposed to the existence or even the use of AI. AI in particular is an essential tool that is going to be widely used in the decades to come and it can have a variety of creative and innovative uses in the entertainment industry, and the implementation of AI we are looking to achieve would result in the technology augmenting human creativity rather than replacing it.

When used ethically and in a manner that recognizes intellectual property rights, AI can help people in their careers and can further opportunities. It can create access to employment for people with disabilities and those who would otherwise be prevented from pursuing work in the entertainment industry. What SAG-AFTRA is eager to do is to channel the benefits of AI into a future that's beneficial to our members, to workers in other industries, and to the public in general. The key is that the companies using AI technology must be required to get the informed consent of any individuals whose voice, likeness, performance, persona, or intellectual property is being used to generate content and companies need to compensate these individuals fairly.

Informed consent and compensation addresses a lot of the most important ethical questions attached to how generative AI works. It also would ensure that the hundreds of thousands of individuals who work in the entertainment industry will be paid for the part they play in creating content for these companies who are using human beings to create their profit.

It's important to understand that all AI generated content originates from a human creative source. No AI algorithm is able to make something out of nothing, and that human generated content that's used in the training data reflects real and substantial work and its intellectual property and it deserves legal protection. There's a bit of a double standard that can be seen happening in a lot of these conversations around AI. After all, if an individual decided to infringe on one of these company's copyright protected content and distribute it without paying for the licensing rights, that individual would face a great deal of financial and legal ramifications.

So, why is the reverse not true? Shouldn't the individuals whose intellectual property was used to train the AI algorithm be at least equally protected? An actor's brand is their voice, as is their likeness and their unique persona, and no company should be able to appropriate that and use it

however they wish without permission. What we are proposing is about keeping our world and our industry human-centered. AI and its algorithms must be here to serve us, not the other way around. Thank you.

Madeleine Varner:

Thank you so much. Next, we're going to hear from John August, a screenwriter and member of the Writer's Guild of America West Negotiating Committee. John?

John August:

Thank you. It's a pleasure to be here.

I'm here on behalf of the Writers Guild of America West, a labor union representing thousands of writers in film, television, and streaming series. Our members and the members of Writers Field of America East have just concluded a 148-day strike where artificial intelligence was a key issue. As Duncan just noted, our fellow artists at SAG-AFTRA are still on strike, with AI as a core issue for them as well, but the fight for protection over our craft and livelihoods doesn't stop at the bargaining table. While we have been able to achieve groundbreaking protection for writers, we need public policy solutions, too.

Obviously, copyright is an area of government scrutiny, both the copyright ability of AI generated work and the degree to which training AI models infringes upon copyright. WGA writers do not hold copyright to most of the scripts we write; those are works made for hire, so studios— our employers— hold the copyright, but through the power of our union, we have over the decades negotiated an assortment of contractual rights in the works we create, including the right to payment for reuse of our work. It's a good reminder that while copyright is important, it's not the end of the story when it comes to protecting artists.

The Guild's new agreement offers helpful guidance in thinking about future public policy on AI. Our agreement defines that AI is not a writer and the material it generates is not equivalent to human writing for purposes of our contract. That means that AI cannot rewrite us, nor can it compete with a human writer for credit and the associated financial benefit of that credit. Further, the studios now have to tell us if they are providing us with material generated by AI and then it cannot require us to use AI tools.

We won these protections because we're a strong union that successfully carried off a nearly five month strike. But we need to remember that most writers and most artists in this country don't have unions to protect them. It's best to think of writers and other artists as tiny businesses, each competing in the marketplace to sell their work. Writers and artists each develop a unique style, voice and brand in order to distinguish themselves. AI fundamentally disrupts that market in ways that could be devastating to the creative economy.

Large language models like the one that underpins ChatGPT have scraped massive volumes of data, including our words and our unique perspectives. This is theft, not fair use. Our works—protected by copyright and our own contractual rights—are being used entirely without our authorization, without any attribution or compensation. Right now, you could ask ChatGPT to write something “in the style of” a particular writer, and they would try to do that—appropriate the unique voice of a writer without that writer's consent.

As FTC Commissioner Alvaro Bedoya recently suggested, this could constitute an unfair method of competition. It is using stolen goods to undercut the price of a seller and create market confusion, and it's not a hypothetical. Right now, authors are finding AI generated knock-offs of their work published on Amazon. They're having to fight to get those fakes taken down, and protect their brands.

This form of AI appropriation may also have consumer implications. From electronics to organic eggs, consumers expect to be told the origin of a product and its authenticity. Consumers make choices based on that information. The same will likely be true with AI.

So, with our strike and this contract, the Writers Guild was able to win groundbreaking AI protections for writers. But it's important to remember that our deal only covers the film and television studios. Most of the real work in AI is being done by companies like Google, Facebook, and OpenAI with which the Guild has no contractual relationship. Public policy will play a crucial role in protecting our members.

In conclusion, AI does have the potential to assist the creative community, but only with the consent of that community. Without guardrails, AI poses a profound threat to writers and the integrity of our creative work. AI and its use raises major questions of intellectual property, transparency of competition, all of which require careful oversight to protect the creative economy.

Thank you for the opportunity to speak on behalf of film and television writers.

Madeleine Varner:

Thank you. Next we're going to hear from Neil Clarke, Founder and Editor of the Award-Winning Sci-Fi magazine, Clarkesworld. Neil?

Neil Clarke:

Hi, I'm the Publisher and Editor of Clarkesworld, a science fiction and fantasy magazine that has been publishing for 17 years. We publish stories not only from established writers, but also new voices from all over the world. These people represent the future of my field, and it is essential to maintain the avenues through which they can be discovered.

Like many of my colleagues, we do this by maintaining an open submissions process in which anyone can submit their stories for consideration. Not long after ChatGPT was released, we started noticing some unusual submissions in our queue and quickly realized that they were generated. It started small. A few in November, around 50 in December, over a hundred in January. In the first 20 days of February, it spiked and we received over 500. On the morning of the 20th alone, we received over 50 and the daily trends were indicating that we would double our normal monthly submission volume of 1100 by the end of the month.

This was unsustainable, so for the first time in over a decade, we closed submissions for something other than a software update. We needed breathing room to process what we had received and time to figure out some way to block, deter, or minimize these submissions. Even before this happened, we had a standing no AI policy. We knew these models were trained without permission on the copyrighted works of others, and it didn't sit right with us. It's also unclear that the people submitting these works actually own the rights they require. There are those who have argued that we should be considering these stories on their own merits rather than rejecting them for being generated. I can confidently state that they were among the worst stories we've ever received.

The problem is not quality at the moment, but rather the sheer speed and volume at which these works can be produced. The best way to describe them is to call them spam, but it's spam at a level we've never seen before. The generated submissions came from outside the science fiction fantasy community, typically from countries with a much lower cost of living and higher unemployment. We traced the source of much of this activity to YouTube, TikTok, and blog content hosted by side hustle con artists that would waive a stack of cash and make false claims

about earning riches from ChatGPT. Their intent was to earn advertising revenue or sell classes to those that tried and failed. They knew these tactics would not get people published. We happened to be on a list of a hundred or so magazines that were used as part of these schemes. Over the months that followed, copycats would refine their approach.

