Copyright + Technology Conference
Copyright Society September 30, 2024

PANEL 2
THE PAPER TRAIL: ATTRIBUTION, AI, AND COPYRIGHT

with DANIELLE BULGER,[1] NICOLAS GONZALEZ THOMAS,[2] JENNI KATZMAN,[3] WILL KRETH,[4] and LEONARD ROSENTHOL[5]

One of the technologies being developed to address copyright issues (among others) in the AI context is attribution – digital "paper trails" that establish the human or non-human authorship of content. Many of the technologies for establishing attribution, such as digital watermarking and online content identifiers, have existed since the 1990s and have been used in specific applications for infringement detection and rights administration. But the need to establish robust attribution trails takes on new levels of meaning in the age of AI; and the emphasis today is on standardization and interoperability of the myriad technologies, chiefly through the Coalition for Content Provenance and Authenticity (C2PA). These were discussed on this panel.

Danielle Bulger: My name is Danielle Bulger. I'm with ArentFox Schiff LLP, where my legal practice focuses on copyright issues related to media, technology, and retail clients primarily. I am based in our Washington, DC office, but I have to admit, it's good to be back in New York, where I lived for a number of years working as a journalist. I was an associate producer and then multimedia journalist. I will say that I think those experiences laid the groundwork for my current practice as outside counsel for approximately the last decade, litigating and counseling on issues related to content. That is why this particular panel is very exciting for me, and I am happy we have such an esteemed panel here today.

---

[1] Danielle Bulger is a Partner at ArentFox Schiff, advising on copyright, anti-piracy, trademark, trade dress, and advertising law. *See Danielle W. Bulger*, ARENTFOX SCHIFF LLP, https://www.afslaw.com/attorneys/danielle-bulger (last visited Jan.. 3, 2024).

[2] Nicholas Gonzalez Thomas is the co-founder and CTO of Musical AI. *See Nicholas Gonzalez Thomas*, THE COPYRIGHT SOCIETY, https://copyrightsociety.org/bio/nico-gonzalez-thomas/ (last visited Dec. 17, 2024).

[3] Jenni Katzman is a leading policy expert on intellectual property and digital safety who currently serves as a Senior Director at Microsoft. *See Jenni Katzman*, THE COPYRIGHT SOCIETY, https://copyrightsociety.org/bio/jenni-katzman/ (last visited Dec. 17, 2024).

[4] Will Kreth is the founder of HAND and is the current working group co-chair at Digital Data Exchange and co-chair of the standards register working group at Metaverse Standards Forum. *See Will Kreth*, THE COPYRIGHT SOCIETY, https://copyrightsociety.org/bio/will-kreth/ (last visited Dec. 17, 2024).

[5] Leonard Rosenthol serves as Adobe System's PDF Architect having been involved with this technology for almost 30 years. *See Leonard Rosenthol*, THE COPYRIGHT SOCIETY, https://copyrightsociety.org/bio/leonard-rosenthal/ (last visited Dec. 17, 2024).

Let's just start off with everyone introducing themselves. Nicolas, if you'd like to go first.

Nicolas Gonzalez Thomas: Hello, my name is Nicolas Gonzalez Thomas. I'm the co-founder and CTO of Musical AI, where we do rights management for the music industry.[6] My background is in computer science.[7] I've been in generative AI since it was called computational creativity back in the early 2000s. I was a researcher in music AI, and then I became an entrepreneur in this space, building AI companies' models, all that kind of thing,[8] until we started to realize the bigger problem was access to training data and seeing that it was either one side we have to just train and then deal with the consequences of that later, or try to get access to data. And then we realized, let's just build a platform for rights management and tackle the challenge of attributing the data that was used for training. So, that's what we're doing now.

Jenni Katzman: Hi everyone, and thanks to The Copyright Society for hosting the event today. My name is Jenni Katzman, I'm with Microsoft.[9] I'm based in DC. I'm a senior director in our US Government Affairs shop. I handle policy for intellectual property, digital safety, and synthetic media.

Will Kreth: Hi everybody, I'm Will Kreth, founder and CEO of HAND Human & Digital,[10] a DOI Digital Object Identifier foundation registration agency and business intelligence platform under the ISO standard.[11] We work with companies like Sony Pictures Entertainment, American Film Institute, and Respeecher, who's doing voice synthesis around unique identification of instances of identity of talent, of legal and natural people, connected digital replicas and fictional characters.[12]

Leonard Rosenthol: Hi everyone, I'm Leonard Rosenthol. I am Adobe's senior principal architect for two of their key technologies.[13] One is for PDF, which I'm sure you're all familiar with. And the other one, which is the one we're going to talk about today, is content authenticity, and that has led me to chair the Coalition for Content Provenance and Authenticity.[14] Shorthand is the C2PA now since 2019 and its founding.

---

[6] *Manifesto: Attribution in Generative AI Music Models*, MUSICAL AI, https://www.wearemusical.ai/Manifesto (last visited Dec. 17, 2024).

[7] BSc, Computer Science, Universidad CAECE; MSc Interactive Arts and Technology, Simon Fraser University.

[8] *Nicholas Gonzalez Thomas, supra* note 2.

[9] *Jenni Katzman, supra* note 3.

[10] *Will Kreth, supra* note 4; For more information on HAND, *see About HAND*, HUMAN DIGITAL, https://handidentity.com/about/ (last visited Dec. 17, 2024).

[11] To access Kreth's ISO standard, *see* ISO 26324:2022, ISO, https://www.iso.org/standard/81599.html (last visited Dec. 17, 2024). .

[12] *Our Story*, RESPEECHER, https://www.respeecher.com/about-us (last visited Dec. 17, 2024) [hereinafter "Respeecher"].

[13] *Leonard Rosenthol, supra* note 5.

[14] *Id.*; For more information on the C2PA, *see About*, COALITION FOR CONTENT PROVENANCE AND AUTHENTICITY, https://c2pa.org/about/ (last visited Dec. 17, 2024).

Danielle Bulger: Thank you, everyone, for those great introductions. As you can see, we have a panel full of policymakers and people who are on the front lines of creating the technologies that we're discussing throughout this panel. So, let's dive right into it and give a little bit of background. As we all know, authors and their works have always been closely connected. The concept of moral rights recognizes the right of attribution and the right of integrity of an author, or the right to be credited, and the right to prevent others from distorting a work. Over the years, we have seen the US acknowledge the existence of moral rights, more or less.

First, we saw it when the US joined the Berne Convention, and recognized moral rights through a patchwork of laws, including state laws related to privacy and publicity and defamation and unfair competition and the like.[15] We also saw it through the passage of the Visual Artist Rights Act,[16] and then later with the passage of Section 1202 of the Digital Millennium Copyright Act, which prohibits parties from providing false copyright management information, including a title's work, author, copyright owner, terms and conditions of use, as well as removing or altering CMI.[17]

The reality is that Section 1202, as you can see, was passed more than 25 years ago. ince then, how we access and share information, particularly with the growth of the internet, has changed, making works of authorship increasingly accessible and susceptible to mishandling and manipulation.

