Silicon Flatirons
KNOW WHAT'S NEXT.

Transcript

Generative AI and Copyright Conference

October 06, 2023

Contents

Published October 27, 2023

Silicon Flatirons Center for Law, Policy, and Entrepreneurship
University of Colorado Law School
2450 Kittredge Loop DR, Boulder, CO 80309

Introduction and Keynote: Maria Strong

https://youtu.be/YTChS7eB2RE

[00:00:00.02] BRAD BERNTHAL: It is my pleasure now to introduce our faculty director and today's lead conference organizer Blake Reid. By way of introduction, Blake researches and teaches in the areas of telecommunications, internet, copyright and disabilities law. For the past decade until this year Blake has directed the Samuelson-Glushko Technology, Law and Policy Clinic, the TLPC at Colorado Law.

[00:00:29.54] For those of you who are not familiar with clinics, they are the labs of the law school. This is where law students do real work for real clients in the real world. Actually I launched the Technology, Law and Policy Clinic, and Blake was a star student himself in that clinic back in the day. Indeed he secured an exemption for security researchers as part of the triennial review in front of the Copyright Office back in the day.

[00:00:58.13] When Blake took over the clinic, he didn't just keep it going, he dramatically improved it from the days in which I led it. His student attorneys conducted technology policy advocacy in the public interest before an array of government bodies ranging from agencies, courts, legislatures, and others.

[00:01:19.88] And along the way Blake found time to write, and he has become a very accomplished scholar in addition to an accomplished teacher. In fact, Colorado Law just put him on the tenure-track faculty this year. His work has been published in the Stanford Law Review, California Law Review, and Indiana Law Journal among others.

[00:01:42.41] He's the lead organizer for today's conference and we are in for a terrific set of discussions. Please help me give a warm welcome to Blake Reid.

[00:01:50.20] [APPLAUSE]

[00:01:56.74] BLAKE REID: Good morning. It's a privilege to welcome you to Silicon Flatirons. Before we begin, I want to thank all of my colleagues on the Silicon Flatirons staff who worked so hard to pull this event together. In particular, I want to thank Shannon Sturgeon and Nate Mariotti for organizing the logistics of the event both for everyone here in person and for the folks who are joining us online.

[00:02:20.56] Christine McCloskey who helped bring in our outstanding group of speakers from all over the country, and Sara Schnittgrund, who has helped in many ways to bring our students into the mix for this event. Please join me in giving them a hand.

[00:02:33.82] [APPLAUSE]

[00:02:37.21] All right. With that, I am delighted to welcome you to our second conference this year on Generative AI. It's an outgrowth of my colleagues' and fellow faculty directors, Margot Kaminski and Harry

Surden's AI Ethics series. My thanks to them for laying such a wonderful foundation for our conversation today. And this is kind of a spiritual successor to the work of our former Silicon Flatirons colleague Christella Garcia, who's now out at Georgetown Law. We're really excited today to focus on the intersection between generative AI and copyright law.

[00:03:09.14] And I want to do a little table setting and then introduce our keynote speaker. So almost as quickly as generative AI tools like ChatGPT and DALL-E burst onto the mainstream this spring and promised radical changes across a wide range of society, they encountered the legal system and the first contact with the legal system, at least here in the United States, was with copyright law.

[00:03:32.71] Now my normal remit here at Silicon Flatirons is leading our telecom and platforms initiative. And you may be used to Silicon Flatirons doing a lot of telecom focused events and wondering what is generative AI and copyright have to do with any of that?

[00:03:48.60] And the answer from my perspective is that the evolution of AI and its intersection with the legal system raises profound descriptive and normative questions about the future of the internet and the flow of digital information that warrant attention, not only through the lenses of copyright law and policy and regulation, but also through Silicon Flatirons' broader lens on the migration to and the evolution of digital broadband networks, platforms and applications, and that's what we're going to try to do today.

[00:04:18.72] So as one commenter has described it, Copyright law is a durable creature of technology regulation that predates the internet by many centuries. We're talking about the Statute of Anne and the Stationers monopoly at our student primer event last night.

[00:04:34.37] Copyright law is often the primary and only functional law of the internet. So unsurprisingly an internet technology taking the world by storm has attracted substantial attention from litigants, policymakers and technologists, who are trying to construct generative AI in its future trajectory through the lens of copyright law.

[00:04:54.51] So in today's conference we're going to try and make sense of the intersection of generative AI and copyright from two different angles, law and policy. On the legal front, we've got high profile law lawsuits from artists, coders, writers, photographers and industrial rights holders, centering on copyright and related claims that are starting to pile up in the federal courts over the past several months.

[00:05:17.87] There's questions about how copyright laws basic rudiments of infringement and fair use and how they map onto generative AI's complex supply chain and mechanics and whether the outputs of AI systems are eligible for copyright protection. And these

questions have roiled scholars, and litigators, government and civil society.

[00:05:37.04] Like our first panel led by Harry Surden is going to explore these topics with a group of distinguished experts and shed some light on the doctrinal state of play. And then from a policy perspective what to do about all of this raises a whole new set of vexing questions about whether and how copyright law should change in response to generative AI's technological, social, cultural, and economic impacts? And we'll discuss that in the second panel.

[00:06:05.40] So more broadly questions about whether copyright is the or even a correct lens for shaping generative AI policy and how copyright intersects with broader agendas about the policy around and the regulation of AI both in the US and internationally, including on issues ranging from privacy, security, transparency, accessibility, competition, consumer protection, discrimination, mis and disinformation, labor and employment.

[00:06:34.10] We'll try and tackle all of that today. And against that backdrop, I could not be more excited to welcome our keynote speaker, Maria Strong, the associate register of copyrights and director of International Affairs for the United States Copyright Office, which is part of the Library of Congress responsible for both administering the registration system for copyrighted works and advising Congress on copyright law and policy.

[00:07:01.37] Maria is going to share the range of activities that the office is undertaking right now in this space. I'm especially excited for us here in Colorado. There's an exciting decision about the winner of the Colorado State Fair's annual fine art competition in 2022, Theatre D'opera Spatial, which was recently denied copyright protection. We'll hear all about that.

[00:07:29.36] In her role as associate register, Maria advises the register on the policy functions of the office including domestic and international copyright issues. She directs the Office of Policy and International Affairs. Maria has also served as the acting register of copyrights back in 2020, has served at the Copyright Office since 2010 following nearly a two-decade stint in private practice, and very close to our hearts here, started her career at the Federal Communications Commission. Without further ado, please join me in welcoming Maria Strong to Silicon Flatirons.

[00:08:05.20] [APPLAUSE]

[00:08:11.89] MARIA STRONG: Thanks, everybody. Thank you so much for the invitation. It's a pleasure to be here in beautiful Boulder. And I almost didn't make it, thanks to the government shutdown drama of last week. And I really want to thank Professor Reid, executive director Bernthal, and an incredible staff here to make sure that I was able to literally catch a last-minute flight and be with you today. It is a pleasure.

[00:08:37.09] As Blake said, we're here today to talk about the many issues that have come out of the intersection between copyright and AI including generative AI. A lot has happened in this space over the past year. Members of the public are now using sophisticated AI, generative AI systems to produce their own images, art, visual works, the classic subject matter of copyright.

[00:09:01.21] And creators have become very vocal in the press, on the Hill, in the courts about their concerns on their work being used and the out-- and how that affects their livelihoods. At the same time, others are also expressing a lot of interest and excitement about the possibility of how these tools might assist them in their creativity. Hence, you can see the opportunities and the challenge.

[00:09:24.53] I will caution everybody it's a very complex ecosystem. It's not binary. And I think that's kind of a challenge for policymakers and lawmakers as we move forward. My goal this morning is to provide an update, sort of set the scene on what the Copyright Office has been doing in response to all these developments.

[00:09:45.77] So let me start at the start, which is to remind everybody the role of the Copyright Office. We are part of the legislative branch of government, and we're situated within the Library of Congress. We are the agency that administers the copyright registration system. And although copyright registration, as you know, is no longer required for protection, registration does confer a number of benefits including the ability to bring a copyright infringement in federal court.

[00:10:11.18] Each year, we issue about a half a million registration claims covering millions of works, but we do more than that. Congress has also told us what it wants us to do. And we are the Congress's primary legal advisor on copyright. It's sort of like having 535 clients at once.

[00:10:30.44] We do work closely with congressional offices on domestic copyright legislation and policy matters including by providing technical advice on their proposals. Some of those proposals might not see the light of the day, and some of those proposals we may or may not be consulted on. So there's-- Washington DC is fast moving. At the same time, it's sometimes slow moving. And we're trying to react and serve our congressional clients.

[00:10:57.03] We also conduct comprehensive policy and legal analysis for Congress for use by the public Congress and often the courts. To be clear, when it comes to advising the president and the executive branch, it's the US Patent and Trademark Office that has that primary lead. When it comes to litigation, we work with the Department of Justice and the Solicitor General as they represent our agency and all agencies before the courts. And they, in turn, work closely with us on those copyright cases that might come before the Supreme Court as we have seen in recent sessions in terms.

[00:11:33.53] On international matters, our office participates as part of an interagency team. And that's usually led by an executive branch agency. So for example, when we review copyright legal developments in another country or developments in something like APAC or OGC, OECD, or the G7, we are part of a team that is evaluating the copyright interests and equities in that development, but-- to be sure who the final decision maker. And that is going to be an executive branch agency, depending upon the issue at hand.

[00:12:06.07] Click. There we go. So let me share a little bit about what we've been doing the past probably almost a year. We have been very active in response to the development of AI and in particularly, generative AI. Earlier this year, we announced a formal launch of our AI initiative, which comprises of a number of pieces of actions. Kept us incredibly busy.

[00:12:30.05] We have provided, for example, and I'll talk about this in a minute, guidance on how to register works that contain AI-generated material. We've held public roundtables with AI developers, artists, academics, business people, and everyone else interested in the intersection of AI and copyright.

[00:12:47.14] We've had meetings with dozens of stakeholders. And we've provided two educational seminars, webinars on both registration and international matters. All of this is on our website. We are a very transparent agency. And that's why we like to gather the information in a diligent, thoughtful way so that we can make sure what we have on the table as we provide the advice that we are going to be doing in the future.

[00:13:14.06] All this led up to at the end of August. We issued a formal Notice Of Inquiry, sometimes called an NOI, where we are seeking public comment on a number of questions that we have identified. And I'll get to that also in a little bit. Already, we've received close to 8,000 comments that have been put in well in advance of the deadline.

[00:13:35.48] So this level of interest is not surprising. These issues are not just technical, but they're actually existential to many creators. These issues touch on the fundamental nature and future of human creativity and how it can coexist and interact with potentially unlimited technological capabilities.

[00:13:55.17] And so the office's foray into the current AI technology has involved copyrightability of the output of AI systems. This is a function of our role in registering copyrights. While we don't investigate the underlying facts set out in an application for registration, we do review the deposit copy of the work to make sure it's copyrightable.

[00:14:16.82] And so a few years ago, we actually started receiving applications to register works creating generative AI. So this is not new

to us. This meant that early on, we had to make decisions and take some public positions before other agencies in our government, before other countries have started to develop things. While there are other countries around the world that do have registration systems, I think ours is probably the oldest and most robust. I'll just acknowledge my little confirmation bias there.

[00:14:47.88] So why don't we take a look at some of the registration decisions that have come down out of our office. So the first one as you see on the screen involved a claim that we received back actually in 2018. And it was essentially a test case. It's a work of visual art that claimed to have been created entirely by AI allegedly with no human involvement. This work is called a Recent Entrance to Paradise.

[00:15:12.83] The applicant, Stephen Thaler, identified the author as the creativity machine and he himself as the copyright claimant by virtue of being the machine's owner. He described the work as autonomously created by a computer algorithm running on a machine. So when we were examining this claim, we weren't working from a blank slate.

[00:15:34.67] As far back as the late 19th century, courts have been looking at the scope of copyright in works created using machines starting back with a camera. Looking at the reference to authors in the Constitution as well as language in the Copyright Act, courts have concluded that human authorship was required for copyright protection. They've rejected copyright claims, for example, in a selfie taken by a monkey and in a book purported to be authored by non-human spiritual beings.

[00:16:05.00] The question in each case was whether the work reflected the human's creative choices even if the machine was used as a tool. So applying these well-established principles to Mr. Thaler's claim, the Copyright Office's review board refused registration. Thaler then sued in federal court because a review board is final agency action. And he sued, arguing that the Copyright Act does not explicitly restrict copyright to human-made works and that any ambiguity should be resolved in favor of protection in order to encourage maximum creation and dissemination of works for the public benefit.

[00:16:46.91] In August, just two months ago, the US District Court for the District of Columbia in the case of Thaler versus Perlmutter issued a decision agreeing with us. It held that human authorship is an essential part of a valid copyright claim. Mr. Thaler has said he's going to appeal.

[00:17:04.38] Now, the Thaler case involves a clear and straightforward set of facts on authorship because he maintained that the work was created entirely by AI. Other applications that we have received are presenting more challenging questions. What about works produced through a combination of human and machine creation? What about

what type and amount of human contribution is enough to merit copyright protection?

[00:17:37.26] So in other case has involved this graphic novel called Zarya of the Dawn, which included text written by the human applicant, Kris Kashtanova, and visual images generated by the AI technology Midjourney. When the office learned that the images had been created by AI, we sought more information from the applicant. And that's called correspondence where our examiners go back and forth with the applicants to better understand the scope of the claim.

[00:18:06.56] As well, we ended up proposing-- canceled the initial registration. In its place, we issued a more limited registration that covered only the human authored elements of this work. And that is the text as well as the selection, coordination, and arrangement of the texts and images. We excluded from the scope of registration the AI images themselves because they lacked human authorship. And I really invite you to read our letter on that so you can really understand the scope of the work we did to understand what was being claimed.

[00:18:47.92] So the rapidly growing use of AI and the attention generated by some of these cases like Recent Entrance to Paradise and Zarya of the Dawn meant that the Copyright Office began receiving requests and questions for further guidance on how do applicants file claims for copyright. And to address the copyrightability and registration issues raised by these questions, on March 16, we issued registration guidance which confirms the requirement of human authorship and instructs applicants to disclaim AI-generated content in their works if it is more than de minimis.

[00:19:27.22] So what does that mean? So what it means is that the question to ask is, would the human-generated material standing on its own be copyrightable if it had been created by a human author? If so, then a very brief description, a statement disclosing the work includes AI material should be included in the application for registration.

[00:19:53.14] This is very similar to our long-standing requirements to disclaim copyrighted material owned by others that has incorporated a new work or you have to disclaim if you're using public domain work. It helps to identify the scope of what is being copyrighted for the public record. So our goal, to be clear, is not to set up barriers to registration, but rather to help people avoid problems such as maybe later questioning about the validity of the registration by third parties or by the courts.

[00:20:23.64] We do expect to further develop and refine this guidance based on the comments we've been receiving over the last couple of months that we've been engaged in this as well as our continued review of applications involving AI. As an aside, I would say maybe you would think, well, wow, with all this, you guys are getting tons and tons of applications. And actually, we're not. And we only

have a little over 200 applications for claims that involve AI-generated material right now. So I think people were expecting a lot more.

[00:20:52.63] Obviously, I think another thing to keep in mind is that technology is constantly changing. So everything I say today is as of the state of technology today. We have also been working on trying to address even more questions we've received from the public about this guidance. So on June 28, the office hosted a public webinar to discuss that guidance.

[00:21:18.22] The head of our registration program went through a variety of potential hypotheticals, provided examples, and answered questions. It's also available online right now. I think a lot of the initial feedback we've heard is that people are finding that webinar very helpful.

[00:21:36.20] So I think there are also going to be a lot of cases where content creators are going to want to register works creating the AI-generated material. I mean, this is-- that is the big question of the day, isn't it? So how it would work would be-- is let's say an author uses AI to brainstorm a plot and characters but then writes her own book herself and omits any of the AI-generated material in that book. In this case, there's no need to disclose or disclaim anything.

[00:22:10.23] But if the author wrote the text herself and included illustrations generated by AI like the Kashtanova case, Zarya of the Dawn, that-- she's going to need to check on the box on the registration form that she is excluding that 2D artwork and/or else provide a very brief statement in the application explaining the role of the AI in her work. And the webinar goes through the examples. I just want to reiterate, we're not looking for an Excel sheet that says every single thing that you did.

[00:22:40.83] We're looking for some basic guidance, so we can examine. And if we choose to, our examiners can engage in correspondence to understand the scope of the claim. I think Blake mentioned a little bit on this very important case. The most recent example of complex registration claim came last month when the office's review board upheld a refusal to register the work of a visual art called Theatre D'opera Spatial, which was created by Midjourney by an artist named Jason Allen.

[00:23:12.71] And this case is the one that is probably familiar to folks here in Boulder as it won the 2022 Colorado State Fair's annual fine art competition. As you can see on the slide, the initial image generated by Midjourney is on the left. And the work that Mr. Allen sought to register with us, the deposit copy, is on the right.

[00:23:33.41] When the registration claim was originally submitted to us, there was no mention of the work having been made or based on Midjourney. But because of the national headlines revealing that AI contributed to this work, our examiners requested additional

information from the applicant. And in response to our questions, Mr. Allen stated that he input numerous revisions and text prompts of at least 624 times to arrive at the initial version of this image.

[00:24:02.70] He also explained that he used Adobe Photoshop to remove flaws and create new visual content. And he used Gigapixel AI to upscale the image, increasing its resolution and size. At the end of the day, Mr. Allen declined the examiners request to disclaim the AI-generated aspects, and he reasserted his claim to the copyright in the entire work.

[00:24:26.46] But the office concluded that the work could not be registered unless the claim was limited to the authorship that the human author himself contributed. And consistent with our registration guidance, the office accepted the claim that the human authored visual edits made with Adobe Photoshop were de minimis and did not need to be disclaimed. But the more extensive features generated by Midjourney and Gigapixel AI needed to be excluded as non-human authorship.

[00:24:57.54] The ultimate-- the application brought the case to our review board and, as I mentioned, the next level of appeal. And the board affirmed that Mr. Allen could not be considered the author of the initial image since his sole contribution was inputting the text prompts that produced it. Because this AI-generated material, which could not be protected by copyright, was more than a de minimis contribution to the work. it needed to be disclaimed.

[00:25:23.94] In the board's view, even though Mr. Allen described the text prompts that he took to produce the image, the resulting image was entirely dependent on how the Midjourney system processed those prompts. The review board opinion points out that Midjourney does not understand grammar, sentence structure, or words like humans. Bottom line is if Mr. Allen were willing to disclaim the AI-generated material, he could file a new application and explain how his modifications to the image rose to the level of copyrightable authorship.

[00:25:59.97] So as I mentioned at the start, our office does do more than registration. We also look at the larger law and policy issues and the policy implications of copyright and AI. And we are asking a lot of questions. And today, I'm going to have a lot of questions. And since the review is open, I'm not going to have a lot of answers to these questions, and I'm looking very much forward to the discussion later.

[00:26:22.62] On the infringement side, we're seeing more and more discussion of the legal implications for the incorporation of copyrighted works into the training of large language models. This is the ingestion issue. And this includes several high-profile lawsuits that have been filed by copyright owners also mentioned I think by Blake early on.

[00:26:43.41] There's a large number of questions. And we're all-- we're trying our best to explore many of these through our notice of inquiry. So for example, does employing copyrighted materials to train generative AI systems constitute infringement? And if so, under what circumstances? Do uses require consent and compensation from the copyright holder, or do they qualify as fair use?

[00:27:06.19] There is considerable debate over how existing case law applies, notably, the Second Circuit's decision in Google Books and the recent Supreme Court case in the Andy Warhol Foundation versus Goldsmith. If consent or compensation for training is required, how can this be feasibly accomplished, given the volume of works ingested? Could this be done or should this be done through direct licensing, collective licensing, or some form of compulsory license?

[00:27:35.48] Our current copyright system incorporates all three models in various other contexts. And I think we're all kind of blown away by the scale of AI. And that is a major factor to consider going forward in any solution. As a practical matter, how can copyright owners find out whether their works have even been used in the machine training? What information about the process is retained by the AI developers, and what level of public transparency is appropriate, feasible, or desired?

[00:28:07.81] If copyright owners are given the ability to either opt in or opt out in having their works included in training materials, how is that going to work, especially for individual creators? Also, is it even consistent with copyright law to require owners to opt out of the unlicensed use of our system? Who should be ultimately responsible for infringement, either during the training process or for the ultimate output, the assembler of the training material, the entity offering the tools, the user of the tools?

[00:28:41.89] Besides authorship and infringement, the final major area of question in our notice looks at the boundary between creativity and personal identity. Performing artists, in particular, have raised questions about AI-generated work imitating their likenesses, their voice, their style. We've all seen reports of deep fakes, AI-generated works that purport to be the Drake and the Weeknd and other examples.

[00:29:08.57] And so while current copyright law does not reach this activity, there are various state laws involving rights of publicity, unfair competition, misappropriation that might apply, but it's very mixed across our country. And what are these-- how are these going to-- laws going to address the new technologies? So take a deep breath. That's a lot of questions. Believe me our notice went through a lot of internal edits just to get to the 34 numbered questions that we have disclosed to the public.

[00:29:38.18] We want to find a way or ways forward that ensure that human creativity can continue to thrive, but also that promotes, rather than inhibits, the development of this exciting technology. So I've spent most of my time right now on what the office is doing. But as we know, AI raises more issues than just in the copyright basket, things. Like bias, access to public data, transparency, cybersecurity, competition, all those very important issues are kind of outside the scope of my lane. And so there are other agencies that are looking into that.

[00:30:15.62] There are also members of Congress who are in various subcommittees that are looking at issues as patent inventorship, competition, security, election safety, national defense, among others. There have been copyright hearings held in both the House and the Senate. Senator Schumer has his Safe AI Act-- it's called the SAFE Act Innovation Framework that he announced in June and his plans to hold some insight forums on various issues over the next couple of months or so. It hasn't really been rolled out yet.

[00:30:49.25] Other members of Congress are also looking at various angles of AI. In addition to Congress, back to what I said at the beginning, there's other agencies looking at AI as well. So for example, the US Patent and Trademark Office has workstreams on inventorship issues as well as standards. The relatively new National Artificial Intelligence Advisory Committee, commonly called NAIAC, which was established by the National AI Act of 2020, they advise the president and the National AI Innovation Office.