Several creating videos that specifically targeted Clarkesworld and even added directions on how to avoid detection. We reached out to YouTube and others in the hopes of having these videos taken down, but no one ever responded. I'm also the developer of the submission system that's used by Clarkesworld and a few other magazines. During the period we were closed, I spoke to professionals with experiencing credit card fraud prevention, network security, spam detection, and even some sympathetic people working on AI research. I evaluated many of the public and commercial detection tools and found their claims significantly overstated. The number of false positives and false negatives made them unusable. To make matters worse, I observed that foreign authors were far more likely to be incorrectly flagged as AI and an experienced eye remains the only reliable method of detection. When we reopened, the tools and techniques we adopted worked for about two months before the spans changed tactics and the levels shot up once again.

It's been back and forth since. Some days are good, many are bad. Unfortunately, we still have to review each of these suspicious submissions manually, but for now, we've been able to stay open. Between the time spent reviewing submissions and maintaining the software to control them, my workload has easily doubled. It's been exhausting, and this is only the early days.

Submissions are still being submitted manually by people. Once they start employing bots to submit these words, what we currently receive in a month may arrive in hours. Quantity and the lack of reliable detection tools will eventually break our ability to do our job. This is not the limit of my troubles with AI. Stories and art posted to our website have been utilized the training data for language models without our knowledge or permission. Anthologies I've edited and published have been pirated and included in the Books3 database used to train various models. This has concerned our authors and some are even withdrawing their work as we can no longer protect them. While it is likely that there are many positive uses for this technology, I don't believe that authors, artists, translators, narrators, and others should be required to sacrifice their work to improve or build these models. Progress can still be made responsibly without stepping on creatives. Regulation of this industry is needed sooner than later, and each moment they are allowed to continue their current practices only causes more harm. Their actions to date demonstrate that they cannot be trusted to do it themselves. Thank you.

Madeleine Varner:

Thank you. Next, we're going to hand it off to Bradley Kuhn, policy fellow at the Software Freedom Conservancy. Bradley?

Bradley Kuhn:

First, I thank the FTC for organizing this panel. It's admittedly humbling to be here among these key individuals from such a broad range of important creative endeavors. Folks will surely notice that I'm not appearing by video today, and I, again, thank the FTC for providing a method for me to join you today without requiring that I agree to Zoom's proprietary terms and conditions.

As a matter of principle, I avoid using any proprietary software, but in this case, it is not merely esoteric principle. Zoom is among the many Big Tech companies that have sought to cajole users into consent for use of their user data as training input for machine learning systems.

If consumers take anything away from my comments today, I hope they remember to carefully read the terms and conditions of all software platforms they use, as they may have already agreed for their own creative works to become part of the company's machine learning data sets.

I admit it may take you a week to read all of those terms, but it's sadly the only way you'll know what rights you've inadvertently given away to Big Tech.

The creative works that I focus on, however, is the source code of software itself. Software is unique among creative endeavors because it is so easy to separate the work that's created by humans, which is the source code, from the form of the work that's enjoyed day to day by consumers, which is the compiled binary. I'm an activist in the area of software freedom and rights specifically because I believe every consumer deserves the right to examine how their software works to modify, to improve, and to change it, be it altruistically or commercially. Free and Open Source Software, abbreviated FOSS, aims to create through licensing and other means an equal field for all software professionals and hobbyists alike and to grant rights to consumers so they have true control of their own tools.

For 30 years, our community has created FOSS and made it publicly available. Big Tech for its part continues to refuse to share most of its software in the same way. So as it turns out, nearly all the publicly available source code in the world today is FOSS, and most is licensed under terms that are what we call copy left, a requirement that anyone who further improves or modifies the work must give similar permissions to its downstream users. This situation led FOSS to become a canary in the coal mine of Big Tech's push for machine learning.

Hypocritically, we've seen Big Tech gladly train their machine learning models with our publicly available FOSS, but not with their own proprietary source code. Big Tech happily exploits FOSS, but they believe they found a new way to ignore the key requirements that FOSS licenses dictate. It's clear Big Tech ignores any rules that stand in the way of their profits.

Meanwhile, Big Tech has launched a campaign to manufacture consent about these systems. Big Tech claims that the rules, licensing, and legislation that is applied to creative works since the 1800s in the United States are suddenly moot simply because machine learning is, in their view, too important to be bogged down by the licensing choices of human creators of works.

In the FOSS community, we see this policy coup happening on every level from propaganda to consumers to policy papers to even law journal articles. I realize that I sound rather pessimistic about outcomes here. I'm nevertheless hopeful sitting here in this panel today because I see that so many of my colleagues in other fields are similarly skeptical about Big Tech's self-serving rhetoric in this regard. And I hope that we can work together to counter that rhetoric fully. Thank you.

Madeleine Varner:

Thank you. Next, we're going to hear from Umair Kazi, director of Policy and Advocacy at Authors Guild. Umair?

Umair Kazi:

Thank you, Maddy. Thank you, Chair Khan, Commissioner Bedoya, and Commissioner Slaughter for centering the voices of creators in this important debate around generative AI. And thank you to the FTC staff for organizing this excellent program.

I'm here today on behalf of the Authors Guild, the nation's oldest and largest professional organization of published authors. Since its founding, the Guild has served as the collective voice of American authors. Our 14,000 plus members include traditionally published and independent authors, novelists in all genres and categories, nonfiction writers, journalists, historians, poets, and translators. The Guild's core mission is to support working writers by protecting free expression and copyright, and fighting for fair contracts and for authors' ability to earn a livable wage.

Generative AI poses a serious threat to the writing profession, and we believe that guardrails around its development and use are urgently needed. AI developers have copied millions of copyrighted works without permission. These works are not only copied many times in the course of compiling training data sets and ingestion, but are embedded in the very fabric of the language models. These works are also used when users prompt the AI system. It is inherently unfair to use copyrighted works to create highly profitable tech, which is also able to produce competing derivative works without the creator's consent, compensation, or credit.

There's a serious risk of market dilution from machine generated books and other works that can be cheaply mass-produced, and which will inevitably lower the economic and artistic value of human created works. We are already seeing that AI is being used to generate low quality eBooks, impersonating authors, and displacing human authored books in the marketplaces like Amazon. In one instance earlier this year, AI-generated books started dominating Amazon's bestseller list in the young adult romance category. We have seen AI-generated books pop up for sale on closely related or very similar topics as those authors have listed for pre-orders, a growth in the number of unauthorized summaries of books, and fake books using names of renowned authors.

Generative AI is being used to create unauthorized derivative works such as a developer using ChatGPT to write the concluding books in George R. R. Martin's, "A Song of Ice and Fire" Series and chatbots like the Dan Brown Chatbot. Beyond the obvious economic impact are the more personal and painful incursions upon authors' voices and personhood. Pulitzer Prize winning author and Authors Guild councilmember Min Jin Lee has likened the experience of having her work used by AI to identity theft. "AI company stole my work, time, and creativity," she recently wrote on X/Twitter. "They stole my stories. They stole a part of me." It's not just well-known and bestselling authors that are suffering. Freelance journalists and professional writers of web and marketing content are reporting losing work at an alarming rate. An Authors Guild member who writes marketing and web content reported losing 75% of their work as a result of clients switching to AI.

A content writer featured in a piece about ChatGPT's impact in the Washington Post stated that he had lost a half of his annual income. There's a widespread and real concern among writers that AI will decimate the profession. In our recent survey, 69% of authors said that generative AI threatens their careers, and an overwhelming 90% said that authors should be compensated if their works are used in training. These concerns are born out of the experience of enduring long-term precarity.