Before we get into the many current laws and developments, including some that happened just this weekend, I think it would be helpful to provide our audience some context concerning attribution and technology.

> *The context for this panel is that AI is the latest in a series of technologies that can be used to manipulate media, including for purposes of deception, misrepresentation (including representing a work as that of an existing artist), and disinformation. These technologies make it all the more valuable to identify the sources of content – including whether those sources are human or synthetic – and the ways in which it's manipulated along its journey through the world. The technologies and standards discussed on this panel are all attempts to address this issue. –Eds.*

---

[15] As well as, for example, Section 43(a) of the Lanham Act (false designations of origin) and § 106 of the Copyright Act (exclusive rights of authors to create derivative works). Since the U.S.'s implementation of the Berne Convention in 1988, Congress joined additional international treaties addressing moral rights and concerning rights management information, such as the WIPO Copyright Treaty (WCT) and the WIPO Performances and Phonograms Treaty (WPPT), obligations for which were integrated into the DMCA. *See generally Authors, Attribution, and Integrity: Examining Moral Rights in the United States*, REPORT OF THE REGISTERS OF COPYRIGHTS (Apr. 2019).

[16] Visual Artists Rights Act of 1990 (VARA), which created 17 U.S.C. § 106A (2015).

[17] 17 U.S.C. § 1202 (2024).

Jenni, I'd like to start with you. At Microsoft, you are on the front lines of addressing risks related to **synthetic content**.[18] Can help us understand how the growth of technology has led to the manipulation of works of authorship with respect to technology?

Jenni Katzman: I think it is sort of important to establish at the outset that before we sort of get to the harms, the reason why we're talking about the harms is because there's a lot of opportunity, and I think that's just really important to say. There's going to be a lot of new expression, a lot of great opportunities, new developments because of AI, we wouldn't be talking about the harms unless that was the case. So, I just want to just lay that out there.

But it is important, obviously, to talk about the harms, to address the challenges that is absolutely necessary. And Media manipulation is not a new concept. It's been around for ages. This has been happening since the 18th century. Totalitarian rulers have used media manipulation. Hitler and Stalin used it for propaganda purposes. It's been used, certainly within the age of technology, and has increased the development of it. There have been products that have been created at companies, including ones that are on this stage.

Leonard Rosenthol: Okay, you can say Photoshop.[19]

Jenni Katzman: Yes, Photoshop. But I would venture to say that people have used Photoshop in ways that have led to great uses, even though bad actors have used it in other ways. And with generative AI, it is different. It is more accessible, easier to use, and in ways, more easier to manipulate for bad actors. And on the other end, it is more difficult to distinguish from real images. In fact, we know in certain circumstances that for adults looking at it, trying to distinguish between real images and AI-generated images, they get it right 60% of the time.[20] We have a tool on our site called Real or Not.[21] You can look at it and see whether or not you can distinguish it. And I think that 60% number is probably an accurate figure. And so, we have to figure out a way to resolve that. But it creates a host of challenges to address different types of **harms**. We have a white paper out[22] on those sorts of harms from things like non-consensual imagery, non-consensual intimate imagery, fraud, and the misinformation, disinformation, other issues

---

[18] Synthetic content is content created by generative AI that is used to train generative AI, as opposed to human-created content used to train generative AI.

[19] *Photoshop*, ADOBE, https://www.adobe.com/products/photoshop.html (last visited Dec. 17, 2024).

[20] Zeyu Lu et al., *Seeing is not always believing: Benchmarking Human and Model Perception of AI-Generated Images, in* 37th CONFERENCE ON NEURAL INFORMATION PROCESSING SYSTEMS (NEURIPS 2023) TRACK ON DATASETS AND BENCHMARKS (2023).

[21] To play this game, *see Real or Not*, https://www.realornotquiz.com/ (last visited Dec. 17, 2024).

[22] *See Protecting the Public from Abusive AI-Generated Content*, MICROSOFT, https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/documents/presentations/CSR/Protecting-Public-Abusive-AI-Generated-Content.pdf (last visited Dec. 17, 2024).

across the ecosystem, including the one we're addressing today, a question of trust and attribution.

Danielle Bulger: Thanks, Jenni. Those are all great points and a reminder that this concept of fake news that we hear so much about did not originate on Twitter or on YouTube. It also actually reminds me of an exhibit on fake news a few years ago at London's Tate Modern art gallery .[23] To your point, Jenni, it did not just start within the last few years. Take for example, these images of Stalin. In this first image, we see Joseph Stalin and four of his colleagues. Twenty-three years later, there are only two. And then finally, Stalin is pictured in the image by himself, showing how those once in his inner-circle disappeared from official images, and how early images even foreshadow this idea of fake news. Leonard, maybe you can tell us a little bit about how you are seeing AI disrupt **content provenance**?

Leonard Rosenthol: Let me start by just setting the stage with the use of that term, content provenance. So, we believe very strongly, and it's why it's core in the name of C2PA--Coalition for Content Provenance and Authenticity--that everything is rooted in provenance. So, just like historically, if you went to an auction at Sotheby's, or the equivalent, you get that nice piece of paper that would tell you the history of the object that you had purchased. We believe that equivalent needs to also exist for all of your digital content. And so, that is the who, the what, the where, the when, the how, and the why of that asset, and not just its creation, but through its entire lifecycle. And so, that is provenance. And everything has to be rooted in provenance: So, whether it's attribution, it's rights, it's all that goes into that asset from the moment that it shows up in the creator's head through the publishing process and everything.

We actually started C2PA before ChatGPT became popular, before the whole idea of generative AI. We'd already been working in this area of provenance and establishing that aspect. So, to us, the use of generative AI or AI in any aspect is just another piece of the provenance chain. So, whether or not that asset is created with generative AI, it's edited at some point with generative AI or vice versa, where is the human in the loop, if you will? I capture an image on my camera. I then use Photoshop and remove some stray hairs or even remove a person, such as we just saw, using AI technology. That's the norm, and it's going to become more and more. And so, that's not a bad thing. We just need to make sure that the information about exactly those pieces of the process are recorded with that asset so that downstream, when someone is viewing that image, that video, that document. Think about this; in all forms of media, you can evaluate exactly where that generative AI shows up, because it's not a yes or a no flag.

My favorite example of this is if I use something like Microsoft Copilot[24] to modify one paragraph of a 300-page brief that I'm working on, is that an AI document now all of a sudden because I changed one paragraph to write it better?

---

[23] Wolfgang Tillmans, *2017 Exhibit from February 15 - June 11*, TATE MODERN EXHIBITION, https://www.tate.org.uk/whats-on/tate-modern/wolfgang-tillmans-2017.