[00:31:23.57] The Federal Trade Commission just held a roundtable public session on generative AI two days ago. The White House has its blueprint for an AI bill of rights. NIST has its AI risk management framework. NTIA is working on its AI accountability policy. So you can see there are-- everyone's interested in AI.

[00:31:46.31] We are also keenly aware about the global nature of generative AI. We can't confine our attention to US borders. And since AI systems are developed and deployed internationally using content from around the world, we need to be aware of developments globally. And so my office is speaking with copyright policy makers in other countries about their approaches, specifically and especially on the authorship issue, registration, copyrightabilities because that's clearly in our bailiwick.

[00:32:16.44] There is a lot of interest on what we're doing in the United States because our office got so far out in front. We have a registration system, and we play a major role, especially in advising Congress. And so I will say why-- there seems to be widespread agreement across borders that copyright requires human authorship.

[00:32:35.19] That is by no mean absolute. There are other jurisdictions like the United Kingdom and similarly situated legal

systems that follow their model, which does provide copyright protection for works created by a computer. And they give that protection to the person who arranged that work.

[00:32:51.36] There are also differing approaches to exceptions and limitations. Some countries have very specific text and data mining exceptions. The pending AI Act in the European Union has transparency requirements mandating the publication of summaries of copyrighted works used in training. Now, that AI Act still has to go through a trilogue negotiating process, but that issue is very live in Europe and elsewhere.

[00:33:18.93] If you're interested in diving deeper on these international matters, as I mentioned, we hosted a webinar in July. It's online. It included leading copyright scholars from around the world talking about many of these issues. Our notice also asks whether approaches developed in other countries should be adopted or should be avoided here. We always have lessons learned to look at.

[00:33:42.64] And some of these are still early days. And it remains to be seen how many different national laws will be interpreted in the context of generative AI. So to wrap up, and I'm sure we're going to be hearing more in today's conference, I want to make a public invitation. Please participate, share your views to our docket.

[00:34:05.15] We have extended the deadlines for October 30 for initial comments and November 29 for reply comments. To be clear, you do not have to answer all the questions. Don't be intimidated. We are going to be reviewing all the comments that we received to this notice. And we are going to be using our human eyes. Our team is going to be doing this.

[00:34:26.18] Right now, we have almost 8,000 comments that are on Regulations.gov. We're also monitoring developments around the world, as I mentioned, but also developments in the marketplace. There's a lot of news that keeps changing and being reported every day on how these companies are either coming up with contractual agreements, changing terms of ingestion, people buying other companies.

[00:34:49.25] There's a lot going on in the marketplace. And that is often instructive as to what can the private marketplace produce and contribute in positive ways to these complex legal issues. And so as for next steps, we're looking to update our compendium of practices, update our registration guidance for applicants. We plan to produce one or more reports with analysis of legal issues and recommendations for any possible desired legislative or regulatory change.

[00:35:17.55] I don't have any expected dates when we're going to release this, but I can say it's not going to be this year. And with that, I want to thank you for your attention. I look forward to your questions. And thank you for having me participate. Thank you very much.

[00:35:30.00] [APPLAUSE]

[00:35:36.66] BLAKE REID: All right, well, a tour de force of the vexing issues at the intersection here. So it is a Silicon Flatirons tradition under the Phil Weiser rule, our former director, that the first question goes to a student. And I see a couple of my students parked down here.

[00:35:56.77] I made all my copyright students come today. So I hope we'll get a volunteer, but cold calling is a possibility, of course. All right, question from a student.

[00:36:12.32] MARIA STRONG: There we go.

[00:36:13.50] BLAKE REID: And if you could wait till the microphone comes down and introduce yourself, so the folks online can catch us, that would be great.

[00:36:23.44] AUDIENCE: Thank you. I'm doing this to avoid a cold call. Thank you so much for being here. I'm wondering if you could speak a little bit if you have looked at the comments that are coming in. Are the majority of those-- what is kind of the breakdown? Are they coming in from lawyers? Are they coming in from creators?

[00:36:44.06] I know lawyers are probably more likely to be familiar with the process, but I'm interested to know who is speaking on this as you're compiling these comments.

[00:36:52.99] MARIA STRONG: Yeah, no, no, thank you for that. So what we do is we use another government agency system called Regulations.gov. And that's where all the comments will get funneled to. And like I said, we have almost 8,000 right now.

[00:37:05.09] Right now, we're going through them to make sure that there's no errors. Like sometimes you can click a box and say your one comment was 10,000 comments. And that kind of makes the number go crazy. We're looking mostly to make sure that we're going to be able to hit Post when the day comes or the day after the deadline.

[00:37:25.50] I will say right now based on our preliminary review, most of the comments we received so far are individuals. If I were to outline a trend, so far, it's mostly creators who are very concerned about AI maybe upsetting their livelihood. I don't want to venture a conclusion, but I will note that many of these folks are also participating in various strikes involving the motion picture industry. And so I think that's part of it.

[00:37:59.64] But to the extent individuals contribute a major part to the creative economy, this is exactly the kind of thing we want to have. We do expect that the number will increase exponentially as we get closer to the deadline with more formal comments from whether it's academics, trade associations, companies, small companies, more individuals. We do get, and I do expect in this docket, a very wide variety of interests.

[00:38:27.00] I think it's really important to keep in mind that, as I said earlier, it's not binary. It's not content, it's not tech. It's never been that. And so it's-- people who are large companies have different views than people who are small companies. And so we expect to see a robust round of comments. And like I said, I don't want to venture against how many more 1,000 more we will get, but we are going to do our best to get them posted so that everyone can review them.

[00:38:51.43] And then if they wish-- provide reply comments, I have spoken to other countries around the world. And they are looking at this very much very closely because it will be a public repository of a state in time at the end of 2023 where the public sees this. So please look. Thanks.

[00:39:10.53] BLAKE REID: I'm tempted to follow up and ask whether you're going to get any comments from artificial intelligences and how they'll be received by the office. I don't want to put you on--

[00:39:18.81] MARIA STRONG: Yeah, we think we've seen a couple, just a very few. They will count. We're not going to be doing any kind of registration examination of comments, but I do expect that. I mean, in prior dockets, we have seen grassroots campaigns, foreign letters, and that's fine. But I do expect probably to see some. We'll find out.

[00:39:43.16] BLAKE REID: All right, open it up to the audience for questions. And someone standing right by a mic.

[00:39:50.06] AUDIENCE: Thanks so much. Dan Murray, Rocky Mountain AI Interest Group. We did have Jason Allen speak at our group, the author of the Theatre D'opera Spatial image you mentioned. So I have a question.

[00:40:04.17] The current approach on copyrightability of AI-generated images, if someone generates the image with an AI tool like DALL-E 3 or Midjourney and then changes it like, let's say, they bring it into Photoshop and they make changes that are not de minimis, is the copyrightability then that you apply for copyright of that image and disclaim the parts that are AI generated? Is that the current approach. That's what I understand.

[00:40:33.50] MARIA STRONG: Yeah, you need to disclaim the part that is more than de minimis. So when we looked at Mr. Allen's original Midjourney thing, we found that there was a lot of AI-generated material there. And that's what he needed to disclaim. He chose not to disclaim it, and that's why we had to refuse his registration.

[00:40:53.56] But we also would note-- and I'm not going to get ahead of our registration department. But if you take a look at certain parts of his-- the image that he submitted for registration, there is arguably some pretty creative things that he did in other sections of that image. And that's why we said if he wanted to get at least the copyright in that

part of the image, he would need to file a new registration so that our examination folks could do that.

[00:41:20.03] So this is the exact problem of what's the difference between AI generated and AI assisted? Because many industries-- I mean, music, just think of music, there's been a lot of artificial intelligence, computer-generated things that assist in the creation of music. Same thing in art, I mean, Adobe Photoshop.

[00:41:41.53] There's a lot of tools out there that are done at the initiation and the control of a human. And so back in 1965, our office actually recognized that we don't have an absolute rule saying no. We are looking for, and I think I can quote, whether the work is basically one of human authorship or whether the computer or the device merely assisting instrument, but the--

[00:42:08.59] What we're looking for is the traditional elements of authorship that were actually conceived and executed, not by man but by a machine. So that is also a question of our inquiry. And that's why we're trying to make those decisions and get better guidance on that because things keep changing. So back to your point, if Mr. Allen wanted to resubmit after he disclaimed the non-AI part, our registration team would take another look.

[00:42:36.43] BLAKE REID: Come down to Professor Bracha in the front row.

[00:42:39.10] AUDIENCE: Thank you. Oren Bracha, the University of Texas here. So again, about authorship, I wanted to ask you the following. There is this argument out there that in some cases, a human can establish authorship simply by virtue of the creativity of the prompt. So to be clear, it's not a case of selection and arrangement of AI-generated materials or adaptation in any way.

[00:43:05.17] The claim is sometimes the prompt itself if rich and creative enough makes the human who just created the prompt the author of the work, whether because of the richness of the prompt or some reiterative process. Does the office have any position on that?

[00:43:21.16] MARIA STRONG: So I will say that is a question in NOI, but I would point you to both Zarya of the Dawn as well as the Theatre D'opera Spatial letters where we took administrative notice of how Midjourney works. And so if you take a look at those opinions, we stayed at based on the examination of those works, realizing other works might be different, that the prompt was merely an extraction. It was not done by the human.

[00:43:51.44] And I think, again, right now if you take a look at the technology, I'll use a Midjourney as an example, there is always going to be an unpredictable nature of the output. You could prompt, prompt, prompt away. And even the folks who are behind the algorithms themselves have said, we're not quite sure how you get to

that last output. So if that's the case, the machine do the output, not the human.

[00:44:17.36] But I will also agree with your question, that we would welcome further input on the prompt question because I think that is something we're going to have to definitely be looking at. Thanks.

[00:44:27.96] BLAKE REID: I think we go to the question in the far back.

[00:44:35.10] AUDIENCE: That was really interesting. Thank you. And I have a-- my name is Steve Heuser. I work for POLITICO, and I have a more of a Washington tech policy type question. Your Copyright Office is part of the Library of Congress officially, right?

[00:44:47.69] So do you-- on these topics, are you all getting guidance directly either from Congress or even from the White House on this? And if you're not, what kind of guidance you would you like to see? If you are, what is it?

[00:44:59.86] MARIA STRONG: So yes. So that's-- we are in the Library of Congress. We are in the legislative branch. Congress, as I showed you, has directed us. We are their primary guidance and advisor on the law. So I can pick up the phone and talk to a staff member, represent my agency and give them impartial legal advice for the question they want.

[00:45:19.58] If you're in the executive branch, let's say the Patent and Trademark Office Department of Commerce, any of these other groups, what they usually have to do is you have to go through your government relations team in order to talk to Congress because that's sort of the procedure, right? Because the executive branch needs to come together. They have a lot more cross-agency coordination because they want to speak on behalf of the executive. And there's a lot of executive branch agencies.

[00:45:45.81] So does the executive talk to Congress? Sure, they do. But I think the initiative that we have right now is being led by the Copyright Office. I have a lot of colleagues who are on the copyright team at the Patent and Trademark Office. And we are keeping them informed of the work we're doing, but the legal analysis and the ultimate conclusions will be the work of our office.

[00:46:09.72] BLAKE REID: All right, we've got a question in the third row over here.

[00:46:16.07] AUDIENCE: Hi. So I'm intrigued by the big difference between the US and the UK's view. And I guess I have two questions. One, does that reflect sort of an historical distinction between the two countries and how they view copyright? And my second question is, who's the outlier here globally, the US or the UK?

[00:46:43.12] MARIA STRONG: On the first question, obviously, our legal tradition comes from a lot of statute. And there you go. But yeah,

the provision that they have in terms of the computer work, that's been there actually for quite some time. So I think they are still trying to find a way to use it because my understanding up until recently, it hasn't really been a big issue for them. They are taking another look at it.

[00:47:08.00] I believe there are similar provisions in-- I want to say Australia, maybe New Zealand as well. In terms of an outlier, I will say that the really interesting part about copyright law, it's national. There's a lot of themes, but the actual way in which laws are written and definitely interpreted varies territory by territory.

[00:47:27.16] I would say that there's some discussion on what if any kind of new right might be needed to address this kind of situation. There's been some in the academic world that have suggested that maybe a sui generis might help solve this AI thing, much like the database work in the EU did. I think that's-- I'd like to hear some more analysis on that.

[00:47:54.71] I think the outlier-- it's not that the UK doesn't respect authors. It's just that they have this additional sub-segment. I think that most copyright offices right now do heavily believe that human authorship is a very important element. I think it goes back to the Berne Convention. I mean, Victor Hugo Les Miserables fame helped create the Berne Convention.

[00:48:18.50] There weren't computers back then. And they were very much interested in protecting the artistic rights of authors. So I think the human authorship element is going to stay pretty firm. I think where the development might be is to the question you meant. And I wouldn't use the word outlier because what if anything, is there going to be another way to address this?

[00:48:39.90] And then that question raises the next issue, is oh, well, if you country over here have a sui generis protection for, let's say, some kind of AI assisted or AI-generated work, and I'm a country over here that doesn't recognize that, this person becomes the outlier because their work can't be protected over here. So there's a big challenge there. If you're looking at trying to find a common understanding across nations, that is going to be a major challenge. Long answer. Sorry.

[00:49:11.48] BLAKE REID: I think we've got time for one or two more questions. Why don't we go right back here.

[00:49:22.05] AUDIENCE: So I think one thing that's interesting is-- so I'm a lawyer. I work at Red Hat. And the lawsuits that are getting the most attention and what everyone obviously has talked about, which is understandable, is like the sort of visual arts, right? It's a lot easier for us to grasp and--

[00:49:40.44] But there is a lawsuit, of course, against Copilot and GitHub. And it just-- I think like with my lawyer hat on, it's like, well,

these are all things that are protected by copyright code and paintings and everything else. However, there's something like visceral from a human perspective an indifference, right?

[00:49:58.95] I mean, arguably, even though source code is protected by copyright is also-- has a lot more usefulness than what we think of as art art. No offense to any of my developer friends. I think they would agree with that. So I was just curious if you-- it sort of cuts I think to that final question, and then it starts to bleed into what I would say maybe even a moral rights kind of protection, which we don't have here.

[00:50:22.85] So I was just wondering if you-- one of the things I've wondered-- and I don't know if you can say this, but is if we will end up with a splicing, if you will, of how much protection is for different types of works that are generally lumped together?

[00:50:36.27] MARIA STRONG: Yeah, I mean, I think that is-- the issue of looking at name, image likeness, and styles is-- and whether rights of a publicity or misappropriation at the federal level is very much a live one. There's been interest on both sides, both the House and the Senate, on that based on the recent hearings. The question would be, could such a law be developed at the federal level, and where would it live, right?

[00:51:05.35] I would bet or guess it would probably not live in the Copyright Act, but [AUDIO OUT] in the US code. I think that's possibly the way things are going to go, but there's a lot of challenges there, right? I mean, copyright also builds on style.

[00:51:22.88] I mean, is style the right word? I think it's different. I think it's going to be more likenesses because there is, as you mentioned, a very visceral image or visceral reaction to an artist or a singer or an actor who wants to be not only credited that their name is associated with their work, but have the right to control people who use that. So I think that's-- that area also is very important for the issue of deepfakes, election false advertising, which is again, being raised in Congress as well.

[00:52:00.06] And so I think-- I'm not going to predict, but I think you can see that I think there's a train going towards that direction on the hill. Yeah.

[00:52:07.22] BLAKE REID: All right, noting the time, I'm sorry because I see a lot of hands up. I'm going to have to invite us to continue this conversation next door at the reception. So come to our first break now. We'll reconvene in 15 minutes, so 10:15 Mountain Time. And we will see you back. But please join me in thanking Associate Register Strong.

[00:52:28.05] [APPLAUSE]

Panel: The Copyright Law of Generative AI

https://youtu.be/XjlhrqK80f4

[00:00:00.14] KAILEY LAUTER: Good morning, everyone. My name is Kailey Lauter. And I'm the current treasurer of Silicon Flatirons Student Group. And as a partner of Silicon Flatirons, we host a lot of events for students, which primes them and gives them a little background for conferences like this. And we also host forums for them to engage with some of the speakers.

[00:00:27.12] So yeah, and we also help students find mentors and jobs and internships. So if any of you have any desire to engage with the student group, I'm happy to let you know of opportunities to get involved.

[00:00:40.55] And it is my absolute pleasure to introduce our next panel, The Copyright Law of Generative AI, moderated by Professor Surden.

[00:00:53.53] HARRY SURDEN: Great. Thank you so much for that introduction. Really appreciate that. So I'm Harry Surden. I'm a law professor here at the University of Colorado Law School, and also the faculty director of our artificial intelligence and law initiative at Silicon Flatirons. And I research issues related to artificial intelligence and law.

[00:01:13.81] I want to thank Professor Blake Reid for organizing such a terrific conference, Professor Bernthal for organizing this, and as well as the terrific Silicon Flatirons team for doing such an amazing job. So please join me in a round of applause for those.

[00:01:28.83] [APPLAUSE]

[00:01:33.66] Thank you also to Registrar Strong for such an interesting presentation. I learned so much.

[00:01:39.87] Joining me today are a terrific panel of experts. In the interest of time, I will only provide a very brief introduction by encouraging you to look at their extensive bios online as a very impressive accomplishments.

[00:01:54.81] To my left, I have Andy Gass, who's a partner at Latham & Watkins, and who runs a copyright practice. And he's been involved in many important copyright cases involving AI. Next is Cynthia Arato, who is a partner at Shapiro, Arato, and Bach, LLP, who is a well known IP litigation practice generally, and specifically, within artificial intelligence itself.

[00:02:21.28] Next is Viva Moffat, professor of law at our friends down the road at the University of Denver, who is a well known expert in intellectual property law and an important commentator on AI and copyright issues. Next, we have Professor Oren Bracha from the University of Texas School of Law. Professor Bracha is a scholarly expert in intellectual property, and in particular, the history of

intellectual property law. And he's also written extensively about AI and copyright.

[00:02:51.94] And last but not least is Katherine Lee, who is a PhD candidate at Cornell University who is a leading voice in AI language models and copyright and privacy.

[00:03:04.76] So we have this terrific panel of experts. These are the right people here at the right time at this issue that is generating tons of interest among the public. And before we get to our questions, I'm going to provide a little bit of background. I know that not everybody is an expert in copyright or AI, or maybe some of you are just joining us and didn't get either access to the primer yesterday or to the introductory remarks. So let me just try and situate the questions generally before we hear some of these ideas from our expert panel.

[00:03:40.93] So let's begin with the basic question, what is generative AI? And I think it's worth explaining that. Generative AI or artificial intelligence is the umbrella term for artificial intelligence uses in areas of human creativity, such as art, music, language and text, video, et cetera.

[00:04:01.34] So what does that mean? Well, artificial intelligence can be used in lots of different areas outside of human creativity, prediction, self-driving cars, robotics, identifying cats in photos, what have you. So when you're using artificial intelligence for particular uses associated in areas of human creativity, we call it generative AI. So how does generative artificial intelligence relate to copyright?

[00:04:28.51] We will see, there is lots of different ways, as Registrar Strong talked about, but there are two main ones that have dominated the discussion. The first is the use of copyrighted works for training artificial intelligence systems. I'll have more to say about that in a second. And the second is whether works produced by artificial intelligence systems should be copyrightable, something that Registrar Strong discussed.

[00:04:56.71] So let me just give you a brief background of the first and the second, and why they're related, and we'll-- to copyright law. And then we'll throw it to our panels for some opinions about how we should solve all of this.

[00:05:10.24] So generative artificial systems like ChatGPT, many of you are familiar with that, produces fluent-like text. In other areas, we have generative AI systems like Midjourney, which can produce art or DALL-E, also generative music systems. These are part of an artificial intelligence approach known as Machine Learning.

[00:05:32.15] And what Machine Learning does is it needs to be exposed to a large amounts of data amount of data, and learns patterns in those-- in that data to be able to produce large amounts of text. So that's called training.

[00:05:45.25] For example, ChatGPT, how did it learn how to produce fluent human-like text? It was exposed to billions of words, a large portion of the internet, and it learned the various patterns of language and thought among those documents on the internet. Similarly, how did Midjourney or DALL-E learn how to make images, novel images, like the one that Maria Strong showed? It was exposed to lots of different photos and paintings and drawings out there on the internet, and learn the patterns encoded in those pieces of art. So that's known as training.

[00:06:26.38] Now importantly, most huge percentage of those things, either text or images that these systems were exposed to, are copyrighted. So that's where we-- copyright comes into the picture. So-- and the question is, is that permissible under Copyright Law to use copyrighted images or texts to train an AI system? Or is it not permissible?

[00:06:51.96] Now-- and then the other question is, what about art that is produced by these systems? Or is the art or the text produced by ChatGPT-- is the text copyrightable? Is the art produced by Midjourney copyrightable? We've already heard an opinion from Registrar Strong.

[00:07:11.36] So where does this fit into copyright more broadly? So Copyright Law is primarily about allowing the creators of creative works to capture the monetary value of their works by stopping others for using their works without paying or paying too little. So that's one main value.

[00:07:29.79] On the other hand, we have something called Fair Use, which is an important aspect of copyright law as well. So Fair Use is permission-free uses of copyrighted works that are not infringement. So think if you're quoting a news article that has a criticism of somebody's book or something like that, that would be your prototypical use of Fair Use. So Fair Use is typically around education, speech, criticism, and also transformative art, and in some cases, allowing technology to work together.

[00:08:06.94] So when we have these two main pillars of copyright clash, we have artists and creators trying to get the values of their works, on the other hand. And the uses of these works for other things that society deems important, whether technology, or art, or speech, or criticism, this is where these things clash.

[00:08:24.36] And that's what we have today in the world of generative AI. We have AI which is deemed as very important and is helping a lot of discoveries. On the other hand, we have creators whose creations are not-- who are being used without their permission, and who feel like they should have some say in whether their creations are used at all or maybe they should be compensated for that.

[00:08:47.20] So with that background, let me dive into our first question for our first panelist, Andy. And let's talk about Fair Use.

