



**U.S. Senate Committee on the Judiciary
Subcommittee on Intellectual Property**

Hearing on

“Artificial Intelligence and Intellectual Property – Part II: Copyright”

Written Testimony of

Dana Rao

**Executive Vice President, General Counsel and Chief Trust Officer
Adobe Inc.**

July 12, 2023

Chair Coons, Ranking Member Tillis and members of the Committee:

Thank you for the opportunity to appear before you today. My name is Dana Rao, and I am the General Counsel and Chief Trust Officer at Adobe.

Adobe and Innovation

Adobe has a long history of pioneering innovation. We are the company that brought the world Acrobat and PDF to enable you to create and share digital documents everywhere. We have a digital experience business that powers small, medium, and large businesses’ websites and e-commerce experiences. And we have a collection of products which enable all types of creative expression, including Adobe Illustrator for graphics, Adobe Premiere for professional video editing, and of course, Photoshop, the world leader in image editing. Since our founding in 1982, we’ve continued to invest in transformative technologies that allow our customers to unleash their creativity, perfect their craft, and power their businesses in a digital world.

At Adobe, our customers represent a vast array of creators – ranging from aspiring digital artists, war-time photographers, fashion designers, marketing and advertising specialists, and the world’s leading enterprises. Our millions of customers use Adobe’s tools to create the billboards you see in Times Square, Academy Award-winning movies, the cover art of albums, illustrations in children’s books, the photographs you see in magazines and newspapers, even the flyers on the windows at your local coffee shop. Throughout our history, we have worked together with our creative community as we harness the power of emerging technologies in the tools we provide to support them.

The Power of Artificial Intelligence

The newest advance is artificial intelligence. Adobe has been incorporating AI into our tools for over a decade to help creators realize their potential. AI is a powerful, disruptive technology – a thinking machine that enables computing to solve problems it never could before, with insight based on data. That next magical step, where our computers complete our sentences, create recipes from a list of ingredients, diagnose diseases, answer questions in insightful ways, and assist in our creative activities is here, now. And generative AI is the latest version of AI technology that is transforming the way we work, create, and communicate. Image-generating AI allows you to simply type in a prompt, and you can generate a completely new image from scratch. You’ve all seen the magic: type in “cat driving a 1950s sports car through the desert” and in seconds, multiple variations of a cat on a retro road trip appear before your eyes.

In content creation, AI will usher in a new age. Creativity will finally become accessible to everyone, regardless of their technical skill level. Billions of new creative voices will be able to express themselves through these new easy-to-use tools. And for the professionals, AI will provide an amazing first step in their creative process and an opportunity to design entirely new experiences bounded only by their imagination, while being more productive than ever before. This is the power and promise of AI.

Innovation and Responsible Innovation

The United States has led the world through past technological transformations like the advent of the internet, mobile phones, and cloud computing, and we have learned it is important to be responsible about the impact these technologies have on society. And now, as the United States finds itself at the forefront of the latest technological advance, to maintain our leadership, we must pair innovation with responsible innovation. Placing thoughtful safeguards around AI development and use will help it realize its full potential to benefit society.

With millions of creators as our customers, and as a leader in AI development, Adobe believes a comprehensive analytical framework for encouraging responsible AI development should consider:

- the importance of ensuring access to data to spur the growth of AI;
- the impact of unfair AI bias on high-risk uses of the technology;
- the impact of AI on creators and IP;
- the economic impact of AI on productivity and job creation; and
- mitigations against the spread of deepfakes.

Today, we are here to talk about the intersection of AI and creators, the need for data to train AI accurately and responsibly, the current state and implications of copyright and AI, and how Adobe is addressing these issues in our own Generative AI technology. In addition, Adobe has specific views on how we can protect the rights of the creators and ensure the United States maintains our leadership in AI innovation, including providing artists a “Do Not Train” tag to allow them to opt out of training, providing artists a way to secure copyright and attribution in a world with AI-assisted works, and establishing a new federal right to allow artists to protect themselves against people using AI to impersonate their likeness or style. Finally, we also

believe we can help advance new economic opportunities for creators enabled uniquely by AI. Pairing these creator protections with support for access to data for AI research will ensure AI innovation continues to develop in the right way, for everyone.

