

FORENSIC TELLS:


The Litigator's Guide to Detecting Deepfakes and Authenticating Digital Evidence



● Authentic Frame



! Suspected Deep Fake



BY CRAIG BALL

Forensic Tells: The Litigator's Guide to Detecting Deepfakes and Authenticating Digital Evidence

By Craig Ball

©2026

Introduction: The Age of Doubt

For decades, I've been telling lawyers that electronic evidence is different from paper evidence—that it carries with it a hidden payload of information about its origins, handling, and integrity. Metadata, I've preached, is the DNA of digital evidence. Now, as we enter an era when any photograph, video, or audio recording can be convincingly fabricated by artificial intelligence, that metadata has become more than a curiosity for the technically inclined. It has become the last line of defense against manufactured reality.

We have arrived at a moment I long feared: the democratization of deception. What once required a Hollywood studio, a team of visual effects artists, and a budget measured in millions can now be accomplished by a teenager with a laptop and a few minutes to spare. Deep fake technology—the use of artificial intelligence to create synthetic media depicting events that never occurred or words that were never spoken—has matured from a novelty to a genuine threat to the integrity of our courts.

This is not hypothetical hand-wringing. Deepfake are turning up in litigation, and allegations of deep fakery are becoming a go-to defense. While reported case law remains sparse—courts don't always identify challenged evidence as AI-generated in their opinions—practitioners report encountering synthetic media issues in custody disputes, insurance matters, and employment cases. Fabricated recordings have been offered as “evidence” of confessions, affairs, and criminal conduct; conversely, authentic evidence is now routinely challenged as potentially fake. Both phenomena will only accelerate as the technology improves.

As Texas lawyers, we pride ourselves on our ability to sniff out falsehood. We cross-examine witnesses, challenge documents, and demand proof. But how many of us are equipped to challenge a video recording that appears to show our client committing a crime he didn't commit? How many of us know what questions to ask in discovery to expose a fabricated photograph? How many of us understand the technical fingerprints that distinguish authentic digital media from synthetic imposters?

This article aims to equip you with that knowledge. We'll explore the technical foundations of digital media authentication, with particular emphasis on the metadata that accompanies—or conspicuously *fails* to accompany—genuine digital evidence. We'll examine the visual, auditory, and logical tells that betray synthetic media and develop practical discovery strategies for obtaining the information you need to challenge suspicious evidence.

My goal isn't to make you a forensic examiner. My goal is to make you a better advocate—one who knows what to look for, what to ask for, and when to call in expert reinforcement. In an age when seeing is no longer believing, the lawyers who understand digital authenticity will be the lawyers who prevail.

Part I: Understanding the Enemy—What Deepfakes Are and How They Work

The Technology Behind the Illusion

Before we can detect deepfakes, we need to understand what we're up against. The term "deepfake" is a portmanteau of "deep learning" and "fake," referring to synthetic media created using artificial intelligence. The technology has evolved at a pace that should give every trial lawyer heartburn.

The first wave of deepfakes emerged from generative adversarial networks, or GANs—a class of algorithms that remain historically significant even as they've been largely superseded. A GAN pits two neural networks against each other: a generator that tries to create synthetic media and a discriminator that tries to spot the fakes. The generator learns from its failures, improving its output until it can consistently fool the discriminator. This adversarial training process produced the first convincing face-swap videos that captured public attention around 2017–2018.

Today's most capable synthetic media systems use different architectures. Diffusion models—the technology behind systems like Midjourney, DALL-E, and Stable Diffusion—learn to gradually add and remove noise from images, eventually gaining the ability to generate entirely new images from text descriptions or to modify existing images in sophisticated ways. Transformer-based models, similar to those underlying large language models like ChatGPT and Claude, now power state-of-the-art video generation and voice cloning. These newer approaches produce results often indistinguishable from authentic media to the naked eye.

For our purposes as lawyers, the technical details matter less than the practical implications. What you need to understand is this: modern AI systems can create photographs that never existed, videos of events that never occurred, and audio recordings of words never spoken. They can place a person at a location they never visited, put words in their mouth they never said, and manufacture documentary "evidence" of conduct that never happened.

Categories of Synthetic Media

Deepfakes come in several varieties, each presenting distinct challenges:

Face swaps transplant one person's face onto another's body in video footage. This is the classic deep fake—taking an actor's performance and pasting the target's face on top. The underlying body movements, gestures, and environment are real; only the face is synthetic.

Face reenactment is subtler. Rather than swapping faces, these systems animate an existing photograph or video of the target, making them appear to speak words or display expressions captured from a different source. The target's face remains their own, but its movements are puppeted by another.

Full synthetic generation creates entirely artificial media from scratch. Using only photographs of a target—often scraped from social media—these systems generate novel images or videos placing the target in fabricated scenarios. No source video is required; the entire creation is synthetic.

Audio deepfakes clone a person's voice from sample recordings, then generate speech in that voice from text input. With as little as a few seconds of source audio, modern systems can produce startlingly convincing vocal performances in the target's voice.

Hybrid approaches combine multiple techniques. A fabricated video might use a real background environment, synthetic face animation, and cloned audio to create a seamless—and entirely false—record of events.

The Litigation Landscape

Deepfakes have already infiltrated our courtrooms, and their presence will only grow. Consider the contexts in which synthetic evidence might appear:

In **family law**, a spurned spouse produces a video appearing to show their partner engaged in domestic violence, substance abuse, or inappropriate conduct with the children. The video is compelling, the children's welfare is at stake, and the accused party insists—truthfully—that the events never occurred.

In **employment litigation**, a plaintiff offers an audio recording of their supervisor making racist or sexist remarks. The recording sounds authentic. The supervisor denies ever making such statements. Without the tools to challenge the recording's authenticity, how does the defense proceed?

In **personal injury cases**, surveillance footage appears to show the plaintiff engaged in physical activities inconsistent with claimed injuries. The plaintiff insists the footage is fabricated. Is it?

In **criminal matters**, the stakes are highest of all. A defendant's alibi may depend on proving that video evidence has been manipulated. A victim's credibility may hinge on demonstrating that an exculpatory recording is synthetic.

These scenarios aren't speculation. They are happening now, in courts across Texas and the nation. And as deep fake technology becomes more accessible and more convincing, they will become commonplace.

The Liar's Dividend

There is another dimension to the deep fake threat that may prove even more corrosive to the pursuit of truth in our courts. Law professors Bobby Chesney (Dean, University of Texas School of Law) and Danielle Citron (University of Virginia School of Law) have given it a name: the Liar's Dividend.¹

The Liar's Dividend is the benefit that accrues to liars from the mere *existence* of deep fake technology. Once the public—and juries—become aware that any video or audio recording *might* be fabricated, wrongdoers gain a ready-made excuse to dismiss authentic evidence of their misconduct. “That’s not me in that video—it’s a deepfake.” “That recording was generated by AI.” “You can’t trust anything you see anymore.”

This isn't hypothetical. We're already seeing it. Defendants confronted with damaging recordings claim fabrication. Litigants caught on video deny the evidence of their own conduct. The existence of deep fake technology provides a ready-made alibi for anyone seeking to evade accountability.

The Liar's Dividend operates as a two-pronged attack on evidentiary truth. First, fabricators create synthetic evidence to frame the innocent or support false claims. Second, the guilty invoke the specter of fabrication to escape genuine evidence of their wrongdoing. Both corrode the fact-finding process, and both demand that lawyers become sophisticated consumers and challengers of digital media.