Should copyrighted-- should AI companies be allowed to use copyrighted works in training? And then the subtext of that is, is it-- should that be considered Fair Use under current law? Is it Fair Use under current law? And what are the pluses and minuses of the situation? And then we'll open it up to the rest of the panel to weigh in on this. So Andy.

[00:09:23.28] ANDY GASS: Thanks, Harry. And thanks so much to all of you for having me here. It's a great privilege to be with all these people. And hopefully, we're going to have a lively conversation about these issues.

[00:09:31.54] I should make clear, at the outset, I'm speaking in my capacity as a private citizen, not on behalf of any clients whom I represent in this area.

[00:09:39.07] So I guess I'll answer that question by starting with a slightly more detailed overview of what is happening in these technologies, and then maybe do a slightly deeper dive into how I see the law working.

[00:09:51.55] So when we talk about companies using copyrighted content to create generative AI models, what we're basically talking about taking either pursuant to some contract with someone or not, unimaginably large quantities of the medium at issue.

[00:10:13.89] Take text for example. Billions or even trillions of words that are found somewhere in a state of nature, whom someone wrote at some point. And what the technology does is it feeds this-- all of these discrete components of text into a neural network of some sort. And the purpose of that is to allow the network, to, as you alluded to, derive almost incomprehensibly large volumes of statistical insights about the underlying correlations between the constitutive elements of the texts that are being fed.

[00:11:02.55] So in a sense, what is being done is like, no human being is seeing these texts. They're being tokenized, broken apart, and studied by this computer program, which is then saying, aha, I have determined to take a radically overly simplistic example. If the first two words that I see are, "Today is," the next two words are much more likely to be, "a Friday" than "a bookcase."

[00:11:31.47] And it's those kinds of insights performed at, again, an almost incomprehensible scale in like, when I say multi-dimensional vectors, these are 75 dimensional matrices of correlations and statistics, that allow the tools to then generate something in response to a query.

[00:11:54.28] And there's a bunch of other stuff that happens between the creation of what's called the large language model and the output, but at bottom, we're focused at the initial phase in the Fair Use question on what's going on with the copies that are made.

[00:12:09.97] And for the base case that I'll describe-- and this may not be all cases, but for the base case that I'll describe, imagine that what's going on is that the company has factually made a bunch of copies, in the copyright law sense of that term, of copyrighted material here, texts, in the example I've been talking about. In the service of teaching this model to, after a bunch of human intervention and fine tuning, spit out things that do not themselves resemble any particular item that was in the training set. So that's the case that I want to talk about in the first instance.

[00:12:51.86] So for example, we're talking about teaching a model to produce an output that says like, hey, Andy Gass' bespoke generative AI model, who was the greatest basketball player on Earth? And the model says, opinions may vary, but I think the leading answer is probably Michael Jordan. That's the kind of thing we're talking about. We're not talking about, hey, model, please regurgitate for me the first chapter of Harry Potter. And it gives you that. That implicates a different set of issues.

[00:13:23.72] So in the base case that I'm talking about, the legal question presented is we have a bunch of copying that has taken place in our system of Copyright Law. The copyright owner is the one who has the exclusive right to reproduce a work in copies. That's 17 USC 106(1), for anyone following along at home. But that's not the end of the story because it's not an infringement if those copies were Fair Use.

[00:13:55.19] So that then brings us to the question of, how do we think about Fair Use in this context? So Fair Use is a doctrine that in the American legal system, goes back to at least the mid 19th century. Justice Story is the one who is originally credited classically with having devised the doctrine. There are traces of it in the English system that predates the ratification of the US Constitution.

[00:14:24.91] And the current embodiment of the doctrine is specifically, if you look at the legislative history of the statutory provision at issue, which is section 107 of the Copyright Act, it's specifically intended to instantiate the common law tradition that predated the codification of Fair Use and the Copyright Act at all.

[00:14:46.64] So what that usually means is while we have these four factors that are written into the statute that I will not burden you with right now, from the text of them, it's probably difficult to figure out how any given dispute would be resolved properly. And instead, what we do as lawyers is think like, well, we've got close to 200 years of history of judges dealing with situations that may or may not be analogous or may be analogous in certain respects but not others.

[00:15:18.89] And if you put together that constellation of precedent, where does this fact pattern line up? So you can think about it from first principles, but the way that we really do it is by reference to these pre-

existing judicial decisions. And so what are the relevant judicial decisions in this circumstance?

[00:15:36.44] So as Harry, I think, rightly alluded to, there is a whole long line of cases in the federal courts in the United States that deal with the question of when and to what extent a technology company can make copies of copyrighted works, not for the purpose of showing those copyrighted works to the public, but instead, for the purpose of creating some new and innovative product that didn't exist anymore. And the courts have grappled with like, do we always allow that? Do we sometimes allow that? What are the circumstances in which we allow that?

[00:16:16.48] And in my view, the general rule that emerges from those cases is a principle that says, in that circumstance where what is happening is the technology company is factually making copies in the-- but not showing them to people, in the service of creating a new product that didn't previously exist.

[00:16:35.50] The courts are going to call that back end invisible to the public copying Fair Use to the extent that the output of the process would not itself be a copyright infringement if it had just come to be without all of that copying that took place.

[00:16:51.49] So what do I mean by that? I mean, take for example, the case of a video game company. And imagine that there is a proprietary operating system that a technology company runs for a platform for operating video games. We'll call that company Sega, OK?

[00:17:11.11] [LAUGHTER]

[00:17:12.58] So Sega has an operating system. And Sega has its own video games that run on the operating system. And imagine that you have a challenger who comes along, and says, I would very much like to create my own video games to run on your operating Sega-- operating system, Sega. And imagine that Sega says, no, thank you. Our system is a closed system. It's a copyrighted product, our operating system. You may not interoperate with us.

[00:17:40.21] And imagine that this challenger video game company comes along and says, well, I take your point, Sega, you prefer that we not do this. But what I want to do is just study your operating system so that I can technologically figure out how to make a video game cartridge that's going to work when someone plugs it into the console that you've created. And I'm going to do that. I'm going to create as many back end invisible to the public copies, as I need to, of your copyrighted operating system in the service of what? In the service of producing my own new video games.

[00:18:17.71] And they're not going to be copyright infringements of your video games. They're going to be different video games. But they are going to compete with your video games. So we're clear.

[00:18:27.16] In that situation the precedent that we have is that the court says, Fair Use. You are allowed to do that.

[00:18:36.25] Similarly, imagine a company that says, I would really like to create a product that makes the texts of books searchable for people. And I'm going to do it in a way that says that when someone inputs their search query and sees the result returned, what they get isn't going to be like the full text of the book, it's just going to be like a little snippet view of the book. And I'm going to be a little thoughtful about it to make sure that a dictionary definition, where the thing itself would be a complete substitute for what the copyrighted work was is not going to be displayed to the public.

[00:19:19.83] Imagine that I want to do that. And imagine further that the authors of all of the books in the libraries of Harvard University, UC Berkeley, University of Michigan, a number of other libraries come together and they say, over our dead bodies are you going to do that. We wrote those books. That content is ours. You can't do this without permission.

[00:19:41.04] In that circumstance what the courts have said is that all of that back end copying, all of that copying in the, again, scanning physically and creating permanent, digital, lasting files of the full text of all these books is Fair Use because it's in the service of creating this non-infringing output, these snippet views.

[00:20:04.86] So what are we, as lawyers, do with this? We think, to what degree is this analogous to the base case of generative AI that I described? To what degree is it disanalogous? And the punch line is that from my perspective, the answer is those cases are very instructive. And what they mean is that the existing law that we have for the base case that I described spits out the result that all of that back end copying in the service of creating non-infringing outputs is Fair Use.

[00:20:34.50] HARRY SURDEN: OK. Great. Thank you so much for that. So what I'm hearing from you is not only should it be Fair Use in your opinion, it is likely Fair Use under current doctrine. Yes. So I can see the fellow panelists very excited to answer. So--

[00:20:49.19] [LAUGHTER]

[00:20:50.22] --let me see.

[00:20:50.82] ANDY GASS: I thought that was unassailable.

[00:20:52.20] HARRY SURDEN: Yeah, no--

[00:20:52.71] [LAUGHTER]

[00:20:54.51] So let me see what Katherine has to say about this. And then we'll turn to the other panelists because I know they've got thoughts as well.

[00:21:00.44] KATHERINE LEE: So I am, first and foremost, a Machine Learning researcher. And my research is predominantly in the memorization of training data and training data extraction from large language models.

[00:21:13.10] So with that in mind, your discussion rests on the premise that the output is non-infringing. My research and my coauthors would be very upset with me if I didn't bring this up, is that a lot of the outputs are duplicates of examples in the training data.

[00:21:30.79] Arguments that I've heard for why that would be Fair Use are that the process itself is transformative, but this argument together feels a little circular because it's saying that because the output is not infringing, because the process is transformative, therefore, the process is transformative.

[00:21:46.61] So obviously, this is not the case with every single situation because not every output coming out of a model is an exact copy. And by and large, the majority of copies-- the majority of outputs are not.

[00:21:59.71] But there is three different cases that we should consider here. One is the situation in which they are. One is a situation in which they are not. And the much harder case, which is the situation in which they're kind of similar, but not really related. And so I would love for you to expand on the different cases a little bit more because we've predominantly focused on just that second one.

[00:22:19.31] ANDY GASS: Yeah, I'd be happy to do that. So--

[00:22:22.07] [LAUGHTER]

[00:22:24.17] So I think we have common ground on the following points. It's not controversial under existing copyright doctrine to suggest that in the scenario where the output of a generative model would not standing alone be an infringement, but that's Fair Use.

[00:22:44.51] I think we have common ground on the proposition that the majority-- you say the majority, I would say, the overwhelming majority of outputs are not copyright infringement in that sense. They are not what the doctrine would call substantially similar to any particular work in the training corpus.

[00:23:06.86] So for that vast majority of output, there shouldn't be much of a question, I think, between us, that the use of the training data to generate those outputs was permissible. The question-- oh, sorry. I didn't want to interrupt.

[00:23:22.14] KATHERINE LEE: This is challenging, right? Because the substantially similar question is really challenging. And to go back to how these models are trained, they're not trained to say like, reproduce the first chapter of Harry Potter, as a response to please

reproduce the first chapter of Harry Potter, but they are trained to produce the next word that in the sequence that they are seen.

[00:23:43.42] So this is a bit-- it's a bit challenging here. And especially when you get training data into large quantities, you get distributions. And distributions are much more concentrated in the areas that you have more data, and less concentrated in other areas.

[00:23:57.02] So the questions of whether or not things are similar, primarily, to my understanding, fall into that third category of where things are maybe similar in some ways and not similar in other ways. And again, not saying that all the ways in which they're similar are problematic because there's a lot of things like facts that we know are not copyrightable. So figuring out and teasing away that part feels like it's, as you said, more of a case by case basis.

[00:24:22.81] But to go back to where this is coming from, I'd just like to say that, I feel like this is more of a case by case basis, and more of a case by case into how the models themselves are also created, because not all models are created equal.

[00:24:34.34] We mentioned before that some models have smaller training data sets. Some models have larger training data sets. Some models contain licensed work. Some models contain copyrighted unlicensed work. And so depending on the way that the models are structured and the way that they're designed, that can impact the copyright analysis as well.

[00:24:51.99] So to lump everything into one category is easy to do, but I think I would like us to tease this apart a little bit more and talk about the different ways in which these scenarios might come up.

[00:25:03.41] ANDY GASS: Totally agree. I was actually specifically trying to tease out the categories. But I actually wanted to ask you a question, and sorry for derailing us a little bit. But--

[00:25:09.80] [LAUGHTER]

[00:25:10.07] --it's-- I don't get an opportunity every day to talk to a machine learning researcher. So--

[00:25:13.82] [LAUGHTER]

[00:25:15.29] So you mentioned that there are-- that not all models are created equal, and that there are models that are trained on less stuff, and there's models that are trained on more stuff.

[00:25:26.36] KATHERINE LEE: Yeah.

[00:25:26.72] ANDY GASS: And that you study memorization, which I think, for the folks who are listening at home, means the circumstance where the model spits out something that's like a verbatim replication of something that it was trained on. So it's not, who was the greatest basketball player of all time? Probably, Michael Jordan. I don't know.

[INAUDIBLE] differ. It's like, it gives you someone's preexisting analysis of that question, right? Not a result that I think virtually anyone wants to see, but a thing that happens.

[00:25:57.36] So my question for you is, as a technological matter, to the extent you can generalize at all, which you may not be able to, does it tend to be the case that that kind of memorization or verbatim regurgitation is more likely the smaller the training data set is or vice versa? Or can we simply not say?

[00:26:23.44] KATHERINE LEE: This is a complicated question. And I'm going to back it up first. So when I say memorization, we have a lot of different intuitive definitions of memorization, one of which is memorization of fact, one could be memorization of style, one could be memorization, as you said, of an exact training data. I'd like to offer that even approximate memorization, approximately getting it right, is also, perhaps, troubling.

[00:26:46.72] The question about whether or not a smaller data set-- a model trained on a smaller data set is more likely to memorize more data, it depends on how long it's been trained for, the architecture and the training dynamics. So what I'm talking about there is learning rate and the learning rate process throughout the entire training. This is just within the model itself.

[00:27:06.90] So I can't answer your question in that framing, but I can tell you that depending on choices that are made, there may be different outputs.

[00:27:16.74] ANDY GASS: If you hold everything else constant?

[00:27:19.71] KATHERINE LEE: I can't hold everything else constant, and I'll tell you why. Because if I want a model trained to a certain quality, that means that I have to adjust a lot of different knobs. So if I have a smaller model and a bigger model, and I want them to get to the same quality, I'm going to have to make a lot of very different decisions, including the training data, including the training time, including the training process.

[00:27:43.39] HARRY SURDEN: Let me just pause for a second because I see some of the other panelists wanting to jump in. Cynthia, did you have something you want to say?

[00:27:50.22] CYNTHIA ARATO: Yeah, I do. So I'm going to start by saying, I do agree with Andy on the fact that I'm speaking just for myself.

[00:27:55.79] [LAUGHTER]

[00:27:58.21] HARRY SURDEN: So we have one area of agreement.

[00:28:00.00] KATHERINE LEE: I'll second that one too.

[00:28:01.76] HARRY SURDEN: We'll take that.

[00:28:02.91] CYNTHIA ARATO: It's the most important area of agreement. So I think I'm just going to speak quickly because we have a bunch of questions to cover. But I think that-- I think the models on the training are copying expressive content. I'm going to just talk about text models right now, large language models.

[00:28:24.30] They're ingesting expressive content for the purpose of extracting that expression so that they can mimic that expression when they do their outputs. And I think that that capturing of the expression is the heart of what copyright is supposed to protect. And that whatever else you want to say about these models, whether the output is infringing, it's not infringing, I don't think there is such a thing as a base case that it-- you can analyze these models saying, the only thing they output is non-infringing because that's just not how they work.

[00:29:00.60] But the capturing of the expressive content means that it's being used for what copyright is supposed to protect, and therefore, the use of those content-- that content without compensation or consent, I think, does violate the copyrights of the content owners, and that it doesn't translate to being a Fair Use.

[00:29:27.83] There's, I think, a big debate out there as to whether the models are copying expression. I know Andy talked about how they are being Fed discrete components of text disaggregated. But I actually would disagree with that. I mean, that they might, in their training and in the process, disaggregate the text.

[00:29:49.34] But if the point is for the model to see what a full sentence is or a full paragraph is so that it can then code those words so that it can mimic human speech when it's producing an output, it has to be ingesting full sentences and full paragraphs in order to accomplish that.

[00:30:10.37] And I think the fact that what it is doing is copying expression can be gleaned from what you see in the output. So memorization is one of those things that if it is possible, through the statistical probabilities, that the model comes up with having looked at the text, to replicate the text, even if that's an undesired result, I think that proves that the expression has been obtained and retained by the model.

[00:30:39.19] There's a phenomenon known as hallucination, where the chat bots of the world are known to provide responses that include facts that are not correct. And I think that that show-- or that they couldn't have of ingested from anywhere because they're made up. And I think that shows that the models are not merely trying to glean facts divorced from expression because they're not able to actually know what's a fact and what's not. They're just giving you a prediction on the next word.

[00:31:13.96] You can-- they also-- they can't extrapolate. Models can't extrapolate. So if you-- there's a recent paper that's come out that if

you ask the chat box, who is Tom Cruise's mom? It can give you the answer. But if you ask the chat bot, who is Mary Lee Phifer's son, it tells you it doesn't know. So I think that all of that together shows that the models are not actually learning or just trying to obtain fact, they're just copying expressions so that they can also be expressive.

[00:31:46.69] And then I'm just going to quickly touch on the back end cases. So Andy was talking about Sega and the famous Google Books case. And I think those are just-- they can be differentiated pretty easily from the generative AI world. So I don't think they're very important precedent to come up with the answer for what is or isn't fair use.

[00:32:09.74] So on the Google Books case, the court there said they already were at the outer boundaries of what Fair Use is. And there, the copying was-- Google explained that the copying was done just to determine information about a book, not the expressive qualities within the book itself. And I think the court was also heavily influenced by the fact that the output was not substitutional, wasn't that the output was not infringing, it was that the output was not substitutional. And here, you see a lot of substitutional effects from what the LLMs are giving you as output, certainly, like for journalism and in news settings.

[00:32:49.74] And then on the Sega case, that is a case about computer code, which the court found had a highly functional aspect, which makes it unique in Copyright Law. And so it's hard to use that as a model for what happens here when you're really dealing with the expression. And it also found that protecting the code in that case would really enable the owner of the code to get a patent-like monopoly on computer code, which it didn't have through the backdoor of copyright. And I think that also is just not a concern that's present when you're talking about copying expressive content or expressive content.

[00:33:33.87] HARRY SURDEN: Great. Well, thanks for that other perspective. And what I heard you say is, you don't think it is or should be Fair Use because of, unlike the Google Book Search case, it's more focused on copying the expressive elements, which goes to one of the aspects of the Fair Use analysis.

[00:33:49.86] Oren, I saw you were-- you had some thoughts.

[00:33:53.61] OREN BRACHA: I do, actually. Thank you.

[00:33:55.05] [LAUGHTER]

[00:33:56.76] So when I listen to those Fair Use debates, I feel a little bit like an atheist in a theological debate, because I have a totally different view I want to lay on the table here, which is, training copies as training copies, ignoring everything else that happens down the road, and maybe I'll get to that, are simply not infringing, period, on subject matter grounds, well before we get to ask questions of infringement or Fair Use at the back end.

[00:34:26.05] And to be clear, I know it's a somewhat unconventional view, but it's not some new strange doctrine that I've just invented. I think that conclusion follows directly from the most fundamental principles of Copyright Law as follows. Basically, when we have physical copies, in which, in some technical sense, we have expressive works embedded, but nobody would enjoy the expression as expression, no human will ever get exposed to the expression as expression from those copies, that does not involve copyrightable subject matter, period.

[00:35:05.08] In fact, what's going on, again, bracketing for a second what happens at the output stage, is it's a classic case of learning. It's a classic case of-- for what, for centuries, people were allowed to do under Copyright Law. To take a copyrighted expression, extract meta information from it in order to produce sometimes, new expressive works. In fact, in the training context, it goes well beyond that because we're not just talking about expressive generative AI with the training process.

[00:35:39.08] The only difference in the generative AI context is that the technology as it exists today has to generate a physical copy. That's the only difference. But look, if nobody enjoys the expressive content, no human enjoys the expressive content from that physical copy, focusing on the physical copy is an act of basically copy fetishism.

[00:36:04.27] It's-- for some reason, we bow to the physical fact, there's a copy there. And we ignore the fact that the basic subject matter of copyright is not implicated, no human enjoys the expression from that, that's just irrelevant for copyright. So what I'm saying, I guess, is that we need to correct the right outcome but the wrong turn taken in Sega versus Accolade all of those years ago, and get back to correct copyright principles.

[00:36:30.25] Now just very quickly, Katherine's very, very pertinent invention. But what happens if at the output, there is substantially similar stuff at the output stage that does infringe copyright? Again, first of all, let's be clear, this is just in the big picture, a small subset of the cases, not just within copyright. There's so much training going on of generative AI that it's not-- that doesn't then generate copyrightable or expressive works. Many, many training instances involve copyrighted works. So we have a lot of cases with no expressive output at all.

[00:37:05.80] But to your point, Katherine, I think most of us, maybe all of us can agree, certainly, those would be cases of potential copyright infringement. The output will be infringing, whether it's substantially similar.

[00:37:19.12] Now there's a whole interesting set of question of whether and how we can trace that back to the person who trained the system and through which doctrines. Maybe we can. But to my mind,

the argument wouldn't be the trainer is the infringer just because they made the physical copies. There might be a line of arguments that connect from the training to the infringing output. But again, not just by virtue of physicalism and copy fetishism.

[00:37:50.05] HARRY SURDEN: OK. That's a really interesting perspective. So just-- you dove into the mechanics of Copyright Law. So just for the audience, we'll explain, there are different ways you can not infringe copyright. One way is to copy something that's a non-protectable element. That's not copyright infringement.

[00:38:05.27] And so you're more in that spirit that what has been copied is not copyrightable at all. Or we've been talking about Fair Use, which is given something that looks superficially like copyright infringing, is there a good defense of Fair Use that would excuse you from doing that? And that's an interesting perspective I hadn't really heard.

[00:38:25.36] Viva.

[00:38:26.35] VIVA MOFFAT: Well, maybe, I'll just jump in very briefly. I know we have a lot of questions to get into. So I mean, I think this is-- I don't know, from maybe a slightly different perspective, in some ways, this-- there is-- this is exciting, like real disagreement on this--

[00:38:39.81] [LAUGHTER]

[00:38:40.76] --on this panel. And I think much of the disagreement at one layer of it is about, is AI like other copying technology that came before it? And do some of these analogies work? And if so, how? Or is it disjunctive and so different in some ways that maybe some of these analogies don't work anymore? And I mean, that's a way, in the way, I've been teaching Copyright Law for years now is, well, what analogies work? And why?

[00:39:14.60] And so I'm not going to answer that question at All but that, I think, is-- I mean, in some ways, I think there are people saying AI is actually different. AI is a different kind of thing. This is not like the copy machine or even digital copying, but is some-- I don't know, something different.

[00:39:33.54] So I think that maybe that thing will-- that perspective threads through all of these questions.