As our most recent income survey found, the median writing related income for full-time authors is just over \$20,000 per year. The consequences of this precarity with generative AI is quickly compounding reach beyond the writing community. Do we really want a world where our books and literature are algorithmically synthesized mimics of the richness of human experience? To safeguard the incentives for creators to continue creating incentives so vital to our democratic culture that they are inscribed in the Constitution, the Authors Guild is lobbying for laws, regulations, and policies that recognize the following and require:

1. Consent and compensation. Require all generative AI companies to seek permission for the use of creative works and to fairly compensate creators.
2. Credit and transparency. Create obligations for all AI companies to disclose what data sets and works they use to train the systems.
3. Permission and payment for use in outputs. Require all AI companies to seek permission and pay compensation when creative works are used in outputs or when names or identities or titles of works are used in prompts.
4. Labeling AI-generated content.
5. No copyright for AI-generated outputs. We oppose efforts to deem AI-generated content protectable under copyright law or through creation of even a limited suite generous right.

Providing copyright or similar incentives to use AI to generate content will exacerbate the threat of AI-generated content flooding and overwhelming market for human works. Thank you for your time, and I will now turn it over back to Maddy.

Madeleine Varner:

Thank you. Next, we're going to hear from Douglas Preston, a bestselling author and a former president of Authors Guild. Douglas?

Douglas Preston:

Thank you. All right, I'm glad to be here. Thank you, everyone.

In addition to being an author and the former president of the Authors Guild, I'm also the plaintiff in a class action lawsuit against OpenAI, along with 15 other authors and the Guild itself. And we're asking for damages for unauthorized use of our copyrighted work and training and building ChatGPT. And going forward, we're asking that OpenAI and other AI developers get permission from authors, properly license our books, and compensate us fairly for that use.

When ChatGPT first came out, like many authors, I was fascinated with it and began playing around with it. And at one point, I asked it to write a poem in heroic couplets about one of my characters. And I was floored at the level of detail it knew when it generated this poem, and that's when I realized it must've ingested many of my books.

And at the same time, many authors were discovering that ChatGPT-3 knew everything about their books as well. And some realized it was even being used to create works that imitated their own. My friend George R.R. Martin, who was already mentioned, was very disturbed when AI was used to write the last book in his Game of Thrones series using his characters, his plot lines, his settings—even his voice.

So we at the Authors Guild investigated, and here are some of the facts we uncovered. ChatGPT-3 used more than 150,000 copyrighted books to feed into its AI system, which led us to the next question. "Where did OpenAI get our books?" They're not just sitting out there on the web unprotected. Well, the Authors Guild found that OpenAI got many, if not most, of these books from pirate websites such as LibGen run out of Russia.

So these pirates had stolen our books, and then OpenAI stole them from the pirates, I guess.

ChatGPT would be lame and useless without our books. Just imagine what it would be like if it was only trained on text scraped from web blogs, opinion screeds, cat stories, pornography and the like.

Now, Sam Altman himself testified that books provide the really high value literary content that large language models require. And he also testified that his goal with OpenAI is to support creativity. But supporting creativity by stealing from creators is like claiming you're supporting the candy store by shoplifting.

This is our life's work. We pour our hearts and our souls into our books. They aren't just products, they're a part of us.

And that's why we joined together, the 17 of us authors, in a class action lawsuit on behalf of all professional novelists against OpenAI. There's nothing complicated about this lawsuit. OpenAI illegally ingested our books to create a product that is currently valued at tens of billions of dollars, and they did this without our consent or compensation. And as Umair mentioned, the average full-time author in America makes only \$20,000 a year. This is a classic case of Robin Hood in reverse, stealing from the poor to give to the already obscenely rich.

In their race to be first, AI developers are swallowing everything they can get their hands on without regard to copyright ownership, intellectual property rights, or moral rights. And they're doing this without the slightest consideration given to supporting the livelihood of America's creative class.

Now, it's been mentioned before. The founders of our country wrote copyright protection into the very first clause of the Constitution—it was that important to them. They believed that their scrappy little country one day would become the creative engine of the world, and that's what we've become, and we can't allow AI developers to ignore copyright protection and injure the entire literary community of our country in their mad rush to succeed. They can succeed and they can also partner with America's authors in a mutually beneficial relationship.

Thank you.

Madeleine Varner:

Thank you. Next, we're going to hear from Tim Friedlander, president and founder of the National Association of Voice Actors. Tim?

Tim Friedlander:

Cool. Thank you for having me here today. I am the president and co-founder of the National Association of Voice Actors, and I'm here to represent the interests of the million strong voice actors in the United States and the surrounding ecosystem of Americans that work with them, ranging from engineers to script handlers to directors, producers, and PAs.

Not all voice actors are celebrities or well-known voices. Most are blue collar, working class voice actors who are working 40 plus hours a week. Over 60% of the voice actors are located

outside of LA and New York. We're not anti-tech or anti-AI, as many have said before. The ability to record audio on our computer at home was game changing. I'm a child of the 80s. I still remember the time that I died of dysentery on the Oregon Trail for the first time. We are not anti-tech or anti-AI. We are pro voice actor.

I'm also a member of SAG-AFTRA and a professional musician for over 30 years, first in the classical arena, and now most recently in hip hop for the last 10 years. We stand in solidarity with our fellow creative industry artists who are in unions. But unlike them, 80% of the voiceover industry is non-union, meaning we lack the protections and contract that organize workers enjoy. Without the intervention of the Federal Trade Commission, AI and synthetic voice clones will place our members at an unfair competitive disadvantage and decimate our profession. It's difficult enough to compete against other humans. Now mom and pop voice actors across the country have to compete against digital clones of themselves and against multimillion and billion dollar tech companies. It's incredibly easy to use AI to capture the voice of an actor from content available on the internet, and to use that sample to create whole works for sale or non-commercial distribution.

Voice actors have been creating and delivering digital audio of ourselves for decades. Contracts we signed years ago are now being used to justify the inclusion of our audio in synthetic voice models. And every time that happens without the consent, control, and compensation of the voice actor involved, the value of that voice actor's product, their voice and the sound of their voice, is diluted and unfairly diminished. We consider this to be a form of theft. We agree to license our voice to clients. The client is buying the exclusive rights to that recording, that performance, but also the sound of our voice. We could potentially now be in conflict with a clone of our own voice, which causes harm to both the voice actor and the company. Pepsi and Coke can't have the same voice providing their commercials. Chevy and Ford can't have the same voice. Currently the only protections that we have from voice actors having their audio uploaded to these sites are the terms of service that people agree to.

Companies can change their terms of service, and currently it's easy as checking a box that says, "I have the right to upload this audio." We've seen damages this year very clearly. A voice actor in New York worked for a company for three years, and year four, they were let go because they were told the company had enough of their audio, and they were going to now create a synthetic version of their voice. In February, Twitter doxxing: Voice actors such as Michael Schwalbe and Abbey Veffer had their voices cloned and had Twitter accounts made that gave out their personal addresses in their own voice and said racist and homophobic things. Thousands of synthetic voices are currently online on websites from video games all over the internet. Fan dubs are being created and turned into adult content, and then sharing those fan dubs on YouTube using the voices of those voice actors.

And most recently, a voice actor in Washington state lost an audiobook job when they decided to take the job in-house. Around that same time, the audiobook company made a press announcement that they were now using ElevenLabs for all of their audiobook productions, effectively replacing all of those human narrators with synthetic voices.

Additionally, an extra burden has now been placed on the voice actors to prove that the audio they're delivering is not AI-generated. Many voice actors have reached out to NAVA asking for some way to prove that they delivered human-generated audio.