This is why the authentication tools discussed in this article matter as much for *defending* authentic evidence as for *attacking* fabricated evidence. When your client has genuine video evidence of the opposing party's misconduct, and that party claims the video is a deep fake, you must be prepared to prove authenticity—through metadata, through chain of custody, through forensic examination—with the same rigor you'd apply to challenging suspicious evidence. The Liar's Dividend can only be defeated by evidentiary competence on both sides of the authenticity question.

Part II: The Metadata Foundation—How Digital Evidence Carries Its Own Authentication

The Hidden Payload

Here's the good news, and it's *very* good news: authentic digital evidence doesn't merely consist of the visible content—the image you see, the audio you hear, the video you watch. It carries with it a wealth of contextual information about its creation, handling, and integrity. This information,

¹ Robert Chesney & Danielle Keats Citron, *Deepfakes: A Looming Challenge for Privacy, Democracy, and National Security*, 107 Cal. L. Rev. 1753 (2019).

broadly termed metadata,² serves as a powerful tool for distinguishing genuine evidence from synthetic fabrications.

Metadata is data about data. When you take a photograph with your smartphone, the resulting image file contains far more than the visual content. It records the camera that captured it, the settings used, the date and time of capture, and often the precise GPS coordinates where the photograph was taken. It may contain a thumbnail of the original image, a history of editing operations, and cryptographic signatures verifying its integrity.

This metadata is the digital equivalent of a chain of custody. It tells us where the evidence came from, how it was created, and what has happened to it since. And critically, when digital evidence is fabricated, this metadata is almost always absent, incomplete, or inconsistent.

Think of it this way: authentic digital evidence has a birth certificate, a family history, and a paper trail. Fabricated evidence is a foundling—appearing from nowhere, with no documented origin and no verifiable lineage.

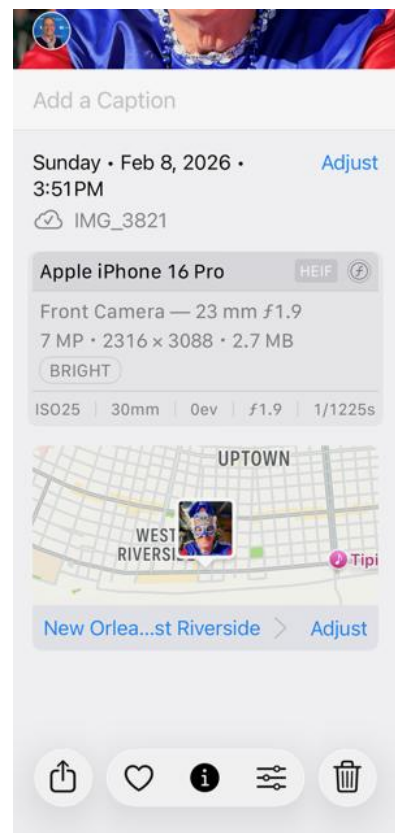
EXIF Data: The Birth Certificate of Digital Images

The most important category of metadata for image authentication is EXIF data—Exchangeable Image File Format information embedded in photographs at the moment of capture. EXIF data is automatically generated by digital cameras and smartphones, and it provides a detailed record of the circumstances of image creation.

A photograph taken with an iPhone, for example, will typically contain the following EXIF information:

Device identification: The make and model of the camera or phone, often including specific hardware identifiers that can trace the image to a particular device.

Capture settings: The aperture, shutter speed, ISO sensitivity, focal length, and other technical parameters used to capture the image. These settings constrain what the image can plausibly depict—a photograph taken at f/1.6 with a short focal length will have different depth-of-field characteristics than one taken at f/2.8 with a telephoto lens.



² For more on metadata: Ball, *Beyond Data about Data: The Litigator's Guide to Metadata*, <http://www.craigball.com/metadataguide2026.pdf>

Date and time: The timestamp of capture, typically drawn from the device’s internal clock. Multiple timestamps may be present, including the original capture time, modification times, and the time zone in which the device was operating.

GPS coordinates: If location services are enabled, the precise latitude, longitude, and altitude where the photograph was taken, often accurate to within a meter or two. This information can place the image at a specific, verifiable location.


Software information: The firmware version of the camera and any processing software applied to the image.

Thumbnail data: A smaller version of the image embedded within the file, which should match the main image but may reveal prior editing if it doesn’t.

Orientation data: Information about how the camera was held when the image was captured.

When you encounter a photograph offered as evidence, the presence of complete, internally consistent EXIF data is a strong indicator of authenticity. Conversely, absent or inconsistent EXIF data is a red flag warranting investigation.

A critical caveat: the absence of metadata is not, by itself, proof of fabrication. Legitimate transmission pathways routinely strip or normalize metadata. When you text a photograph through iMessage or share it via WhatsApp, the messaging platform typically removes EXIF data for privacy reasons. Social media uploads routinely strip GPS coordinates and device identifiers. Screenshots capture only pixels, not the metadata of the underlying image. Cloud synchronization services may alter timestamps. Email



attachments may be re-encoded.

This means metadata analysis is most powerful when you can obtain the *original* file from the source device, *before* it has passed through any transmission pathway that might strip or alter metadata. Absent metadata is a question to be answered, not a conclusion to be drawn. The question is: *is there an innocent explanation for the missing metadata, or does the absence—combined with other factors—suggest fabrication?*

Why Deepfakes Lack Authentic Metadata

Here's the critical insight: AI systems that generate synthetic images don't create authentic metadata. They aren't cameras. They're software programs running on computers, and the "images" they produce are mathematical outputs, not optical captures.

When a deep fake image is generated, several telltale metadata anomalies typically result:

Missing EXIF data: The image may contain no EXIF data at all, or only the minimal metadata added by the software that saved the file. There will be no camera identification, no capture settings, no GPS coordinates—because no camera was involved.

Inconsistent EXIF data: A fabricator may attempt to add EXIF data to make the image appear authentic, but creating consistent, plausible EXIF data requires technical sophistication most fabricators lack. The added data may contain internal inconsistencies, impossible combinations of settings, or identifiers that don't correspond to real devices.

Software signatures: Deep fake generators often leave their own fingerprints in the metadata. The image may be tagged as having been created by specific AI software, or it may bear telltale characteristics of particular image processing pipelines.

Mismatched thumbnails: If the fabricator modifies an authentic image to create the fake, the original thumbnail may remain embedded in the file, revealing the true original content.

Impossible timestamps: The metadata may show creation dates that are impossible or inconsistent with the purported events depicted.

This doesn't mean a sophisticated fabricator can't create convincing metadata. With sufficient effort and expertise, they can. But the effort required is substantial, and most fabricators don't bother—or make mistakes when they try. The absence or inconsistency of metadata thus provides a valuable first-line screen for identifying suspicious evidence.

Video and Audio Metadata

The principles that apply to photographs extend to video and audio evidence, though the specific metadata formats differ.

Video files contain extensive metadata about their creation and encoding. Authentic smartphone video will include information about the recording device, the codec used to encode the video, the frame rate and resolution, and often GPS coordinates and timestamps. Video files also contain structural metadata about how the video was encoded—information revealing whether the file is a direct camera output or has been re-encoded, which would be necessary to incorporate deep fake modifications.

The container format of a video file (MP4, MOV, AVI, etc.) carries its own metadata distinct from the video stream itself. Examining both the container metadata and the stream metadata can reveal inconsistencies suggesting manipulation.