[00:39:40.65] HARRY SURDEN: Yeah, that's a great point. And that is actually my perspective. I think generative AI is a new enough thing in many dimensions that we haven't seen a close analog in the past, even though superficially, it seems similar.

[00:39:54.15] All right. Let's turn it. Wow. We're only on question 2. You got a lot of questions.

[00:39:57.69] [LAUGHTER]

[00:39:58.44] And I am surprised that lawyers--

[00:40:00.96] ANDY GASS: I didn't finish my answer to question one, Harry.

[00:40:02.41] HARRY SURDEN: Yeah.

[00:40:02.55] [LAUGHTER]

[00:40:04.20] Well, you'll get a chance to weigh in on question 2.

[00:40:06.39] So this one's for Cynthia. Should purely or mostly AI-generated works be eligible for copyright protection? And part of the background is this idea that maybe, we're seeing for the first time, users of things like Midjourney can enter very minimal effort and produce high quality outputs-- high quality artistic outputs?

[00:40:30.20] So I can create art on Midjourney. I don't have any artistic skills. And we-- usually, in the past, time, effort, and skill have been linked to high quality. And now we're in a world where anybody can generate high quality output.

[00:40:44.75] So let me throw this over to Cynthia.

[00:40:46.81] CYNTHIA ARATO: Sure. So I think there's a major difference here between can AI-generated works be subject to copyright protection? And should they? And so the registrar gave you a background earlier this morning about their copyright office's idea or position that it can't be copyrightable. And that comes from the text of the US Constitution and the Congressional Copyright Act. And as she explained, there's one district court decision that has held that that is correct.

[00:41:20.02] But the question of should actually turns this from a legal discussion to a pretty deep philosophical question that is really about humanity and are machines supposed to be treated like humans or not. And so the more I read about this, the more I was really surprised at how philosophical the discussion can get.

[00:41:46.04] So I'm just going to go through some of the arguments that I've seen pro and con for why AI-generated content should or shouldn't be given copyright protection. And I'm talking here about the purely AI-generated content, not these-- not the questions about the hybrid type of material that the Copyright Office is allowing to be registered, although that comes into play.

[00:42:13.75] So the-- there are a bunch of reasons that people have given why you should or shouldn't recognize copyrightability in generative AI work. And they range from moral to economic to the essence of humanity. So people who feel that these works should not be given copyright protection say that it's not a moral result if you let a machine own something where the machine can't take responsibility

for, whether it's doing something right or wrong. There's a whole economic argument.

[00:42:58.23] The purpose of the copyright clause in the Constitution is to give economic incentive to private actors so that they will continue to generate works that are valuable to society. But the goal is to benefit society, not to benefit the private actor. And so there are questions about whether a computer needs any kind of incentive to create a new work. And if it doesn't, why do we need to give the copyright to the work that's generated by a computer?

[00:43:31.99] There are arguments about human exceptionalism that just the creative process by humans is just fundamentally different than the creative process by a machine, in ways that people may not actually really be able to articulate, but I think feel very strongly. And that the benefit of copyright protection should only be given to that special human creative force.

[00:43:58.30] And then there's just the functional questions about, if the output of the machine is indistinguishable from the output of a human being, and the goal of copyright is to increase the production of these creative works, why shouldn't we give copyright protection to a functionally indistinguishable product generated by a computer?

[00:44:24.02] And then there's some practical questions which is, how can we really determine whether the output of a computer is original? Is there a way to do an audit trail to the original data that the computer has been fed? Or how do we know-- what do we do when the computer generates the same output multiple times to different people?

[00:44:44.53] So those are some of the, I think, interesting questions that get raised when you talk about whether you should provide the protection. And then there's a separate question of the works that are generated by the machines, but with the guidance of a human hand. And if there's a sufficient guidance of the human hand, then it should be protectable.

[00:45:08.83] But then you get into complicated questions about, well, but who's the right. copyright owner there? Is it the person who put the prompts in? Were the prompts expressive enough to warrant the prompter getting the copyright? Is it the machine itself? Is it just the owner of the machine? Is it the human being who wrote the computer code that enabled the machine to function? And I don't think anyone has perfect answers to all of those questions, but those are the questions that I see.

[00:45:40.16] HARRY SURDEN: That was a terrific summary. Thank you for that. And just-- can I put you on the spot and ask you your opinion?

[00:45:46.73] CYNTHIA ARATO: I'm not here to give my opinion on that one.

[00:45:48.85] [LAUGHTER]

[00:45:50.87] I'm the-- the lead person can just set the stage.

[00:45:53.95] HARRY SURDEN: OK. Excellent. Excellent. So let me throw this out to the panelists, either for opinions or whether you think under-- we've heard what the current state of the law for purely AI-generated works is. Yes, Katherine.

[00:46:08.80] KATHERINE LEE: I'll just jump in here and say that the technology is changing a lot. And the ways that we interact with the technology is changing a lot. And the questions about how much creative control someone has also depends on what the user interface looks like for the model.

[00:46:23.86] So we talk a lot about the generative AI models, but there's this entire system that's built around them. And so the ways that we interact with the model influences the outputs that we get out, and influences the ability that we have to modify certain parts.

[00:46:36.01] So for example, some image generation models, you can circle a bit of the image and say, put a dog there. And in some of them, you can even draw a little picture of the dog, and use that as a prompt, and it'll fill in the rest of the dog for you, maybe in that little pose.

[00:46:49.87] So the types of the ways that we get to interact with the system is going to continue to change as we-- as more products get developed and as we play with it more, and people imbue their own creative process to the way that they interact with the systems. So when we're talking about AI generative models-- sorry, generative AI models, we should really be talking about generative AI systems. So I'll just leave that there.

[00:47:12.01] HARRY SURDEN: Yeah, great point because we're just thinking about it the way they are right now. But five years from now, there could be fine levels of control that might add the human hand more. And we might be talking about something different. Viva.

[00:47:25.70] VIVA MOFFAT: Maybe consistent with my prior point on the last topic. I mean, one thing I've really wondered about-- I'm not going to answer your question either.

[00:47:32.08] [LAUGHTER]

[00:47:33.53] One thing I've wondered about is whether this is-- this question of whether AI-generated work should be protectable is really going to push to almost its breaking point copyrights current very wide net of any even modestly original thing that's fixed in a tangible medium of expression is protectable. I feel that we are already close to a breaking point when every photograph on my iPhone is a protectable work and every tweet that everybody sends is a protectable work.

[00:48:10.94] Again, is this a new enough and a big enough change that it's going to push us to thinking, well, maybe we should change the place where the filter is. The copyright filter now is there's no hurdle at the outset of protection. Now talk about taking us out of step with the world, I think, talking like that would be.

[00:48:33.71] But I think this question really pushes that issue of whether we really should continue to have this extremely broad initial coverage for Copyright Law.

[00:48:45.68] HARRY SURDEN: Yeah, that's a great point. And it shows another way in which our copyright code developed in the 1970s, primarily, did not anticipate such new technology.

[00:48:55.89] And I think lurking on the back end is this idea that you-- for human authors, you get 70 years plus life. So if you have your AI system today generate 10 million outputs-- 10 million pictures, and we were to make them copyrightable, each one of them could have 140 years of copyright protection with $150,000 worth of statutory damages, potentially. So-- and that's a different world we lived in when copyright was created at fixation back in the '70s, they didn't anticipate that idea.

[00:49:31.77] Did I see another hand from the panel on this issue?

[00:49:35.98] OREN BRACHA: It depends on you. If you want to get to other questions, I'll skip.

[00:49:38.76] HARRY SURDEN: Yeah, well, do you want to give your quick--

[00:49:40.64] OREN BRACHA: OK.

[00:49:41.01] HARRY SURDEN: --viewpoint on [INAUDIBLE]?

[00:49:41.89] OREN BRACHA: Since you insist, I'll do it very quickly.

[00:49:43.75] HARRY SURDEN: This is a big topic. Sure.

[00:49:44.67] OREN BRACHA: OK. This is a very-- I think, very useful discussion and setting up of the stage. So bottom lines. I think we should actually, at this point, bracket, put aside the deep philosophical, metaphysical questions, although, it's an endless source of student notes that will keep coming on the stage. But we're really not at the point where we really reach the point where we do things that machine creates like humans, and we should treat them like humans, either in the kinds of interest they have in the creation or in the legal rights they get.

[00:50:20.37] So let's put all of that aside. What are we left with? With the policy questions. And actually, personally, I don't think the policy questions are limited to economics. It's questions about human interests implicated, whether economic or otherwise.

[00:50:35.29] And I think, just to keep it very short, when we look at works that at least get close to being machine created, and there's a deep question of when that happens, there's very little argument for copyright protection.

[00:50:53.19] So duration is just an example. I think that no category of works should be protected for a century. But certainly, not that. Why? Because what's the big brouhaha about machine creation? The one main thing is exactly the slashing of costs. It's so much more cost-effective to generate stuff that we couldn't in the past for that cost. And cost is the main case for copyright in this area.

[00:51:21.58] So bottom line, on the merits, I think-- and again, the interesting questions, of course, come with the fusion of human creativity and machine creativity. Those are the hard questions. But with regard to pure machine creation, there's no reason really.

[00:51:39.39] So bottom line remains to be seen with all the developing models for user creation. But I think at the moment the direction actually chartered by the Copyright Office, as described by the registrar, it's a great reaction to the current situation, as described. And then we can wait and see how the technology develops.

[00:52:01.17] HARRY SURDEN: Yeah, that's a great point. So just as a bit of background to summarize what you're saying here. The historical view of copyright before technology was, in order to incentivize creators to spend the time, effort, and skill to create a painting or a song, you give them several years-- many years of exclusive rights on the back end. But now we're at a moment where we can create high quality outputs at near zero cost. You just type it in. So do we need 140 years of exclusive rights on the back end to incentivize that?

[00:52:35.35] OREN BRACHA: Thanks for translating.

[00:52:36.59] HARRY SURDEN: Yeah.

[00:52:36.85] [LAUGHTER]

[00:52:38.29] You're welcome. And that was a great point. All right. Now this is a point that I think is near and dear to many of us, which is, what about the policy impacts of generative AI on creators?

[00:52:51.25] So on the one hand, we have these artists out there who are painting, musicians, poets, writers. This is how they earn their living. And this is very valuable. And we're moving to a world where AI can generate pretty good versions and maybe even great versions of some of those things. So we might be concerned.

[00:53:14.38] On the other hand, we have seen this before in technology where artists and creators say, here's a new technology. Oh, no, here comes photography back in 1900. Portraiture is going to go away and then painting.

[00:53:27.76] And then when it came to digital photography, they said, oh, no film photography is going away. And then Photoshop, and then so-- there's a lot of new technology brings a lot of new opportunities. And artists often find new uses for new technological tools in ways we can't anticipate.

[00:53:47.33] So we want to be both concerned, but also wary of the past. So with that in mind, let me throw it to Viva.

[00:53:54.19] VIVA MOFFAT: Right. So you stole my thunder a little bit.

[00:53:56.88] HARRY SURDEN: Oh, sorry.

[00:53:57.35] VIVA MOFFAT: Because that-- so one piece of this, certainly, is that there has-- artists have-- I think we've seen in the past with each of these new kinds of copying technologies that artists and creators have managed to adapt and in various ways.

[00:54:17.78] And in fact, Harry, you said, before tech, there was this copying. And in fact, I don't-- there was no copyright before technology. So literally, really, we only have Copyright Law because of disruptive technology.

[00:54:31.65] So if the history of Copyright Law is really the history of disruptive technology, there's always been this concern, legitimate, I'm not undermining it, about, well, what's going to happen to the creators, the very-- the fundamental point of copyright law is, how do we make sure we get enough creativity? How do we get enough of these things that we like, which are music, and paintings, and poetry, and novels, and video games, and everything, all of these things?

[00:55:02.10] And I think that, Harry, as you point out, the history of this has been that we do seem to keep having human creation. And I'm actually, I'm a big optimist on this point. I think that human beings are going to keep creating things. And I actually also think that human beings are going to continue to want human-created things.

[00:55:22.37] I think it's possible, this is maybe my overly optimistic view, is that there may be even more of a premium put on human-created or authentic experiences, going forward, given that I think we will certainly have massive, massive quantities of non-human-created or AI-assisted works.

[00:55:46.94] It may lead us-- I think we're already in a world, I'm going to offend somebody here, where we have lots of low value works out there in the world. And I think it's-- I mean, I put all the photographs on my phone in that category.

[00:56:03.05] [LAUGHTER]

[00:56:04.64] But-- and lots of-- I read a bunch of stuff before this about, oh, TikTok people creators using AI to aid them in their TikTok creations so you can tell how old I am right now. I don't even know how

to use the words right here. I would put a lot of that, probably, in the low value, again, I'm going to make somebody mad, space.

[00:56:26.84] I think we're going to keep getting lots of that, lots of that stuff. And on and on the high end-- higher end, I absolutely am an optimist. I think the-- so when we're talking about creators and livelihoods, though, that involves creators are one thing, livelihoods involve business models. And the technology that comes along often does really either seriously undermine or completely destroy incumbent and legacy business models that do enable creators to make money from their creations.

[00:56:57.20] And this is one where I think the-- I just put it out there is like, I don't really know. And I'm not sure anybody knows what's going to happen to people who write marketing copy, to web designers, to graphic designers, journalists, screenwriters. These are all people who I think will keep wanting to create and may have less opportunity, because of AI, to make money from that kind of creation.

[00:57:26.02] And so now I don't even know what question I was supposed to answer exactly. But I think that those are-- that's-- I'm optimistic about humans, but I am very deeply concerned about, in some ways, humans' ability to make money off of their creations going forward. And I think AI has a possibility of maybe being different than previous technologies in terms of disrupting that. So maybe I'll just put that out there as a question.

[00:57:57.15] HARRY SURDEN: Cynthia.

[00:57:58.33] CYNTHIA ARATO: Yeah, I think that you touched on a lot of the different industries. And I think that AI might impact them in different ways and in ways that matter differently to the United States. So I'm just going to talk about journalism because you mentioned it.

[00:58:14.61] I think in terms of the industry as a whole, if generative AI is allowed to copy indiscriminately from source material and generate outputs that substitute from that source material, you're going to find it very challenging, as you pointed out, for the journalism industry to continue and flourish.

[00:58:38.03] And I think as important as fine art is, and as much as we all appreciate it, journalism is a pillar of our Democratic workings. And without the economic incentive to continue to be able to function as a business, the loss of journalism, I think, is a very dangerous outcome to what could be and what seems to be very substitutional use from the gen AI.

[00:59:04.99] So I think that-- for that perspective, creators are quite concerned. And I think they're rightfully concerned.

[00:59:12.46] HARRY SURDEN: That's a great point because creators aren't just necessarily serving a creative role in society, they're serving other institutional roles. I love that point. Andy.

[00:59:21.02] ANDY GASS: So I mean, I certainly share the vision that it would be bad if newspapers ceased to exist. That is common ground. I think I have a little bit of a different perspective on the likely effect of AI on all of that. Really, just two points to add to the conversation, which has been very illuminating to this point.

[00:59:42.54] One is just another angle that has always struck me about the potential of these tools to affect creative endeavors is the degree to which they actually unlock the possibility of people creating who, without this technology, couldn't do it for whatever reason, either because they physically were disabled or otherwise, didn't have the wherewithal to pull it together.

[01:00:07.51] So it's another consideration to throw into the hopper as we evaluate all of these competing implications that there is, I think, unassailably a dimension of this where we're going to democratize, in some respects, the universe of people who can create really cool stuff that wouldn't previously exist.

[01:00:23.68] I guess, the other comment that I would like to throw out there, which I don't mean in a cynical way, I just mean in a realpolitik pragmatic way, is that there's an international dimension of all of this that we should bear in mind, which is that the US is not the only jurisdiction that matters here.

[01:00:44.44] And these tools exist. They are probably going to continue to exist, at least, somewhere on the planet. And as we think about what the right rules of the road are to protect all of the legitimate interests that people have articulated in response to this question, I think it's important to do so in the context of understanding that, well, we probably don't want to make it so that the way that things happen is just everyone goes to another country that is very hospitable to these tools. And they're the ones who are going to be generating the things that are going to take the jobs that people don't want to see taken.

[01:01:18.34] So it's a complicated set of interrelated policy levers to think about. And that's really the sum total of my thinking on those questions.

[01:01:26.82] HARRY SURDEN: Yeah, those are some great points. And I want to emphasize one point you said, which is, new technologies unlocking new creators and democratizing. And we-- 15 years ago, we didn't have the term YouTube Creator. And now people spend their livelihoods using these new technologies. So AI will probably unleash new types of creativity even as it disrupts existing ones. Kath.

[01:01:52.53] KATHERINE LEE: Yeah, just really briefly. I'm going to keep beating the drum that there are many different types of generative AI models and many different types of generative AI systems. And I do not think that our future is guaranteed to be one way or another.

[01:02:04.98] It really depends on what types of products are built with these systems, how they're integrated, and how people are able to use and leverage them. So I can imagine tools for journalists that could be built with AI that would empower them and make their lives easier. Will that necessarily have growing pains? Definitely. But I could also imagine much more sinister versions of our future.

[01:02:24.70] And so I think that as we're sitting here, thinking about it, I share your optimism for humanity. And I share the optimism that we're going to figure some things out. It does absolutely depend on broader and institutional and governmental interests. But this is not a foregone conclusion. Our future here is not a foregone conclusion. It really depends on how we structure the products that we build, how we structure those models, and how we integrate them, and teach people to use them.

[01:02:51.64] HARRY SURDEN: That's a great point. Did anyone want to respond to that? Well, I love your point, which is, and this is why conversations like this are so important, that the technology-- we shape the technological future, all of us. And it's important not to let the defaults shape it. So I love that point.

[01:03:10.84] So I'm going to throw this over to Oren here. So as Professor Reid said, for better or for worse, copyright is often the first on the scene when it comes to technology law like AI. But it's not necessarily the best or the most equipped set of legal regimes to deal with complex issues. What are the-- some of these pitfalls for allowing Copyright Law to regulate AI by default?

[01:03:39.68] OREN BRACHA: So the answer is, that's correct.

[01:03:41.24] [LAUGHTER]

[01:03:43.00] And I think the premise in the question is correct, which is, a lot of those big issues that are arising now, those social concerns and complaints about AI in the realm of culture will be laid at the doorstep of copyright. They are already being laid at the doorstep of copyright in hips, as we know.

[01:04:04.69] And the reason is, in our society, Copyright Law is the main tool we have for dispensing cultural policy. It's not the only one, but it is the main one for better or worse. So that's going to happen. It is happening.

[01:04:19.24] Now two things to notice about copyrighting. First of all, copyright is and was designed and is built around a particular policy problem and as a particular solution to that specific problem in the field of expression. Copyright is about the problem of undercompensation, attributed to the fact that once you create something and publish it, it's very hard to exclude others from it. And the fact that almost always, it's much cheaper to copy than to create.

Copyright is tailored to be a solution, at least, a partial solution for that problem.

[01:04:57.65] Second thing to note is copyright solution, which by the way, it's not the only one, but copyright solution is market-based. That's what copyright is about. The solution it offers is basically, it enables creators to sell their creation at marked up prices. So it relies on market transactions. It's market-based. And there are known advantages and disadvantages to that strategy, but that's what copyright is about.

[01:05:28.81] Now enter generative AI in the field of expression. Again, two things to notice about the complaints, and arguments, and concerns, specifically about the broad legal arguments being made.

[01:05:45.10] So the broad legal arguments, when you really examine them, are not about the specific policy problem that copyright is designed to solve. Well not all of them. The classic arguments about output of generative AI that is substantially similar to copyrighted works, of course, it's about that. But the broader arguments like training copies, and copying of style, and some other arguments that are percolating up there, the real concerns that drive them are different concerns.

[01:06:16.62] And actually, I think there are legitimate and grave concerns. There are-- those concerns are actually reflections of more general concerns about generative AI in society in general, and not just the expressive sphere, just to identify three concerns like that in the expressive area.

[01:06:35.56] First of all, concerns about the disappearance of sources of livelihood and human income in the creative industries. Second, concerns about opportunities and potentials for people to engage and enjoy the inherent value of creativity. Not the market value, the inherent value of creativity. Third, concerns about the sources and potentials for disruptive innovation, the kind of innovation that doesn't only deliver very well within given patterns, but rather, break the patterns.

[01:07:08.62] Those are very important, legit concerns, I think. Those are exactly the concerns that copyright is not built to deal with, it was not designed to deal with.

[01:07:19.81] And moreover, if we really examine closely, which of course, I cannot do in my remaining two minutes, each of those concerns, what we would see is that each and every one of them is extra market. It's really external to market logic. And that's not a criticism, I think that's a great thing.

[01:07:39.32] But those are the things at the bottom of those concerns. Why? Because the logic of the market is best market efficiency. From that perspective, what on Earth could be wrong with generative AI? It

delivers wonderful results. It satisfies market demand at great cost efficiency. What's to complain about?

[01:08:00.08] All of those concerns, in one way or another, challenge the logic of the market from this or that perspective. And remember, copyright is market-based. That's the logic it adopts. What is the outcome we're getting here? The outcome is an incredible misfit, basically, between the two that people naturally go to copyright in order to try to solve all of those problems and the characters of the problems we're trying to solve.

[01:08:32.51] And therefore, it's much like trying to use a hammer to do the job of a screwdriver or indeed, an entire toolbox. What are you going to get in those cases is a very bad job. And you're going to smash a lot of things.

[01:08:46.97] HARRY SURDEN: Those are really great points. I mean, a lot of the issues that are coming up are cultural issues, issues of economic policy, trade policy, competition policy. So for example, just one-- imagine that the courts say that training data is not Fair Use. As a practical matter, it would be-- it would make training AI systems either more expensive, maybe slow them down. And that's really a competition issue.

[01:09:15.60] On the other hand, it might give new sources of revenues to creators whose works are being used and livelihoods are being disrupted. Again, economic issue. And we're using copyright as a tool to deal with all these other issues that are outside of the area of expertise.

[01:09:31.76] CYNTHIA ARATO: I'm just going to say really briefly. I think that those economic issues, though, are what Copyright Law is about. I don't actually view that as being outside of Copyright Law. There are other extra copyright issues that I think generative AI implicates, but a lot of the market economic points really are at the heart of copyright.

