May 24, 2024

The Honorable Mark R. Warner
United States Senate
Washington, D.C. 20510

Dear Senator Warner,

I write in response to your May 14, 2024 letter expressing interest in the measures signatories are taking to implement the *Tech Accord to Combat Deceptive Use of AI in 2024* and expressing concern about the use of generative AI by malign actors to undermine trust in public institutions, markets, democratic systems, and the free press.

The power and promise of artificial intelligence (AI) are upon us. At Adobe, we believe AI can supercharge and democratize creativity to transform how we work, play, and learn. With generative AI models, like Adobe Firefly, our first generative AI foundation model, if you can dream it, you can make it.

However, we recognize that the power of anything good can also be used for bad. We have all seen AI-generated deepfake images, audio, and video insidiously insert themselves into our everyday lives, whether intended to help bad actors reputationally disparage a candidate in an election, commit sophisticated financial fraud, or intentionally inflict emotional pain. The bad actors are finding ways to disrupt our lives by undermining our ability to trust what we see and hear online.

The implications for democracy are profound. We recently conducted the Adobe Future of Trust Study, and one of our findings was that with misinformation becoming more prevalent, most respondents (84% U.S., 85% U.K., 84% France, 80% Germany) express concern that the content they consume online is vulnerable to being altered to fuel misinformation. The study also found that most consumers (83% U.S., 88% U.K., 84% France, 79% Germany) believe that governments and technology companies should work together to protect election integrity against the detrimental effects of deepfakes and misinformation. As we saw recently with AI-generated audio of President Biden being used to steer voters away from the polls in New Hampshire, even shoddy digital fakery holds the power to sow confusion. And the danger of deepfakes is not just the deception, it is also the doubt that they cause. Once we know we can't trust what we see and hear digitally, we won't trust anything, even if it is true. And when you begin to doubt everything, when you can't tell fiction from fact, democracy itself is threatened. Addressing this issue has never been as important as in 2024, with more than four billion voters participating in over 40 elections around the world.

The Tech Accord's seven principal goals represent important actions signatories can take, and many already are. Provided below is a discussion of the actions Adobe has already taken to uphold the principles of the Accord and fulfill its goals. We look forward to continuing to engage with you, your colleagues, and your staff on this critically important issue.

Sincerely,

Dana Rao

Dana Rao
EVP General Counsel and Chief Trust Officer
Adobe Inc.

1. **What steps is your company taking to attach content credentials, and other relevant provenance signals, to any media created using your products? To the extent that your product is incorporated in a downstream product offered by a third-party, do license terms or other terms of use stipulate the adoption of such measures? To the extent you distribute content generated by others, does your company attach labels when you assess – based on either internal classifiers or credible third-party reports – to be machine-generated or machine-manipulated?**

Adobe is committed to advancing Content Credentials across our tools and platforms and maintaining Content Credentials on the content we display. Content Credentials are available in popular Adobe creative applications like Photoshop and Lightroom — where creators can attach information such as date, edits made and tools used. This past year, Adobe announced additional Content Credentials availability for select generative AI features, including Generative Fill and text-to-image in Adobe's generative AI model Firefly. Also in 2023, Content Credentials came to more Adobe applications including Illustrator and Express. We are committed to continue to add Content Credentials to all of the relevant Adobe products in our portfolio, from video, to audio, to documents.

Adobe co-founded the Content Authenticity Initiative (CAI) in 2019 and in just five years, the CAI, and its accompanying open standard, the Coalition for Content Provenance and Authenticity (C2PA) has grown to more than 3,000 members made up of media and tech companies, NGOs, academics, and more including NVIDIA, Qualcomm, OpenAI, Google, Microsoft, TikTok, Associated Press, Wall Street Journal, Sony, Nikon, Leica and many others. The CAI actively promotes the widespread adoption of Content Credentials, which are based on the open technical standard defined by the C2PA. Content Credentials function like a "nutrition label" for digital content. They allow creators to attach information to a piece of digital content such as date, location and edit history. This information is cryptographically embedded into the metadata of a content and is designed to travel with it wherever it is used, published or stored. In this way, Content Credentials allow the creators to establish levels of authenticity with their audiences, giving them a way to be trusted in this digital age. They also give consumers a way to see important context about the digital content they are consuming. Through this increased level of transparency, consumers can make more informed decisions about whether to trust digital content. While it has been rewarding to see the engagement across the industry, for this solution to work, we need to see provenance attached at the moment of creation by the smartphones and other devices consumers are using to capture content, to the ultimate platform where they see it.