But it's not all terrible. I personally am working on a synthetic voice that I have consent, compensation, and control for. There are some things that humans can't physically do, such as narrate the New York Times cover-to-cover every morning, or provide a realistic voice for someone who is nonverbal. But this tech should enhance and not replace voice actors.

We're asking for a few things.

First, the three C's. For us, this is consent, control, and compensation. Consent to have our voices in the models. Control over where that model and that voice can be used. And compensation to be paid fairly. A federal right of publicity.

If the FTC is not able to do this through rulemaking, we would request the FTC's support for federal legislation to establish that right. Transparency of ingested content on which these foundational models are trained in order to know if our voice is present. Protections prior to the generation of any AI-created content that might include voices of professionals that have not provided consent and are not being compensated. Protection of our voices' biometric data for privacy and commercial purposes. An independent third party to verify that audio files are ethically sourced. And finally, clear labeling of any AI-generated content to ensure the consumers are fairly informed.

And as such, NAVA recently endorsed the [US Senators] Schatz and Kennedy's AI Labeling Act of 2023, which is Senate Bill 2691 that we are fully endorsing.

Thank you for your time.

Madeleine Varner:

Thank you. We're going to hear now from Sara Ziff, founder and executive director of The Model Alliance. Sara?

Sara Ziff:

Thanks, Maddy.

My name is Sara Ziff, and I'm the founder and executive director of the Model Alliance, a nonprofit organization that advances workers' rights in the fashion industry. I come to this having worked as a model myself for many years and experiencing the pitfalls of what remains a largely unregulated industry. Models are typically hired as independent contractors through management companies which, unlike talent agencies, are held to very few legal standards. So when we talk about how generative AI is impacting workers, we need to consider the context of an industry that is truly like the Wild West— where workers have few protections at baseline and also cannot collectively bargain here in the US.

Our community of models, content creators, and other fashion creatives are concerned about the impact of generative AI. We recently conducted an informal poll to get a better sense of their concerns, which generally fall into two key areas:

The first is around the use of 3D body scans in connection with generative AI, and the second is around the creation of AI-generated models – particularly AI models of color.

Increasingly, companies are asking models to undergo scans that generate a 3D model of their body or face. In our poll, nearly 18% of models who responded had been asked to undergo a scan by a brand or a management company.

When a model signs a management agreement, she typically hands over a power of attorney, and thus rarely, if ever, sees her own contracts with the brands. As a result, models have little transparency into their own business dealings.

Those who had been scanned described not being given information about how their scans would be used, unknowingly handing away rights to their image, and not being fairly compensated. For people whose livelihoods are their image, this is particularly troubling in light of the rise in deepfake technology, specifically deepfake pornography.

The second concern is around the creation of AI models and influencers, which are digitally created, fictitious representations of human models.

Fashion workers are worried about the threat of these AI models replacing jobs – not only for models, but also photographers, stylists, and hair and makeup artists among others.

Members in our community have expressed particular concern about companies using AI-generated models as part of their diversity and inclusion initiatives.

For example, Shudu, a digital model who was created through AI in 2017 by the world's first all-digital modeling agency, has appeared as a face of high-end brands such as BMW and Louis Vuitton. Critics have called this a form of “digital blackface” since Shudu is a Black woman, and the creator who profits off her image is a White man.

And earlier this year, Levi's announced that they are creating AI-generated models to increase the number and diversity of their models. In an industry that has historically been discriminatory, creating digital representations of models of various ages, ethnicities, and body types rather than hiring and paying a diversity of real models is concerning.

I appreciate the FTC's mandate is to protect both competition and the consumer, and that the agency is set up to prevent unfair methods of competition and deceptive practices affecting commerce. With that in mind, there's a real risk that AI may be used to deceive investors and consumers into believing that a company engages in fair and equitable hiring practices and is diverse and inclusive, when they are not.

To address these concerns, we first aim to pass the Fashion Workers Act, our signature bill, which would establish basic labor protections for models and content creators working in New York's fashion industry. This would help address the lack of transparency that leaves models in the dark about how their digital image is being used, and establish a necessary foundation for regulation around generative AI in the fashion industry.

In considering regulation, it's essential that we center the experiences and expertise of those who will be directly impacted. And to that end, we're also developing a research study in partnership with the Worker Institute at Cornell University to better understand the impact of generative AI on fashion workers, particularly workers of color, and develop policy recommendations. So if anyone is interested in learning more or getting involved, we welcome you to reach out.

At The Model Alliance, we believe now is a critical time for solidarity between workers across creative fields who contribute heavily to our culture and economy. Unfortunately, it's not enough to win protections through collective bargaining agreements. There are many workers, including members of our community, who cannot engage in collective bargaining, and so we have to ensure that they are included.

As many others have said, we're not anti-technology. We're anti-exploitation. And we believe this round table is an important step toward ensuring our creative community is afforded the protections we need and deserve.

Thank you for the opportunity to speak with you today.

Madeleine Varner:

Thank you. We're now going to hear from Karla Ortiz, a concept artist and illustrator.

Karla Ortiz:

Hi. FTC Chair Lina Khan, Commissioner Slaughter, Commissioner Bedoya, Maddy Varner, and all esteemed officers at the FTC, thank you for allowing me to be a part of this amazing panel to discuss generative AI and its impacts on the creative economy.

My name is Karla Ortiz. I have worked as a professional concept artist, illustrator, and fine artist for the past 14 years and have been training for that all of my life. My work has helped shape the world's big-budget films and TV shows including Marvel Studios Loki 1 and 2, Avenger Civil Wars, Guardian of the Galaxy 3, and most known for my design of Dr. Strange – look in the first movie, which you can see him right there.

I deeply, deeply love what I do. Making a living as a professional requires a whole life of practice and study. The creative economy only works when the basic tenants of consent, credit, compensation, and transparency are followed. The country's creative economy, represented here by notable figures from each sector of our industry, is great because it adheres to those tenant. That success is gravely threatened by generative AI. For the first time in my life, I am worried about my future as an artist.

Generative AI technology that uniquely could use to consume and exploit the hard work, creativity, and innovation of others, no other prior technology is like this. I first encountered generative AI around a year ago. I'm still shocked by its blatant exploitation. My work and the work of almost every artist I know was stolen without consent, credit, or compensation. Without transparency, we can't even know the extent of how much of these companies have taken. They took our work and data to train for-profit technologies that then directly compete against us in our own markets using generative media that is meant to mimic us.

AI companies also engage in deeply unethical practices. For example, data laundering, where a company outsources its data collection to a third party under the pretext of research to then immediately use that for commercial purposes. An example of this is LAION. Specifically, it's more popular LAOIN-5B, a dataset that contains 5.8 billion text and image pairs, which again includes the entirety of my work and the work of almost everyone I know.

Beyond intellectual property, datasets like LAOIN-5B also contain deeply concerning material, like private medical records, non-consensual pornography, images of children. Even social media pictures of our actual faces end up in that dataset. Many of the largest and most prolific sellers of generative AI use these datasets. These companies describe their stolen training material as publicly available data or openly licensed content to disguise their extensive reliance of unauthorized use of copyrighted and private data. These companies even break their own promises not to sell users information.

These are just few examples of the unprecedented volume of copyright infringement and violation of other rights. No matter what these sellers say, most, perhaps all of these models rely on stolen works. Not just this, but many of these highly profitable companies encourage users to

use artists full names to generate imagery, exploiting our very identities and reputations. Some artists have had their names used in props hundreds of thousands of times, maybe more. And these numbers come from incomplete records, by the way. My own name, Karla Ortiz, has also been used thousands of times. I never gave consent. I never got credit. I never got compensation.