[01:09:52.65] HARRY SURDEN: Can you expand a little bit more on that?

[01:09:54.51] CYNTHIA ARATO: Well, just the-- it's all-- it goes back to the original rationale behind why we have copyright rights to begin with, and why the Constitution gave Congress the power to grant exclusive rights to the creators of those works, it's to incentivize their creation.

[01:10:10.50] So the more you take-- like on the training data example, if training data were to be found to be a Fair Use or the use of training data to be a Fair Use, you're depriving the creators of an important revenue stream. And that's going to-- and that's an economic consequence and a market consequence. And it goes to the heart of the incentive policy behind copyright.

[01:10:36.73] HARRY SURDEN: Yeah.

[01:10:37.05] OREN BRACHA: May I respond quickly?

[01:10:38.64] HARRY SURDEN: Please.

[01:10:39.33] OREN BRACHA: I totally agree that copyright-- by the way, again, even if the interest you care about at the end are not economic interest, is about the economic dynamics and making sure that we get what we want there.

[01:10:50.83] I do think, though, that it is focused on a particular problem within the economic dynamics, especially when we talk about incentives. And it is that problem of not being able to recoup your investments because of the impossibility to exclude others, and the gaps between creating and copying.

[01:11:10.56] And to some extent, that's, of course, relevant to the field of generative AI. But I do think that the broader problems are not limited to that at all. For example, the problems of losing-- even the problems of losing sources of income.

[01:11:24.06] Much of the issue there, let's be honest, does not come from the fact that the AI free rides on the efforts. It's less the expressive efforts of creators that became before. And much of the issue derives exactly from the fact that unknowingly, perhaps, it's so much better, so much cost effective in creating something new out of what's out there in the comments.

[01:11:49.44] It's really troubling on some level, right? And it will displace the livelihood of people. It is a troubling issue. It's just not the kind of economic issue that copyright was designed around.

[01:12:01.83] HARRY SURDEN: Good point. Andy.

[01:12:04.02] ANDY GASS: I know we've got to get to questions shortly, but I appreciate your comment about competition in this space and the valence of different perspectives you could take on the IP issues for what that portends for competition in various markets.

[01:12:16.41] The Federal Trade Commission actually held a session on Wednesday, in which a number of people aired anti Fair Use perspectives, I think, on the use of copyrighted content to train models. It struck me, my background is as an antitrust lawyer, that if you really care about fostering competition in this environment in a way that it will ultimately redound to maximize consumer welfare, which has historically been the gravaman of competition policy, at least for a while now, what you would see the FTC doing, which they probably won't do, would be filing amicus briefs in all of these copyright cases, saying, the use of this training data should be Fair Use because that will, more than anything else, allow for an explosion of competition among the technology side that's innovating in this space.

[01:13:06.32] I'm not holding my breath, but from a theoretical perspective, that's how that debate would play out if people were thinking about it in that way.

[01:13:14.11] HARRY SURDEN: Yeah, that's a terrific point, linking it to competition policy.

[01:13:19.22] So our last question before we open it up to the audience for comments, goes to Katherine. So Katherine, how should we address the potential for AI to inadvertently infringe on existing copyrights, especially when they're trained on vast data sets?

[01:13:36.80] KATHERINE LEE: So I'm going to answer this from a technical perspective. And then maybe people can chime in with other perspectives. I've been saying this throughout, but there's a lot of choices that you get to make in this process of creating a generative AI system. There's a lot of choices that go into the system. There's a lot of choices about who you work with, what you work with, and how you deploy the system. All of that feeds into how much, or how many, or what looks like the copyright infringement looks like.

[01:14:03.98] So I am not-- I am neutral, looking forward to this. I think there's a lot of work that can be done on the technical side to make some of this a little bit better. So one example is using retrieval-based models.

[01:14:18.16] Retrieval-based model is a different type of architecture than your typical generative AI model. It takes-- it goes and looks inside a data set for relevant works, and then uses that work to inform the generation that it produces. So the advantages of that are that it's much traceable-- it's much more traceable. So you can say like, oh, yes, if it's going to be similar to something, maybe it's similar to that work that was retrieved.

[01:14:45.10] Obviously, this is not perfect, and it's not a perfect solution, but I'm just trying to give you some sense about the creativity in the space that we have to work with here. There's other ways of designing models that can make these issues easier to deal with or harder to deal with.

[01:15:00.11] There are also ways of controlling the inputs and the outputs to these systems. So regardless of whether or not the model itself is producing potentially infringing outputs, again, if it's maybe not shown to anyone and it's stopped by an output filter, maybe our case here is a little bit easier.

[01:15:15.32] So I started this process talking about-- I started talking today about exact memorization of training data. That's a really easy case to catch as an output filter. You can just say, hey, does this generation look exactly the same as any other piece of training data? But it's actually not that easy because how much of that generation needs to be exactly the same as something in the training data. And

that's a question that I'm not going to answer because as I believe, there are entire volumes of work written on substantial similarity.

[01:15:47.16] And for the record, I think that there is no easy, simple technical solution here. But that doesn't mean that technical methods can't help improve this problem. So don't want to be technical solutionist and say that, yes, there is one function that I can give you. And if you call this function, it will be fine. I do not think that's the case here.

[01:16:08.57] But that does mean that we can produce things that like, look at the lighting of different images and compares that. Looks at bits and pieces of the images and compares that to the training data. Looks at different ways of capturing style-- capturing and measuring style. And all of that can inform the decisions that we make downstream.

[01:16:26.58] So I'm going to reiterate this point again. I'm pretty optimistic. I think there's a lot of different choices that we have in this process and in this pipeline. People are getting really creative about how they design their model architectures. People are also exploring other ways of creating and collecting training data that may change this entire analysis, whether or not someone uses their own model that they have trained themselves or uses a model off the shelf. And how they interact with that may also change the copyright analysis here.

[01:17:01.93] So there's a lot of different actors, a lot of different choices being made. I would love for us to design good incentives so that the choices in the future that we have align with our values.

[01:17:12.18] HARRY SURDEN: That's terrific. And I think, especially the focus on the input and the output filters is a good idea and something we're starting to see happen today. So just last night, GPT-4 recently came out with DALL-E 3 integration. And I typed in, hey, would you produce a version of this copyrighted character and this copyrighted work? And it refused. It said, no, I know these are copyrighted. And I'm not going to do that. So I thought that was fascinating.

[01:17:40.03] Anybody else want to weigh in on this? Yeah, Oren.

[01:17:42.70] OREN BRACHA: So I just want to quickly provide the bookmatch version on the legal side to what Katherine is describing on the technical side, which I guess, I totally agree or just bow to her on the technical side. I would just describe it not as endless set of choices, but rather, what arise out of those choices are different models.

[01:18:03.31] And I think, similarly, on the legal side, what we get is different legal models that should deal with those different models on the technical side, specifically, on a very general level, to the extent that we go down the food chain of using the AI, up to the level of the users who put in the prompt these days.

[01:18:28.96] Doctrines-- legal doctrines that will be more relevant is direct infringement doctrines. To the extent we're going up the food

chain, and the food chain could be very complicated with many links in between, we probably have to go to existing secondary liability doctrines as we know them. And it's not at all it's fine in that area, but that's what we'll have to turn to.

[01:18:50.54] And then interestingly, possibly, we'll have to add more. We might need to think about some other innovations, either through some sort of internal regulation or adding-- and some people won't like that some sort of an idea of a negligence standard on the secondary liability level. But we'll have to match the different technical models and food chain with different legal structures.

[01:19:16.12] HARRY SURDEN: Yeah, great point. Cynthia.

[01:19:17.89] CYNTHIA ARATO: I just wanted to say something about the nomenclature. So I think that the idea of inadvertent infringement, I'm not sure inadvertent is the right word. I mean, in Copyright Law, for things that are registered before the infringement began, at least, you talk about willful infringement, innocent infringement, not really inadvertent.

[01:19:41.98] And so I think that there's-- and you can be a willful infringer if you do something that's in reckless disregard for a copyright owner's rights. So I think that depending on the stages at which you're looking at the use and the alleged infringement, some of what people might use the word inadvertent might actually be willful from that reckless perspective because I think people ingesting huge quantities of data, they are doing it knowing that they're ingesting things that are copyrighted.

[01:20:14.98] And if they have a genuine argument that they think it's Fair Use, that might impact those findings of willfulness or recklessness. But I just think inadvertent is a difficult word to use when you're talking about copyright infringement.

[01:20:27.94] ANDY GASS: You're suggesting I didn't make a genuine argument?

[01:20:29.86] [LAUGHTER]

[01:20:31.80] HARRY SURDEN: Well, thank you for that. It's important words, and their meanings matter in law. So let me toss this out to the audience for some questions of particular student volunteer. Oh, we got an eager student right upfront.

[01:20:50.90] AUDIENCE: Thank you for this panel. This was very insightful. So I'm thinking of how to phrase my question. And going to Oren's point of there might be a reduction in economic costs, essentially, associated with using generative AI models to generate artwork, stuff like that. Given that the whole reason behind Copyright Law is to provide individuals the economic benefit derived from their work, how do we account for those decreasing transaction costs, if you

will, when applying copyright law to AI-generated works of creative works, if you will?

[01:21:50.89] OREN BRACHA: I'm happy to respond.

[01:21:52.17] HARRY SURDEN: Yeah. Yeah. Go ahead and--

[01:21:52.40] OREN BRACHA: I just want--

[01:21:52.96] HARRY SURDEN: --and take that.

[01:21:53.29] OREN BRACHA: --to take-- Look, so again, I'm going to answer on a pretty abstract level. But hopefully, it will be responsive. Remember my basic take on certain economic costs or at least, worrisome economic effects. My basic take is that the strongest reasons why those effects are concerning are not market efficiency or anything like that.

[01:22:23.82] There's something really concerning if a large amount of people lose their livelihood in creative industries. Now the question becomes, again, what's the best way of addressing that? And you remember, I don't think that classic Copyright Law. Property rights is the best way of addressing that, among other things because of transaction costs.

[01:22:47.25] But that's not the only reason. The market, for example, compensates people according to market success. I'm not sure that that's what we want to do if we have other concerns.

[01:22:59.41] So the abstract answer is, again, there's a whole set of institutional choices that with a continuum, all the way from classic, private property rights, copyright to a really public general tax and transfer scheme. And in between, we-- that-- somebody mentioned it, but we haven't really discussed it yet, you might find other institutional options like general compulsory license for a flat rate or something like that. So there's a whole menu in between market solutions and completely public tax and transfer solutions.

[01:23:39.04] The only thing I would still advocate is, even if we go for a solution, which is not extreme on the continuum, not all the way to some general tax and transfer, let's keep in mind what we're trying to do here. And again, I think what we should be trying to do is not necessarily mimic the market and its outcomes. Therefore, if we-- even if we have a compulsory license, I would argue that a compulsory licensing fees shouldn't be calculated by some hypothetical market transaction formula, but rather, by something else.

[01:24:12.31] Hopefully, that's responsive, but it's pretty abstract, I guess.

[01:24:15.82] AUDIENCE: Yes. Thank you so much.

[01:24:17.15] HARRY SURDEN: OK. We've got a question in the back.

[01:24:22.62] AUDIENCE: My name is Vicki Mandel. And I'm a consultant. I have my own law firm. And I have a question for Kathie Lee because it is quite of an opportunity to be able to talk to a Machine Learning researcher.

[01:24:37.34] And Kathie, you've talked about all the decisions that are being made while the machine is learning and being trained, human decisions that are guiding the growth of the machine and the learning. And so I was wondering if you could give some examples of those decisions that are being made. And I'm particularly interested in examples that have to do with teaching the machine to be ethical.

[01:25:11.72] And I realize, this is a little bit outside of the copyright question, but I'm very interested in AI regulation generally. And that's one of the big issues is how to teach the machines to make ethical decisions. So any examples at all, but those that focus on ethical value, moral choices are especially interesting. Thank you very much.

[01:25:33.71] KATHERINE LEE: Yeah. So this is a really interesting question because from my perspective, copyright is only one of the many issues that generative AI raises. And on a technical side, like copyright and privacy, and potential IP leakage all have similar technical roots.

[01:25:48.90] And if you pull back even further and say like, what's the goal here? The goal is to make useful models. What is useful? Sometimes, useful is creating copyrighted content. Sometimes, useful is not creating copyrighted content. And the process of creating a useful model requires alignment.

[01:26:06.33] Alignment is a word that is very overloaded in the technical community. So to give you a very brief summary of it, it's basically constraining or structuring or guiding models to do things in accordance to the way their human creator wants them to do. So there's many different ways of implementing alignment. And to generate or to not generate copyrighted content could be thought of as one subset of this alignment problem.

[01:26:33.80] There are other subsets including, like I said, privacy, which is like, don't reveal anyone's private information. That could be a value that we hold that we want to imbue into a model. So some of these techniques include things like reinforcement learning by human feedback is something you may have heard, that means we're collecting a bunch of human-labeled data, data that somebody has said OK, and somebody has said not OK. And then giving that as part of the training process of the model. So that's one thing that you could choose to include or choose to exclude.

[01:27:02.74] The data collection process of collecting the human alignment data is also a factor in this. And different people may ask different people-- sorry, different actors and different model trainers may use different communities to respond to that information.

[01:27:20.78] So that's just one choice that's like a data curation choice. There are other choices that we've already talked about data set size, and scale, and exactly what's contained in it, like licensed, non-licensed data. You, can ask the same questions for private, not private data. Again, the word "Privacy" is also really overloaded both on the technical side and on the legal side, I'm given to understand. And so not getting into that.

[01:27:40.52] But yeah, there's a lot of curated curatorial decisions here that all play into that. I gave another example, which is retrieval-based models. Another example for alignment is using the output filters or choosing parts of your training data to use or to exclude to change the biases of the model. That's another version of this.

[01:28:06.11] Another one is prompt tuning, which is a way of designing prompts, the prompt that you give to the model, to guide it in one direction or another. So you've probably seen some of this before, which is to say like, some people have been jailbreaking these models. And what that means is they try to go around the alignment that model creators have put on the model by giving it prompts that are unexpected, and put it into different scenarios.

[01:28:37.14] So there's, again, a whole plethora of different choices that could be made here. And also of different actors who have different responsibilities to different people, and who may or may not have responsibilities to each other.

[01:28:52.58] HARRY SURDEN: Great. Thank you. And we have time for one minute, one quick question. So if we can ask quickly and answer quickly.

[01:29:02.55] AUDIENCE: So this is a very simple question but a tough answer. What do the panelists think about giving humanity itself time to transition over to the adaption of the generative AI because it's literally going to disrupt every single area of the life and put livelihoods at stake?

[01:29:22.84] HARRY SURDEN: I mean, if you're referring to the pause, I'm not a fan of it because I don't think it's particularly effective. I think there are other ways to do things more directly. But I'll open it up to the other panels.

[01:29:36.16] KATHERINE LEE: I think the technical community is very split on this. There are people-- and again, it matters what AI you're talking about. If we're talking about creating a photo-- sorry, this is a bad example. Let's-- let me roll back.

[01:29:55.34] OK. There are different levels that we're rolling out products in. And some of those things are general purpose products, and some of them are much more specific and more narrow. So I think maybe thinking about it in that frame and thinking about how we can engage with folks who are part of those communities or those

industries to create products that they are actually excited about and are useful for them could be one different approach for that. But the technical community is also extremely divided about how exactly to approach this question.

[01:30:24.77] HARRY SURDEN: Well, thank you for a very interesting panel. Please join me in thanking our panelists here.

[01:30:29.90] [APPLAUSE]

Panel: The Future of Generative AI and Copyright Policy

https://youtu.be/c6vmVuCwZtA

[00:00:00.23] MARLAINA PINTO: Good afternoon. Welcome back. My name is Marlaina Pinto. And I'm the editor in chief of the Colorado Technology Law Journal. We're a student-run technology law and policy journal that publishes scholarly articles on topics like telecommunications, and space law privacy, entrepreneurialism, and most relevant to why we're here today, AI and intellectual property.

[00:00:32.57] We're fortunate enough to have Silicon Flatirons host these conferences, sparking thoughtful discussions that can ultimately turn into articles for our journal. In fact, our spring 2024 issue will include articles covering views from the investment crowdfunding symposium that happened last month.

[00:00:53.46] So keep us in mind as you're thinking of these topics. And if you are so inclined to write something yourself, please consider submitting it to our journal for publication. With that said, I'm happy to get back to today's program and introduce the next session, which we'll discuss the future of generative AI and copyright policy. So I'll pass the mic to our moderator for this session, Professor Blake Reid.

[00:01:22.95] BLAKE REID: Thanks so much.

[00:01:23.55] [APPLAUSE]

[00:01:26.86] All right. Let's get started. And I think we're going to pick up in just a second with Professor Bracha's provocation from the last panel about copyright serving as an incredible misfit for the range of policy issues. I think we're going to-- there's a lot to unpack with what Professor Bracha teed up.

[00:01:50.25] But before we do that, I want to introduce our panel. And I want to do this in a way that doesn't go through people's bios, but rather gives folks a chance to introduce their view on this discussion. And the question I want to ask everybody to grapple with for no more than a minute is, how do you construct generative AI? How do you see generative AI from your professional perch?

[00:02:20.24] And this morning, through the lens of Copyright Law, we constructed AI in a few ways. We constructed it through the lens of authorship. Is AI an author or not an author? Through the lens of

infringement and Fair Use, we ask, is AI a copy machine or a legitimate copy machine or not?

[00:02:38.91] But there are lots of other ways that we might construct AI. We might construct it as a consumer product for search. We might construct it as a tool for creators to aid in their work. We might construct it as a replacement for creators. We might construct it as an object of research and development. And I'm sure you all will come up with some other ways.

[00:03:03.84] But I'd love you to reflect, what brings you to this conversation? How do you think about generative AI? And I want to start with Sarah Jeong, the deputy features editor at The Verge. Sarah.

[00:03:15.30] SARAH JEONG: Well, in my day to day work, we see it as a thing to cover. And that means something that is misinterpreted, lumped in, and really, a lot of different things that operate very differently, things that are not even necessarily AI by even an expansive definition of AI are lumped in together. And as part of our work, we try to be very careful. And we try to interrogate what the technology is that we're looking at.

[00:03:45.67] But yeah, it's-- I think that the thing that I'm primarily interested in is that it is a source of content. Notice I didn't say like a content creator. I think that there's been a really interesting linguistic shift like in the last even five years-- 5, 10 years.

[00:04:06.34] At first, we started talking about the creative industry as the content creators, the content industry. And then eventually, we shifted to talking about creative work as content. So people start referring to their art as, oh, I make content, which I think is really fascinating because you're basically admitting that your artistic output is ultimately ad inventory.

[00:04:30.27] And I think that generative AI is really interesting because it is creating a thing not necessarily art, not necessarily culture, not necessarily information. It's creating content. And that's-- but that's what's relevant right now is ad inventory.

[00:04:48.98] BLAKE REID: Perfect. All right. I'm going to shift us up to the screen to introduce another speaker who's, despite being on the screen, is not an artificial intelligence, although possesses the power of one. And that is Amanda Levendowski, who is an associate professor of law at Georgetown Law. Amanda.

[00:05:08.74] AMANDA LEVENDOWSKI: I think that generative AI is a great source of billable hours for a number of lawyers.

[00:05:13.44] [LAUGHTER]

[00:05:16.42] As a clinical professor myself, it's definitely a source of opportunities for my students to learn about where sociotechnical harms, law, and technology intersect. But I would also just say that to

go off Sarah's really astute point that it generates content, a lot of that content is bullshit. It's really good at generating beautiful bullshit.

[00:05:40.06] And I think that's something that we have to hold in our minds at once, which sometimes, it's really compelling, and interesting, and engaging, and yet, a lot of what it's really doing is reflecting and amplifying human biases back at us. And keeping both of those things in mind, I think, is really key to the conversation.

[00:05:57.23] BLAKE REID: Thanks so much. Moving down the row, we have Ellen Goodman, who is a distinguished professor at Rutgers Law School, but also wear some other hats that I'll let you explain to whatever extent you want. Ellen.

[00:06:12.35] ELLEN GOODMAN: Right. Let me just dispose of that. So I wear another hat where I'm on loan to NTIA, the Department of Commerce, working on AI accountability policy. And the only other thing I'll say about that, before I tell you that everything I say is in my personal academic capacity, is that Colorado lent us two amazing interns, one from the law school and one from the Department of Computer Science, Xelef Botan, and they're sitting here, and Christine Chang. Can you raise your hands?

[00:06:43.81] [LAUGHTER]

[00:06:46.52] Without their-- I mean, it is absolutely true that students from here get amazing opportunities and it is absolutely true that they have been and are amazing assets to the work we're doing. OK, that's it for NTIA.

[00:07:02.72] So I come-- and I think the interesting thing is-- and we'll talk about copyright and the hammer analogy, and I think probably, we all look at generative AI and AI in general, depending on what tool sets we're bringing to it. And so I come from a media policy background.

[00:07:21.41] And my first interest-- my first contact with algorithmic decision systems and with AI was really about Democratic self-governance and human autonomy, and the question of whether these tools empower people to take more control and exercise, more authority in their lives, or whether they do the opposite by obscuring decision making, producing bullshit, and otherwise, disempowering people. And so that's-- it's at that-- it's that fulcrum that I'm most interested in.

[00:07:53.63] BLAKE REID: Awesome. All right. Moving down the line to our colleague from across campus, Casey Fiesler, who is an associate professor of Information Science here at CU. Casey.

[00:08:06.23] CASEY FIESLER: So for me, generative AI has been a source of TikTok views. [LAUGHS]

[00:08:16.12] One way that I've been thinking about it is How-- is this tension between, in what kinds of tasks is it going to be doing to

replace us versus augmenting us? And I think the answer is both in different contexts, where the replace-- any kind of replacement is probably going to be bad. But in the end, when people say like, oh, AI is stealing our jobs, AI isn't doing anything. People are firing you.

[00:08:46.75] So it's not-- it's-- so it's not-- that's not really a feature of AI, but more of a consequence. But I was trying to think of a way that I would use AI that I-- this kind of AI that I would consider appropriate in terms of thinking about what it could do for me.