[24] *See* COPILOT, https://copilot.microsoft.com/ (last visited Dec. 17, 2024).

The normal answer, one would assume, is no. But yet knowing that is still extremely important. And that's why we need this deep expression of provenance, both for humans as well as for AIs.

Danielle Bulger: And to that point, when we're talking about provenance in the digital era, we are really talking about **metadata**.

Nicolas, tell us a little bit about the consequences of removing a work's metadata or altering it.

Nicolas Gonzalez Thomas: Earlier on in the development of AI, when you trained on data, there were ways in which you could keep a trace of the metadata, and you could trace when you're learning something from a data set you can then, when you generate, have some sense of where in that data set you are getting most influence from. But with the larger neural networks and the larger data sets,, that data is getting stripped, and it's getting lost. All of the IP is getting taken into the model, all of the knowledge, or even if it's images, if it's text and we are completely losing all of that metadata. And it's very easy for a company doing that, or anybody doing the training to argue that it'll hinder the development of the technology if you are required to have some tracing of the metadata. That is the challenge now, where you can very easily advocate for the need to have no tracing of the metadata, because if you do that, then you are hindering the development of the AI. But there are ways in which that can be done.

This is part of the work that my company does, where we look at the model as something that's a black box, and you ignore the training of the model, and all you look at is the input and the output. Then you can find some links from the input data and the output data, and there are paths forward. There are solutions that the rights holders whom we are talking to find acceptable. That's basically the path forward that we are advocating for. The rightsholders do want to license their content, and there are ways in which they can keep the tracing of the metadata.

Danielle Bulger: It sounds like the removal of this metadata could be to the detriment of artists and their ability to profit from their works?

Nicolas Gonzalez Thomas: It's completely destroying the ability for any way in which you can attribute to the input. So, for example, in the panel earlier, we were talking about, if you prompt for, let's say, New York versus Tokyo.[25] Now, that picture from New York had an owner, and so, if you strip the owner of that picture, you will no longer be able to attribute to specifically the owner of those pictures that that model learned on. Whereas if you are able to do the attribution, you are able to say, "These specific images that came out with the prompting of New York now receive more compensation than the one that was Tokyo, because that prompt was New York and not Tokyo."

The AI companies are advocating for a pro rata split across all of the metadata, all of the data sets, so that if you are Beyoncé, you will get paid equally to my music. And so, for me, it's in my favor, but not for Beyoncé. All the rights

---

[25] *See* Dave Davis remarks, *in My Blanket and Me: Blanket Licensing for Generative AI*, 71 J. COPYRIGHT SOC'Y (forthcoming in this present issue).

holders who have the most valuable asset will not want to license their content. And that's what we're finding, that all the rights holders and the major labels we're speaking to, they will not license their content unless the attribution is specific to certain data points and not just across the board on the dataset.

Danielle Bulger: Will, I want to turn to you, and I will put this graphic on the screen. Last year, the survey firm YouGov asked Americans how concerned they are about various potential consequences arising from artificial intelligence.[26] 85% of respondents said that they are very concerned or somewhat concerned about the spread of misleading video and deepfakes.[27] When we are thinking about copyright and generative AI, how are chain of custody rights related to the licensing of one's name, image, likeness, voice essential or foundational with respect to your work?

Will Kreth: Well, part of our partnership HAND has with C2PA is to provide that element of the chain of custody, that attribution trail, that source verification as to where an image came from.

I was part of an accelerator project at the International Broadcasting Convention just two weeks ago in Amsterdam, where it was called Design Your Weapon Against Disinformation and Deepfakes.[28] And the example we used was a video of President Joe Biden from public television, from PBS, backing out of the electoral campaign, out of the election this year, and making that announcement that when that video went out on their YouTube channel, someone took it, and then, with a voice model, replaced his voice with things he would never have said, words not suitable for work in shared company.

But that story of providing that attribution trail of like saying that the individual, in that case, we provide a unique identifier for Joe Biden as a public figure, was part of the cryptographic metadata, of the three things that C2PA talks about are those cryptographic metadata, watermarking, and fingerprinting of the original media object, in that case, it was a video file. And to protect that through its travels through the workflow and the supply chain, so that it would be detectable whether it was embedded in the file type or whether it was not present, which then would potentially tip the HAND to say, "This is not a legitimate or consent-based usage of that asset, of that media file."

So, it's one thing to provide deepfake detection tools - there are many companies, and you probably have heard about folks like Reality Defender and others that are providing deepfake detection, kind of essentially AI fighting AI.[29]

---

[26] Taylor Orth & Carl Bialik, *Majorities of Americans are concerned about the spread of AI deepfakes and propaganda*, YouGov (Sept. 12, 2023), https://today.yougov.com/technology/articles/46058-majorities-americans-are-concerned-about-spread-ai.

[27] *Id.*

[28] *2024 Accelerator Project: Design Your Weapons in the Fight Against Disinformation: Final Showcase Session*, IBC, https://show.ibc.org/accelerator-project-design-weapons-fight-against-disinformation.

[29] *Disinformation*, Reality Defender, https://www.realitydefender.com/solutions/disinformation (last visited Dec. 17, 2024).

And I know those folks, and they're good people. But there's also the story around provenance and that story of, do you know where it came from? Can you provide attestations of the authenticity of the work and of the individuals and the components inside the file? And it's a story where there's a lot of work to be done, but we're making steady progress on that.

And I think that's part of the story, that there's not just throwing our hands up in the air and saying nothing can be done and looking at the data with regards to trust in the society, how to engender trust, and to foster that, and to kind of make resiliency and harden that story from source to play out, from origination to distribution endpoint as part of that work we're doing.

Danielle Bulger: When you were just talking, it reminded me of the news story a few weeks ago with James Earl Jones, him assigning or licensing his AI rights in his voice.

Will Kreth: And the estate, now he's passed away, rest in peace James, is the custodian, if you will, of the rights for licensing with Lucasfilm and future Darth Vader instantiations.[30] And the folks at Respeecher, who we work with, they are working on that voice modeling.[31]

So, we have a video of our project. We did two projects there.

> **Video:** A bill aimed at protecting performers. The actors' union wants "explicit consent" to use the digital likeness of performance. Fans rallied to drown out deep fakes. Deep fakes of notable people are becoming indistinguishable from reality. We are alarmed by the circulation of false images, to be more exact. My name is Evan Shafraner. I am an actor, a comedian, and a rapper, and I'm here to get scanned. Welcome to the Scan Truck spaceship. Come on in. Happy to be here. Imagine a world where artists can protect their identity and effortlessly track the usage of their consent-based digital replicas across platforms.
>
> To have a digital replica made of myself, which I own, is amazing. I'm a performer, and I'm an artist, and we're putting our likeness out there. Somebody is making it work for us, but it's still our image at the end of the day. I am excited because I think that a lot of people aren't aware that this is being created right now for our protections. When I found out, it immediately sparked in me, "Oh, wow. I can create my own short film or animated series using my own image." I'm excited for everybody to see this new opportunity in ways that we can work together.
>
> A world designed to support the four C's, consent, control, credit, and compensation of talent.