Perhaps most harmfully, these exploitative products compete directly with artists and are already replacing us. That may be the most disturbing harms of generative AI. Not vast murdering sci-fi nonsense, but one built on works taken without credit, consent, compensation and transparency, and marketed and used as a replacement for the creators of those works at a fraction of the cost. The whole process is rotten.

As a side note, due to all of this, I am also a plaintiff in a class action against generative AI image companies as well. Their plan is simple, to go as fast as possible, promising promises of progress and innovation while normalizing the exploitation of creative professionals, hoping that by the time anyone tries to stop them, it'll be too late to protect us American, or “median humans” as a prominent AI executive likes to call us.

But with help of the FTC and others looking out for American rights, we hope that game will not succeed. I think this panel is a great step in that direction. Regulatory agencies should act now to protect artists, consumers, and other Americans from this unconscionable exploitation. Regulatory agencies should demand full transparency from generative AI companies and opt-in-only practices.

Lastly, regulatory agencies should strongly consider seeking algorithmic disgorgement on products built on data acquired without consent, credit, or compensation, regardless whether that company is transparent or not. Urgent measures like these will be needed to avoid, in my opinion, the diminishing or outright destruction of most, if not all creative professional livelihoods and the protections of all of our rights.

Thank you.

Madeleine Varner:

Thank you so much. Next, we'll hear from Steven Zapata, a concept artist and illustrator speaking on behalf of the Concept Art Association.

Steven Zapata:

Thank you so much and a huge thank you to the FTC for putting on this event and to all of the attendees. I have to say the representation here is exemplary and I'm sure we are all so thankful to see it.

My name is Steven Zapata. I'm a concept artist and illustrator, and I'm here speaking on behalf of the Concept Art Association, an advocacy organization for artists working in entertainment, particularly concept artists, those artists who generate striking visuals in the early design process of an entertainment property.

CAA is committed to elevating and raising the profile of concept artists, their art, and their involvement in the entertainment industries. We're here today because a future with unregulated AI will hurt concept artists and all other sorts of artists across many fields. We need regulation, intervention, and oversight. The massive datasets that text-to-image models are trained off of contain the copyrighted works of artists, whether it be work done for ourselves or work done for companies in a work-for-hire capacity.

To advertise our work, most of us put our art online, on social media and our personal websites. This leaves it exposed to unethical scraping practices, which can result in AI models being created that can mimic the signature styles of particular artists. The companies offering these models often encourage users to request work in the styles of particular artists by name, and many of these generations end up online with our names attached to these pieces that we didn't make. These pieces have subject matters and qualities that we have no control over. And the negative market implications of a potential client encountering a freely downloadable AI copycat of us when searching our names online could be devastating to individual careers and our industry as a whole.

Also, the nature of the typical work-for-hire contract has put us in a bad bind. Almost every commercial artist today works under all-encompassing work-for-hire contracts that surrender all rights to a company. And unfortunately, visual artists don't have strong union representation to push back on this. As it stands, you can work hard for a company like Disney for 25 years and they can take all your work, train a model on it, and then have a free and tireless replica of you to use in perpetuity. That may currently be legal, but it certainly is not just. It is clearly unfair that our work should be used without our permission or even our notice to train products that we will have to compete directly against in our very marketplace.

Given these industry pressures, artists may be coerced by clients to utilize these systems themselves to stay up to speed with the market, thus normalizing the exploitative practices and foundations of these models. This will inevitably damage the perception of our field and art in general, as it will contribute to the idea that beautiful art is made easily. I can assure you, it is not. And insofar as it is made easily by an AI, it is because it has been trained off of the beautiful work of thousands upon thousands of artists who had to invest time and effort into creating their art.

So we need regulation, intervention, and oversight. We as creators should have complete control over how our work is used, but we need help. Some of the potential actions and remedies that we hope to see include, first and foremost, ensuring that all commercial AI models utilize only public domain content or legally licensed datasets acquired in an opt-in capacity. Opt-out is completely insufficient here. This could mean current companies shifting to the public domain and possibly destroying their current models in the process so that opt-in becomes the standard.

We also need transparency on datasets, and divulging your dataset should be compulsory.

Mainstream models like DALL-E 3 don't reveal their training data and don't let you search it, but they do offer an inefficient one by one opt-out system that you can use if you think maybe your art is in there. But because these AI systems can't unlearn, this will only remove the images from future training datasets used by this one company and it's already too late to get out of the most current model. Future tools that would verify compliance with future regulations will also depend on this transparency.

We should also have AI companies pay a fine for their past practices and pay all affected artists a fee per generation. This is to compensate artists for utilizing their works and names without permission, should be retroactive for as long as the company has been for-profit. We must close research to-commercial loopholes, interpreted or actual, that allow for-profit companies to monetize the results of non-commercial research.

To close: consent, credit, compensation, control. This is what creators reasonably seek in this new era where our work will be used to add tremendous value to these new technologies. We need oversight. We need compulsory transparency and tools to verify compliance.

Thank you.

Madeleine Varner:

Thank you. We're now going to hear from John K. Painting, director of the Electronic Media Services Division of the American Federation of Musicians. John?

John K Painting:

Thank you, Maddy. Good afternoon, everyone. It is an honor to be here.

As the labor institution which represents and protects the interest of musicians who prepare and perform instrumental music in recording studios for sound recordings, film, television and streaming services, as well as in live theater, symphony, opera, ballet, clubs, festivals, and more all over the US and Canada. The American Federation of Musicians is certainly no stranger to dealing with the encroachment of technology on our profession.

But the rise of generative artificial intelligence yields a more existential fight than we have faced before, as we approach the potential disappearance of performers livelihoods and by extension the disappearance of a component of humanity and culture.

From our history dealing with technological advancement, like our sister entertainment unions and guilds, the solutions sought have been traditionally approached in two ways: collective bargaining with industry and legislative lobbying. Both paths tend to seek secondary income to those performers whose work has been diminished by advancing technology.

For example, in 1942 with live performance and radio orchestras threatened by recorded music, the AFM struck the record labels and eventually succeeded in establishing the Music Performance Trust Fund, supported by label revenue that sponsors free live concerts around the US and Canada.

More recently, the AFM has had to combat the encroachment of prerecorded music into live engagements, such as preventing virtual orchestras from replacing pit musicians and musicals.

We've been here before and we will be here again. Make no mistake, generative AI is coming and it will change everything.

Back in 1927, in the face of the technological advancement of adding sound to motion pictures, then-AFM President Joseph Weber said, "Nothing will destroy the usefulness of an organization surer than to set its face against progress. No matter how unfavorable we may at present see the same to our interests, we must face these changes head on rather than stick our heads in the sand."

But if I'm a regional live theater producer, for example, with a six-figure annual budget to spend on pit musicians, how much do you think I would be willing to pay to have a piece of technology that replaces that cost entirely? And in that scenario, should the consumer still be paying the same ticket prices to see that artificial musical? Is disclosure of that fact enough? What if every theater in the country had the same technology?

Compared to previous technological shifts, the job losses here could be steep: recording and scoring of all types can be wholesale replaced if composers and songwriters can replace the output of live performing instrumentalists with an AI-generated facsimile. Producers may replace the composers and songwriters themselves. Live touring across all genres would be impacted if bands and artists could tour with AI-generated backing musicians, holographic orchestras, AI backing tracks.