Access to Data is Critical to Enable AI Innovation to Develop, and Develop Responsibly

AI is only as good as the data on which it is trained, which is why the question of data access is important to companies and organizations building foundation models. Like the human brain, AI learns from the experiences or information you give it. And like the human brain, the more information you give it, the better it will perform. An AI system trained on a small dataset is more at risk of producing wrong or unsatisfactory results, or reproducing harmful biases that exist within the dataset.

Again, think of it like the human brain. If you'd never been taught what a car is, it would be hard to accurately depict one or answer a question about what it is or what it does. To produce accurate results, AI needs a large dataset representing the universe of possible answers to learn from. Additionally, a narrow dataset can lead to unfair bias. If you've only ever been taught that lawyers are men, you are likely to conjure up an image in your head of a man when someone is talking about a lawyer, even though over half of the graduates of law schools are women. AI works the same way. Training on a larger dataset can help ensure you capture a broader set of perspectives in the data itself, so that when you type in "lawyer," you will see a result set that reflects the society in which you live. Given those technical realities, governments need to support access to data to ensure that AI innovation can flourish both accurately and responsibly.

Protecting Creators' Rights in the Age of AI

One of the important issues that arises from AI's need for data access is its implication on creators and copyright. Top of mind for the creative community is whether an AI can be trained on the copyrighted or branded materials that exist on the web without permission. Understanding this has implications for the policy positions we should consider taking.

Copyright is a critical intellectual property protection that has been part of our legal system since the founding of our nation. It is enshrined in our Constitution as an important lever to "promote the progress of science and useful arts" by giving creators a right to protect their interests by protecting against unauthorized copying of their work. Over the past two and a half centuries, the fair use doctrine has played an equally important role in the application of copyright law, codified in the 1976 Copyright Act as a statutory exception to copyright, designed to allow copyrighted work to be used for purposes such as parody, research, and transformative uses. Fair use ensures that copyright law does not "stifle the very creativity which it is designed to foster."

How AI Models are Trained

Whether fair use applies to the use of content to train AI requires an understanding of how an AI model actually works. Training a generative AI model for image creation is a multistep process. To start, hundreds of millions of images are ingested for their underlying facts to train

the AI model. These facts (and the labels associated with the image attributes) are used as part of a feedback loop to create “weights” (mathematical constructs that are built into the decision-making points of an AI model), which help the AI learn what attributes to include in response to a text prompt. The weights are iteratively adjusted until the model generates a satisfactory output based on the prompt. Once the model has been trained, a user can type in a text prompt and the trained model generates a new image or text based on its weights, or its learned understanding of what attributes it thinks would be most likely associated with the words in the text prompt.

From a copyright perspective there are two core questions: Is the output image a copyright infringement of an image that was used to train the AI model? And is using a third-party image to train an AI model permissible under fair use?

For the first question, as noted above, the current technical understanding is that an output image is a new image “hallucinated” from the user’s input text prompts and is not reusing copies of the images that were part of the training dataset to simply assemble a “composite” output image. The input images are used to extract facts for training the model and its weights. This question is currently the subject of several ongoing litigations and will eventually be decided in court or by Congress.

For the second question, whether it would be fair use to use an image to train a model, using an image to train an AI model would typically be considered a transformative use because an AI model, on its own, is a software program, which is very different than the original image. However, if the output of the AI model is substantially similar to a copyrighted work that it had ingested, the question remains whether fair use would be applicable, even though training the model itself may not have been considered a copyright infringement. This question is also the subject of several ongoing litigations and will eventually be decided in court or by Congress.

Adobe’s Approach

At Adobe, we considered all of these questions and developed a comprehensive approach to AI – from training our AI models to economic solutions for artists – that aims to empower creators and enable AI innovation to thrive.

Adobe Firefly: Designed to Be Commercially Safe

Adobe recognized the various unanswered legal questions around access to data in designing our own family of generative AI models, [Adobe Firefly](#), which we launched in March 2023. We chose a path that supports creators and customers by training on a dataset that is designed to be commercially safe.

