

MOTIVE. METHOD. MEANS.

I would like readers to take seriously the possibility that federal and state law enforcement agencies now or will soon be able to fabricate video surveillance evidence using generative AI.

The question is not hypothetical. We live in a world where video is treated as secular truth: more reliable than testimony, more objective than memory, more persuasive than a first-person account. That faith was always somewhat misplaced. Cameras have frames. Files have metadata. Recordings can be cropped, compressed, muted, stabilized, color-corrected, slowed down, narrated, selectively released, or stripped of context. **Generative AI adds a new possibility: not only presenting real footage in a misleading way, but fabricating or recreating footage that appears to have been captured by a real device at a real place and time.**

This possibility matters because the justice system increasingly relies on surveillance technology at precisely the moment when public trust in institutions has been weakened. **Many people no longer trust police departments, prosecutors, news organizations, or even juries to distinguish truth from institutional convenience.** That trust cannot be restored by more cameras, more dashboards, or more AI products. It can only be restored by **due process**, adversarial testing, transparent evidence-handling procedures, and a renewed commitment to the presumption of innocence.

The issue is not whether every video should be presumed fake. The issue is whether courts, defense attorneys, journalists, and jurors should continue to treat video as uniquely authoritative when the technical conditions that made that assumption plausible are rapidly disappearing.

MOTIVE

The motives are not mysterious. The incentives are plain enough to name.

First, prosecutors and law enforcement agencies have an institutional interest in making charges stick. Video evidence can strengthen the state's negotiating position, encourage plea agreements, and make defendants less willing to risk trial. In the real world, most criminal cases do not end with dramatic jury verdicts. They end with bargaining, delay, risk assessment, and fear.

Second, video can shape public opinion before a jury is ever seated. Surveillance footage circulates through television, internet news, social media, and edited police releases. This is especially dangerous in political contexts: demonstrations, marches, sit-ins, labor actions, campus protests, or other mass gatherings where public sympathy may determine whether official force is viewed as necessary or abusive.

Third, police departments and municipalities have a strong interest in avoiding civil liability. A recording that supports an officer's account can reduce reputational damage,

disciplinary exposure, insurance pressure, settlement value, and the risk of a wrongful-death or excessive-force judgment. I am not implying that police departments falsify evidence as a matter of course. I am saying that any system in which careers, budgets, public reputation, and money are at stake requires independent checks.

“Those who serve” are as human as those who serve time.

The problem is not confined to criminal law. Much of the surveillance infrastructure in the United States resides in the private sector. It belongs to fast-food franchises, malls, banks, retail chains, corporate campuses, warehouses, plants, hospitals, universities, apartment complexes, residential HOAs, and private security vendors. A recording can influence an insurance dispute, an employment decision, a workplace investigation, a premises-liability claim, or a civil lawsuit. If video evidence is becoming easier to manipulate, the problem extends into corporate governance, risk management, insurance, and civil liability.

Falsifying evidence is already a crime. In federal matters, 18 U.S.C. § 1519 provides that knowingly altering, concealing, falsifying, or making a false entry in a record or tangible object with intent to impede, obstruct, or influence an investigation or federal matter can be punished by fines, imprisonment of up to 20 years, or both. State laws likewise generally treat evidence tampering, obstruction, perjury, and related conduct as serious criminal matters. The point is not subtle.

No one is above the law. That principle matters most when those handling evidence are themselves agents of the state.

METHOD

Generative AI has matured by leaps and bounds.

In 2024, OpenAI described Sora as a text-conditional diffusion model trained jointly on videos and images, using a transformer architecture that operates on “spacetime patches” of visual data. OpenAI stated that its largest model could generate a minute of high-fidelity video, and framed this research direction as a step toward systems that can simulate aspects of the physical world. Google DeepMind’s Veo line now emphasizes realism, prompt adherence, audio, dialogue, ambient sound, and “real world physics” as part of the model’s appeal. These systems are not perfect. They still make mistakes with hands, reflections, object permanence, causality, and physical continuity. But surveillance footage is often blurry, compressed, occluded, low-resolution, dark, and shot from awkward angles. The very defects that make surveillance video hard for humans to interpret also make synthetic manipulation harder to pinpoint.