A machine-generated creation which utilizes a performer's voice, face, or in our case instrumental sound without the actual participation of that performer in the creation is tantamount to replacing the performer entirely. It removes the human from the creative process and the performing process. When that happens, the work is gone, the wages are gone, the secondary residuals are gone. But it's more than just economics. The humanity contained within artistic and cultural expression will be gone too.

Solutions cannot merely be the bargaining of replacement remuneration if the job opportunities are replaced wholesale, or establishing minimum hiring requirements for producers who choose to be union signatory, or extensions of copyright law to establish a financial value to the performers material. That may be sufficient for those whose material is being used to train the system, but what does that mean for the young cellist who's practicing today? What will that mean for their job prospects in five, ten, fifteen years if recording work and orchestra work dries up?

The solutions need to be wider than the traditional paths we've all taken owing to the cultural damage that this problem yields. As soon as it becomes broadly accepted that art no longer requires the involvement of human creativity, we'll have crossed the point of no return. Consumers should not be expected to accept such an alternative at the same price points.

Musical expression transcends culture and time. It will always be a part of our society, but that doesn't guarantee it as a viable career. The end game must be the protection of the profession.

Thanks so much for your time.

Madeleine Varner:

Thank you. And finally, we'll hear from Jen Jacobsen, executive [inaudible 01:06:37] Rights Alliance. Jen.

Jen Jacobsen:

Thanks, Maddy. And thanks to the FTC commissioners and staff for the opportunity to participate today. We're so grateful to the commission for your work on this issue. And I'm honored to be here alongside all these other representatives of the creative community.

I'm Jen Jacobson, Executive Director of the Artist Rights Alliance, which is an artist-run nonprofit that fights for the right of musicians, songwriters, and performers to be treated and paid fairly in the digital marketplace. Like many of the groups represented here, ARA is a proud member of the Human Artistry Campaign, a coalition of over 150 groups in 30 countries spanning creative professions and advocating for responsible and ethical AI.

Musicians are not against AI. We don't want AI to be banned. In fact, musicians have been using AI-driven tools for years to auto tune vocals, generate beats, assist with studio production, and in other parts of the creative process as well. What's new today, though, are the expansive AI models that ingest massive amounts of musical works and mimic artists voices without obtaining creators consent or compensating them.

Musicians want to be treated fairly, whether that means enforcing existing laws designed to protect creators and consumers, or in some cases enacting new laws where there are specific gaps to be filled.

Unfortunately, in today's reckless, careless rush to launch new generative AI products, we are seeing what is euphemistically referred to as AI training or learning, but which is in fact illegal copying of artistic works on a massive scale without consent or compensation and often without the artist even knowing.

Of course, this is a clear infringement of creators copyrights, but it is also an “unfair and deceptive act” that impacts both artists and consumers – and leads to unfair competition in the music marketplace.

Musicians’ work is being stolen from them and then used to create AI-generated tracks that directly compete with them. For example, we might see dominant streaming platforms packing playlists with AI music that they obtain free of charge or at a massive discount, which then lowers their own royalty obligations and diminishes artists wages.

The increasing scale of machine-generated music dilutes the market and makes it more difficult for consumers to find the artists they want to hear. It makes it harder for artists to connect with their fans, and it devalues human creativity.

And perhaps even more disturbingly, AI models are now using artists’ faces, voices, and performances without permission to make digital impersonations that not only create consumer confusion, but also cause serious harm to both fans and artists. These deep fakes have depicted a band canceling a concert that wasn't actually canceled. They've shown artists selling products that the artists never endorsed. We've seen false depictions of musicians badmouthing their own fans. This isn't a hypothetical harm. This type of consumer deception and fraud are happening right now. It's hard to imagine anything more personal to an artist or to anyone than being depicted as doing or saying things that they would never do or say. It's not only confusing to fans, but humiliating to the artists themselves and undermines their public image.

This conduct may violate “right of publicity” laws in several states, but it needs to be recognized universally as a misappropriation that causes real harm not only to the artists, but to the entire market by confusing consumers and creating unfair competition. And no one is as well-positioned to protect consumers and fans in all fifty states than this agency.

The final point I want to make is about the importance of choice. Many AI companies who have illegally vacuumed up hundreds of thousands of musical works and recordings now say that artists can simply contact the company and “opt out.” This is essentially proposing a new form of uncompensated labor that musicians and composers have to perform if they want to avoid exploitation. It's also completely impractical given the proliferation of new services ingesting work without permission. Such burden shifting is not only unfair, it is morally wrong and antithetical to basic principles of artistic integrity. Artists have the right to control whether and when their work is being used, especially when these uses compete with their own livelihood or violate their own values.

Art is about a human-to-human connection, about sharing emotions and lived experiences. Machines can't share emotions or lived experiences because they haven't had any. Only humans can do that. For musicians, like all the other creators represented here today, the unethical use of AI poses an existential threat to our livelihood. But for all of us, the very foundations of human creativity and culture are at stake. It's hard to imagine anything more important than that.

Thank you again for including ARA in this discussion, and I look forward to answering any questions.

Madeleine Varner:

Thank you, Jen. And thank you everybody for your remarks. So now we're going to move into a Q and A portion. And I just wanted to note too that Duncan Crabtree-Ireland of SAG-AFTRA had to leave a little early to return back to the bargaining table.

So to get started, I'm going to open with the first question.

Recently there's been some attention paid to the idea of offering opt-outs to artists who don't want to particular company from using their works to train generative AI models. Do such measures offer sufficient protections for creators? I'm going to go first to John August from the Writers Guild of America West. John?

John August:

Thank you. As Chair Khan said in our opening remarks, the advent of AI doesn't change fundamental ways that the world is supposed to work, and that world works in permission first. That's why we think it's critical that we require artists have affirmative consent before the work can be used to train generative AI models and that they have to be compensated fairly when they do so. The same should be true for all artists, including artists like us who do work for hire and don't hold the copyright on our work.

And this system needs to be opt-in and not opt-out. As Jen just said, there are so many companies out there developing and training AI models, to be forced to continually track all of them down to opt out is an enormous administrative burden on individual artists. It's not practical. It has to be opt-in rather than opt-out.

Madeleine Varner:

Thank you. Karla, I'm curious about your thoughts on this.

Karla Ortiz:

Yeah, so opt-out is an ineffective and inappropriate standard for commercial use of copyrighted works including a generative AI. Once a model is trained on data, it cannot be deleted unless the whole model is retrained from scratch. By the time a model is made public, it's already too late to opt out. Number two, most AI companies keep that training data secret, preventing artists from even knowing if their works were used to train a model. Number three, existing opt-out procedures often ask users to list works used to train the model they own, but as we just mentioned, that training data is secret, so it's an impossible task. And four, there are hundreds of AI models already in the market and more.

Does that mean we have to opt out on each and every one of them? That's a full-time job. What about if those models update? What about if they don't publicize and they use third parties? What if those models in the opt-out forms are not an artist's native language? What about artists who never spend time online or don't even know this is happening?

Basically, tech companies must respect artists ownership rights. These seeking to profit from others works should have the burden of obtaining permission. Explicit opt-in is the only way forward. It's really how we ensure generative AI models exclude unauthorized works from the beginning. Thank you.

Madeleine Varner:

Thank you. Neil Clarke of Clarkesworld, I'd love to hear your thoughts on this as well.

Neil Clarke:

I wholly agree with all the things that have been thrown on the table there. This is what you get when the fox designs the chicken coop. The wrong people are making the decisions when you have something like this happening. And there's no standard, there's no consequence for them ignoring these things at the moment either. So we really need to have them seeking consent, explicit consent, to include these works.