We trained our first Firefly model only on licensed images from our own Adobe Stock photography collection. If needed for accuracy or bias reduction, we intend to expand that dataset to include openly licensed content and public domain images where copyright has expired. To help ensure copyrighted or branded materials are not created as part of Firefly’s

output, we have a content moderation team that performs extra filtering on the images before they become part of the Firefly dataset.

Having a dataset designed to be commercially safe makes Adobe Firefly more attractive to both the creative community – who are eager to use the power of generative AI, but in a way that respects their fellow creators – and to enterprises who want to use the model to generate brand campaigns and marketing materials without fear of infringing on someone’s IP. This approach provides companies with assurance that, even in an evolving legal landscape, they can leverage the power of this new technology.

Enabling Do Not Train

In addition, as we developed Adobe Firefly, we believed there were important steps we could take to empower creators in the age of AI. At Adobe, we developed a technology called Content Credentials. Based on our work with the Adobe-led Content Authenticity Initiative, Content Credentials allow creators to securely attach a “Do Not Train” tag in the metadata of their work. This credential will travel with the content wherever it goes, so if a web-crawler is scraping the web to build a dataset, they will be able to recognize that credential and prevent companies from training on works whose owners don’t want them used in training. This technology is part of the open standard behind the Content Authenticity Initiative, and anyone can join to implement this credential. With government support, we can make this standard industry-wide so that everyone can enable and respect these credentials to allow an artist to opt out of training if they choose.

Ensuring Copyright Protection for Creation with AI

As discussed above, many of Adobe’s customers are creators. Ensuring the creative community could obtain IP protection for their work was important to us as we considered the question of whether the output of an AI can be protected. We believe that with the help of attribution technology, creators should be able to obtain IP protection over AI-generated works, with some modifications.

As US Copyright law states, you can’t copyright an idea. What you can copyright is the expression of an idea. For example, you cannot copyright the word “book” – that is an idea. But you can copyright a painting of a book because that’s an expression of the idea. When you apply this principle to image-generating AI, it means that a prompt may not be copyrightable because the prompt represents the idea, and the output is based on the AI’s interpretation of that prompt. When you type in the “cat driving a 1950s car through the desert” from earlier, the AI decides whether it’s a Siamese cat or a Tabby cat; it decides whether to make the car a convertible or a pickup truck; it chooses the color of the sky, the number of cactuses in the desert and so on. The AI is expressing the idea, not the human. Copyright law is designed to protect the rights of human creators – therefore an AI output (the AI’s expression of the idea) may not be copyrightable.

But most creators will not use the raw, unmodified AI output as their final creation. Many creators are ideating and brainstorming in a generative AI tool and then adding their own style

and expression to an image. Some artists are using AI just to quickly change the sky color of their artwork, rather than laboring through a tedious pixel-by-pixel manual task. Take our cat in a car in the desert example. Maybe the artist starts with the AI-generated output but then uses other non-AI tools to change the color of the car, add some mountains in the desert, add another cat in the passenger seat, and so on. Or, maybe they already created the image of the cat in the desert using non-AI tools and they used AI to simply swap out one of the cactuses for a desert tree. In these instances, the human creativity and expression they are adding to their work should be copyrightable, but they will need to prove what was AI-generated and what was human-created. Fortunately, there are tools (such as any that implement Content Credentials) that will allow a creator to distinguish the portions of a work created by AI versus the portions of a work expressed by a human.

Enabling Creator Attribution in a Digital Age

Another aspect of creator rights is their need and desire for attribution. In the digital world, it is very easy to reproduce digital content in different forms and in different outlets, and it is often hard to maintain the source attribution for the original work. With Content Credentials, creators can cryptographically associate their identity with their work (for example, an image or a piece of digital art), and then if the work is reused or repurposed, their identity will travel with their work across all platforms that support Content Credentials. To support this approach, Congress should require that all platforms that receive such attributions maintain them, so the attributions are not stripped away and artists can receive credit for their work.