[00:09:04.88] And I would not want to use AI to help me build a world for a Dungeons and Dragons campaign because I can do that for myself. That's the creative thing that I would be really excited about doing. But what I would love to do is write 100 words about an NPC that I want to build and then have ChatGPT give me a character sheet for it, because that seems not fun to do and a lot of work. And so-- but then I would take it. And if I didn't like it, I would change it.

[00:09:33.71] So I think that there are ways that it can be like a helper in the most generous way there.

[00:09:43.34] BLAKE REID: Fantastic. All right. Last but not least, over to Mike Fricklas, who is the chief legal officer and corporate secretary for Advanced Media. Mike.

[00:09:53.42] MICHAEL FRICKLAS: Thanks. So just to give a little context, Advance is a company that's in a bunch of different things, from venture capital to owning media companies. We own the Conde Nast magazines, we own a bunch of newspapers around the country. And our roots are really in journalism.

[00:10:10.79] And so even though the journalism from an economic point of view is a relatively small part of the overall enterprise, it's one that I'm really interested in. And before, I was at Advance, I was general counsel of a company called Viacom, which was in media in many, many different ways. We were in books. We owned Paramount Pictures, lots of cable television, lots of things like that.

[00:10:30.35] My whole career, I started here as an engineer at the University of Colorado. My whole career has been at the intersection of business and technology, and there have been lots and lots of transformations.

[00:10:41.93] There's two kinds of transformations from technology. One is the good stuff. And that could be some things become obsolete because you're able to do them better. And you lose some jobs because the people who are doing the old stuff don't have a thing to do anymore. Maybe you were the person who made brass belt buckles in 1810, and fashion changed, and there's no brass belt buckle in industry anymore. And so your jobs are being destroyed.

[00:11:07.76] But there are also those things that are being damaged in the way that are really necessary things. And so in my case, the thing

that I'm spending more of my time on is thinking about journalism, because the AIs ingest a lot of our work, and that's news reporting, but it's also creative work, like at The New Yorker, Vogue, Vanity Fair, think of the newspaper work, and providing people with an alternative way to access the same information, but without putting the effort in to create that stuff. And it's not stuff that AI creates. So you need those things.

[00:11:44.12] AI can spit back to you the story that was written about the Ukraine yesterday, but it doesn't go to the Ukraine and walk among the bombs and try to figure out the personal injury stories that are going on there. And so we need it.

[00:11:55.47] There's also work that's being done on-- well, I'm getting ahead of myself, but just this particular one. That's how I'm spending my time. And somebody made the comment, it's keeping a lot of our lawyers busy because we have a business to run. And we have to figure out not just what we think the theory should be about an alternative to the copyright system, but how are we going to make our business run and the journalism work over the next three months, six months, two years in addition to trying to figure out how to work these systems into making that happen.

[00:12:27.71] BLAKE REID: So I want to reflect some of these back, just so we get a sense of the breadth here. So I heard, AI as a source of content, I heard AI as a source of sociotechnical harms, a source of bullshit, a source of big effects on Democratic self-governance, a potential threat to human autonomy, a potential cause of disempowerment, a replacement for us or a replacement for people to replace other people, maybe an augment for us, maybe a destructor, displacer of labor, maybe as a consumer of other content, and maybe as an output and an imperfect replacement for journalists.

[00:13:13.15] So this is a pretty wide range of social impacts that we're dealing with here. And as Professor Bracha teed up in the previous panel, we are filtering this all through copyright, and filtering it through pretty arcane doctrinal questions about authorship, and about infringement, and Fair Use, and so forth.

[00:13:34.10] So I want to go meta for our first question here. And I want to turn to Sarah because your work has explored historically how copyright has shown up first on the scene. This is not the first time that a technological disruption has come along, and copyright has been there to meet it, and with lots of consequences, good and bad. So I wonder if I could turn to you to break down, how is copyright succeeded and failed in this endeavor in the past? And then we'll turn to AI.

[00:14:08.95] SARAH JEONG: Yeah. So I mean, I'm not going to go all the way back to the printing press because I actually, I don't remember when that happened. I was probably too young.

[00:14:19.43] [LAUGHTER]

[00:14:20.84] But you brought up this question yesterday, and it was primarily to a room full of one else, and so I had to explain what Napster was. But--

[00:14:28.78] [LAUGHTER]

[00:14:29.81] --I don't think I have to explain Napster to you guys. What happened with the internet is like in the '90s, let's go back, there was like a genuine question of whether the law applied to the internet. And it seems really absurd for us to think about it now. But there was like there were actual books written by well-regarded people, Jack Goldsmith, Tim Wu, one of my favorite books called Who Controls The Internet, where they had to explicitly make the argument that the government, in fact, controls the internet in specific ways. And that was-- they had to actually approach that question in a serious way.

[00:15:06.45] And what happened was that in the '90s, the first to rush into the breach into this wild frontier, it was Copyright Law. It seems really weird now. It's like, the fact that copyright got there before national security, before privacy, before child exploitation, not really, but also, it got there first. And it was the most successful in bringing the internet to heel, to the point where when startups were spinning up very early on and flying by the seat of their pants, and they didn't have a lawyer. They'd go like, no, you really need a lawyer. I was like, for what? It's like, all you need a DMCA agent. You need a copyright lawyer.

[00:15:49.70] Obviously, these days, they're aware, I hope, that they need many more different kinds of lawyers. But yeah, copyright, for whatever reason, became the first thing. And so what you see is, immediately, whenever you have a problem on the internet, people run to copyright to try and fix it.

[00:16:09.05] So you see, there was a case in 2014, super fascinating one. Garcia v. Google, where a woman who has been tricked into playing a part in the Innocence of Muslims, that's the film that may or may not have sparked Benghazi. She was receiving a lot of death threats on the internet because of this upload to YouTube.

[00:16:34.34] And so she sues Google for-- based on her copyright interest in this film because she appeared in it, which is not how copyright works. Just because you appear in a film or in a photograph doesn't mean that you own that work or have an interest in it, whatsoever.

[00:16:54.24] And so it goes through. It's a really convoluted case. And it's really interesting to go and look at the Ninth Circuit argument there or the second one, the en banc one. Because the lawyer for Cindy Garcia starts out by reading out the death threats. And immediately,

the judge interrupts and goes, what does this have to do with copyright law?

[00:17:15.29] And yeah, the audience was full of like, it was MTA lawyers. It was Hollywood lawyers. They were here because they were really invested in how this was going to go down for Copyright Law. And meanwhile, Cindy Garcia just wanted to stop being afraid.

[00:17:28.19] And so you see this pattern play out again and again where people run to copyright to try and get things done because there is a robust structure in place for taking content down from the internet for judging whether or not content should stay up on the internet, and also for reaching through the screen and finding the person who did the thing, and getting them. Because the RIAA was incredibly successful in subpoenaing IP addresses, finding individual infringers, going after them, basically, finding who did this, this anonymous person who did this, and getting them.

[00:18:03.45] And so-- and that's copyright is an incredible ill fit for when you're dealing with privacy concerns, being stalked, being threatened, domestic violence. And yet, it comes again then again to the fore.

[00:18:16.90] I think that the law has become a little more robust in dealing with a lot of these other issues over time. It is the internet is certainly much more regulated than it was when I was in law school, which is not that long ago, honestly. But still, it is-- maybe copyright in the internet is still probably one of the most robust bodies of internet law that we have.

[00:18:40.90] BLAKE REID: So I want to open this up to the panel. And I think the life span of the commercial internet is probably the perfect technological period to tap in our inquiry to here. We're 25 years into the life of the commercial internet. We're at the midlife of the internet, call it.

[00:19:00.69] How, as we take measure of that, has copyrights stacked up in being that first law of the internet? We've got some examples of it not doing so well. Other perspectives?

[00:19:13.48] ELLEN GOODMAN: Can I just-- well, I just want to-- I want to challenge this framing that copyrights first contact with the internet because in some ways, the reason why we think of it that way is because there was a recession and 230. 230 was actually first contact or non-contact.

[00:19:31.42] And so 230 precedes all of that and says, we're not going to-- we're not going to make any laws around here. And I think that as a result, when I look back on those 25 years, it's been-- was it 25 you said? Was it--

[00:19:46.15] BLAKE REID: Yeah. If we go back to '96. And then we're even more than that.

[00:19:49.57] ELLEN GOODMAN: Yeah. That it's the way pine trees create a dead space underneath them, this monoculture. And copyright has operated that the lack of any laws. And so we don't have this kind of taut muscle around-- that we might ordinarily have around the internet. And I think we're paying the price for that in AI because copyright is the only kind of muscle that's been exercised.

[00:20:19.58] BLAKE REID: So maybe it's a little different in that copyright at the introduction of-- the inception of the commercial internet was coupled by an intentional abdication of every--

[00:20:28.22] ELLEN GOODMAN: Well put. Yes.

[00:20:28.91] BLAKE REID: --corner of law. Right. Other thoughts? Casey.

[00:20:31.70] CASEY FIESLER: Well, I think this is so interesting. And I think that one of the reasons that we got to this place where copyright has this power, especially in online spaces is because of the systems of power that underlie it. And if you're talking about something like YouTube's content ID, what is YouTube going to do to make Universal Music happy?

[00:20:53.39] But now that everyone is a content creator, there are ways in which copyright law becomes the only kind of recourse for normal people. We see this for things like revenge porn. But also very recently, a content creator, a friend of mine who's extremely popular on TikTok and Instagram did not have a Facebook page because like, why would she want to be on Facebook?

[00:21:20.54] And someone was taking all of her videos and reposting them on a Facebook page that had 100,000 followers. And they all thought that this page was her. And she messaged the person running it. And they were like, well, I'll hand it over to you if you want it. I'll be your social media manager.

[00:21:37.01] And she tried for a long time to get Facebook to take down this page because it was impersonating her. And finally, I said, file a DMCA notice. Gone, hours later, because that's the only thing that often works.

[00:21:52.35] So now if we're thinking about the copyright issues underlying generative AI, people's content from the internet, for example, being used, I still think we're seeing an interesting power dynamic here because the lawsuits that we're seeing that are getting attention are things like Getty Images and famous authors and that sort of thing, and not indie artists on DBR or whatever.

[00:22:17.35] BLAKE REID: All right. I promise, we'll get to AI, but we'll get Michael's reactions. And then Amanda, if you want to jump in on this one, or if not, we'll jump to the next question.

[00:22:24.55] MICHAEL FRICKLAS: Yeah. I do think that one of the things that's really interesting when the internet first became a thing. There was a lot of writing about read-write culture and the like, about everybody becoming a creator and it not being concentrated among a few companies who had the power to exploit copyright. And I think a lot of those early cases had to deal with those big companies.

[00:22:42.50] But I think one of the things I've really seen recently is this whole point that you're just raising, which is now that everybody is a creator. They're also starting to see the threat. Neil deGrasse Tyson was on Stephen Colbert's Late Show on Monday night. And he was asked this question. And he said, I feel really strange about it. AI is a really great new tool. It has all this capability. But I'm really uncomfortable with the idea that my book was taken and is being used for this new product.

[00:23:07.76] And so I think one of the things that's changed is that-- well, and this point that you raised, we see-- we put up videos, we see other people take our videos, download them, and re-upload them to YouTube, and generate thousands-- hundreds of thousands of views for themselves. We've seen-- we see lawyers get-- now getting copies of articles. People are complaining about the articles. And the articles weren't generated by us. They were faked.

[00:23:34.67] And so bringing some law in order to this whole universe, I think, is something that involves everybody and not just the large companies.

[00:23:43.31] BLAKE REID: Amanda, any thoughts on this one or should we move to our next topic? I thought AI was going to solve our mute problem.

[00:23:50.74] [LAUGHTER]

[00:23:51.92] Not happened yet.

[00:23:52.80] MICHAEL FRICKLAS: I hope so.

[00:23:53.33] AMANDA LEVENDOWSKI: It's still a work in progress, which is exactly what I think this panel is about is the fact that there's still a lot of bugs in the system. And I want to go back to what Sarah was saying about copyright being one of the first on the scene, that's because the internet is really made out of copies. And so when you're looking for a law that relates to copies, even though section 230 is also there, looming as well, if you're dealing with copies, it makes a lot of sense.

[00:24:17.61] And then you have these two competing parties thinking about fairness. You have people who want to use copies, thinking, well, it's fair for me to use this copy so that I can expand knowledge, or share music, or create art, or do whatever they want. And then there's other people that say, hey, this is really unfair to me because I want to be getting credit compensation and consent to use my works. And

that's the kind of tension we've been seeing between these two sides for as long as we've been talking about copies on the internet.

[00:24:48.81] And I think that reveals a real FU in Fair Use from the perspective of authors and artists who want to have another C word, "Control," over how those works are used. And that comes into tension with the whole doctrine that governs fairness and copyright because Fair Use allows for uses without any of those things authors and artists want. No consent, no compensation, no credit, all totally fine in the right circumstances under Fair Use.

[00:25:15.85] And so it makes a lot of sense that now that we're seeing another conversation involving really, a debate about what's fair on the internet, copyright is there as it always has been.

[00:25:28.32] BLAKE REID: Well, Amanda, why don't we take this opportunity and shift into AI? We've got this backdrop of the internet and maybe contestable how copyright has played its role. But let's shift into thinking about copyright as a tool for governing and regulating AI.

[00:25:48.36] So you've been a proponent, at least in some of your scholarship, of analyzing AI problems through the lens of Copyright Law. Take us down what the pros and maybe if you feel like it's some of the cons of using copyright as a lens for AI regulation and governance, what does that look like.

[00:26:09.69] AMANDA LEVENDOWSKI: Yeah. So I think one of the cons is that Copyright Law is supposed to promote a very particular type of progress of the useful arts and sciences. It's not to promote social justice. It's not to promote equity. It's not to promote fairness in the colloquial sense, not in the sociotechnical sense. And sometimes, it happens to do those things.

[00:26:31.08] We can look at some of the Supreme Court cases over the years, some of the circuit court cases over the years, some of the district court cases over the years, and see where those values happen to align. But that isn't really copyright's purpose, necessarily.

[00:26:44.08] And so in my scholarship, I've talked about how Fair Use may be necessary to mitigate some of the biases we're seeing in AI systems, particularly around race, gender, sexuality, class, and disability, but at the same time, that doesn't mean that it's promoting fairness for the authors and artists who would really prefer that their works not be used to debug those systems for whatever reason.

[00:27:04.24] And so in the other context, on the other side of things, we have very specific uses that may not be Fair Use, like I think certain types of face surveillance that are using profile pictures for the exact same purpose of recognizing faces, which I call particularized identification. But that's cutting it incredibly fine to make a distinction between these various types of AI uses of copyrighted works.

[00:27:31.45] And so I think the challenge that we're really seeing now is, like I was saying, this FU and Fair Use, this lack of agreement about what actually is fair from a fair use perspective. And I think that's what a lot of this litigation is going to get at as we see it unravel or ravel. I don't know if it's going to ravel or unravel.

[00:27:50.30] [LAUGHTER]

[00:27:52.29] BLAKE REID: So I want to open it up to the panel on, I heard at least two fronts. And maybe let's start with this notion of interest convergence between copyrights, interests, and broader questions of whether we think about social justice or think back to the previous panel about goals of the market, and how those might diverge sometimes. Curious for folks' reactions on that.

[00:28:20.53] SARAH JEONG: Yeah. Actually, so previous applications of copyright to thing-- things that aren't really quite copyright, I've been pretty critical of. In this instance, I actually think this is in copyright's wheelhouse.

[00:28:35.17] The downside, as Oren stated in the next panel, is that copyright in the United States, at least, is very much market-oriented. And the harms-- the most urgent harms that are posed by generative AI, which by the way, I don't think it's 100% bad, but there are some serious problems looming, they're not really market problems because here's the thought experiment.

[00:29:04.09] It's that, let's pretend we lived under full communism. Fully socialist society. No one has to worry about going hungry tomorrow. All of the artists have health care. They live in stable housing. They don't have to worry about getting sick or taking care of their children or sending their children to school.

[00:29:24.74] Do we still have a problem here? And it's yes, actually, we do, because there is more going on here than compensation, than the systems at play for paying for work. It's that when you make it so cheap to generate content-- and again, I'm not saying art, I'm not saying culture, I'm not saying information, because none of this stuff is necessarily informative, true, valuable, or even contributing to human flourishing in any way whatsoever. When you have so much of it, it becomes harder to find the information you're looking for.

[00:30:07.88] And this is already a thing we're seeing with search engines. We're seeing search engines decay in their reliability over time, because I mean, if you just look the first page of Google results, a lot of that's already AI-generated, a lot of those links. A lot of those links are not necessarily reliable information.

[00:30:24.88] You might notice that it's a lot of really wordy repetitive paragraphs, like it's-- and that it's not necessarily true. It's not necessarily even addressing the question that you asked. It is becoming harder and harder to find the answers to the questions that

we have. It's harder for us to figure out where supposed answers came from, who said things, whether it was actually said by the person.

[00:30:49.69] And then we're just getting-- there's also sort of much harder to measure this impact. But when you get the culture flooded with stuff that's just rehashed images that are just mixed up in between pieces of everything that's existed before over and over again, what does that do to us as people, as a culture that wants to progress? I don't know. I actually really don't know. But I don't love it.

[00:31:21.17] I know as journalists, there is a problem when consumers cannot actually find out what's happening or cannot trust what it is that they're reading, and there's no reliable way to distribute and disseminate information because there's too much information.

[00:31:37.78] Historically, the problem has been that there's not enough access to information, that there's not enough information, that we aren't incentivizing the creation of information or making it accessible to people at a price or for free in a way that is good for society. But right now the problem is now and increasingly going to be that there's too much content and that it's not possible to navigate through all the content.

[00:32:04.45] So under full communism, it's still a problem. And so yes, there is an issue there with copyright being market-based that it's maybe not a good fit.

[00:32:15.01] On the other hand, copyright has been, since even before the statute of Anne, the primary way in which the state has regulated the dissemination of information for the purposes of strengthening our culture, strengthening our bonds with each other through how we communicate.

[00:32:35.36] So why not? We have this robust body of law. Why not turn it to the purpose of this big problem? Other part is that we have the robust body of law, but the institutions already exist. We already have scholars.

[00:32:49.85] I mean, imagine, if we went at this from a privacy perspective, which there's a lot of privacy problems. There really are a lot of bias problems. I'm glad that someone is working on this stuff. But it's like, if we went primarily to that first and foremost and leaned heavily on it, we'd find out really fast that the US has very much under-invested in regulating anything from these perspectives. And that's a problem, for sure, but in the immediate right now, what framework do we turn to? I think copyright is such a terrible fit, honestly.

[00:33:30.01] BLAKE REID: Mike, head down to you.

[00:33:31.62] MICHAEL FRICKLAS: I want to-- I think those are a lot of really interesting and important problems. But I think one of the ways that we think about some of the solutions to some of those problems is that the reputation of the speaker is tremendously important, and

maintaining that connection with the truthful understanding about who's speaking, and relying on sources that you know to be trusted.

[00:33:55.63] So one way you do that is through brands. The New York Times is a great brand. And if you go to The New York Times, hopefully, you have the experience that whether you agree or not with their opinions, you trust the facts that are generated, and you know that they're thoughtful.

[00:34:10.65] The same thing with a professor at the University of Colorado. It's-- they have a name, they have-- you know you're going to get a certain amount of useful information. I think we do rely--

[00:34:20.41] BLAKE REID: A lot of useful information.

[00:34:21.85] MICHAEL FRICKLAS: I think--

[00:34:22.27] [LAUGHTER]

[00:34:23.17] I think we rely a lot in navigating this expansive amount of information more and more on knowing, on finding somebody, finding a source, finding sources that deliver to you on a consistent basis the kinds of information that you want to achieve.

[00:34:36.83] And so staying away, thinking more about the policy front than the copyright front, I think sets of rules that allow people to have trusted connections with sources of information is very important aspect of the ecosystem.

[00:34:53.02] BLAKE REID: Mike, can I follow up just a really quick follow-up directed at you, and then open it up to the rest of the panel to keep going on this convergence question? Do you feel like copyright accomplishes that goal? Or in what ways does copyright succeed on that front? In what ways does it fail?

[00:35:10.06] MICHAEL FRICKLAS: Fair enough. Both, it succeeds in some ways and fails in others. So in terms of success, what it does is it provides-- and I'll use The New York Times again as my example, provides as The New York Times with a source of revenue.

[00:35:22.63] The people who like The New York Times, they read The New York Times. They trust it. They appreciate what it does. They engage with it in an economic manner. So they pay for subscriptions. They see advertising. The New York Times gets traffic that's helpful to their advertisers. And so there's that market transaction where copyright enables the journalism and enables the brand, it enables the trusted source to continue to do the very hard work through hundreds and thousands of journalists to give people something that they want.

[00:35:55.67] And so if copyright is functioning correctly-- and there's lots of issues around the edge, if it's functioning correctly, that works.

[00:36:02.89] I think copyright, the bill, I won't go back to the Statute of Anne, but we're pretty much operating under a law that was mostly

written in 1909 or in the years leading up to 1909, with some tweaks in 1976. And Fair Use is a judge made concept that's really a policy set of rules, more than it is a legal set of rules.

[00:36:25.66] You can have cases that go absolutely opposite directions, depending on the situation. And that's true for other kinds of concepts built into copyright, like what is a copy to what is a derivative work. These are words that are not really that filled with meaning, and they leave a lot of discretion.

[00:36:45.74] And so if you're a journalist trying to figure out how to make a living, and you're trying to say, OK, what can people do with my work that's going to compensate me or not compensate me for the effort that I did? It's really hard to have a level of predictability.

[00:36:59.59] And as well, I think, really big and well-funded actors have pushed some of those concepts beyond any recognition. Those concepts weren't written-- to start up, they weren't written for technology. They were written for normal expressive kinds of works. Does this play looked too much like that play? Did they just change the name of the characters? Those are the kinds of issues we were dealing with in the first place, not whether generating billions of page views in a particular aspect are-- undermines the incentive to create or not. They're just not designed.

[00:37:33.01] So all that's judge created, judges, they have two litigants in front of them. They don't have all of society. They don't have hearings. They don't have major policy debates. And I think our Congress has been very-- in situation that makes it very difficult for them to get into constructive policy debates as well.

[00:37:53.00] So it's really strained what copyright can do. And it's left just a high degree of unpredictability.