Audio files similarly contain metadata about their recording. Common formats like WAV, MP3, and AAC include information about the recording device, sample rate, bit depth, and encoding parameters. Voice recordings from phones often include call metadata, carrier information, and timestamps verifiable against carrier records.

More sophisticated audio analysis can examine the acoustic characteristics of a recording—the background noise profile, room acoustics, and compression artifacts—to determine whether the audio is consistent with its purported recording environment.

The Digital Chain of Custody

Beyond creation metadata, digital evidence carries information about its subsequent handling. Every time a digital file is copied, modified, or transmitted, metadata may be created or modified. File system metadata records when files were created, accessed, and modified on particular storage devices. Email headers record the transmission path of attachments. Cloud storage services maintain access logs and version histories.

This handling metadata can establish—or undermine—the chain of custody for digital evidence. If a photograph is offered as having been taken on a particular date but the file system metadata shows it was created months later, the discrepancy demands explanation. If a video is claimed to be an unmodified original but the metadata shows it has been re-encoded, further investigation is warranted.

For attorneys, understanding this metadata ecosystem is essential. The metadata is there, embedded in the files and recorded in the systems that handled them. We simply need to know how to find it and what it means.

Part III: Discovery Strategies for Digital Authenticity

The Fundamental Principle: Ask for the Original

The single most important discovery strategy for authenticating digital evidence is deceptively simple: obtain the original file, in its original format, as it came from the source device.

This sounds obvious, but it's routinely neglected. Lawyers too often accept printed photographs, compressed video clips, or transcribed audio recordings without demanding the underlying digital files. In an age of synthetic media, that's malpractice-adjacent. The printed photograph has been stripped of its metadata. The compressed video clip may have been re-encoded in ways that destroy evidence of manipulation or impair resolution. The transcription tells you nothing about whether the original audio was authentic.

Under Texas Rule of Civil Procedure 196.4, parties may request production of electronic data in the form in which it is ordinarily maintained or in a reasonably usable form. For authentication purposes, specify that you want the data in its original format—the format created by the recording device, prior to any conversion, compression, or processing.

Your production request should seek:

The original digital file in its native format (JPEG, HEIC, MP4, MOV, WAV, etc.)—not converted or re-encoded.

All application metadata associated with the file, including EXIF data, container metadata, and any embedded information.

File system metadata from the device or storage medium on which the file was originally stored, including creation dates, modification dates, and access dates.

Transmission records if the file was sent via email, text message, or other electronic communication, including complete message headers and any cloud storage records.

The source device or forensic image thereof, if the authenticity of the evidence is genuinely contested and the stakes warrant it.

Interrogatories Targeting Authenticity

Interrogatories provide an opportunity to lock down the opposing party's authentication claims before the metadata speaks for itself. Consider:

For each photograph, video, or audio recording you intend to offer as evidence, state:

(a) The make and model of the device used to create the recording;

(b) The date, time, and location of creation;

(c) The identity of the person who created the recording;

(d) Whether the recording has been edited, modified, enhanced, or altered in any way since its original creation, and if so, describe each modification;

(e) The complete chain of custody of the recording from creation to the present, including every device and storage medium on which the recording has been stored;

(f) Whether any artificial intelligence, deep learning, or synthetic media generation tools were used to create or modify the recording.

These interrogatories serve multiple purposes. They establish a baseline account against which the metadata can be compared. They force the opposing party to commit to specific claims that can later be challenged. And they create a record for impeachment if the evidence is ultimately shown to be fabricated.

Requests for Admission

Texas Rule of Civil Procedure 198 permits requests for admission, which can be powerful tools for challenging suspicious evidence:

Admit that the photograph marked as Exhibit A was created by a digital camera or smartphone, not by artificial intelligence software.

Admit that the photograph marked as Exhibit A has not been modified, altered, or edited since its original creation.

Admit that the EXIF data embedded in Exhibit A accurately reflects the circumstances of its creation.

Admit that you do not possess the original device used to create Exhibit A.

The beauty of requests for admission is that they demand a definitive response. The opposing party must admit, deny, or explain why they can't do either. A failure to respond results in deemed admission. A false response has serious consequences.

If the opposing party admits the evidence was created by a camera and not by AI, they're locked into that position. When your expert examines the metadata and finds telltale signs of synthetic generation, you have a devastating impeachment.

Deposing the Custodian and Creator

Deposition of the person who allegedly created digital evidence, or who serves as its custodian, is essential when authenticity is contested. Your deposition outline should cover:

The alleged circumstances of creation: Where were you when you took this photograph? What device did you use? What were you doing? When and why did you take it? Was anyone else present? Were other photos taken at the same time? These questions establish a narrative that can be compared against the metadata.

The device used: What is the make and model? When did you acquire it? Where is the device now? Can you produce it? Is it associated with cloud accounts that might have backup copies? Have you factory reset it since the alleged creation date?

Technical knowledge probing: Have you ever used image editing software? Photoshop, GIMP, or similar tools? Are you familiar with AI image generation tools like Midjourney, DALL-E, or Stable Diffusion? Have you ever used face-swapping or filter applications?

The chain of custody: After creating this recording, what did you do with it? Did you send it to anyone? Upload it anywhere? Where has it been stored? Have you made copies? Edited or modified it in any way?

Confrontation with metadata: If you've already obtained and analyzed the metadata, confront the deponent with any inconsistencies. "You testified that you took this photograph on June 15th at your home in Houston. The metadata shows it was created on June 20th in Austin. Explain that." Let them try.

Forensic Examination of Devices

In high-stakes cases where authenticity is genuinely contested, you may need forensic examination of the source device—a complete, bit-by-bit forensic acquisition of the device's storage.

Forensic acquisition is powerful because it can recover:

Deleted files: Photos and videos deleted from a device are often recoverable from unallocated storage space, though the ability to do so is quite limited on modern phones. If the opposing party claims an image is original but forensic examination reveals an earlier, different version, you have strong evidence of fabrication.

Editing history: Some devices and applications maintain logs of editing operations. Forensic examination can reveal what modifications were made and when.

Application data: If the device user installed deep fake applications or AI image generators, that fact will likely be visible in the application data.

Browser and AI history: Most deepfakes are now created using online tools and large language models. That activity may appear in browser history, cache, and AI application history.

Timeline reconstruction: File system metadata, combined with other device activity, can establish a detailed timeline of when files were created, modified, and accessed.

Under Texas law, obtaining forensic access to an opposing party's device requires demonstrating specific, credible evidence that the device contains relevant information not obtainable through less intrusive means. *In re Weekley Homes, L.P.*, 295 S.W.3d 309 (Tex. 2009), established the framework for court-ordered forensic examination in Texas. The requesting party must show that

the responding party has somehow defaulted on its discovery obligations, such as by failing to search for relevant evidence or producing documents inconsistent with its claims.

If you've already obtained metadata showing anomalies suggestive of fabrication, that evidence can support a motion for forensic examination. The inconsistencies themselves may constitute the "specific, credible evidence" that *Weekley Homes* requires. Consider seeking a neutral forensic examiner to balance the need to know against the right to safeguard privileged and confidential content.

Third-Party Discovery

Don't neglect third-party discovery. Digital evidence often touches third-party systems that maintain independent records:

Cloud storage providers: Apple iCloud, Google Photos, Dropbox, and similar services maintain extensive metadata about stored files, including upload dates, modification histories, and access logs. Subpoenas to these providers can yield valuable authentication evidence.