Adobe does distribute content created by others in its Adobe Stock platform – a site that hosts stock photography videos, 3D content and other digital content. Adobe Stock offers generative AI images for license and does require contributors to label applicable generative AI images as generative AI. Adobe Stock attaches Content Credentials to any generative AI content downloaded from Stock to inform Stock customers that generative AI was used in the creation of the image.

2. **What specific public engagement and education initiatives have you initiated in countries holding elections this year? What has the engagement rate been thus far and what proactive steps are you undertaking to raise user awareness on the availability of new tools hosted by your platform?**

Adobe, together with the CAI, created Media Literacy Resources available for free on its website. These curricula, developed in collaboration with education experts, are specifically designed to help middle school, high school, and college/university students develop critical media literacy skills to better navigate the ever-changing digital information landscape.

To date, roughly 600 education leaders from across the world have accessed the media literacy lesson plans via the Adobe Education Exchange (EdEx). In addition, educational institutions in Ukraine, India, as well as across the United States, have leveraged the resources. We also take proactive steps to promote the resources using our EdEx newsletter which has a global reach.

As we look ahead, Adobe understands how crucial media literacy is to address the challenges posed by deepfakes and other forms of deceptive synthetic media. People need to know that there are deepfakes out there that could deceive them, and to use provenance tools like Content Credentials to help them verify the authenticity of digital content. As such, we are working on several initiatives to create public awareness campaigns. We welcome the opportunity to work with you and your colleagues in the Senate to share and implement these resources in education and with the public at large.

3. **What specific resources has your company provided for independent media and civil society organizations to assist in their efforts to verify media, generate authenticated media, and educate the public?**

Independent media and civil society organizations are an important part of the digital content ecosystem, and they have a strong interest in promoting trust in their content. Through the CAI, we have welcomed them as original participants in the effort to ensure we heard from all viewpoints. In addition, we have been working with many of these organizations to help them implement Content Credentials by offering engineering support. CAI also makes publicly available on its website CAI's open source software development kit – a set of tools and libraries that enable developers to create, verify, and display Content Credentials based on C2PA standards. These tools have already been deployed by active human rights defense organizations like the Guardian Project.

4. **What has been your company's engagement with candidates and election officials with respect to anticipating misuse of your products, as well as the effective utilization of content credentialing or other media authentication tools for their public communications?**

Adobe has been focused on working with both policymakers and elected officials to help them understand how they can use open-source Content Credentials to help deliver campaign content that can be trusted by citizens. To this end, Adobe has engaged directly with the Federal Elections Commission (FEC) to recommend publication of national standards and best practices that encourage the use of Content Credentials and related technologies in digital campaign content. In addition, Adobe has directly engaged with the White House to have a requirement that provenance be used in the Federal Government's content creation be added to its Executive Order on Artificial Intelligence last Fall. We hope the leadership of the Federal Government using this technology will help proliferate the use by candidates and elected officials.

As part of the FY2024 Further Consolidated Appropriations Act, Congress appropriated $55,000,000 to the states and U.S. territories for the improvement of election administration and election security.[1] Given the rise of AI-generated misinformation in our elections space, states should make use of this funding to integrate content provenance technologies, like Content Credentials, in their election administration processes to bolster trust and transparency. In fact, the U.S. Election Assistance Commission (EAC) paved the way for such use case and issued grant guidance earlier this year approving the use of election security funding to counter disinformation generated through the use of AI.[2] In addition, Adobe has been engaging directly with policymakers to encourage every stage of the digital content supply chain to maintain provenance wherever content goes to ensure the public can easily inspect Content Credentials when present wherever they are consuming online content.

[1] Election security grant information from the Election Assistance Commission.
[2] "HAVA Grants Guidance: Using HAVA Funds to Combat AI-Generated Mis- and Disinformation" publication, page 2.