[00:37:59.56] BLAKE REID: Sarah, I want to get into institutions down the line. But before we leave these questions about copyright doctrine and convergence, I want to make sure to invite Casey and Ellen on those. Any reactions on that?

[00:38:17.01] CASEY FIESLER: So I actually think that your point about the importance of the source of information is really-- is going to be really, really important. I've been thinking a little bit about how mechanisms for attribution could potentially help with that. When you see a screenshot of a tweet from the CDC, if there was some way in a technical manner that you knew for certain that that came from them.

[00:38:42.51] But even then, I feel like there's this very generous assumption that people care. And one of the most pressing things I saw recently was that someone photoshopped a tweet from the CDC. And then tweeted, oh, the CDC tweeted this and then deleted it right away. But I got a screenshot. And it was about there being an Ebola outbreak at Burning Man. And--

[00:39:04.97] BLAKE REID: [LAUGHS]

[00:39:05.28] CASEY FIESLER: --this is what started--

[00:39:07.62] BLAKE REID: That's not funny. Sorry, go ahead.

[00:39:08.97] [LAUGHTER]

[00:39:10.42] CASEY FIESLER: It-- but so many people believe-- like 100% believed this, just because they saw a screenshot of a tweet. And everyone knows how easy that would be to fake, I would like to think.

[00:39:25.60] And now the issue is that you're going to see a press conference with the director of the CDC, a video of it. And you're not going to be able to know that that's real. And so the only thing that you can do is trust the source of the information. Did it come from a professor at CU Boulder? Or did it come from some random TikTok user with a cartoon avatar?

[00:39:46.42] And so how do we get people to care about the difference? I don't know. [LAUGHS]

[00:39:52.27] BLAKE REID: Well, Casey, I want to ask you the same follow up I asked Mike, which is, does copyright give us any tools at all to effectively solve that problem?

[00:40:01.55] CASEY FIESLER: See, I think actually, this is one of the examples of where it does. I mean, I think the kinds of things that have been brought up, privacy, bias, misinformation, I don't know the copyright does give us the tools to deal with those things. And you're absolutely right, we just-- we don't have-- we're going to end up with these AI laws that are mishmashes of all of these kinds of things that then only apply to AI, which I think is going to create a weird regulatory framework.

[00:40:35.48] We watch congressional hearings with all of these complaints about where data on social media is going. And yet, we have almost no meaningful federal data privacy laws. And we might get something related to AI, and then it won't help us with everything else, which is a slightly different issue.

[00:40:53.47] BLAKE REID: All right. I think we're going there. So let's go there. Let's talk about--

[00:40:57.49] ELLEN GOODMAN: Wait, wait. I want to--

[00:40:58.33] [INTERPOSING VOICES]

[00:40:58.75] BLAKE REID: Please. Please.

[00:40:59.75] ELLEN GOODMAN: OK. So just take a step backwards. It's not surprising that copyright is being used to deal with fears about job displacement, human dignity, and autonomy, concentration of power. It's a property rights system. Property rights have always been used and abused for all kinds of things. And I mean, as you pointed

out, it's not just the little creators, it's also Disney. On both ends of the copyright spectrum, there's been abuse of the copyright system. All right. So it's not surprising.

[00:41:34.12] Second of all, I think there-- and to Casey's point, yes, if we're concerned about our epistemic system, and truth, and attribution, and we're concerned about the demand side, not just the supply side, that's a problem that copyright is not going to solve. And that's a bigger societal problem.

[00:41:54.38] So I think we can stipulate, copyright has a piece of-- has a piece of this, but not all of this. But-- and maybe we can-- in the next round, we can get to this, I'm interested in the plumbing things that we can do that will help on the copyright side, but also help more generally with some of the other problems.

[00:42:13.16] So for example, with attribution, there are provenance systems. And there are ways in which they can be fortified. And there are ways in which that can maybe-- without expanding copyright liability, there are ways in which maybe that can address some of the credit, the needs for credit that copyright holders want. It doesn't-- and it may-- and it may also advance compensation and control and consent if at least, we can determine what the provenance of both the training data is, the inputs, and also, the outputs.

[00:42:49.06] And I think the same is true for-- and Oren-- Professor Bracha said that this is not free-riding. And I just-- I-- maybe we can talk about that a little bit because I think it is free-riding.

[00:43:03.28] SARAH JEONG: Exactly.

[00:43:04.51] ELLEN GOODMAN: And I certainly think it's free riding on the journalism. And so to the extent that it's-- that that's what we're concerned about, is copyright the right tool if we don't really have the right copyright doctrine to deal with that? But we do have other ways of having-- I think-- who said tax? Who said tax and-- yeah, yeah, tax and redistribution. OK.

[00:43:27.86] So we have ways to deal with that. All right, I'll shut up. But yeah, so copyright is definitely a piece of this.

[00:43:33.69] BLAKE REID: So I have two directions I want us to go. And I think we'll go first to the question of, what redeeming qualities does the copyright system have specifically? What are the-- what are the set of problems that we can solve? And then I want to shift us over to the broader range of problems that we have a pretty good sense it can't solve. And Ellen, I'll come to you for the latter question, but let's start with that first one. What is copyright going to be really good at? Amanda.

[00:44:04.32] AMANDA LEVENDOWSKI: I love-- that was just met with silence. Everyone's like--

[00:44:07.92] BLAKE REID: No. OK. Ellen pointed at you, so you got it.

[00:44:11.58] AMANDA LEVENDOWSKI: Yeah. I think they're-- even though it is not meant as a privacy enforcement tool, and I think Sarah is very right to be skeptical of that, even though a hilarious amount of my scholarship is about why we should do that. Anyway--

[00:44:25.77] [LAUGHTER]

[00:44:27.75] --I think that it is actually given the interest over time that have shaped Copyright Law, which are very powerful. We do have a really robust hammer to hit some nails. And some of those nails may be related to privacy.

[00:44:42.99] We've seen this play out. Casey alluded to non-consensual intimate imagery earlier. The DMCA is a really powerful tool for taking some of that stuff down when simply asking nicely to a platform doesn't actually make any difference. And we saw that with Celebgate in I think, 2014, which I know Sarah has also written about.

[00:45:02.61] We may see it happen with tools like face surveillance, in the context of, say Clearview AI. There may be entrepreneurial plaintiff's lawyers who are looking for in the same way that they've used right of publicity in this sense, which is really not exactly what right of publicity is for, as Jason Schultz has written about. It can be appropriated for some of these privacy-related, justice-oriented solutions.

[00:45:27.85] And I think that in the hands of creative, competent lawyers, there may be an opportunity to push copyright law into the space where it doesn't cleanly belong. But Warren and Brandeis looked to Copyright Law to get some guidance in their hallmark privacy law review article. I mean, to the extent that we're talking about scholarship not by CU professors, which I honestly don't think we should get into too much.

[00:45:51.09] [LAUGHTER]

[00:45:52.86] But it's been there all along. And there's not necessarily a reason not to think that it can't do some of this work here.

[00:45:59.56] The other thing that I think is so important is, one of the factors or-- I guess, I should say, sub-factors of a factor of Fair Use that some courts look at, including the Supreme Court, is whether a technology serves a public benefit.

[00:46:12.17] And one of the things I'm exploring right now is, well, how can you dub something of public benefit without truly grappling with those countervailing public harms? And is there a place in copyright law to force courts-- force lawyers to grapple with those public harms, and come to a fairer resolution than Fair Use necessarily requires, that may, as Ellen was saying, put some credit, put some compensation, put some consent back where artists, and authors, and

journalists want it, not necessarily where Copyright Law requires it? And that's something that I'm looking at in some of my current work.

[00:46:49.21] BLAKE REID: Other thoughts on copyright as a hammer, Mike?

[00:46:51.97] MICHAEL FRICKLAS: I already made my argument about why the economics and underpinning really matter. One of the questions I have-- well, really two. One is the whole tax and redistribution thing. I want to address that for a second. One of the things that copyright does is it keeps journalism, it keeps authorship apart from dependent on government subsidies.

[00:47:12.88] And I think that government subsidies are a really dangerous thing when it comes to-- when it comes to content. We see-- if everybody here watches what happens in politics, and you see attacks on the media all the time.

[00:47:26.92] One of the simple ways to attack the media is to threaten the economic foundations. Then people get very cautious about challenging government. And that's one of the fundamental things that we do.

[00:47:38.92] And so I think that's a weak system for that reason. And I think charitable sources of-- are weak systems for that as well. And I also think that we shouldn't operate from the assumption that if there's a licensing model that AI can't do everything that AI wants to do.

[00:47:56.02] There are dozens, hundreds of collective rights organizations that have built up, once you have a system of rights, the market's been able to generate solutions to that. There's the copyright clearance center for print. There's dozens of music licensing organizations from composers represented by CISAC and ASCAP and BMI. There's photography with Shutterstock and Getty Images as another solution.

[00:48:20.08] But there's lots of ways to organize when you have a need and you set yourself a desire to solve a particular problem of rights. There are ways to license in bulk, and even vast bulk, that can be organized when people have an interest in doing it.

[00:48:38.72] I actually-- so I actually don't think copyright funds journalism. I think that journalism has been funded largely as an accident for the last 100 years. If you look at the decline of the Alt Weeklies, a lot of that has to do with Craigslist because it turned out that everyone was picking up and reading the Alt Weeklies, and putting ads in the Alt Weeklies because they needed to sell their used car or sublet an apartment, find an apartment, find a job.

[00:49:07.16] All of that economic activity went somewhere else. It was an accident that that part of the private sector was bundled with journalism. That has nothing to do with journalism. You could take any

random business, bundle it with a newspaper, and suddenly, the newspaper is viable.

[00:49:23.96] I don't think The New York Times makes its money-- its reporters, who do an incredible public service-- trying to decide how much to say about my former employer. The thing is that good reporting is very expensive, and it does not make that money back. You have to make that shortfall up with other ways.

[00:49:46.60] In the case of The New York Times, it is not Copyright Law, it is Trademark Law. It is the swag that they sell. It is the fact that there is cachet involved in you having a New York Times subscription. It is the crossword puzzle because now you are erudite because you have a crossword subscription. It is a New York Times Cooking. New York Times Cooking makes a lot of money compared to what its footprint is at The Times.

[00:50:12.77] And part of that is the fact that it's very hard for people to find recipes online without being blasted in the face with a lot of ChatGPT generated AI garbage that has a lot of ads on it. To cycle back, Wirecutter also makes a ton of money because people don't know how to dredge through the garbage. And they trust The New York Times because their trademark is extremely strong.

[00:50:37.01] And the fact that their trademark is extremely strong actually doesn't have that much to do with the reporters. The reporters, they don't care about their trademark being strong. They're just there to do journalism. Meanwhile, the Advertising Division is going to the advertising conferences and saying, hey, give us multi-million dollar deals because we're The New York Times, and this is what we do.

[00:51:00.53] And the image that they present of The New York Times being, we're-- I mean, the Washington Post can eat out on Watergate for the next 100 years. But New York Times has to go with like, we tell the truth. They had an advertising campaign that was just "The truth" dot. And they were able to sell a lot of ads on that.

[00:51:23.64] And it's like, there is this aspect of news, news-making, and doing the public service that is a little bit accidental, a little bit savvy, but it doesn't have to do with the market value and the demand for this public service. And I don't think that Copyright Law is doing a lot of the work here. I do think that shifts-- especially in copyright actually, shifts in the marketplace have resulted in a radical decline in how expensive news can be, and also, how many people you can put in a newsroom on how much money.

[00:52:10.09] You-- I talked to someone who's a generation ahead of me. And they'll casually throw out, oh, I wrote this essay. This is the backstory of how I wrote this essay. And then I bought a summer home with the royalties from it.

[00:52:23.29] And I'm just like, dude, I love my job, but I'm not going to get a summer home out of it. That's not happening. That's not in my future.

[00:52:33.04] The things have changed very, very quickly. And I think a lot of that is accidental. And it's like, it's not like-- it's not like my copyright got suddenly super invalid compared to my colleague from another generation. There have been just seismic shifts because we chose to align ourselves with specific kinds of business interests to fund this thing that was ultimately a loss-making proposition. And we bet wrong in 1,000 different ways.

[00:53:04.94] BLAKE REID: Mike, I saw you itching there. Do you want a two-finger rejoinder to anything before we move on?

[00:53:10.22] MICHAEL FRICKLAS: I didn't mean it that way. But I absolutely think that those brands are built by their content. And I think that every reporter for The New York Times and every editor sits there and says, does this make my brand stronger? Does it make it weaker? Is this who we are as a journalistic organization? And they care really, really deeply about that.

[00:53:28.31] I spent a lot of time with journalists. I spent a lot of time with people like Sarah, who cares really, really deeply about her name and her byline being associated with quality journalism. And so I do absolutely think that that's an important facet of all of this.

[00:53:42.77] Certainly, the markets have changed. The positions have changed. The economics of journalism changed. But I think that fundamental connection between a user of news and information services and knowing the source--

[00:53:57.62] I'll make one other point which is that, I think Casey is right that there's a lot of misinformation out there. And I think a lot of people are gullible about it. I also think it's remarkable to me how people I talk to who are younger and digital natives have a pretty good source of what to rely on and what not to rely on, or at least, what to check on and what not to check on.

[00:54:19.30] When my kids go to Amazon just to pick on somebody, they quickly point out that doesn't look like it's real, and that does look like it's real, or that best of list is paid for by an advertiser, it's not legitimate. And I think there are a whole bunch of signals that they rely on for that.

[00:54:37.94] I think it's really important as we think about AI to keep those signals in place so that people have ways to tell what's true and what's not. It's not going be perfect. And there's going to be plenty of disinformation and people who are maybe laws, and hopefully, laws that make it harder to create deepfakes and to create mis- and disinformation, and not identify its source.

[00:54:59.96] But I also think just that attribution piece. It's also important from an attribution point of view that people not just take the attribution. So coming back to copyright for a second. There are some experiences that people have designed that show attribution or something as akin to attribution as a way to create validity for the output of an AI.

[00:55:24.20] And those get you the attribution piece, but what they don't do is anything about the economics. So in other ways that they're even trading more on those trademarks that are valuable and saying, oh, you can believe in this. It came from The New York Times, but it's not generating any revenue for The New York Times. It's not generating a link, like they used to get from search, it's not generating an ad revenue at all. And these are issues that Copyright Law may not be perfect, but it's going to have to address.

[00:55:49.37] SARAH JEONG: One really quick addition that occurred to me. I really want to go back to tax and redistributions law, I think, off the table because there is a very private, public framework for monetizing music.

[00:56:04.43] We don't have this moral panic around the fact that the streaming royalties are very much intertwined with Washington. This is-- it's so arcane that I really don't want to get into it. Let's not-- let's just not do it. But there is no moral panic around that. There is no moral panic around the fact that we have a quasi redistribution framework there. We-- but-- and yet, for some reason, tax and redistribute has never been considered in a serious way for news, art, et cetera. Maybe we should put it on the table. It's like, I mean, music is imperfect, but it hasn't failed.

[00:56:44.84] BLAKE REID: I mean, the obvious answer there is, the government has terrific taste in music.

[00:56:48.31] [LAUGHTER]

[00:56:50.78] Ellen, I want to shift to our next part of the conversation and come to you. And we previewed a bunch of this already. But let's start talking a little more explicitly about the problems that copyright does not cover. And/or maybe to frame it a little more positively, the overarching regulatory agenda for AI. Give us a sense of what that landscape looks like. And we've got bits and pieces of it already, but bring it together.

[00:57:18.82] ELLEN GOODMAN: So let me-- the take off point that I'll use is this word "Trust." And I think, so Michael, you used the word brands or sources that people trust. And trust has become-- I think it was the OECD that first used this trustworthy AI. And I think it would be really interesting to interrogate whether that's a good framing or not. But that is what Europe and the US have come to use as the goal. And so trustworthy AI.

[00:57:51.18] And it doesn't have a ton to do with copyright, although, as I said, it touches on copyright to the extent that we view credit attribution. Those are all attributes of trustworthy AI, safe, reliable, privacy respecting, fair, all of these abstract nouns, are those nouns, adjectives, are-- that's what makes trustworthy AI. OK.

[00:58:18.63] So very few of those, whether or not an AI system, and let's just take it out, let's say just narrow AI, so a hiring, an AI system that's used for hiring, whether or not it actually does what it says it can do, which is to sort the wheat from the chaff in terms of the applicants, whether it does that in a way that's better than the humans-- the human judgment that it replaced, that is a question of efficacy and validity. And that is a question of trustworthy AI. I don't think copyright has anything to say about that. And we could go down the list.

[00:58:55.59] So I think each of those, in the US structure, is dealt with-- those attributes are dealt with in verticals, in sector-specific verticals. And we could talk about whether or not the pros and cons of dealing with them that way.

[00:59:12.81] The only other thing-- the last thing I'll say about copyright is that copyright is a speech regulation. It's like one of the only speech regulations that we have. And I think that's another reason for this overweighting of it because now we've got AI or before that. Just like ones and zeros and our digital world is all mediated by speech. And then once we breach the digital meatspace divide with autonomous agents, our whole world is going to be mediated by speech.

[00:59:49.26] And we have, in the US, all of these, for better and for worse, obstacles-- First Amendment obstacles to regulating speech. And so that is-- in answer to your question, what can copyright do or not do? Copyright is our only speech, one of our few speech regulations. And it's a question about how much we can regulate these other things, including the production of deepfakes with any other tool.

[01:00:16.17] BLAKE REID: Although, I would observe with copyright that where it giveth on speech, it taketh away because we have to deal with all those gnarly questions about fair use, which at least, Justice Ginsburg claims is a creature of the First Amendment.

[01:00:30.09] But I want to keep going on this list, and make sure we've got the full list. Any other contributions to the list of problems that we need to be worrying about that we haven't mentioned already? Any others that folks want to throw on the pile? And then I want to get to this question about verticals and how we do this, and what institutions. But any other policy concerns that we ought to have on the table here?

[01:00:56.28] All right. Pin drop, we've covered it all. Well, let's talk then about how we go about doing this. And one, I think, question

that's come up throughout the conversation today is, where's the institutional home for doing this?

[01:01:13.57] So with copyright, we've got-- the Copyright Office is doing some policy making, some adjudication. Interesting just to point out that after all the conversation about authorship this morning and AI authorship, that didn't come up in this conversation as a real policy concern at all, interesting that no one mentioned that. But we've got the copyright--

[01:01:35.16] ELLEN GOODMAN: It's just copyright.

[01:01:36.48] BLAKE REID: That's just the copyright thing.

[01:01:38.76] CASEY FIESLER: They talk to them a lot.

[01:01:39.52] BLAKE REID: So we've got the Copyright Office. We've got courts. We've got Congress. How do we feel about those institutions for policy making? And then let's turn to the institutions that might be able to deal with some of the other problems that you teed up on. How do we feel about copyright institutions?

[01:02:00.46] CASEY FIESLER: You don't think that tech companies should just be regulating themselves?

[01:02:03.76] [LAUGHTER]

[01:02:05.32] BLAKE REID: We always have that option.

[01:02:08.92] MICHAEL FRICKLAS: In fact, the Google-- not to pick out on one company, but Google in their privacy policy, of all places, has said that they will train their AI on anything that's publicly available, whether it violates terms of use, whether it violates copyright or anything else. So it is an example of the big tech companies just trying to establish the rules.

[01:02:31.90] BLAKE REID: All right. So we could add self-regulation in the mix. Let me zero in, though, let's talk about courts. So I think a lot of what we were talking about this morning were hard questions about infringement, about Fair Use that are ultimately going to be decided, at least in the first instance, by courts.

[01:02:49.24] How do we feel about courts handling these big questions about AI? Are they a good venue, a bad venue?

[01:02:58.41] SARAH JEONG: I mean, there's problems with every single venue. It's not-- I mean, courts trust in the courts, probably, never has been lower. Courts also not great on technology issues, necessarily.

[01:03:15.67] But on the other hand, one of the weird things that happens when courts meet bizarre technologies and bizarre setups is that sometimes, they'll just put the brakes on things. And is that necessarily a bad thing? I don't know. Was it a bad thing that the

Supreme Court put the brakes on torrenting? I actually don't know. I'm not sure. I don't think it was good per se, but maybe it wasn't bad.

[01:03:44.50] So I'm not sure about the courts. I'm like ah. Congress--

[01:03:49.33] [LAUGHTER]

[01:03:51.22] I think that DMC is a pretty decent law, honestly, for what it does, the fact that it was created well in advance of the Web 2.0 that it came to govern. But on the other hand, the Communications Decency Act, which happened before that, terrible law. Awful. Just not-- no good, whatsoever. And so that's a mixed record. And also, that was a Congress that was much more functional and probably, had less of the brain drain problem that we face.

[01:04:21.41] The admin state and the executive branch, pretty good under the Obama administration, but there's been a period in which it's not been quite as good as it was before. It was-- it's been gutted. So I don't know. It's-- there are no good solutions here. And then of course, then there are the tech companies. And-- no.

[01:04:46.94] Right.

[01:04:47.60] MICHAEL FRICKLAS: Yeah, so-- I'm sorry. I just thought I'd-- I just thought I'd add a little bit to your list. So there was a recent summary judgment decision relating to copyright, actually, in the legal industry about copyrightability of West law notes, if I remember the case right. And the court said, I really have no idea how to apply the fourth Fair Use factor and the one about the effect of the use. And so I'm just going to kick it to a jury.

[01:05:19.44] So the juries on your list.

[01:05:21.42] [LAUGHTER]

[01:05:23.43] In the case of situations where Congress hasn't acted, we've seen the states act. In California, there's legislation pending that would provide compensation for journalists from large tech platforms as a competition remedy. So that's been picked up by the state because it's a little stuck in congressional level.

[01:05:42.99] We've seen state AGs get into the act. We've seen international bodies, certainly, the EU and other countries get into the body, as well as a variety of international treaty organizations. All of these organizations have a variety of influences and pressures on them. Some of them are very parochial. Very few of them, actually, are trying to consider the system as a whole free of those kinds of influences.

[01:06:07.85] So I think it's a really-- and everybody wants to jump in on the act right now because AI is a hot topic. And so it gets a lot of attention. And people are really not that knowledgeable about it. They're not that expert on it. We can't really expect courts who might hear criminal trial one day on some thing having absolutely nothing to

do with technology, in the next day, they become an expert on artificial intelligence and the way it works.

