How photographers can protect their photographs—and democracy—from generative AI

Sight was the first of our senses to be technologically shared in a world we did not witness with our own eyes. Photography—writing with light—has historically meant a one-to-one relationship between what was before a camera (defined as a lens focussing light on a recording media) and what came out the other end, created by a human.

For nearly 200 years, democracy and photography evolved together, to their mutual benefit. We could agree on what we were seeing in an authentic photograph, and often what it meant. The photographic still serves as a hallmark of visual reliability. Photographs have often been the catalyst for social understanding—and social change—because we could trust what photographs represented.

Photography can still represent "truth." But a photographer is the key to establishing visual trust. There is a well-developed but heretofore implicit ethic ("tell the truth") to establish this trust relationship. First and foremost is, "I was there," i.e., the photographer is a witness.

But photography as we trusted it in 2022 no longer exists. At the end of that year, OpenAI released DALL-E, the first publicly available artificial intelligence tool that generates images based simply on typing a text prompt. Other companies followed suit, and soon the number of AI "phictions" that were made to look like photos exploded. It has been estimated that by mid-2023 AI had already generated more images than all photographers had taken until the invention of the first portable digital camera in 1975. AI-generated images are now so realistic that most people today can't tell the difference between them and an authentic photograph.

Because digital photographs (and video, Latin for "I see") are the primary way we learn about our world in open societies, a lot is at stake. Of course, bad faith actors will use generative AI to repeatedly lie to us with their propaganda when running for public office but there's more. Imagine another scenario: a family member documenting a wedding manipulates photographs taken on a smartphone—erasing a transgender child—to deny what others who could not attend understand about their family. With no camera-recorded provenance, there is no record of how that happened. Another liar's dividend is paid-in-full many years later: who will know what to believe in their digital family album?

What are photographers—from amateurs to professionals—to do? Copyright is no protection when an AI company scrapes your images into the database from which it derives its images. Congress is unable to legislate or regulate generative AI, especially when AI companies spend untold amounts to make sure they can continue monetizing our attention. The same is true for social media companies, whose profits rely on "free" (and increasingly fake) content. There are other options, such as poisoning photographs before publishing them online, but these "solutions" all put the onus on every photographer to protect their practice from an existential threat. And though public skepticism about photography may now be necessary, it's not sufficient to protect the trust relationship that photography has developed with its audiences. Seeing is no longer believing.

Fortunately, solutions exist that could help protect the trust relationship that photography (and more recently videography) has developed with democracy. According to Hany Farid, a professor at UC Berkeley's School of Information specializing in the analysis of digital images and the detection of digitally manipulated images such as "deepfakes," there could be an alignment of interests between authentic photographers and AI companies: when lens-based photographs are digitally and indelibly documented as authentic at the point that the image is captured (with no AI applied anywhere in the process) and permission is granted to scrape them, the resulting generative AI models produce better results. As it stands now, the models

scrape a rapidly increasing number of previously AI-generated images, leading to overall (and irreparable) degradation of their models.

Today, there are more than 57,000 photographs taken every second, and 92% of them are taken on smartphones. Photographers should demand that their cameras—from capture to editing to publishing—cryptographically protect the original digital provenance, as Leica has already done. Apple, which commands more than 60% of the smartphone market in the U.S., could distinguish itself by building the same open-source authentication technology into its camera hardware. (On the other hand, Samsung, heavily invested in built-in generative AI for its smartphones, believes "There is no such thing as a real picture.")

Photographs unaltered by generative AI should be credentialed as such. The Content Authenticity Initiative (CAI) is coalescing companies to address some of these issues, but their content credentials do not visibly differentiate authentic photographs from AI-generated or -altered images. The CAI, convened by Adobe, itself a purveyor of the leading suite of generative AI and photo editing tools, needs to help the public easily identify images that *aren't* generated or edited by artificial intelligence.

Absent any progress on alignment facilitated by CAI, content credentials will soon serve mainly to provide free advertising for AI companies when applied to images. And the effects of disinformation will proliferate: it's just too much to expect that the public will drill down into the credentials of any image to determine if it is fake or not. Anyone should be able to trust—at a glance—that a credentialed photograph was not generated or edited by artificial intelligence.

We're at an inflection point in the power of photography. But we don't need to resign to the future generative AI envisions for us. These solutions will reinforce the trust relationship that has existed since the invention of the photograph, that a human witness—using a camera—was the maker of the photograph. We should be able to trust *and* verify what we see recorded by a camera.

Unless we act soon, images generated by AI will irreversibly outnumber published authentic photos, and we won't know when society can no longer distinguish between authentic and generated images.

There are some things that AI should not do. One of them is to take photographs for us.

<div align="center">###</div>

Marshall Mayer was awarded an MFA in Photography from UCSD, adopted digital photography in 1994, and takes notes with an iPhone at http://take-note.com. He is also the producer of the Writing with Light Bibliography and founder of trust.photography, a Discord community.