Madeleine Varner:

Thank you. I'm going to move on to our next question. Like "opt-outs," the idea of licensing out work for training purposes either by individual creators or others who may retain rights to said work has garnered attention. What interests or concerns do you have with this kind of approach? And I'm going to kick it to Jen Jacobson of ARA first. Jen?

Jen Jacobsen:

Thanks Maddy. ARA does believe that it's possible to have a market for licensing our work to responsible, ethical AI developers. The entire music streaming market is based on licensing, so we know it can work. But there are three things I would say about it.

First, artists need to have complete opt-in choice, as we've just heard, about whether to license their work or not.

Second is artists need to have enough power to negotiate fair license terms with these gigantic AI developers. There's actually a bill in Congress sponsored by representative Deborah Ross that would give small and independent musicians an antitrust exemption so they can come together and negotiate collectively, both with AI developers and streaming platforms, which is something they do not have the leverage to do currently.

And the third thing I would say is there is not a one-size-fits-all licensing system that will work for all creators or even for all musicians. Assuming there is a level playing field for negotiating, we think the best way for musicians to license their work is in the free market, which may look different for every use, every artist and every company.

Madeleine Varner:

Thank you. Steven Zapata. I'd love to hear your thoughts on this.

Steven Zapata:

Without a doubt, licensing will be essential in the future, but we must accomplish that through an opt-in system, otherwise there would be no real negotiating leverage for creators. And the focus of licensing, I think should go towards new opt-in foundation models, not the fine tuning of existing unethical models. As to when companies hold the rights to work done for hire and want to license or train off of that, we need regulation.

I mean, visual artists, for example, lack the union representation to push back against contracts that claim all ownership. And without regulation, I think predatory contracts will just run rampant in this sector. And collective licensing is also troubling. The early experiments we are seeing with "contributor funds" from companies like Shutterstock are paying out less than pennies. I mean actual fractions of a penny per used image. That's all they want to pay a creator for their life's work and to create a tool that will directly compete against them forever.

If that's not unfair, I don't know what is.

Madeleine Varner:

Thank you. John K. Painting of AFM. I'm curious about your thoughts.

John K Painting:

Sure. This goes back a bit to part of the solution I mentioned in my remarks. Let's say I'm the drummer for Taylor Swift and her catalog is used to train a generative AI system that can pump out an album on a prompt, like give me an album of Christmas songs in the style of Taylor Swift,

then the drums on that generated album must have used my sound to mimic the style. So not only should I have had the option to opt in there, of course, but I should see some form of benefit or compensation for that because those new parts are clearly copying mine.

But this is still only a bandage on the problem, because if this scenario works really well, it likely means that I'm not getting hired to record any new albums anymore because this system can just pump this stuff out. So do these new albums have the same value? Should they cost the same to buy and stream? Does the art have the same meaning and the same impact? That's all the existential crisis that we're facing right now.

Madeleine Varner:

Thank you. So I'm going to move on to our next question, which is: What kind of insight do you feel like you have now into how your work or likeness is being used by generative AI systems, and what kind of transparency do you feel is needed?

And Umair from Authors Guild, I'm going to start with you.

Umair Kazi:

Thanks, Maddy. Our members are keenly aware that their works are being used by AI systems. We get reports. Our members try out different prompts and AI systems and chat box reveal details about their works. And of course, several of our members have also reported finding their books in Books3, which is a data set of containing 200,000 books that's downloaded from a pirate source, which was used to train Meta's LLaMA, Bloomberg's GPT, and others.

But Books3 is kind of an anomaly, is that it was publicly downloadable and its contents were visible and searchable. There is a lack of transparency from AI developers about training data sets, which makes it very difficult to ascertain which works were actually used to train the models and how.

Much of the information about ingestion comes from the prompting experiments that I mentioned. We and other authors have been able to prompt, for instance, GPT to produce extensive detailed summaries of works and text in the style of and even incredibly compelling outlines for possible derivative works like sequels using settings, characters.

As far as what we want, we want AI companies to be required to fully disclose the complete lists of copyrighted works, books in particular is the medium that we deal with most often, that are in the training data sets or provide specific links to where the data sets were obtained from. And anyone compiling a training dataset should be similarly obligated to disclose the sources. And in the case of copyrighted works, a complete list of works that have been included in the dataset.

We also think authors and artists should have the right to say that they don't want their identities, works, voice or style used in outputs. And we believe that there should be a requirement to conspicuously label fully or substantially AI generated words in online marketplaces.

Madeleine Varner:

Doug Preston, I'm curious to hear your thoughts on this question.

Douglas Preston:

Yes, thank you. As a plaintiff in this lawsuit, one of the things that I did was to go to ChatGPT-3 and ask it many detailed questions about my books, my characters, very minor characters that have never appeared in a book review or in Wikipedia, and it was absolutely extraordinary how knowledgeable ChatGPT was. It was able to write a whole paragraphs in my style, using my characters, using my settings, using everything in my books, and yet it was all stolen from me and fake and AI generated. And that was very disturbing. And I also used that Books3 database, that dataset that Umair was speaking about, and found that that dataset has all my books, all 40 of them on it, including many different editions. That was another shock.

But as far as how we know our books are being used in AI training, we have absolutely no idea. It's a black hole. OpenAI is training ChatGPT 5 right now, is building it. They refuse to answer any questions from the Author's Guild about what data sets they're using, where they're getting their books, and how they're being used. There's no transparency at all. It's an absolute black hole. All we want is control of our creative work, fair compensation and opt-in consent.

Madeleine Varner:

Sara Ziff, I'm curious about your thoughts on this question.

Sara Ziff:

Yeah, thank you. Models have very little insight into how their work or likeness is being used in general, let alone in the context of generative AI. Normally they don't see their contracts with the brands and often don't know how their image will be used, whether how much they'll be paid. So generative AI introduces the potential for further exploitation in an already exploitative work environment. In terms of the transparency we want, I think some key elements would include requirements for explicit consent, notification of use, compensation, and liability for misrepresentation.

Madeleine Varner:

Thank you. Bradley Kuhn, I'm curious about your thoughts on this question.

Bradley Kuhn:

Well, first of all, there's now really no question that the body of Copyleft in FOSS is a huge part of the software assisted development machine learning systems such as Microsoft's GitHub CoPilot. We've exposed many of these examples, but the companies are also playing cat and mouse by simply excluding, on the generative backend, the most egregious examples of copyright infringement when they're publicly called out.

In my opinion, there's no reason that big tech shouldn't be regulated to make these systems transparent, completely end to end. In my view, the public should have access to the input set, have access to the source code of the software that does the training and generation, and most importantly, access to the source code that does these forms of backend generation exclusion, the latter of which I think would expose the duplicity of big tech's policies here.

Finally, I expect that once we have real transparency, it will bear out what many of the other speakers have noted today, that the issues of machine learning systems can't be solved merely with a financial compensation model to creators. FOSS really shows this explicitly since most FOSS authors are doing the writing of the software altruistically and the compensation that they're seeking is the requirement for future improvements to the comments, not a financial compensation. So we really need full transparency in these systems to assure that essential non-monetary policy, licensing term and consumer's rights are upheld.

Madeleine Varner:

Thank you. So I'm going to ask our final question for the day, which is, you've all described the harm to creators and their professional interests when companies use generative AI tools in certain ways, such as when those tools rely on the use of protected works or people's voices and images. Can those uses also wind up harming consumers? And if so, how?