The relevant capability is not necessarily the creation of a cinematic deepfake from scratch. A more plausible evidentiary threat is constrained generation: taking existing footage, replacing a few minutes or seconds, altering a body position, changing the apparent timing of contact, removing an object, inserting a gesture, matching camera motion, and rendering the result back into the degraded visual language of a bodycam, doorbell camera, parking-lot camera, jail transport camera, or store security camera.

A video does not need to be entirely fake to be misleading. A few missing minutes, a muted audio track, altered color levels, stabilized motion, cropped context, shifted timing, missing frames, or a “cleaned up” version substituted for the original can change the perceived meaning of an event. Generative tools magnify that risk.

This is where generative video intersects with 3D and virtual reality. Video is becoming less like a fixed record and more like a model of an event that can be re-rendered. Modern systems increasingly reason across image, video, depth, pose, camera movement, audio, lighting, and scene geometry. That does not mean they possess human understanding. It means they can produce outputs that preserve enough spatial and temporal consistency to persuade a human viewer under ordinary conditions.

One technical bridge is neural rendering. Earlier methods such as NeRFs—neural radiance fields—showed that a system could learn a 3D representation of a scene from multiple 2D views and synthesize new viewpoints. More recent techniques, including Gaussian splatting, can turn camera captures of real-world locations into navigable 3D environments.

Virtuix, a publicly traded VR gaming and education company, describes its Virtual Terrain Walk system as using AI-driven 3D reality capture to generate geo-specific environments for military planning and rehearsal. Its own materials say the system can support 12 or more users, simulate conditions, and allow trainees to “walk the terrain before you fight on it.” Using existing drone footage, virtual terrain environments can be generated in as little as two hours. As of March 30, 2026, the company had signed an agreement with the U.S. Navy, and sold test units to the U.S. Military Academy at West Point, the U.S. Air Force Academy, Yokota Air Force Base, and U.S. Marine Corps. A video demo, produced in partnership with Florida-based **Hatalom**, is available online on their website for public viewing.

Consumer, enterprise, and defense markets now openly advertise tools that reconstruct real-world terrain into navigable synthetic environments. Once an event can be represented as a scene, the question becomes what else can be added, removed, replayed, or exported.

A training system designed to let soldiers rehearse a battlefield is not the same as a system designed to falsify police evidence. But it points to the underlying capability. If a real location can be captured, reconstructed, annotated, populated, modified, and walked through, then the evidentiary world needs to grapple with the difference between a recording of an event and the rendering of a plausible event.

Generative AI also changes the economics of manipulation. Traditional video forgery required time, expertise, and a tolerance for visible defects. But the trend is toward lower cost, faster workflows, better automation, more plausible motion, and easier integration of multiple inputs: text prompts, reference images, audio tracks, depth maps, pose estimates, masks, scene descriptions, and existing video. Agencies, vendors, or contractors may increasingly rely on opaque workflows that blur the line between enhancement, reconstruction, simulation, and fabrication.

This is especially dangerous because the visual errors of generative media overlap with the ordinary errors of surveillance media. Distorted hands, smeared edges, missing details, shifting shadows, and strange body geometry can indicate AI synthesis. But sometimes the same effects may result from mundane causes. The more technically complex the media pipeline becomes, the harder it is for an ordinary defendant or overburdened public defender to challenge the result.

Detection is an arms race. Authentication must happen upstream. The law needs to know where a recording came from, which device captured it, whether the original file still exists, who accessed it, how it was copied, what was done to it, and whether the version shown to the defense is the same version that left the camera.

MEANS

The means are the surveillance infrastructure already in place.

The United States is covered by public and private cameras: police bodycams, license-plate readers, traffic cameras, doorbell cameras, store cameras, transit cameras, school cameras, workplace cameras, warehouse cameras, hotel cameras, hospital cameras, and cameras installed by homeowners' associations. These systems do not merely record. Increasingly, they index, search, classify, correlate, and share.

Flock Safety is one of the most visible examples of this trend. In his 2025 feature with Forbes Magazine, Flock Safety CEO Garrett Langely is pictured with a 1950s style flat-top haircut and a light blue polo shirt, drone hovering above his left shoulder. **He describes the company's mission as preventing all crime in the U.S.** According to the company website, its platform connects communities, businesses, and public safety so incidents can move from capture to "clear, searchable evidence."