New Federal Rights to Provide Creators Protection Against Impersonation

Finally, Adobe considered one of the core questions raised by our customers: what happens when someone uses an AI model to replicate their style, in direct economic competition with their original work? We believe that it is important that the law offers artists protection against this type of economic harm, and accordingly, we propose that a federal anti-impersonation law be established specifically to protect artists from someone using AI to impersonate their work. Such a law would provide a right of action to an artist against those that are intentionally and commercially impersonating their work through AI tools. This type of protection would provide a new mechanism for artists to protect their livelihood from people misusing this new technology, without having to rely solely on copyright, and should include statutory damages to alleviate the burden on artists to prove actual damages, directly addressing the unfairness of an artist's work being used to train an AI model that then generates outputs that displace the original artist. In addition to style, a federal right of publicity could be created to help address concerns about AI being used without permission to copy likenesses for commercial benefit. The potential for AI to be used for economic displacement is a critical problem to solve, and we believe this Committee should support a legislative solution for it.

New AI Business Models to Benefit Creators

As I shared earlier, we believe AI holds vast potential to unlock new opportunities for creators. First, we believe that Generative AI will increase productivity by automating tedious or repetitive tasks, allowing creators to do in seconds what used to take hours and increasing the creative output of every creative professional. In addition, by making creating easy and fun,

billions of latent creators can now become creative, because the tools are easy to use. As more people create, the demand for professional content will only continue to grow. And leading more people into creative careers will help safeguard jobs from AI and automation in those creative functions that require true creativity and which are most difficult for an AI to replicate. In the creative world, we believe art will always require soul and a story along with it in order to create the real connection and meaning that viewers, readers, and listeners want from it. Helping the U.S. transform into a creator economy will ensure that we maintain a workforce ready to take advantage of where the world is going using the tools (including AI) that are available.

Second, we believe the AI business model itself can unlock new economic opportunities for creative professionals, and Adobe is exploring ways to help creators monetize their work in the age of AI. One approach would be to enable creators to license their style directly to consumers. In this approach, people could subscribe to a particular artist, who has worked with an AI model to have it replicate their style. The consumer could then create their own work in that artist's style, but for a fee. This would be an entirely new revenue stream for artists, and we are currently testing an AI model with creators to do just this.

The Content Authenticity Initiative and Content Credentials

In 2019, Adobe founded the [Content Authenticity Initiative](#) (CAI) to bring transparency to online content to help creators get credit for their work and fight harmful deepfakes. The CAI uses the Content Credentials provenance technology (described above) to function like a nutrition label for content. Content Credentials can tell you the name, date, time, and edits that were made to that content. That nutrition label travels with the content wherever it goes – a stamp of authenticity that gives people more context about the content they are consuming. Without this transparency in digital content, people are at risk of believing lies and of no longer believing the truth. This poses a significant threat to our democracy.

In just four years, the CAI has grown to 1,500 members from across industries, including: Microsoft, Qualcomm, Nikon, Leica, Wall Street Journal, Washington Post, BBC, and more. Recently, Universal Music Group became the first big music label to join, which shows the importance of transparency and authenticity in all types of digital media – whether it be images, videos, or audio. The CAI uses an open standard organization called [C2PA](#) to ensure that this technology is free to anyone to implement, and free for everyone to join. Adobe believes these problems are common to all of us, and we all need to work together to solve them.

The Content Credential technology will be very useful to address the deepfake issue. However, this technology can also serve as the backbone for many of the pro-creator proposals discussed above. Do Not Train, artist attribution, and identifying which parts of a work are AI versus human to help secure copyright protection are all features that can be (and are) included in the Content Credential technology and in the C2PA standard. But in order for this solution to truly empower creators and the public alike, we need it everywhere. We encourage Congress to take

a leading role by including provenance requirements in AI legislation for everyone that makes, uses, and distributes content.

A Collaborative Approach to Bring AI to the World Responsibly

With every technological transformation, we open the door to new opportunities, experiences, benefits, and conveniences. AI has already begun to transform our lives and we are at a pivotal moment. It is critical we take this opportunity to address – in advance, and thoughtfully – the implications of AI as we build our future together. Today’s discussion is an important step in the ongoing collaboration across industry, government, creators, and the public that is needed to bring this technology to the world responsibly.

Thank you, Chair Coons, Ranking Member Tillis, and members of the Subcommittee, for holding today’s hearing and underscoring the need to consider these important intellectual property issues as part of the Senate’s ongoing work on artificial intelligence.