And for this question, we're going to go to Tim Friedlander from the National Association of Voice Actors. Tim?

Tim Friedlander:

Cool. Thank you for this question. The ability to create a synthetic voice from anyone who has recorded audio is easy and simple and dangerous. Currently now it only takes three seconds of source audio to create a realistic voice clone. And this synthetic content can be used to deceive consumers into believing that a trusted voice is communicating with them. This can lead to relying on false and misleading information and potentially even implicate the human whose voice has been used to harm people. In the last couple of days, we've seen Tom Hanks and MrBeast release statements saying that there is synthetic versions of their image and voice cell in ads they did not authorize. And currently there's a show streaming on a major media streaming channel that is using AI generated voice dubs for their English dubs.

At a minimum the consumers should know when AI is used to generate voices and receive a warning that the information they're going to receive may not be accurate. We're also coming into an election, and we know that synthetic voices will be used, they have been used, and it's possible for anybody to use a synthetic voice for disinformation. And no one, a voice actor or anybody in general, wants to be the recognizable voice of disinformation in this coming election.

I want to leave just with a personal thought that we all know the sound of our friends and family and our loved ones' voices. You call somebody and say, "Hi, it's me." We know who we're speaking to. I spoke with a friend of mine, his name's Patrick, last night. Two months ago or a few months ago, he received a phone call from somebody claiming to have kidnapped his daughter, was demanding ransom, played a recording or a voice of his daughter screaming, "Papa, help me. Help me. Papa, save me." He said, "I knew deep in my bones, that was my daughter's voice." His daughter was safe. It was a scam, it was a fake. But he said "The voice that I heard on that line was my daughter's voice. It was the same voice I heard when I dropped her off at school. And it was the same voice that I heard when I picked her up when I found out she was safe." The danger to consumers is real. It is tangible. And it is here now.

Madeleine Varner:

Thank you. And that concludes our Q&A portion. Now we're going to turn it over to Commissioner Bedoya for some closing remarks. Over to you, Commissioner.

Commissioner Alvaro Bedoya:

Hi everyone. Can you hear me and see me? Wonderful. Tim, what a note to end on. As a parent, that is chilling. And I just want to say before I begin, I want to touch on two issues that I think hopefully will compliment some of the conversations that have been had today. I'm profoundly worried about fraud and AI use. I'm profoundly worried about what's happening to your work, your property, the fruit of your minds, quite literally.

I just want to offer two thoughts.

The first doesn't really have to do with the law. It has to do more with first principles and really what's at stake. And my hope is that this audience in particular, might think that important. The second does have to do with a very specific legal application of generative AI that concerns me profoundly and that I recently wrote about in the LA Times on Labor Day. But let's first talk about what's at stake here.

There is this foundational moment in American letters. I think it was 1837, Ralph Waldo Emerson gives this talk called the American Scholar. He goes on about how for way too long American writers, American poets have looked to Europe. He basically calls out writing for being boring. He says, he has this wonderful line, "Genius is always the enemy of genius by over influence." And he says that for 200 years, English poets have been Shakespeare-ized, basically saying that people have been copying the work of this genius rather than creating new work. And he said, "We need an American voice. We need an American scholar. It needs to be rooted in the soil, it needs to be free of these European strictures." And sitting in the audience that day was a journeyman printer from Brooklyn named Walter Whitman.

One of the great mysteries of American letters is that basically out of nowhere in 1855, a basically middling journeyman printer and occasional journalist creates what is arguably the greatest work of American letters, "Leaves the Grass." This collection turned poetry itself on its head. Previously, poetry had restricted itself to certain strictures around rhyme and meter. Previously, yes, poetry maybe touched upon love and human sexuality, but it was always demure to a certain degree, think Keats' "Ode to a Grecian Urn." And suddenly everything was possible. Suddenly Whitman made it acceptable for people to write in free verse. Whitman made it acceptable to write about love and sexuality in a way that was previously considered completely not on the table.

There is no Allen Ginsburg without Whitman. There is no Gabrielle Garcia Lorca without Whitman.

When I think about the work all of you do and I think about the challenges we're confronting, and you think, what can we expect from generative AI? I think there are people who basically expect everything, expect the world, expect there are no limits whatsoever on what is possible. And what I think it's important to remember is that because of the fundamentally recursive nature of what we call artificial intelligence, where it takes what it's read, what it's seen and returns it to you, I think that it is wildly unrealistic and frankly dangerous to think that it will break every bound of human creativity.

I think that, to paraphrase Emerson, what can we expect? We can expect a lot of bad copies of Shakespeare, but good luck trying to create the next Walt Whitman. I don't think it will ever extinguish the genius of human creativity. And it's important we disabuse people of the notion that it remotely will. And so, as a question of first principles, I think we need to remember that the only place genius comes from is people, is humans. And any effort to persuade the world otherwise will fail. But we need to make sure to remind people of that.

Let me get off my rant for a moment and speak about one specific legal application, which is something that concerns me, which is the capture of the art of actors and the art of writers. A couple weeks ago I started digging into this and I saw two things.

First of all, I was reading all these industry analyses about what AI promised for industry, and they were all over the place, but the one thing they have in common is they say, to the degree that work requires social and emotional skills, artificial intelligence is going to struggle.

I was reading these reviews from very fancy think tanks about that. And then Justine Bateman posted a note on social media. She said, "Hey, background actors, I hear some of you have been scanned on the lot. Why don't you share your stories with the world?" And these stories started coming in. "Yeah, hey, I was on the lot. And on day three with no notice, I was told to go to that tent over there. I walked into that tent and I was scanned from 360 degrees by 80 cameras. They paid me \$25. They told me that if I said no, I'd get fired." Someone else writing about, same story, surprise request to get scanned, surprise capturing not just in many instances of people's bodies, but also their emotions.

Hearing these stories reminded me in Congress perhaps of 1914, when Congress passed the law that created this commission that all of you are speaking before today, the Federal Trade Commission. In 1914, Congress had the choice of passing a law that specifically enumerated acts 1, 2, 3, 4, 5, 6, 7 that would be illegal. One House of Congress actually considered that, actually voted on it. Then they stopped and said, no, we need to consider the fact that there will be innovation in unfair methods of competition. We need to create an institution that is free to meet the innovations of large, powerful entities that will stifle competition in American industry in whatever corner it may be found.

And when I hear about writers who worry, new writers, young writers worry that the moment I arrive, I'm going to be asked to feed my scripts in, to train a new AI. When I hear about background actors, young actors, how lots of future actors are discovered, but who are the least powerful, least experienced, least savvy of all actors being forced to get scanned in the nude sometimes or in other really uncomfortable situations, it strikes me as more than innovative and it fills me with concern.

I will say, because this is a law enforcement agency, these are allegations. I did not investigate these things, but the shape of what I'm seeing concerns me profoundly. And you have to know that in myself, in our chair, in my colleague, Commissioner Slaughter, you have public servants who are keenly interested in your stories and what's happening to you and what's happening to your work. So with that, I will end and I'm not sure, actually, I did read the script. I don't know if I wrap everything or I pass it back to you, Madeline. Let me pass it back to you.

Madeleine Varner:

Thank you so much, Commissioner Bedoya. On behalf of the FTC, I'd like to thank everybody for a robust discussion, for sharing your perspectives and experiences and for raising more awareness on these issues. We're sincerely so grateful.

We will be posting the video and transcript of this round table on the event webpage. And in the meantime, thank you again so much to all the participants as well as the folks who are watching right now.

We hope you all have a great rest of your day. Thank you.