AI service providers: As many deepfakes are now fabricated using large language model services, the AI service provider may have a record of the fabrication.

Social media platforms: If the evidence or related files were posted to social media, the platforms maintain records of upload times, source IP addresses, and sometimes original file metadata.

Communication providers: If the evidence was transmitted via email or messaging, those providers may have records of the transmission, including original file hashes.

Telecommunications carriers: For audio recordings, carriers maintain call records that can verify whether calls occurred when and where claimed.

Surveillance systems: If the evidence purportedly depicts events at a particular location, independent surveillance footage from that location can corroborate or contradict the evidence.

These third-party records are powerful because they're created and maintained by neutral parties with no stake in the litigation. Inconsistency between the offered evidence and third-party records is strong evidence of fabrication.

Part IV: The Tells of Synthetic Media—What to Look For

Proving Physical Impossibility

While metadata provides the most reliable indicators of authenticity, visual and auditory examination of the evidence itself can reveal telltale signs of synthetic generation. The fundamental principle: AI systems learn to mimic the *appearance* of real media, but they don't understand the

three-dimensional physical world that real media depicts. Deepfakes often contain artifacts that are physically impossible in the real world.

Your eye has spent a lifetime learning the physics of reality. Somewhere in your visual cortex, you understand how light falls on faces, how fabric drapes on bodies, how reflections work in eyes, and how shadows correspond to light sources. Deep fake detection, at its core, is about systematically applying that intuitive physical knowledge to suspicious media.

Facial Anomalies

Faces are the most common subject of deepfakes, and they're where synthetic media most often fails.

Unnatural blinking: Early deepfakes notoriously failed to reproduce natural blinking patterns because training data often consisted of photographs—images in which subjects have their eyes open—rather than video. Modern systems have largely corrected this, but blinking anomalies remain worth examining.

Eye reflections: This is a powerful tell. When a person is photographed, light sources appear as reflections in their eyes. In an authentic photograph, these reflections will be consistent across both eyes—the same light sources, in the same positions, with the same appearance. Deepfakes often fail to maintain this consistency. The reflections in the left eye may differ from those in the right, or be absent entirely, or depict impossible light source geometries.

Examine the eyes closely, under magnification if necessary. Compare the reflections in each eye. Ask yourself: are these reflections consistent with a single, coherent lighting environment?

Facial boundaries: Where the face meets the hair, ears, and neck, deepfakes often exhibit subtle artifacts. The boundary may be unnaturally sharp or unnaturally blurred. There may be color discontinuities or texture mismatches. The face may appear to “float” slightly above the underlying head, with a visible seam.

Asymmetry failures: Human faces are naturally asymmetric—the left side differs slightly from the right. Deepfakes often produce faces that are too symmetric, or with asymmetries in the wrong places or of the wrong character.

Expression inconsistencies: Authentic facial expressions involve coordinated movement of dozens of muscles. A genuine smile involves not just the mouth but the eyes, cheeks, and forehead. Deepfakes sometimes produce expressions inconsistent across different facial regions—a smiling mouth paired with neutral eyes, for instance.

Teeth anomalies: Teeth are surprisingly difficult for AI to render convincingly. Examine them for impossible geometries, unnatural coloring, or inconsistent numbers. Missing teeth, extra teeth, and teeth that shift position between frames are all warning signs.

Temporal Inconsistencies in Video

Video deepfakes must maintain consistency not just within a single frame but across thousands of frames. This is technically demanding, and temporal inconsistencies are common tells.

Flickering: Watch the face closely across multiple frames. In deepfakes, the face may flicker or shimmer subtly, particularly around boundaries, resulting from frame-to-frame inconsistencies in the synthetic generation process.

Unnatural motion: The face's motion should correspond naturally to the head and body. In deepfakes, the face may seem to slide or drift relative to the underlying head, or fail to track properly with head rotation.

FORENSIC TELLS: SPOTTING FAKE MEDIA

AUTHENTIC FRAME vs **SUSPECTED DEEP FAKE**

Two images may look equally persuasive. Only one arrives with a traceable evidentiary lineage.

VISUAL ARTIFACTS
Look for unnatural reflections, lighting and shadows, ear blurring, facial asymmetry, unnatural neck blends, irregular hand poses, and image compression/hallucination.

AUDIO ARTIFACTS
Be alert to monotone delivery, lifeless expressions, off-sync lip movements, unnatural pauses and intonation, and some incorrect background noises.

METADATA INCONSISTENCIES
Inspect internal dates, times, GPS locations, device data; check for alterations, gaps, or lack of expected timestamps and identifiers.

HASH MISMATCH
SHA-256 Hash 1: c910cf48...
SHA-256 Hash 1: c6b27fb...
DO NOT MATCH!
For baseline authenticity, hash the received version and compare it against the original/counterpart provided natively.

DEEP FAKE FOOTPRINTS
Look for production tags in metadata from software like diffusion models "GAN", or names of known synthetic media tools, and "explainer" images typical of deep fake datasets.

LACKS PROVENANCE
Scrutinize transfers for discontinuities, such as a file that's a screenshot, or shows irregular transfer, editing, and export history! native source.

Lighting inconsistencies: As a subject moves through a scene, lighting on their face should change consistently with the environment. Deepfakes may fail this test, with the face appearing to have its own independent lighting that doesn't correspond to the scene.

Audio-visual synchronization: In videos with speech, lip movements should correspond precisely to the audio. Deepfakes often exhibit subtle sync errors—lips that move slightly before or after the corresponding sound, or mouth shapes that don't match the phonemes being spoken.

Environmental Artifacts

The environment surrounding a deep faked subject often contains tells:

Inconsistent shadows: Every object in a scene casts shadows consistent with the light sources. Deepfakes may introduce faces or figures whose shadows are inconsistent—pointing in the wrong direction, having the wrong intensity, or being absent entirely.

Reflection failures: Beyond eye reflections, consider reflections in windows, mirrors, water, polished floors. A synthetic figure may fail to appear in reflections where they should, or appear differently than expected.

Perspective errors: Synthetic elements may not conform to the perspective geometry of the scene. A face may be slightly too large or too small for the body, or oriented at an angle inconsistent with the figure's posture.

Edge artifacts: At the boundaries of manipulated regions, look for halos, color fringing, or unusual blurring resulting from the compositing process that blends synthetic elements into authentic backgrounds.

Audio Deep Fake Tells

Audio deepfakes—synthetic voice cloning—have their own characteristic artifacts:

Unnatural prosody: Human speech has natural rhythms of emphasis, pitch variation, and pacing. Synthetic speech often has a subtly mechanical quality, with prosody too regular or failing to correspond naturally to the emotional content.

Breathing anomalies: Real people breathe. Their speech is punctuated by inhalations, their sentences constrained by lung capacity. Synthetic speech may lack natural breathing patterns or insert breaths at unnatural locations.

Background inconsistencies: The acoustic environment—room tone, background noise, reverberation—should be consistent throughout. If the voice was synthetically inserted into an authentic recording, there may be subtle mismatches.

Clipping and artifacts: Synthetic audio generation can produce subtle digital artifacts—clicks, pops, or moments of distortion—not characteristic of natural speech.