---

[30] Emma Roth, *James Earl Jones lets AI take over the voice of Darth Vader*, THE VERGE (Sept. 24, 2022), https://www.theverge.com/2022/9/24/23370097/darth-vader-james-earl-jones-obi-wan-kenobi-star-wars-ai-disney-lucasfilm.

[31] *See* RESPEECHER, *supra* note 12.

Danielle Bulger: This is an amazing concept, and it's really the future. You all are on the front lines of new technologies, especially when it comes to digital replicas. When we think about content provenance, what risk do deep fakes pose to your business?

Will Kreth: Well, I think there's an opportunity here to get it right. And I mentioned when I was on my panel at IBC that I want companies in the industry to realize that we can be on the right side of history, not the wrong side of history - regarding the rights story, regarding to the licensing story, regarding to the utilization of notable public figures, of which we're identifying with unique identification, a registry, and cryptographic metadata to protect it, protect those instances, protect those as new objects in the media supply chain. Or we could just let the horses run free and people create whatever they want willy-nilly. And I think a story around cooperation, collaboration across the industry to provide systems and methods to instantiate ways and means to protect those valuable rights, those name, image, likeness, and voice rights of people is the right thing to do.

Danielle Bulger: Great. Leonard, I want to turn back to you, because you have been working with PDF technologies for the last three decades, and, as the chief architect of the Content Authenticity Initiative for Adobe, you really have seen this issue grow. I'm curious why you think, if you do, content provenance technology is the best solution for identification of artificial intelligence and human involvement in the creation of an asset.

Leonard Rosenthol: Yeah, absolutely. It goes back to Nicolas mentioning AI, fighting AI. You can look at any number of other approaches over the years. The one I always go back to is the copy-protection wars in the 1980s.[32] You can't win it. One will get better and then better and better. It simply won't work. The other thing to consider is scale. Even if you had an AI detector that was 99% effective, if you consider that Facebook, for example, gets in excess of 100 million images a day, that's still a significant number of images that are going to not get caught by that detector. And I, for one, would not want the legal responsibility of missing the one that's the most important on that given day. So, it's an unwinnable war.

But provenance, on the other hand, is just facts, as it's always been. If you state what you did or why you did it upfront, and that gets, as you heard, carried along throughout the lifecycle of the asset, there isn't a question. And then you tie that, as Will mentioned, to personhood. Not just famous people, but all of us in this room, any individual, you want to tie that to your ownership. We talked in the last panel about creators, about ownership. I always think of this as a stack. You've got identity, you have authorship, who authored the content, you have ownership, as was mentioned earlier.

---

[32] At that time, commercial software vendors were experimenting with hardware-based copy protection technologies, such as holes in floppy disks and "dongles" for printer ports on PCs. These techniques were rejected as ineffective, inconvenient for users, or both.

Who owns that asset? What rights? You build rights on top of that, what are the rights that you're establishing across the board for the usage of that asset? And then has been mentioned, remuneration. Pay me for those rights if you want, either collectively, as was just talked about in the last panel, or individually. I want to buy a couple of seconds, and I want to do it right now as a machine-based transaction. How do we enable that sort of thing going forward? And all of that sits with provenance as the root whether it is human-based, AI-based, or again, most likely the combination of the two, you need all of that information. You need all of that data somewhere. Again, whether it's embedded, which we prefer, or it lives out on the cloud or on a blockchain, we don't really care necessarily. But it all has to come back to that single provenance record, which we call a **content credential**.

Danielle Bulger: To that point, because you are on the technical side, let's talk about some of the technical steps that the industry generally is taking to address the needs of rights holders over controlling artificial intelligence training usage.

Leonard Rosenthol: Technologically, this is an area that is still very, very fresh. I've actually attended, I think, three panels in the last week and a half in different cities around the US, and previous to that in Europe and other countries on exactly that topic, which is how do you assert these rights? How do you state them? It's not that we don't understand what they are, but we don't have a standard or even a set of standards today. So, that work is ongoing.

But I will give you some sort of a 10,000-foot view of directionally where the industry is going, which is that there are two different models for asserting those rights. We call them *location-based* and *unit-based* or sometimes *asset- based.* And the reason we need two is that in one case you're asserting that based on the fact that it lives on my website or somewhere at a specific location on my website, because I own the website, I own the domain technically. Therefore, I have the rights to assert AI training usage. That's location-based. And that works in many scenarios. It's great for professional publishers like the BBC or the New York Times or someone like that. But if you're a scientific publisher, for example, I'll use Elsevier, for example, they are one of my favorite examples, they're not the creator of that content. They may or may not have the rights to assert that they can tell someone else whether or not you can train on that. So, location-based doesn't work in a scenario like that, and that's where you need unit or asset-based rights that carry along with the asset.

Provenance, like we've been talking about already, one of the things you want to put in the provenance of an asset are the rights, as I mentioned previously. And so, that's where asset and unit-based come into play. So that regardless of where that scientific paper lives, it doesn't matter whether Elsevier is selling it, if I put it on my site, I put it up on a Dropbox directory somewhere, it doesn't matter. The rights, the information about that I've associated with it go along with it.

Danielle Bulger: And you touched on the C2PA earlier, but can you just give us a brief overview? The Coalition for Content Provenance and Authenticity is a new concept to some in this room. If you could give us a brief overview, especially

given how involved you are with C2PA, then we can go into the level of detail the C2PA provides.

Leonard Rosenthol: Yeah, I think we've talked about a lot at the 10,000 view for a few minutes, so I'll drop down a little bit. So, we talked about the general concept of provenance. And what that represents technically, for those of you who think about it in this regard, is that we take all of these facts, the who, what, where, when, how, and why--we call those *assertions* in our grammar, in our lingo--and you can have as many of these as you want. There are *standardized assertions*. You can have *custom assertions*. All these assertions get bundled together as part of something we call a *claim*. They're all cryptographically hashed, so we know if any of them get modified, and then that claim and all those assertions are *digitally signed.*

And as you heard, since I was involved in PDF, it should come as no surprise it's the exact same technology we use to digitally sign PDF files.[33] It's the same technology used on the web for that lock icon in your browser. So, it's a well-established digital signature technology that signs everything, makes it all one big tamper-evident package, which we call a *manifest.* That's the technical term. The end user term is *content credential.* We have this nice little logo called our CR pin for the credential.[34] And so, that's the "I" in the image, and now we've got this little CR. But the idea is that's how somebody, as they're scrolling through their social media feed, for example, would say, "Oh, wait a minute, this asset has a credential. I'm probably going to treat that one differently than one that does not." And so, that becomes very significant. Much like you look for a UL for Underwriters Lab or a nutrition label. We like to compare it to if you go shopping, you look at the nutrition labels on all the foods you buy. And that just keeps getting added to over time. So, as I mentioned, as each part of the process, creation, modification, publishing, distribution, et cetera, everything just continues to build on that content credential.