The concerns are not abstract. In October 2025, U.S. Senator Ron Wyden accused Flock of failing to protect Oregonians from searches by federal and out-of-state law enforcement and of misleading state and local customers about data sharing with the Department of Homeland Security and other federal agencies. **Citing Wyden's letter, 404 Media reported that Homeland Security Investigations, the Secret Service, and the Navy's criminal investigation division had access to Flock's nationwide network.** Local reporting said Flock confirmed that CBP, HSI, the Secret Service, and NCIS had access as part of a pilot, and that CBP and HSI conducted about 200 and 175 searches respectively.

Amazon Ring further complicates the picture. In October 2025, Ring announced an intended expansion of Community Requests to local public safety agencies using Flock Safety's Nova platform and FlockOS. On February 12, 2026, the planned integration was canceled and Ring later stated that no customer videos were sent to Flock. Nor is there any known relationship between Virtuix and Flock Safety. However, the termination of Amazon's connection to Flock Safety certainly would *not* rule out pre-launch testing and proof-of-concept validation.

An abuse scenario does not require a single all-powerful system. It requires only a chain of access, incentives, and technical opacity. A private camera records an incident. A cloud system stores it. A police department obtains it. A prosecutor reviews it. A vendor enhances it. A contractor exports it. A defense attorney sees a compressed derivative. A news outlet airs a clip. A jury later watches something labeled “video evidence” without ever seeing a complete provenance record.

Before generative AI, this was already a problem. In 2020, KTVU reported on Bay Area police departments editing body-camera footage before public release. **Vallejo, California settled the Ronell Foster wrongful death case for \$5.7 million while the department had defended the shooting and released what critics described as heavily edited body-camera footage supporting its narrative.** That case did not involve generative AI. That is precisely why it matters. The evidentiary problem predates synthetic video. AI simply raises the stakes.

There are at least three categories courts must learn to distinguish.

First, there is authentic footage that has been selectively edited. This includes clips that are shortened, annotated, narrated, slowed down, muted, cropped, or released without surrounding context.

Second, there is authentic footage that has been processed or enhanced. This includes stabilization, sharpening, noise reduction, color correction, frame interpolation, audio filtering, exposure adjustment, and format conversion. Some enhancement is legitimate. But every transformation should be documented, reproducible, and disclosed.

Third, there is substantially synthetic footage: footage in which visual or audio content has been generated, inserted, removed, reconstructed, or replaced. This is the category most people think of when they hear “deepfake.” But in evidentiary settings, the boundary between category two and category three may become blurry.

That boundary is where the legal system is weakest. Courts are accustomed to asking whether a photograph or video is “what the proponent claims it is.” But machine-generated and machine-altered evidence complicates that question. In 2026, Reuters reported on proposed Federal Rule of Evidence 707, intended to address machine-generated evidence by requiring reliability scrutiny similar to expert testimony, and on proposed Rule 901(c), aimed at fraudulent or inauthentic digital evidence such as deepfakes. The fact that such proposals are being debated should tell us something: existing evidentiary rules were not built for a world where machines can generate litigation-ready media at scale.

Procedural and technical rigor offer a partial solution.

The evidentiary value of a recording depends not only on what it depicts but on confidence that the original file has been preserved intact. Courts should distinguish rigorously between an authenticated evidentiary master and derivative copies that have been enhanced, cropped, color-corrected, stabilized, compressed, annotated, or otherwise processed for presentation. We cannot rely on an MP4 traveling from an

officer's body camera to the police department, then to the district attorney's office, then perhaps to the prosecutor's whiz-kid video editor cousin to be "cleaned up" and "have the levels adjusted" before the defense has a chance to see it. That is not a chain of custody. That is a trust fail.

The technical safeguards are not mysterious. Cameras can cryptographically sign footage at the moment of capture. Files can carry tamper-evident provenance metadata. Evidence systems can preserve immutable audit logs. Access, copying, enhancement, export, and viewing can be logged. Original files can be preserved in read-only storage. Defense counsel can receive authenticated masters rather than compressed derivatives. Expert review can be funded when liberty is at stake.