Vocabulary and phrasing: A softer tell, but a person's speech patterns, vocabulary, and habitual phrases are as distinctive as their voice. A synthetic recording may capture the voice accurately but use words or constructions the purported speaker would never naturally use.

The Limitations of Visual Detection

A word of caution: the visual and auditory tells I've described are useful but not definitive. Deep fake technology improves rapidly. Artifacts obvious in early deepfakes have been largely eliminated in current systems. What constitutes a reliable tell today may be undetectable tomorrow.

Moreover, human perception is fallible. We see what we expect to see. If a juror believes an image is authentic, they may overlook anomalies. If they believe it's fake, they may perceive artifacts where none exist. Even when they learn an image is fake, they can have trouble dismissing it in deliberations.

For these reasons, visual examination should complement, not replace, metadata analysis and forensic examination. The metadata provides objective, verifiable evidence of authenticity or fabrication. The visual tells provide supporting evidence that helps explain to a jury why the metadata matters.

Part V: Authentication Under Texas Law

The Rule 901 Framework

Texas Rule of Evidence 901 establishes the authentication requirement: the proponent must produce evidence sufficient to support a finding that the item is what the proponent claims it is. This is a relatively low bar—the proponent need not conclusively prove authenticity, only provide sufficient evidence from which a reasonable juror could find the evidence authentic.

Rule 901(b) provides an illustrative list of authentication methods, several relevant to digital evidence:

Testimony of a witness with knowledge (Rule 901(b)(1)): A witness who perceived the events depicted can testify that the evidence accurately depicts those events. “I was there when this happened, and this video accurately shows what I saw.”

Distinctive characteristics (Rule 901(b)(4)): Evidence can be authenticated by its appearance, contents, substance, or internal patterns, taken in conjunction with circumstances. Metadata can serve as such a distinctive characteristic: the EXIF data in a photograph is an internal pattern that, taken together with testimony about the circumstances of creation, supports authenticity.

Evidence about a process or system (Rule 901(b)(9)): For evidence produced by a technological process, authentication can be established by showing that the process produces accurate results. This is relevant to demonstrating that EXIF data accurately reflects capture parameters or that forensic acquisition tools reliably preserve digital evidence.

The Sufficiency Dispute: Laying the Predicate

When you offer digital evidence, opposing counsel may object that you haven’t laid a sufficient authentication foundation. When you challenge opposing evidence, you may argue the proponent hasn’t met their burden.

In *Tienda v. State*, 358 S.W.3d 633 (Tex. Crim. App. 2012), the Court of Criminal Appeals considered social media authentication and articulated a totality-of-the-circumstances approach. Authentication doesn’t require conclusive proof, just sufficient evidence from which a jury could reasonably conclude that the evidence is what its proponent claims.

Under *Tienda*, relevant circumstances for digital evidence authentication include: (1) testimony from someone with personal knowledge of the evidence’s creation; (2) metadata or other

embedded information indicating origin; (3) distinctive content identifying the creator or confirming the depicted events; and (4) circumstances foreclosing the possibility of fabrication.

For deep fake challenges:

When offering evidence: Establish a clear chain of custody from the recording device to trial. Provide testimony from the person who created the recording. Offer metadata as corroboration. Present any forensic examination.

When challenging evidence: Attack the weakest links. If the proponent can't produce the original device, note that. If metadata is missing or inconsistent, highlight those deficiencies. If visual examination reveals artifacts suggestive of manipulation, present expert testimony explaining them.

The Best Evidence Rule and Digital Originals

Texas Rule of Evidence 1001 defines "original" for electronically stored information as "any printout—or other output readable by sight—if it accurately reflects the information." This somewhat awkward formulation means that a printout of a digital photograph can qualify as an "original" for best evidence purposes.

However, for *authentication* purposes as distinct from best evidence, the original digital file is far more valuable than any printout. The printout lacks metadata. It can't be examined for synthetic generation artifacts at the pixel level. It has lost information that can't be recovered.

When you demand original digital files in discovery, you're making an authentication argument, not a best evidence argument: you need the original digital file because it contains information necessary to evaluate whether the evidence is what the proponent claims. That information is simply not present in any derivative.

Part V-A: Federal Court Practice—A Companion Guide

Many Texas lawyers practice in federal court or handle matters that may land there. While Texas and federal evidence rules share common ancestry, federal practice offers some distinct tools.

Federal Rule of Evidence 901 and the Conditional Relevance Framework

Federal Rule of Evidence 901 mirrors its Texas counterpart: the proponent must produce "evidence sufficient to support a finding that the item is what the proponent claims it is." Federal courts apply the same conditional relevance framework under Rule 104(b)—the judge determines whether a reasonable jury could find the evidence authentic, and if so, the jury makes the ultimate determination.

Federal courts have generally followed the same totality-of-the-circumstances approach to digital evidence authentication. The leading federal articulation appears in *United States v. Vayner*, 769

F.3d 125 (2d Cir. 2014), emphasizing that authentication requires evidence that the item is what its proponent claims—not merely that it appears genuine on its face.

Self-Authentication Under Rules 902(13) and 902(14)

Federal Rules of Evidence 902(13) and 902(14), added in 2017, provide a streamlined authentication pathway for certain electronic evidence. These rules allow self-authentication of electronic records when accompanied by a certification that the records were generated by a reliable electronic process (Rule 902(13)) or that their integrity can be verified through a hash value or similar digital identification process (Rule 902(14)).

For deep fake challenges, these rules cut both ways. A proponent may invoke Rule 902(14) to authenticate a recording by showing its hash value matches one created at capture—powerful evidence of integrity if the original hash was reliably created. Conversely, the absence of such verification, when it could have been provided, may support an authenticity challenge.

The Federal Judicial Center’s guidance on these rules emphasizes their application to evidence whose integrity depends on the reliability of electronic systems—precisely the territory deep fake challenges occupy.³

Federal Discovery: Rules 26, 34, and 37(e)

Federal discovery rules provide robust tools for obtaining digital evidence in native format.

Under **Federal Rule of Civil Procedure 34(b)(2)(E)**, a party must produce ESI in the form requested, or if none is specified, in the form ordinarily maintained or in a reasonably usable form. Specify native format production to preserve metadata.⁴

Rule 26(f) requires parties to confer about ESI issues early in the case. Use this conference to establish protocols for digital media evidence, including native format production and metadata preservation requirements.

Rule 37(e) addresses ESI spoliation. If a party fails to preserve digital evidence that should have been preserved, the court may order proportionate remedial measures or, if the party acted with intent to deprive, may presume lost information was unfavorable, instruct the jury accordingly, or enter default judgment. When challenging evidence authenticity, consider whether spoliation of metadata or original files has occurred and whether Rule 37(e) remedies are available.

³ Paul W. Grimm et al., Federal Judicial Center, *Authenticating Digital Evidence* 4–10 (2017) (explaining authentication of electronic evidence through proof of reliable systems, hash values, and certifications under Fed. R. Evid. 902(13) and 902(14)).

⁴ For more on forms of production: See Ball, *Forms That Function* (2026); http://www.craigball.com/Ball_Forms_That_Function_2026.pdf.

Federal Expert Testimony: Daubert

Federal courts apply *Daubert v. Merrell Dow Pharmaceuticals, Inc.*, 509 U.S. 579 (1993), to expert testimony, including testimony on digital forensics and synthetic media detection. While *Daubert* factors (testing, peer review, error rates, general acceptance) overlap substantially with the Texas *Robinson* framework,⁵ federal courts sometimes apply them with greater rigor.