Danielle Bulger: Let's narrow this down even further and talk about practicalities. We're looking at this image. But practically, let's just think about a piece of video. If a piece of video has been edited with AI to clean up some minor blemishes, for example, will the watermark say that it has been edited with AI, or what exactly does the C2PA provide in that instance?

Leonard Rosenthol: So, the answer is that we give you a lot of flexibility. We don't mandate exactly what to say, but what we do is we describe this vocabulary where you can describe as much or as little as you like. So, you could just say, "This big hammer made with AI." But chances are you probably don't want to do that. It's not as useful as I mentioned before. What you probably want to do is you want to say, in the case of a video, "This range of frames represents a clip that I took from here. This range of frames represents this clip I took from

---

[33] *See, e.g., Overview of Digital Signatures in Adobe Acrobat Sign*, ADOBE (Sept. 1, 2024), https://helpx.adobe.com/sign/config/digital-signatures/overview.html.

[34] *Introducing Official Content Credentials Icon*, COALITION FOR CONTENT PROVENANCE AND AUTHENTICITY, https://c2pa.org/post/contentcredentials/ (last visited Dec. 17, 2024).

here, this audio I licensed from somebody else over here." Because that's usually how you build a video. Now, let's say you go and you wanted to use a little AI. One of my favorite examples these days is things like removing the "ums" and the "ahs" and the pops that you usually find in voice. Well, great, okay, you can ask an AI to do that automatically. You want that identified in that stream as exactly where those things were removed. And if it actually changed the images of the video, then identify which actual frames or maybe even which rectangles on individual frames were impacted.

And again, that's all about having as much data in that provenance as possible so that downstream somebody can then understand, "Do I want to trust this? " 'Cause at the end of the day, that's what we're trying to help people decide when it comes to, as you mentioned, deepfakes. Do I want to trust this?

Danielle Bulger: Right. That is the core question. Do I want to trust this? And can I trust this based on the information presented to me? Jenni, you have a unique position working at Microsoft. I'm curious from the AI developer and service provider position, what are companies like Microsoft doing to help consumers? Because again, we're going back to this question about trust. What's Microsoft doing? And I understand these are your opinions, so not official statements, but just keeping in mind your experience, how can service providers like Microsoft and AI developers help consumers better understand the source of digital content? Because that's essentially what provenance is.

Jenni Katzman: Yeah, we'll give a plus one to C2PA, which Microsoft co-founded.[35] I think C2PA is really essential in this space, and it is for building trust in the ecosystem. There's no silver bullet here, but that is a really critical step. And it's not just about the AI content. It's also about *authentically captured media* as well. I want to put that out there. It's not just when we talk about one type of content, it should be about all content too. But since the end of 2023, we have automatically attached those content credentials to images generated through our services, such as the DALL-E 3 model,[36] our Microsoft Designer,[37] and Microsoft Paint.[38] When you're using those services, it'll automatically be attached there. That is, again, to create trust in the ecosystem. That way, you can look in the content credentials, see when that information was created, and any other information that the person creating it is putting in there.

And then also, LinkedIn, which is part of Microsoft, has now adopted C2PA. So, anything that's carrying that content, or that technology is automatically included in LinkedIn, particularly in the feed. You could click on the content credential, see the information on it, including whether it was created AI-

---

[35] Along with Adobe, Arm, BBC, Intel, and Truepic. *C2PA Founding Press Release*, COALITION FOR CONTENT PROVENANCE AND AUTHENTICITY, https://c2pa.org/post/c2pa_initial_pr/ (last visited Dec. 17, 2024).

[36] *DALL·E 3*, OPENAI, https://openai.com/index/dall-e-3/ (last visited Dec. 17, 2024).

[37] *Frequently Asked Questions*, MICROSOFT DESIGNER, https://designer.microsoft.com/FAQ.pdf (last visited Dec. 17, 2024).

[38] *Paint*, MICROSOFT, https://www.microsoft.com/en-us/windows/tips/paint (last visited Dec. 17, 2024).

generated or partly AI-generated. And that's really important to people using LinkedIn to see it creates transparency. It's going to be expanded on LinkedIn. It's included in ads as well. So, that's, I think, incredibly important to see on platforms, on social media. That's pretty big, and I don't think we're seeing that in places like LinkedIn. We're seeing things more on the services I mentioned beforehand. So, it's a pretty big step on LinkedIn, and I'd like to see that in more places.

Danielle Bulger: It sounds like Microsoft is doing a lot, and, in recent months, there have been a number of technology companies getting involved. We've seen Google increase its involvement with the C2PA, joining the steering committee and implementing content credentials into its core services like Google Search and ads, and likely eventually, YouTube.[39] About two weeks ago, it was announced that Amazon likewise joined the steering committee, announcing that it would attach content credentials to its foundation AI models like Titan Image Generators.[40]

And these steps were followed by Meta's decision to join the same committee in early September.[41] On the screen, these are other steering committee members of the C2PA, and there are also a host of other general members and contributor members.[42]. Will, back to you as a contributing member of the C2PA with HAND Human & Digital, I think we are all curious, has this standard been widely adopted? Should it be?

Will Kreth: I say yes to the second part. It should be widely adopted. You saw those two companies on that last slide. Hard to imagine them agreeing to much of anything in the world, knowing they all have different viewpoints and agendas and different lines of business that they operate, and spaces they play in, but quite remarkable, I'd say the credit goes to Leonard and the team at C2PA, who have done an amazing job of pulling together a really diverse and broad coalition here. And I think that is something very noteworthy and should be recognized.

Our little participant member role in C2PA is one of a cast of thousands when you look at the website and the farm of logos there. But we're glad to be part of it, because we did work with a company for our proof of concept. You saw Evan Shafran and the actor in the video, that example of his replica. We worked with a C2PA implementation through a third-party company called EZDRM that

---

[39] *Google to join C2PA to help increase transparency around digital content*, COALITION FOR CONTENT PROVENANCE AND AUTHENTICITY (Feb. 8, 2024), https://c2pa.org/post/google_pr/.

[40] *Amazon Joins the C2PA Steering Committee*, COALITION FOR CONTENT PROVENANCE AND AUTHENTICITY (Sept.. 12, 2024), https://c2pa.org/post/amazon_pr/.