[01:06:37.15] ELLEN GOODMAN: So--

[01:06:37.50] AMANDA LEVENDOWSKI: [INAUDIBLE].

[01:06:37.71] ELLEN GOODMAN: --when-- oh, sorry, Amanda, you go.

[01:06:40.53] AMANDA LEVENDOWSKI: I would love to hear from you.

[01:06:42.00] [LAUGHTER]

[01:06:43.86] ELLEN GOODMAN: I'm just going to say, when Kathie Lee was talking in the last panel, so much appreciated what she had to say about how we have choices. And-- but while she was talking, and she was talking about ways in which technologists could move in one direction or another, and they need incentives. OK, so how do you get incentives? You get big judgments that focus the mind.

[01:07:05.92] And so I like the-- and that we've had this desert because of 230 and we haven't really developed those. So I like that idea. But I think courts are going to need help because if you think about-- if they have to decide- I'm talking about outside of the copyright context, if they have to decide in a product liability case or in a defamation case, because we're going to have defamation cases, well, what's the-- what's-- is there a duty of care? What is reasonable?

[01:07:33.81] And actually, God forbid if it's an intent-based standard. AIs don't-- they t-- whose intent? And so-- but on those first two, what's-- or especially on the reasonableness, this is where I think researchers and technical standard setting like NIST, and this kind of what I think of as the plumbing here, as that develops, that will assist courts in actually coming-- having inputs into what best practices are.

[01:08:05.77] BLAKE REID: Amanda.

[01:08:06.69] AMANDA LEVENDOWSKI: I think before we get to any of these institutions, there's another very important institution that we haven't talked about that is positioned in all of the ones we have, which are lawyers. Hopefully, all of these-- well, certainly all the major developers and big tech companies have lawyers inside and outside that are advising them on copyright and a whole host of other AI related issues.

[01:08:28.78] But I think that's really an opportunity for two reasons. The first is, if lawyers are taking a client-centered approach, their goal is to present their clients with all possible consequences of their actions. And to pair that with their ethical obligations and empowerments under Rule 2.1, Advisory, they can advise on moral social, political, and economic issues, not just legal ones. And a lot of the harms we're talking about are straight at that intersection.

Sociotechnical harms are moral harms. They are social harms. And they are political harms.

[01:08:59.92] So if you have a client-centered lawyer advising some of these clients, using their full toolbox, they may be able to come forward and say, listen, you asked me to look at Fair Use, but I'm really seeing a bigger issue, which is, should this technology exist at all?

[01:09:13.72] Now that might not be something that the client wants to pay $500 an hour to hear, but that's something that the partners can write off at the end of the day. That doesn't mean that lawyers get to shirk their ethical obligations to have those difficult conversations with clients.

[01:09:26.77] And I think that we've seen-- I don't know what those memos look like behind the scenes. I would like to think that there is a number of really competent lawyers, really competent firms, or solo practitioner offices giving this kind of really excellent advice to their clients, but I think that this also is a fundamental rule 1.1 issue competency, which is, a lot of lawyers are not really competent to advise on these complex sociotechnical issues.

[01:09:53.08] And that is a bigger challenge for, I think, all of the different institutions we've talked about because they are built largely of lawyers. If there's a competency issue, we're actually not having a robust thoughtful conversation about how to address harms. We are running straight past that and going into the law that we're most familiar with, and bypassing a lot of the issues that really deserve attention.

[01:10:16.97] So I would add lawyers to the frame. And I would say, particularly, that client-centered lawyering that really embraces that 2.1 ethical obligation and pairs it with taking competency seriously can be a really powerful tool in shaping better technologies.

[01:10:32.47] BLAKE REID: Unfortunately, I think the reality is that lawyers are going to be using AI to do their research and--

[01:10:38.38] [LAUGHTER]

[01:10:39.19] --and finding hallucinated cases. All right. I want to take the prompt on ethics and turn to one last topic before we open it up to Q&A from the audience. And that's ethics, not in the legal ethics sense, but in the broader ethics sense.

[01:10:52.45] So Casey, this is the focus of your work. And a question, is law even the right vehicle to broach this array of problems that we've thrown out on the floor today? Or should we be thinking about ethics?

[01:11:06.33] CASEY FIESLER: Well, I actually-- I was going to add one more thing to the list, which is the court of public opinion, which I do actually think can have a legitimate impact here. You can be concerned about Fair Use, and you might get sued. But also, with enough people

who are willing to very loudly critique a technology, a company might come up and say, our model is ethically sourced. And then suddenly, you have a market-- it's organic, a market differentiator.

[01:11:39.66] BLAKE REID: Is it gluten-free?

[01:11:42.03] CASEY FIESLER: Because I think that a lot of the concerns that are being framed as copyright, if you ask the average person on the street, they might actually frame it as an ethical issue. A lot of my work has been around understandings of Fair Use or like misunderstandings of Fair Use. And the average person doesn't know what the four factors are or what makes something Fair Use or not. But they have this intuition that maybe this doesn't seem fair-- fair in an ethical sense. And I've been using the example of fan fiction writers.

[01:12:18.99] So I'm on the legal committee for the Organization for Transformative Works, which I mentioned because I'm not about to speak for that organization or make assumptions about the prevailing attitude of its members. But a lot of fan fiction writers who exist because of Fair Use, that is-- it is the underlying-- this fan works are Fair Use is why it exists.

[01:12:41.47] A lot of people are very upset that their fan fiction has been used to train these AI models, very upset. And this might seem counterintuitive but it actually makes sense because the norms of that community around copyright are so deeply entrenched in attribution and credit because it's a gift economy, no one makes any money. And so the worst things that you could do as a fan fiction writer would be to steal someone else's work without credit or make money from your fan fiction.

[01:13:15.52] And so the idea that all of this free labor that was given with love to other people was used to create this technology that A, they might not be directly benefiting from because I think there's a lot of problems with like, are the benefits of this technology being seen by the people who are being harmed? But also, it's being commercialized. They're not being credited.

[01:13:43.07] And people might say like, oh, well, I think-- I don't know that they would argue that it's not Fair Use, but it doesn't feel right. And I think a lot of people feel that way. And there's definitely-- there's a good number of people in the AI ethics world who feel that because this technology was built off the backs of content without people's consent, and without attribution, without compensation, that it is an inherently unethical technology. I don't know if that's the prevailing opinion, but it's a very legitimate way of thinking about it.

[01:14:20.79] And that's not something that is a legal issue. That's not going to be solved with the court decision in Getty Images. Even if a judge says this is Fair Use, that's not going to change people's opinions about whether it was ethical.

[01:14:36.12] BLAKE REID: All right. I have so many questions I want to go deeper on. Lightning round reactions on the ethics question. And then the clock is telling me that we must turn to the audience for questions. Any final thoughts on this question about ethics?

[01:14:50.40] ELLEN GOODMAN: Copyright doesn't care about your feelings.

[01:14:51.99] [LAUGHTER]

[01:14:53.80] CASEY FIESLER: Yeah.

[01:14:54.24] ELLEN GOODMAN: That's what I heard in the last panel.

[01:14:56.21] MICHAEL FRICKLAS: But I think--

[01:14:56.77] SARAH JEONG: Except in Europe.

[01:14:57.09] MICHAEL FRICKLAS: I do--

[01:14:57.70] ELLEN GOODMAN: Except in Europe. In the United States. I mean, I just have-- I care about ethics as much as the next person. But I really don't like the framing of these trustworthy AI issues in terms of ethics because ethics doesn't really have bite. And I would rather have something else have bite.

[01:15:17.70] I do like the idea-- do you think-- well, that was lightning. So I'll stop.

[01:15:23.01] MICHAEL FRICKLAS: I do believe that ethics matters, though. Because I do believe-- I mean, copyright-- the copyright wars of 15 years ago, I think, are a really good example, which is, you can have really good laws. And it could be absolutely clear that something's copyright infringement. And if 100 million people decide that it's not, there's really nothing you can do about it.

[01:15:38.29] And so I think it's really important that these laws are framed on a base of what people think is the right thing.

[01:15:46.79] BLAKE REID: Sarah.

[01:15:47.26] SARAH JEONG: I think these ethical considerations have led to a folk copyright law that is bubbling up. And usually, you just go, oh, that's not how it works. In this case, the reason why we didn't go into the AI authorship question is because I think it just demonstrates how much of this is up for grabs.

[01:16:07.12] Some-- a lot of the questions around generative AI and copyright, copyright literally doesn't make sense anymore. And if that's the case, then that means this is up for grabs because a court is going to come in and either make sense of it by pretending like this was the law all along but actually, making up something entirely new. Or we get a new statute, or we get some kind of a mega-disaster court ruling where they try to make sense of the thing, but actually don't. And we deal with the fallout from that.

[01:16:38.90] So yeah, I guess, what I'm saying is that the folk copyright that's bubbling up is actually relevant now because we're on the precipice of big change being possible at all.

[01:16:54.24] BLAKE REID: Amanda, final words.

[01:16:56.25] AMANDA LEVENDOWSKI: I'm going to let it go to our beautiful audience.

[01:16:59.31] BLAKE REID: Beautiful. And speaking of our beautiful audience--

[01:17:01.74] AMANDA LEVENDOWSKI: I can't see them, but I know they're beautiful in my heart.

[01:17:03.93] ELLEN GOODMAN: They are.

[01:17:04.61] BLAKE REID: They are. Just plug it into DALL-E, beautiful audience. I'm sure there's very accurate picture. Make sure to add Boulder. That's a little different.

[01:17:14.16] [LAUGHTER]

[01:17:15.24] As always, the first question will go to a student. And I'm scanning the room for students. Fantastic. Right back here.

[01:17:34.21] AUDIENCE: Hello. I think we heard a lot about the inadequacies of the system that we have currently and politicians and/or courts who don't fully understand AI. But I'm not a legal scholar, but I do know that the decisions will be made soon and now, whether it is by Google and/or our representatives currently.

[01:17:55.99] So what piece of legislation would you like to see enacted, if it could just be one right now, that you think would most benefit us in the future with generative AI? Obviously, we have limited information, and we don't know what things will look like, but decisions do need to be made.

[01:18:12.51] BLAKE REID: I'm so glad you asked because I wanted to ask you about political economy. We get-- we only get the political economy to do one thing, what's it going to be? Sarah, can we start with you?

[01:18:22.98] SARAH JEONG: I would release funds to probably, the FTC, depends. There would be one of the agencies to establish, basically, a commission. Similar to how the FDA has an inside line on the pharma companies, how the FTC has a close relationship with the telecoms. One have an agency or subagency that is plugged into the AI companies in a way such that they get transparency into what is happening to the data in the pipeline. Where the data is coming from? What is happening with the data? And start-- and empower them to do rulemaking, and also, sue companies based-- for judgments that are larger than the cap that the FTC currently has.

[01:19:18.13] All right. Transparency and monitoring. Ellen, let's go down the line.

[01:19:21.79] ELLEN GOODMAN: This is on legislation?

[01:19:22.96] BLAKE REID: Yeah.

[01:19:23.29] ELLEN GOODMAN: OK. So massive more resources in the federal government. I'm not sure if FTC is one possibility, but across the federal government. And also, additional subpoena power, additional technical competence so that there can be regulation.

[01:19:41.18] BLAKE REID: Casey.

[01:19:42.86] CASEY FIESLER: I like the AI Bill of Rights framework. I think it has challenges, in part, because rights are kind of reactive because it has to be violated, and you can do something about it. But I do like that framing that this is a human rights issue. And so I might like to see some actual legislation come out of those ideas.

[01:20:03.18] MICHAEL FRICKLAS: There's also-- I like all of those. There's legislation pending that we created expert agency at the federal government level that would not take over all the regulation of AI, which is, I think, AI is going to be part of everything, but would provide expertise to the various regulatory agencies when they deal with issues that are framed in an AI context.

[01:20:20.57] BLAKE REID: Fantastic. Amanda.

[01:20:23.31] AMANDA LEVENDOWSKI: Everyone already picked the big pie in the sky like major stuff. But one small thing that I would welcome is codifying public benefit as a genuine fifth factor of Fair Use. And saying that as part of that factor, courts really need to grapple with the countervailing public harms.

[01:20:39.57] I think all the time about Perfect 10 versus Amazon, which is the case that established that Google image search was lawful Fair Use. And in that particular case, the copies in question were pirated images of nude women who had previously had their images behind a paywall. And the court didn't grapple with the privacy invasion, and dignity invasion, and autonomy invasion there.

[01:21:03.06] It didn't grapple with Safiya Noble's incredible work about the biases that are now ingrained in Google image search, particularly, around searches for women and people of color. And it certainly didn't deal with what is now an issue which they could not really have predicted then, that Google is becoming a site of promoting misinformation, given all of the generative AI nonsense.

[01:21:25.78] And while I don't necessarily think that a public benefit factor would overwhelm factors one or four, which, as we've talked about, the market harm is key, and we didn't really spend a lot of time on transformativeness, but factor one has always been very important. I

do think that forcing courts to grapple with those real sociotechnical harms, as well as the legal harms, would give some breadcrumbs to other legislators or agencies who are looking to create some laser-focused legislation or laser-focused enforcements to address specific issues rather than having the world be quite so broad and big.

[01:22:00.94] BLAKE REID: Terrific question. Thanks so much. Let's go to the audience. And I'll come right back about five rows from the back here.

[01:22:16.33] AUDIENCE: Well, thank you very much for the discussion today. There have been legal and economic arguments that AI platform provider model really resembles a natural monopoly, which would seem to complexify the property rights or the licensing aspects of dealing with copyright. If this natural monopoly characteristic is, in fact, true, does that change your approach to the issues?

[01:22:39.34] BLAKE REID: Yeah, so we haven't touched too much on the competition policy dimensions here. Thoughts on that?

[01:22:44.80] SARAH JEONG: Sorry, what resembles a natural? National or natural?

[01:22:49.42] AUDIENCE: Natural monopoly.

[01:22:50.35] SARAH JEONG: What is a natural monopoly?

[01:22:52.03] AUDIENCE: In the sense that--

[01:22:52.93] SARAH JEONG: No, no, no. The-- what--

[01:22:54.22] BLAKE REID: AI models.

[01:22:54.88] SARAH JEONG: AI is a natural monopoly.

[01:22:56.68] AUDIENCE: AI platform.

[01:22:58.77] SARAH JEONG: OK.

[01:22:59.55] BLAKE REID: All right, just for folks online who didn't hear that. The notion is, could the provision of AI platforms be regarded as a natural monopoly. And what consequences does that have for policy? Mike, you want to take the first stab?

[01:23:12.33] MICHAEL FRICKLAS: Sure. I think it's a little too early to tell on the question of whether it's a natural monopoly. So I probably wouldn't address the issue at the beginning as that kind of an issue. There are lots and lots of work being done on AI large language models from all kinds of places. And even though some of them are enormously compute-intensive, there are others that seem to be less so. And there's a major movement to develop a public cloud that might be funding research at universities and the like, and provide an alternative.

[01:23:41.52] I do believe that if they do become a monopoly, it has a lot of consequences for lots of these other things. Markets don't function well when a party on one side is a monopolist. And we're seeing that play through, for example, in the existing internet now with a number of different competition challenges from the federal government.

[01:24:00.49] BLAKE REID: Other thoughts on competition, Casey?

[01:24:02.05] CASEY FIESLER: Beyond just monopoly, I mean, we haven't talked a lot about the FTC. I feel like a lot of AI regulation right now would be coming in terms of unfair trade practices and deceptive practices and that sort of thing, which is going to be the FTC or like states. So we haven't talked a lot about that, but that's the structure that we have right now that would probably be dealing-- might end up dealing with some of the stuff beyond copyright.

[01:24:30.28] BLAKE REID: Ellen, any reactions?

[01:24:32.62] AMANDA LEVENDOWSKI: I mean, just an observation that there's open source AI, open source models. And that suggests that they're not a monopoly.

[01:24:44.60] And in fact, the whole-- the voluntary-- White House voluntary commitments on frontier models, if you followed that and the representation that the companies that might be thought of to be-- you have to be at least an oligopoly, OpenAI, Anthropic, Microsoft, Google. They have made agreement. They've agreed, in part, to be licensed, as if they were monopolies. And then there's a critique that that will surely make them monopolies because only they will be able to comply with these licenses.

[01:25:18.33] So I think this question of how concentrated are the large language model purveyors and how we treat that is really important.

[01:25:29.01] MICHAEL FRICKLAS: But if you go back to 1996, and you think about all the number of search engines, Yahoo, names like people in this room don't even remember, like AltaVista, and even AOL--

[01:25:38.47] ELLEN GOODMAN: Jeeves.

[01:25:39.20] MICHAEL FRICKLAS: --it's like-- Ask Jeeves, there you go. So I was like--

[01:25:41.89] ELLEN GOODMAN: OGAI.

[01:25:42.43] MICHAEL FRICKLAS: I think it--

[01:25:43.07] [LAUGHTER]

[01:25:43.79] I think it's really hard to predict the way that these markets are going to evolve or in the really early stages of having these APIs available. And it's not just the search engine functionality, but it's

all the things that got added onto it that began to create these big economic forces.

[01:25:58.52] BLAKE REID: Sarah, have something to add?

[01:25:59.23] SARAH JEONG: I have no idea whether it is a natural monopoly. I think that you gave some really good points there. The other part being that I would contest natural for reasons what you stated, the monopoly over the data would exist in part because of Copyright Law, that we live in a post-Viacom, YouTube, post-HathiTrust world. And so there are specific actors that have data and access to data because of how things shook out in court and how we decided to approach Fair Use.

[01:26:34.32] So is it a natural monopoly? I don't know. I don't know if it's natural. I don't know if it's a monopoly. Certainly, if it turns out to be, it would make everything easier because then it's easier to regulate. It's easier to regulate one or two entities rather than like 20, with a new one coming out of the bushes every five seconds.

[01:26:53.46] So yeah, I don't know. It is an interesting question.

[01:26:59.45] BLAKE REID: Amanda, any reactions?

[01:27:02.05] AMANDA LEVENDOWSKI: In the interest of making sure we end on time, no, I do not have any reactions.

[01:27:05.41] [LAUGHTER]

[01:27:06.25] BLAKE REID: Terrific. All right. I think that gives us time for one last question, bearing in mind that we're short on time. And I want to go to the person who's right next to the microphone holder in the back.

[01:27:16.48] AUDIENCE: Very convenient. Hi. My name is Morgan English. And I am really happy that I got to participate in this panel, or not participate but be in the audience for this panel.

[01:27:26.77] BLAKE REID: You're on this panel now. Welcome.

[01:27:29.32] AUDIENCE: Personally, I work with the Concept Art Association. They represent concept artists. And they were some of the first people to sue these platforms. So it's really great to hear from having the artists perspective represented.

[01:27:42.43] I had a question about another way some of these platforms are being regulated, which is through self-censorship, and whether-- especially for this type of technology. And maybe Sarah, you could talk about this since you know about the history of a lot of technological advancements, where usually, that expands the marketplace, where I think, in this instance, it might be a little bit more unprecedented where we're shrinking the marketplace.

[01:28:09.89] And if that's something that can really be dealt with through technological regulations or through legislative regulations, is this going to get solved or are we going to really see a minimization of creative artists in the industry?

[01:28:31.64] SARAH JEONG: One of the things that I tell journalists when we start talking about AI is going to come take our jobs is that, AI already took our jobs. This is-- it already happened. You just-- we weren't calling it AI.

[01:28:48.13] In a broad sense, search engines are AI, always have been. Astronaut meme. It's-- we started writing headlines for search engines a long, long time ago. Revenue became diverted to-- long time ago, newsrooms began shrinking. We already lost our jobs to AI. It already happened.

[01:29:10.67] In the '70s, if you look at old stats of The New Yorker, you'll see that they're much longer. All of the articles are not well-edited, honestly. They're like about 150% as long as they should be. And it's because they had too much ad inventory. They needed the text in order to be able to run all of the ads in The New Yorker. That is not the market anymore.

[01:29:37.47] I think that we are all in for a hard time. Anyone who wants to do anything that promotes culture, and human flourishing, and the progress of the useful arts and sciences is in for a hard time. We've been in a hard time. And the money is going to dry up. And things are simply not going to be diverted to us.

[01:30:00.69] And the thing that I'm going to point out is that Silicon-- quickly because I know we're running out of time. Silicon Valley Web 2.0 came about in the wake of the financial crisis and in the wake of the Microsoft antitrust action. So what you have is you have to market interventions. The intervention being the antitrust action, and also, monetary easing by the Federal Reserve.

[01:30:28.29] Interest rates were very low for a very long time. And that period coincides with the Silicon Valley boom, the second boom.

[01:30:36.54] In essence, we subsidized the tech industry. That is money that could have been going to something else, instead, we chose to give it to the tech industry. And that is the world we live in today. So I'm just going to leave it there.

[01:30:52.47] BLAKE REID: All right. Looking at the time, I'm going to defer further conversation to the reception because we're a couple of minutes.

[01:30:59.46] All right, Mike, you're twisting my arm here. Really quick thought.

[01:31:02.78] MICHAEL FRICKLAS: Really quick. I have two quotes for you. And I just thought you'd like these. One is from Sam Altman, one's

from Brad Smith. Brad Smith's the president of Microsoft. Sam Altman is obviously the CEO of OpenAI, has been running around on AI issues.

[01:31:15.20] Brad Smith said in front of Congress, "It's critical for authors to retain control of their rights under Copyright Law and earn a healthy return on their creations in talking about AI." Sam Altman said, and I'll cut down the longer quote but he says, "We believe it's really going to deliver that the content owners likenesses, people totally deserve control over how it is used and to benefit from it." So we have two of the biggest players in AI saying that we're going to work this out.

[01:31:40.13] BLAKE REID: On that note, we're going to work this out. Before we go, it comes to me to do a little housekeeping. October 17, we hope you will come join us again for the Startup Variety Show. If you liked today's programming, on October 24, we're going to do the next event in our AI ethics series. And please save the date for our annual flagship conference, February 4 and 5.

[01:32:04.83] We'll be moving next door for a reception. The earlier admonitions about speaking to students still stand. With that, please join me in thanking all of our speakers today for a wonderful day.

[01:32:15.93] [APPLAUSE]