For deep fake detection testimony, be prepared to establish: (1) the expert’s detection methodology has been tested and validated; (2) the methodology has been subjected to peer review; (3) the known or potential error rate; and (4) general acceptance in the relevant scientific community. The field of synthetic media detection is evolving rapidly, and not all detection methods have the established track record *Daubert* favors.

Expert Testimony on Authentication

Texas Rule of Evidence 702 governs expert testimony admissibility. An expert may testify if specialized knowledge will help the trier of fact understand the evidence or determine a fact in issue.

Digital forensic experts provide crucial testimony on deep fake authentication:

Metadata examination: A forensic expert can examine the metadata and testify about what it reveals regarding the file’s origins, handling, and integrity—explain missing or inconsistent metadata and compare it to the proponent’s claims.

Forensic acquisition and analysis: An expert can describe the forensic device examination process, testify about findings, and establish the reliability of forensic tools.

Deep fake detection: Experts in synthetic media detection can examine questioned evidence and testify about whether it exhibits characteristics consistent with deep fake generation, in terms a jury can understand.

General technical education: An expert can provide jurors the background needed to evaluate the evidence—explaining metadata, how deepfakes are created, and why certain artifacts are inconsistent with authentic media.

For deep fake challenges, expert testimony is almost always necessary. The technical concepts are beyond common knowledge, and jurors need guidance. If you anticipate challenging digital evidence as potentially fabricated, budget for expert assistance.

⁵ E.I. du Pont de Nemours and Co., Inc. v. Robinson, 923 S.W.2d 549 (Tex. 1995).

Beyond Authentication: Hearsay, Completeness, and Prejudice

A brief but important reminder: authentication is necessary but not sufficient for admissibility. Even authenticated evidence faces other hurdles.

Hearsay: A video offered to prove the truth of statements made in it is hearsay, subject to the usual exceptions and exclusions. Authentication establishes the recording is genuine; it doesn't make the statements admissible.

Completeness: Texas Rule of Evidence 107 (and federal Rule 106) may require admission of additional recording portions when fairness demands context. A selectively edited clip, even if authentic, may be excludable.

Prejudice: Under Texas Rule 403 (and Federal Rule 403), even relevant, authenticated evidence may be excluded if its probative value is substantially outweighed by the danger of unfair prejudice, confusion, or waste of time.

These doctrines interact with deep fake challenges in important ways. A fabricated recording is not merely inauthentic; it's also potentially excludable as unfairly prejudicial and misleading. When challenging suspicious evidence, deploy the full evidentiary toolkit.

Challenging Admissibility Versus Challenging Weight

An important tactical distinction: you can challenge digital evidence at the admissibility stage under Rule 901, or allow it in and challenge its weight before the jury.

Each approach has tradeoffs. Challenging admissibility, if successful, keeps the evidence from the jury entirely. But the authentication threshold under Rule 901 is low, and judges are often reluctant to exclude facially relevant evidence on authentication grounds. A failed admissibility challenge may telegraph your concerns, giving opposing counsel time to shore up their authentication.

Challenging weight allows the evidence in but attacks its credibility before the jury. This approach lets you present your full case for fabrication—the missing metadata, the visual anomalies, the expert testimony—to the fact-finder. Texas juries are capable of evaluating competing claims about evidence authenticity, and a well-presented challenge can be devastating.

In most cases, I recommend a hybrid approach: raise authentication objections to preserve the record and educate the court, but prepare primarily for a weight challenge before the jury. The admissibility objection may succeed—particularly if the metadata evidence is stark—but the weight challenge gives you a second bite at the apple if it doesn't.

Part VI: Building Your Deep Fake Challenge—A Practical Checklist

Phase One: Identification and Preservation

When you first suspect evidence may be fabricated, your immediate priorities are identification and preservation.

Identify all versions: Determine what versions exist. Is there an “original” file? Has the evidence been produced in multiple formats? Has it appeared in third-party contexts—social media, news coverage, other litigation?

Preserve everything: Send a preservation demand immediately.⁶ The opposing party has a duty to preserve relevant evidence once litigation is reasonably anticipated, but fabricators have powerful incentives to destroy evidence of fabrication. A clear, documented preservation demand strengthens later spoliation arguments.

Document your receipt: When you receive the evidence, document exactly what you received, when, and in what format. Create forensic hash values (MD5 or SHA-256) for every digital file. These hash values serve as fingerprints revealing any subsequent modification.

Secure chain of custody: From the moment you receive the evidence, maintain a clear chain of custody. Store files securely. Document every access. Use only working copies for examination. This protects against claims that anomalies you discover resulted from your handling rather than fabrication.

Phase Two: Metadata Extraction and Analysis

With the evidence preserved, proceed to metadata examination.

Extract all metadata: Use forensic tools to extract every piece of metadata from the digital files. ExifTool⁷ is a powerful and free command-line utility for extracting metadata from image and video files. Commercial forensic tools like EnCase, X-Ways Forensics, or FTK provide more comprehensive capabilities.

Document your extraction: Record the tool used, its version, and extraction parameters. This documentation supports reliability of your process under Rule 702. And again—*use only working copies*.

⁶ For more on preservation demands, see Ball, *The Perfect Preservation Letter* (2020); http://www.craigball.com/Perfect_Preservation_Letter_Guide_2020.pdf.

⁷ ExifTool by Phil Harvey, <https://exiftool.org> (last visited Feb. 26, 2026).

Analyze for consistency: Compare extracted metadata to the proponent’s claims. Does device identification match the claimed source? Does the timestamp match the claimed date? Do GPS coordinates match the claimed location?

Look for anomalies: Is EXIF data missing where it should be present? Is there evidence of image editing software? Are capture parameters plausible for the depicted scene?

Research device signatures: Different devices produce characteristic metadata signatures. A photograph from an iPhone has different metadata characteristics—and even a different storage format—than one from a Samsung Galaxy. A synthetic image differs from either. Research expected characteristics for the claimed source device and compare.

Phase Three: Visual and Technical Examination

While metadata examination proceeds, conduct or commission visual examination.

High-resolution examination: Obtain the highest-resolution version available and examine under magnification. Compression artifacts and synthetic generation artifacts are often more visible at close examination.

Facial analysis: If the evidence depicts faces, conduct the facial examination described earlier—eye reflections, facial boundaries, asymmetry, expression consistency, teeth.

Environmental analysis: Examine the broader scene for shadow consistency, reflection accuracy, perspective correctness.

Temporal analysis: For video, examine frame-by-frame. Look for flickering, temporal inconsistencies, and audio-visual synchronization errors.

Consider AI detection tools: Several software tools claim to detect deepfakes using AI. These can serve as screening mechanisms, but they aren’t infallible—they have both false positive and false negative rates. Never rely solely on an automated detection tool; use it as one input among many.

Phase Four: Corroboration and Contradiction

Digital evidence doesn’t drop like manna from Heaven, nor does it exist in isolation. Investigate the broader evidentiary context.

Independent verification: Can the depicted events be independently verified? Other witnesses? Other recordings? Physical evidence?

Contradiction search: What evidence might contradict the offered evidence? If it purportedly shows the plaintiff at a particular location, do cell phone records, credit card receipts, or testimony place them elsewhere?

Device investigation: Does the allegedly-used device exist? Can it be examined? Does its forensic state corroborate or contradict authentication claims?