[41] *Meta Joins the C2PA Steering Committee*, COALITION FOR CONTENT PROVENANCE AND AUTHENTICITY (Sept. 5, 2024), https://c2pa.org/post/meta_pr/.

[42] *Membership*, COALITION FOR CONTENT PROVENANCE AND AUTHENTICITY, https://c2pa.org/membership/ (last visited Dec. 17, 2024).

provided a way to provide an attestation.[43] You saw the content credentials from Leonard for the metadata about the provenance of his replica and how it was then detectable at a distribution endpoint.

So, it was kind of a proof of concept from a paper learning standpoint of saying, "How could a workflow and a supply chain story work to provide attestation of authenticity once someone views his replica in a different place?" So, that Evan might be compensated in a story around [inaudible] [00:40:20] mentioned chain of custody earlier in an automated and a kind of programmatic way, as opposed to friction in a manual process, a lot of reconciliation and spreadsheets going back and forth and back and forth.

And that machine-readable data of metadata is so essential in any scaling story when we look at the opportunity for how one's going to provide a way to make it easy for folks to do business and all the ways to automate the different elements of the story. So, I would say our role in C2P, while minor, is just one of the companies trying to look at interesting use cases, interesting case studies, ways to use it.

Leonard Rosenthol: Danielle, can I jump in for one second? Just because I want to extend beyond some of Will's examples. And it's not just technology companies. I mentioned before, media companies. So, BBC, CBC, the Canadian Broadcasting Company, New York Times, Reuters are all publishers who've jumped on this. But the most recent one, which I think has some really interesting potential and I look forward to it, is the US government, in two examples. One is the Department of Defense (DOD). They have a branch which is responsible for publishing photographs from the DOD for things that they do. Every image that's coming out of the DOD has a content credential associated with it. And the next one, which I think is going to be very cool, is that NASA has announced that they would be using C2PA as a way to label their content as authentic, for all future photographs coming off their satellites and other things similarly. So, no more fake moon landings. So, I look forward to that.

Will Kreth: I knew that was coming.

Danielle Bulger: And the C2PA is just one solution in this broader discussion about content provenance and ensuring that the information we view online is accurate, or stated otherwise, that consumers have the tools they need to make the decision for themselves, whether the actual picture or video, whatever content they're viewing, is accurate. But Nicolas, I think you also bring a unique perspective here, particularly with your company. What role do rights management platforms like Musical AI play in this discussion with respect to provenance solutions?

Nicolas Gonzalez Thomas: I just want to give some context. Historically you would sell an asset, and it was a physical asset in music or in books, and then all of that became digitized, and then there was piracy. We needed legislation to

---

[43] Tommy Flanagan, *EZDRM, Microsoft demo takedown at IBC based on C2PA, CMCD*, RETHINK RESEARCH (Sept. 19, 2024), https://rethinkresearch.biz/articles/ezdrm-microsoft-demo-takedown-at-ibc-based-on-c2pa-cmcd/.

enable the technology that would be part of enabling a consumer to pay and a creator to get paid. That kind of worked together, legislation with technology in order to enable that. And now it's happening all over again with AI, where now you train on all content.

And you could argue, "Well, I don't know what it's generating from, it's as if I'm collecting all of the valuable assets, all of the gold in the world, but I don't know who owns it, so I won't pay anybody. And because I can't prove who owns it, you can't enforce me to do that." And so, that's where my company, Musical AI, is building a platform, or we have a platform. Right now, we're collecting assets from rights holders in order to be able to offer blanket licensing and bundling for AI companies to train. And what we do is something that all the rights holders we are talking to are very happy with the solution.

So, for them, they would be happy to license the content in this way because they are being approached by AI companies to train on their content. But the way they are being compensated is not good enough for them. And so, they are asking for gigantic lump sums of upfront payments so that I can give you my content, but the AI companies don't have that money. And so, there's this issue where the rights holders of major labels, smaller labels, and independent artists are not willing to give up their content because they won't be compensated in a fair way. And so, the AI companies don't have the funding in order to pay, let's say, a billion dollars to Sony for all of their content. And so, what we are doing is because of the situation where all of the metadata, once you put everything into a model–large language model, large music model, large video model–all of that is lost, you can't trace back the original. And through the model, there are inventions that are happening at the model level where we think that's not a sustainable solution.

The solution that we currently have, which we are advocating for the media's industry to adopt, is that you can actually trace in a very, not 100% accurate way, but it's the same way in which the industry works today. If something is played on the radio, there isn't an exact calculation on who gets paid, but it's good enough for the consumers to pay the artist.[44] It's good enough for some transactions to happen along the way. There's monitoring, there's estimations, and the same thing can happen with a model.

You could look at what came into the model, you could look at what came out of the model, and you can have a pretty good statistical calculation to assess, "These are the data points that most influence this output." And so, now it's something that you don't care what the model is doing. All you look at is what came into the model, what came out of the model. And that could be something that can be used across the board, all kinds of models, all kinds of media.

And that solution we currently have with music, and we are collecting a lot of content from rights holders who are saying, "I would license this for you. I would license it through this technology." The one thing that they do care about

[44] It is generally considered to be accurate for contemporary pop hits, if perhaps less so for obscure or "long tail" content.

is security, because there's so much uncertainty. They want to be able to say, "I want my license to end in a year, and this model has to be sunsetted. It needs to be able to get out of the market."

Danielle Bulger: It sounds like there's an aspect of detection as well. So, the AI can be used to detect whether there was use of AI in creating the output.

Nicolas Gonzalez Thomas: There's a lot of angles there. So, for example, if something is put on a distribution platform, was that generated through AI? That's what we could call detection. Was it human or AI-generated? Or, for example, is something that came out of a model, which model? If something's on Spotify, and I created it, which model was used? So, that's a different kind of detection. So, was it a model that was able to do that, or was it not legally able to do that? But which model was used?

But also, you don't know what the training data of this model is, and you can detect also with a certain degree of certainty, not 100%, was it trained on Sony versus Warner material, for example? In that kind of situation, there are different kinds of problems, and there are different ways to address them in different ways. And their companies are focusing on very specific aspects of this.

Danielle Bulger: As a litigator, these new technologies are interesting to me, and I'm sure a lot of people in this room, because essentially we'll be able to better detect whether AI was used, and vice versa, in the creation of works of authorship. But let's switch gears really quickly. On July 31st, the Copyright Office published part one of its report on copyright and artificial intelligence. Here's an image of this great report that's available online. In it, they address two congressional proposals intended to address the unauthorized use of digital replicas, the No AI FRAUD Act[45] and the NO FAKES Act.[46]

As mentioned earlier, there have been so many developments in this area, including just this past weekend. We now are seeing various states look into their own legislation around AI.[47] Leonard, can you identify any clear differences between these many laws or proposed laws?