Adobe has also directly engaged with campaign organizations, and engaged both the DNC and RNC, to promote the use of digital content creation tools that currently support Content Credentials, including exploring the use of media creation products to create campaign content that have Content Credentials built into them. Adobe recognizes that to enable campaigns to leverage the power of Content Credentials, the barrier to entry has to be low and the technology pervasive.

5. **Has your company worked to develop widely-available detection tools and methods to identify, catalogue, and/or continuously track the distribution of machine-generated or machine-manipulated content?**

We appreciate that companies, government bodies, and other organizations are exploring deepfake detection tools. However, our view is that detection capabilities available today are limited, and accurately detecting machine-generated content will become harder over time as generative AI tools advance and get better at creating photorealistic digital content. In addition, detectors will always struggle with human nuance such as satire or parody. As such, we have decided to focus our efforts on developing content provenance tools and solutions and driving widespread adoption of Content Credentials in an effort to help good actors prove what is real. However, we certainly support the research and development of detection technology as another tool in combatting deepfakes. We also believe Content Credentials themselves can send an important signal to an AI detector about the trustworthiness of content, to help improve the accuracy of such detectors, if and when they become useful and available.

When we saw the dangerous rise of deepfakes and misinformation five years ago, we knew we needed a way to help people understand how a piece of digital content came to be and whether to trust it. In our view, technological solutions like provenance, allows us to more effectively build transparency and trust online and combat misinformation. By building an end-to-end chain of trust and content authenticity, we undermine the power of bad actors to spread misinformation by giving good actors a way to be believed. Our aim is to give good actors the tools they need to prove what's true, so that the public will have a verifiable way to distinguish fact from fiction and we will collectively restore a more trustworthy digital ecosystem.

6. **(To the extent your company offers social media or other content distribution platforms) What kinds of internal classifiers and detection measures are you developing to identify machine-generated or machine-manipulated content? To what extent to these measures depend on collaboration or contributions from generative AI vendors?**

While Adobe is not a social media company, we do offer content for license through Adobe Stock. Adobe Stock takes several steps to identify generative AI content, including requiring contributors to acknowledge (via a check box) that digital content was "Created using generative AI tools," introducing specific human and generative AI moderation systems to efficiently categorize and review content for any potential policy violations and refining Stock's audit systems to monitor the existing collection more efficiently for content that may violate Adobe Stock's generative AI policies. Non-compliant generative AI content may be removed from Adobe Stock and the contributor's account may be terminated. Currently, Adobe Stock does not collaborate with generative AI vendors on these measures.

7. **(To the extent your company offers social media or other content distribution platforms) What mechanisms has your platform implemented to enable victims of impersonation campaigns to report content that may violate your Terms of Service? Do you maintain separate reporting tools for public figures?**

While Adobe is not a social media company, to the extent that users can share digital content, we offer users the ability through our products and services to report potential issues that may violate our Terms of Service.

8. **To the extent your company offers generative AI products) What mechanisms has your platform implemented to enable victims of impersonation campaigns that may have relied on your models to report activity that may violate your Terms of Service?**

Adobe offers several methods for users to report abuse or misuse of our products and services. We have a "Report Abuse" feature that users can leverage to flag possibly violative content to us. Users may also contact us at abuse@adobe.com if they are unable to utilize in-product reporting features. Our Adobe Content Policies and product-specific Community Guidelines make clear that we do not tolerate behavior or content that violates the rights of others (learn more at the Adobe Transparency Center).

9. **(To the extent your company offers social media or other content distribution platforms) What is the current status of information sharing between platforms on detecting machine-generated or machine-manipulated content that may be used for malicious ends (such as election disinformation, non-consensual intimate imagery, online harassment, etc.)? Will your company commit to participation in a common database of violative content?**

While Adobe is not a social media company, we think this is a great idea, and Adobe would welcome the opportunity to work with industry and government to address these issues through a workable structure. Adobe primarily works through the CAI and the C2PA to drive awareness around content provenance technologies. With 3,000-plus members, the CAI provides robust collaborative opportunities to share information between members. Additionally, because the C2PA is a smaller subset of companies working on developing an open technical standard, information sharing opportunities amongst member companies via the C2PA's steering committee and its technical working groups are more robust.