The C2PA standard already provides a framework for establishing the origin and edit history of digital content through Content Credentials. Its FAQ states that cameras can generate and sign Content Credentials offline using locally stored cryptographic keys. That is promising. It is not enough. Metadata can be stripped. Signing systems can be poorly implemented. A camera can sign a recording of a screen displaying synthetic media. Credentials can document provenance without proving semantic truth. Cryptography can tell us that a file came from a device. It cannot tell us that an officer's description of the event is honest.

That is why technical safeguards must be paired with legal safeguards. Original recordings should be preserved automatically. Any derivative version should be labeled as derivative. Enhancement steps should be disclosed and reproducible. The defense should have prompt access to originals. Judges should be trained to treat machine-generated, machine-enhanced, and machine-authenticated evidence as separate categories. Juries should be instructed that video, like testimony, requires context and scrutiny.

Polygraph (lie detector) tests may also provide a counterweight to video evidence. But they illustrate the same basic problem: no technical instrument deserves blind trust. The Supreme Court has upheld rules excluding polygraph evidence in some contexts, and the Department of Justice has long treated polygraph admissibility as limited and contested. If even lie-detector evidence remains controversial after generations of use, courts should be cautious about treating opaque AI-mediated video as self-authenticating truth.

Surveillance feeds voyeurism and a false sense of omniscience. It trains citizens to watch one another, trains institutions to doubt human testimony, and trains jurors to believe that the image is more real than the person standing in front of them.

Women often oppose intrusive surveillance, including workplace monitoring. But this is not only a matter of personal discomfort. It is a matter of institutional legitimacy. ***A justice system cannot survive if people believe that the state can choose its preferred version of reality, render it as video, and call it evidence.***

Nor can the problem be solved by saying that people who have nothing to hide have nothing to fear. People have bodies, homes, jobs, children, immigration status, health histories, political beliefs, enemies, ex-partners, and bad days. Surveillance footage can be used to misread all of that. It can be used to punish people who are poor, politically inconvenient, frightened, angry, or simply unlucky. Those who live with disabilities and physical and mental health conditions are particularly at risk.

We are all embedded in a global surveillance culture where electronic recordings are considered more reliable than eyewitness accounts or testimony. In fact, **no single source of information is reliable**. Even conventional video editing techniques—adjusting lighting and color balance, muting or heightening sound levels, changing playback speed—can influence a jury, or more likely a district attorney, judge, or public defender. Nobody wants to serve prison time or wait years for a day in court, so plea bargains and probation are the name of the game. People who find themselves unexpectedly caught up in the criminal justice system usually do not have the option to choose a private defense attorney. A 15-minute closed-door discovery session where video evidence is shared and the defendant is not even present to confirm or deny what happened can alter strategy and tactics for both sides.

This is why we have the jury system, and why presumption of innocence is the law of the land in the United States. Generative AI video already possesses the capability to distort and falsify evidence. My hope is that if this clear and present danger to our justice system is understood and shared widely enough to reach the general public, it can begin to make a difference in what are always life and death decisions.

SOURCES:

<https://virtuix.com/vtw>

<https://www.barchart.com/story/news/1041755/virtuix-vtix-signs-navy-agreement-expanding-vr-tech-to-a-fourth-military-branch>

<https://www.youtube.com/watch?v=YuxKjAcRaLI>

<https://www.flocksafety.com/>

<https://www.flocksafety.com/products/flock-nova>

<https://www.404media.co/ice-secret-service-navy-all-had-access-to-flocks-nationwide-network-of-cameras/>

<https://www.eff.org/deeplinks/2025/12/effs-investigations-expose-flock-safety-surveillance-abuses-2025-review>

<https://blog.ring.com/about-ring/ring-expands-community-requests-to-additional-community-safety-partners/>

<https://www.forbes.com/sites/thomasbrewster/2025/09/03/ai-startup-flock-thinks-it-can-eliminate-all-crime-in-america/>

<https://www.hatalom.com>

https://scholarsarchive.byu.edu/cgi/viewcontent.cgi?params=/context/sigma/article/1327/&path_info=05_The_Impact_of_Gender.pdf

<https://techforhumanitylab.clahs.vt.edu/window-into-the-intimate-the-undermining-of-womens-personhood-through-digital-surveillance/>