Leonard Rosenthol: Yeah, so there are about 20-some-odd states that are currently looking at laws, or some are further along than others, as has been mentioned, as well as the federal one here in the US. And then you heard about the EU AI Act.[48] The UK, Japan, China even, everybody, Australia is working on something. But if we go back just within the US, the states started out in very

---

[45] No AI FRAUD Act, H.R. 6943, 118th Cong. (2024).

[46] NO FAKES Act, S. 4875, 118th Cong. (2024).

[47] *See, e.g.,* California AI Transparency Act, S.B. 942, 2024 Leg., Reg. Sess. (Cal. 2024) (requiring companies to provide detection tools to ascertain whether content has been AI generated); Ensuring Likeness, Voice and Image Security (ELVIS) Act, H.B. 2091, 113th Gen. Assemb., Reg. Sess. (Tenn. 2024) (prohibiting use of AI to mimic a person's voice without permission).

[48] Regulation (EU) 2024/1689 of the European Parliament and of the Council, 2024 O.J. (L).

different places. For example, California, you've seen the 12311 bill, which did not make it, which was around AI transparency and labeling.[49]

The one that did was around transparency and in documenting it.[50] We talked about that in the previous session. What you see in the EU is more on the stuff that was talked about before, which is copyright ownership, rights holders, et cetera. But what was interesting is that my state, New Jersey, went completely off the rails when they started the process, which is that the original version of the New Jersey law was that they wanted to mandate that anyone selling a camera or smartphone into the state of New Jersey had to have that camera include labeling technology that it was made, that every picture was made by the camera.

So, they were not even looking at AI. They were completely focused on humans and cameras, which is interesting, especially something that, as far as I know, there are no camera manufacturers actually in the state of New Jersey. But it just shows you that everybody's starting from some point on this very large continuum and that companies, organizations are helping them all come together. And one of the things that isn't there today, although the Partnership on AI has a wonderful vocabulary, which is a huge help, they're still missing pieces in that vocabulary.

And so, we got to get them on the same page with the vocabulary and then get them all moving, hopefully, too, even if they're not exactly the same law, that would be kind of strange anyway. At least they're aligned on where they want to go. And I'm hopeful that will happen sooner rather than later.

Danielle Bulger: I want to make sure we have enough time for questions. But, before we do that, Jenni, since we have you, I am curious if your company endorses any requirements on it with respect to labeling synthetic material and any challenges with provenance in watermarking it. It sounds like a great solution, but there have to be some challenges.

Jenni Katzman: Yeah, no, we do endorse some things on labeling. So, we think that any company that can create sophisticated audio or visual content should be required by law to use this state-of-the-art provenance tooling, things like C2PA, so that people can understand whether something is AI-generated or not. There obviously should be some caveats around that. It should be technically feasible within the realm of something that's cost-feasible, available, that's accessible for people with disabilities.

And obviously, there's things that work for certain types of content. Not everything works in text or audiovisual or depending on what the content is. And also, we've seen legislation that has every sort of list of things from tamper-evident to tamper-resistant to indelible. And you can't just label. There's no standard out there that meets every single thing. So, to the extent that the legislators want to future-proof things, there's just no standard that is perfect. And

[49] Cal. S.B. 942.

[50] California Digital Content Provenance Standards, A.B. 3211, 2024 Leg., Reg. Sess. (Cal. 2024).

we just can't live in that world right now. And I think people just need to understand that.

And then in terms of challenges, in that same vein, basically every standard out there, and Leonard can attest to that, is subject to adversarial attack. We need to address that, but we also need to live in that reality, too. So, I think that we all need to figure out a way to be part of a system that tries to address it. And one of the ways that Microsoft is advocating to do that is to prohibit the intentional and deceptive attempts to tamper or strip provenance metadata. And there's a bill out there that's included in a larger package by Senator Warner that would do that.[51] We really do emphasize the "intentional" and "with the intent to deceive" piece of that. And we think that's really important. And we think platforms play a really important role in that because if we're going to be labeling this and including this information, it really doesn't serve the purpose to then have that information stripped away.

This has been done in other contexts, physical assets like ID numbers on automobiles, where you can't do this. We see this in the copyright context. We've mentioned the DMCA where that's included.[52] You don't want to have your copyright information taken away. So, there is precedent for this, but we just need to do it in an intelligent, meaningful way.

Danielle Bulger: Great. We have a host of other questions for our panelists. But I also want to make sure, especially since we have a great turnout today, that we open this opportunity to the audience. I see we have a question in the back. Thanks, Kate.

Questions from the Audience

Speaker 1: Hey, guys. First of all, great panel. Thank you very much. And I think this is probably a question best for Leonard, but others feel free to jump in. So, you guys talked about, I think one of you mentioned removing a blemish, maybe removing hair from somebody's face, whatever. And obviously, you do that to an image of Kamala Harris, no big deal. You start changing words to make Joe Biden, fake Joe Biden say something that he didn't actually say, much bigger deal. So, my question is, if C2PA is widely implemented, as widely implemented as you'd like it to be, is there going to be context for us end users as to what was changed as opposed to that something was changed? Because if it's just a flag that says, "Yes, this was changed through AI," that's somewhat helpful, but it's a lot

---

[51] *See, e.g., Senate Intel Chairman Pushes Companies to Follow Through On Commitments to Combat Deceptive Use of AI*, MARK R. WARNER, U.S. SENATOR FROM THE COMMONWEALTH OF VIRGINIA: PRESS RELEASES (May 14, 2024), https://www.warner.senate.gov/public/index.cfm/2024/5/senate-intel-chairman-pushes-companies-to-follow-through-on-commitments-to-combat-deceptive-use-of-ai; https://www.warner.senate.gov/public/index.cfm/pressreleases?ID=7EAF94B2-11B8-44DF-B062-7D84421738B7.

[52] 17 U.S.C. § 1202 (2024).

more helpful, obviously, to an end user if you have context around that. So, you know the details of what actually is fake and what's real.

Leonard Rosenthol: A hundred percent. So, there's two parts to that problem. So, one is having that data, having a way to reference and include that information in the file. Like, "This rectangle was removed and it represents a blemish," to use your example. So, today, C2PA can incorporate that information already. So, there's a place to stick the information. So, that was Problem 1. We needed to have a place to stick it and to carry it along. The part we're still working on, because it's in some ways the harder part of the problem, is now what and how do you present that to you or to anyone as an end user?

And more importantly, how do you do that in the few seconds as someone is scrolling very quickly through their social media feed? So, we don't have an answer to that yet. That is a problem we as a group are actively working on. We have, in fact, just had a recent influx of cash to do real user research, to go out into the field and get answers to that question and around the world, because what we've already found is that it differs. People in the UK versus the EU and the US see things and understand things very differently. So, we're working on it as the best answer I can give you. But yes, we know.

Danielle Bulger: Great. Bill [Rosenblatt] in the back.