Digital forensics: Consider whether forensic examination of the proponent’s devices might reveal deleted evidence, deep fake software, or other indicators of fabrication.

Phase Five: Expert Engagement

For any serious deep fake challenge, expert assistance is essential.

Forensic examiner: Engage a qualified, certified digital forensics expert to conduct or validate your metadata examination and provide testimony.

Deep fake specialist: Consider engaging an expert specifically skilled in synthetic media detection. Academic researchers and specialists can provide cutting-edge analysis and compelling testimony.

Prepare for *Daubert/Robinson*: Under Texas Rule of Evidence 702 and *E.I. du Pont de Nemours & Co. v. Robinson*, 923 S.W.2d 549 (Tex. 1995), expert testimony must be based on reliable principles and methods reliably applied. Prepare your expert for reliability challenges and ensure their methodology is documented and defensible.

Phase Six: Presentation

Plan how you’ll present your challenge to the court and jury.

Demonstratives: Technical evidence is most effective when presented visually. Prepare demonstratives showing metadata anomalies, highlighting visual artifacts, and explaining technical concepts in accessible terms.

Narrative: Develop a narrative explaining why the evidence is fabricated and who fabricated it. Juries find technical evidence more persuasive within a comprehensible human story.

Anticipate rebuttal: Consider how the proponent will respond and prepare counter-arguments. They may claim metadata was stripped by innocent transmission. They may claim visual anomalies result from compression. Have answers ready.

Part VII: Emerging Authenticity Technologies

The Content Authenticity Initiative

The deep fake problem hasn’t escaped the technology industry’s notice. Several initiatives are underway to build authenticity verification into digital media from the moment of creation.

The Content Authenticity Initiative, led by Adobe in partnership with camera manufacturers, news organizations, and technology companies, is developing a system called Content Credentials. Under this system, cameras and editing software cryptographically sign media at each stage of creation

and modification. The resulting credential—embedded in the file as metadata—provides a verifiable chain of provenance from camera capture through editing to final publication.

If this technology achieves widespread adoption, authentication becomes simpler. Evidence carrying valid Content Credentials can be traced to its source; evidence lacking credentials will be immediately suspect. But adoption isn't yet universal, and retrofitting billions of existing devices and software tools will take years.

C2PA and the Provenance Standard

The Coalition for Content Provenance and Authenticity (C2PA) has developed technical standards for content provenance that multiple manufacturers are adopting. Under the C2PA specification, digital media carries cryptographically signed “manifests” recording its creation and modification history.

Apple, Google, Nikon, Sony, and other major manufacturers have announced C2PA support. As this technology proliferates, expect provenance information to become a routine component of digital evidence.

For practitioners, the advent of provenance technology creates new discovery opportunities. If the claimed source device supports Content Credentials or C2PA, demand the provenance manifest. If the manifest is missing, ask why. If it's present, verify its cryptographic signatures.

Blockchain and Distributed Verification

Some systems use blockchain technology to create immutable records of media provenance. When media is created, a cryptographic hash is recorded on a distributed ledger. Any subsequent modification changes the hash, making tampering detectable.

These systems aren't yet common but are gaining traction where authenticity is paramount—journalism, legal evidence, insurance documentation.

The Computational Detection Arms Race

AI tools designed to detect deepfakes are locked in an ongoing arms race with tools designed to create them. Detection systems learn to identify artifacts; generation systems learn to eliminate them; detection systems adapt; the cycle continues.

Current detection tools are useful but imperfect. They should be employed as one component of a comprehensive authentication strategy, not as definitive arbiters. A negative result doesn't prove authenticity; a positive result doesn't conclusively prove fabrication.

Metadata remains one of the most valuable authentication tools—but it's a heuristic, not a rule. Unlike visual artifacts that can be eliminated by improvements in generation technology, metadata anomalies are harder to eliminate because authentic metadata requires an authentic source. A

sophisticated fabricator can strip metadata, alter it, or attempt to transplant it from an authentic source—but each manipulation creates its own forensic artifacts and inconsistencies. Stripped metadata raises questions; altered metadata often contains internal inconsistencies; transplanted metadata rarely survives close scrutiny against the purported source device. The absence of a coherent metadata story isn't proof of fabrication, but it's a powerful indicator demanding explanation.

Part VIII: Ethical Obligations and Practical Considerations

The Duty of Competence

Texas Disciplinary Rule of Professional Conduct 1.01 requires competent representation, including the legal knowledge, skill, thoroughness, and preparation reasonably necessary for the representation. Comment 8 to ABA Model Rule 1.1 (on which the Texas rule is based) specifically notes that competence includes understanding “the benefits and risks associated with relevant technology.”

In an age of deepfakes, competent representation requires awareness of synthetic media technology and the ability to either challenge suspicious evidence or engage appropriate expert assistance. Lawyers who remain ignorant of these issues do so at their clients' peril—and their own.

This doesn't mean every lawyer must become a digital forensics expert. It means every lawyer must know *enough* to recognize when deep fake issues may be present, what questions to ask, what discovery to pursue, and when expert consultation is necessary. This article is a step toward that competence, but only a step.

Candor to the Tribunal

The corollary to challenging fabricated evidence is the duty not to offer it. Texas Disciplinary Rule 3.03 prohibits making false statements of fact to a tribunal and offering evidence the lawyer knows to be false.

If you discover that evidence supporting your client's case is fabricated, you have an ethical obligation not to offer it. If you've already offered evidence and subsequently discover its falsity, you must take reasonable remedial measures, including disclosure to the tribunal if necessary.

This obligation extends to situations where you don't know the evidence is false but have serious doubts. The rules require knowledge of falsity before they prohibit offering evidence, but prudence and professionalism counsel against presenting evidence you suspect is fabricated. At minimum, investigate before offering.

The Danger of False Accusations

A word of caution: accusing an opponent of fabricating evidence is a serious matter. If you're wrong, you risk sanctions, professional embarrassment, and damage to your client's case—not to mention the unwarranted harm to your opponent. The mere existence of anomalies in digital evidence doesn't prove fabrication; there may be innocent explanations.

Pursue deep fake challenges when the evidence warrants suspicion, but do so responsibly. Base your challenge on documented technical evidence, not speculation. Engage qualified experts. Consider alternative explanations. Present your challenge as raising questions that require answers, not as foregone conclusions.

And if investigation demonstrates that initially suspicious evidence is authentic, accept that conclusion. Persistence in a baseless fabrication claim is itself a form of advocacy misconduct.

The Jury's Common Sense

Never underestimate the jury. Jurors may not understand EXIF data or generative adversarial networks, but they understand when something doesn't add up. They understand that photographs come from cameras and carry digital fingerprints. They understand that authentic evidence has a history and fabricated evidence doesn't.

Your job is to translate the technical evidence into terms jurors can grasp and relate to their common sense. "This photograph has no birth certificate. It appeared out of nowhere, from no device, with no record of its creation. The metadata that every authentic photograph carries is simply not here. Why not?"

If you've done your technical homework, the jury's common sense will do much of the heavy lifting.

Part IX: Case Studies and Practical Examples

Case Study 1: The Custody Video

Consider a hypothetical custody dispute in which Mother produces a video appearing to show Father striking the child during a visitation exchange. The video is compelling, the children's welfare is at stake, and Father insists—truthfully—that the events never occurred.

How should Father's counsel proceed?