Bill Rosenblatt:[53] Technologies like watermarking and fingerprinting and so on have been around for a long time, and one of the barriers to adoption has been patent thickets and royalties. It's great to see all those big tech companies participating in C2PA. That's going to go a long way towards getting it adopted. But what is C2PA doing to make sure that these solutions are easy to adopt from a liability standpoint and from a royalty standpoint? Do you have a patent pool? Are there FRAND [Fair, Reasonable, and Non-Discriminatory] licensing policies, et cetera et cetera?

Leonard Rosenthol: No, it's an excellent question. So, the core C2PA technology is licensed under the same terms as the W3C. It's under the W3C patent policy. But our core technology, there are no patents registered, there's no licensing, there's no royalties, nothing on that core technology. Done. The way we have approached watermarking and fingerprinting is by not standardizing the technologies, because, as you said, there's lots of stuff out there. Many of them are already standardized by other organizations, SMPTE[54] is a good example. But regardless, we are not going to reinvent a new watermarking technology.

As you said, many of our members already do that. What we are standardizing on is how to interoperate between them so that if you have an image with a Microsoft watermark and another image with a Digimarc watermark, and another

---

[53] Founder and president of GiantSteps Media Technology Strategies; Adjunct Professor of Music and Performing Arts Professions, New York University. *See Bill Rosenblatt*, NEW YORK UNIVERSITY STEINHARDT SCHOOL OF CULTURE, EDUCATION, AND HUMAN DEVELOPMENT, https://steinhardt.nyu.edu/people/bill-rosenblatt (last visited Dec. 17, 2024).

[54] *Homepage*, SOCIETY OF MOTION PICTURE AND TELEVISION ENGINEERS, https://www.smpte.org/ (last visited Dec. 17, 2024).

one with an open-source watermark, there is a way for the client to retrieve that information and use it to look up provenance, because, as I said, at the end of the day, it's all about the provenance. Watermarking is just the way to get you the provenance. And so, you need a standard to do that. And that's what we're working to standardize.

Danielle Bulger: Our clock up here is showing two more minutes, so we'll take one more.

Speaker 2: Hello. One thing I was wondering is that the open-source models are so good and so close to the state of the art that what efforts are being taken to kind of bring the use of the open-source models into the fold for the kind of adoption of these standards?

Leonard Rosenthol: Nicolas, do you want to say anything about that?

Nicolas Gonzalez Thomas: Well, in music, there aren't any really good open-source models. That's where I could speak with most confidence.

Leonard Rosenthol: Yeah, I mean, the short answer is policy. I think this is where people are either going to have to do the right thing or they're going to be legislated to do the right thing.

Danielle Bulger: Right here.

Speaker 3: Thank you. Sorry to pick on you, Leonard.

Leonard Rosenthol: It's okay. It's the game in town. I'm good.

Speaker 3: I appreciate everyone has been talking about new technologies and possible forthcoming technologies. I'd like to focus on something that exists right now and the programs that exist right now. And I'll go with three of yours, Photoshop,[55] Lightroom,[56] and Acrobat.[57] So, in the news business, when it comes to the taking and publishing of breaking news photographs, the general instruction, at least here in New York City, is to photographers, "Send me the JPEGs. Don't touch them, don't crop them, don't tone them, just send them to me."

Now, just over the weekend, I received an ad from your company and it showed a picture of a woman who has a bicycle. And, well, she's leaning against a chain fence. And the ad was something like, "See how easy it is to do this with Photoshop AI. With three clicks, you remove all the chain links from the fence." So, I'm assuming that there is an audit trail of that, because, after all, Photoshop does offer the ability to click undo. So, the first question I have is, can you embed that audit trail into the metadata of the picture so that an editor, let's say, receiving that picture will see there has been some manipulation of it?

And the second thing is you mentioned about using PDFs. Now on Acrobat, at least Acrobat Professional, Acrobat Professional CC, there is an option under protect and encryption, and it says, "Encrypt with password, restrict editing." So, can you take in my original example, the JPEG, which has been edited or has not

---

[55] *Photoshop*, *supra* note 19.

[56] *Lightroom*, ADOBE, https://www.adobe.com/products/photoshop-lightroom.html (last visited Dec. 17, 2024).

[57] *Acrobat*, ADOBE, https://www.adobe.com/acrobat.html (last visited Dec. 17, 2024).

been edited, and lock it so that nothing can be done with it? And it will also reveal the edit trail to show people so those consumers of it, in this case, editors, will see whether or not something has been done because even though they make that rule, people do break it.

Danielle Bulger: That's what the standard provides.

Leonard Rosenthol: Yeah, absolutely. That's right. What you just described is provenance. That's exactly describing that audit. What you call as an audit trail, we call as provenance. It's the exact same thing. So, to your question though, so, in Photoshop, yes, Photoshop today does that, it records all of that information. However, including it into the actual image is an opt-in decision by a user. So, for privacy purposes, we do not do it automatically. A user has to opt into it.

But yes, if you go up and you turn on the option, and it's right in front of them actually at multiple times, and you say add content, credentials, yes, you will get every single action that the user did. We shorten something like you don't get every single stroke of the brush, but you're told the user used the paintbrush. So, I mean, it's not every single thing because that would get very large. But yes, at the end of the day, yes, it does that in our case, and yes, you can opt into it.

Danielle Bulger: All right, we have time for one more. We'll go all the way in the back.

Leonard Rosenthol: And I'm around, and we all are for lunch as far as I know. So, please come see all of us. We'll just pick one more.

Mark Traphagen: [58] Hi, my name is Mark Traphagen. I had a question about interference with the provenance information. And that is, would the C2PA participants, what kind of an intentionality standard are you talking about? Is it a single intentionality or is it a double intentionality like existing section 1202(b), where what is prohibited is an intentional alteration or interference with copyright management information, knowing that it would induce or encourage infringement?

So, with respect to provenance information, would the interference be just intentional interference with the provenance information, or would it be intentional interference, knowing something bad will be done?

Jenni Katzman: Yeah, yeah, I'll take that. Yeah, no, I think it's the latter. I mean what you're asking, what we're advocating for, what the bill is, what we are. Yeah, we'd advocate for the latterpiece on it. I'd have to look at the legislation again to be able to tell you what the legislation says. And again, it's not a standalone bill. But my understanding is that they may be incorporated into other bills as well. But it is two pieces, too. It's not just the intentionality, but it's with the intent to deceive. Yeah.

Danielle Bulger: Well, let's keep this great discussion going during lunch. Thank you all for your time. And a round of applause to our panelists

---

[58] Mark Traphagen is a Professorial Lecturer in Law at GW Law and is a contributing author to the legal treatise Copyright Throughout the World. *See Mark Traphagen*, GW LAW: ADJUNCT FACULTY, https://www.law.gwu.edu/mark-traphagen (last visited Dec. 17, 2024).