Discovery: Demand the original video file in native format (likely .MOV for iPhone). Demand all metadata. Demand Mother's iPhone for forensic examination by a neutral examiner or, alternatively, a forensic image of the device. Propound interrogatories asking Mother to identify the precise device, date, time, and circumstances of recording, and whether any editing occurred.

Metadata examination: Does the file contain expected QuickTime metadata for an iPhone recording? Does the creation timestamp correspond to the alleged incident? Do GPS coordinates match the exchange location? Does the device identifier correspond to Mother’s actual iPhone?

Visual examination: Does Father’s face exhibit tells associated with face-swapping or reenactment? Are there temporal inconsistencies, lighting anomalies, or edge artifacts? Is audio properly synchronized with lip movements?

Corroboration: Were there other witnesses? Independent recordings—security cameras, dashcam footage, bystander videos? Do Father’s text messages or calls around that time suggest any disturbance?

Expert engagement: Engage a digital forensics expert for metadata examination. If analysis suggests fabrication, engage a synthetic media detection expert.

Case Study 2: The Workplace Recording

In an employment discrimination case, Plaintiff produces an audio recording purportedly capturing racist comments by her supervisor during a one-on-one meeting. The supervisor denies making any such statements and suggests voice cloning.

Discovery: Demand the original audio file, all metadata, and production or forensic imaging of Plaintiff’s phone. Demand call logs and recording application logs. Propound interrogatories identifying the recording application, date, time, and meeting circumstances.

Metadata examination: Does the audio file carry expected smartphone recording metadata? Does the timestamp correspond to a time when the alleged meeting could have occurred? Can calendar records, badge access logs, or other business records verify both were present?

Audio analysis: Engage an audio forensic expert. Is the acoustic environment consistent throughout, or are there discontinuities suggesting splicing? Does the supervisor’s voice exhibit characteristics of synthesis? Is background room tone consistent with the alleged location?

Voice comparison: Obtain authenticated samples of the supervisor’s voice—voicemails, recorded meetings, deposition testimony. Engage an expert to compare acoustic characteristics and speech patterns.

Corroboration: Can anyone verify the meeting occurred? Are there calendar entries, emails, or other documentation? Is there independent evidence of discriminatory animus that would make the alleged statements plausible?

Case Study 3: The Surveillance Video

Defendant produces surveillance video showing Plaintiff engaged in physical activities inconsistent with claimed disabilities. Plaintiff insists the video must have been manipulated.

Discovery: Demand original video files, all metadata, the investigator’s work records and chain of custody documentation, and the recording equipment for examination.

Investigator deposition: What equipment was used? Technical specifications? How is footage stored and transferred? What processing was applied? Can the investigator identify Plaintiff with certainty?

Metadata examination: Do timestamps correspond to claimed surveillance dates? Is device information consistent with claimed equipment? Has the video been re-encoded inconsistently with the investigator’s account?

Visual examination: Is Plaintiff’s face consistent throughout? Are there edge artifacts, lighting inconsistencies, or temporal anomalies? Is camera perspective consistent with claimed recording circumstances?

Geolocation verification: Can surveillance locations be independently verified? Do environmental details—signage, vehicles, weather—correspond to claimed locations and dates?

Part X: Looking Forward—The Evolving Landscape

The Technology Will Get Better

Let me be blunt: the deep fake technology we face today isn’t the technology we’ll face tomorrow. The artifacts I’ve described—facial boundary anomalies, reflection inconsistencies, temporal flickering—are artifacts of *current* technology. Future systems will reduce or eliminate these tells.

The rate of improvement has been startling. What was obviously fake two years ago is convincing today. What is detectable today may be undetectable in six months. We’re in an arms race between generation and detection, and generation has the structural advantage: generators only need to eliminate the tells that current detectors look for, while detectors must anticipate what’s coming next.

This is why metadata remains the most reliable authentication tool. Metadata can’t be authentically fabricated because authentic metadata requires an authentic source. As visual tells become less reliable, the importance of metadata examination will only increase.

The Courts Will Adapt

The legal system isn’t built for rapid technological change, but it does adapt. Courts are beginning to recognize the deep fake challenge and develop responsive doctrines.

Expect developments in several areas:

Authentication standards: Courts will likely demand more robust authentication foundations for digital media, particularly in high-stakes cases. The minimal foundation that suffices for uncontested evidence won't cut it when fabrication is credibly alleged.

Presumptions and burdens: Courts may develop presumptions related to metadata. Evidence carrying complete, consistent metadata might benefit from an authenticity presumption; evidence lacking expected metadata might face heightened scrutiny.

Expert evidence: As deep fake challenges become more common, courts will develop clearer standards for expert testimony on synthetic media detection.

Discovery practices: Courts may impose enhanced disclosure obligations for digital evidence, requiring production of original files with metadata rather than derivatives.

Rules amendments: The challenge of synthetic media hasn't gone unnoticed by the rulemakers. Thoughtful scholars and jurists—including my good friends Judge Paul Grimm (retired, U.S. District Court for the District of Maryland) and Professor Maura Grossman (University of Waterloo and Osgoode Hall Law School)—have been at the forefront of efforts to consider whether amendments to the Federal Rules of Evidence are needed to address AI-generated evidence. While formal rule changes take time, the conversation is underway. Practitioners should monitor developments from the Advisory Committee on Evidence Rules. In the meantime, the existing rules—properly understood and vigorously applied—provide substantial tools for challenging synthetic evidence.

The Practitioner's Imperative

For practitioners, the imperative is clear: education and preparation. The attorneys who understand these issues today will be better equipped to serve their clients tomorrow.

Familiarize yourself with metadata and how to extract it. Learn to use basic forensic tools like ExifTool. Develop relationships with qualified, certified forensic examiners. Follow the evolving research on synthetic media detection.

When you encounter digital evidence, cultivate healthy skepticism. Don't assume authenticity simply because the evidence *looks* real. Ask the foundational questions: Where did this come from? What device created it? What does the metadata show? What should be here that isn't?

And when you suspect fabrication, pursue that suspicion rigorously. The discovery tools are available. The authentication challenges are cognizable. The experts exist. The only question is whether you have the knowledge and the will to deploy them.

Conclusion: The Metadata Is Not a Lie

We stand at an inflection point in the history of evidence. The assumption that photographs, videos, and recordings depict reality—an assumption that has undergirded evidentiary practice for a

century and a half—is no longer safe. The technology to fabricate convincing synthetic media is accessible to anyone with a computer and an internet connection. Deepfakes are here, they're improving, and they're turning up in litigation.

But this is not cause for despair. It is cause for adaptation. The same digital technology that enables fabrication also enables authentication. Every authentic photograph carries within it the fingerprints of its creation. Every legitimate video bears the metadata of its recording. Every genuine audio file has a provenance that can be traced.

The synthetic imposter typically has no such lineage. It springs into existence fully formed—as if from Zeus' forehead—from no camera, at no location, bearing no authentic record of its creation. When we know what to look for—and now we do—the absence of that lineage is as revealing as any physical tell at the poker table. This is not an absolute rule; sophisticated fabricators may attempt to spoof provenance, and legitimate evidence may lose its metadata through innocent handling. But the probability calculus favors the prepared advocate: *authentic evidence usually has a coherent origin story, and fabricated evidence usually does not.*

The tools for challenging deepfakes exist: metadata examination, forensic analysis, expert testimony, robust discovery practice, detection tools. What remains is for the bar to develop the knowledge and skill to deploy them effectively.

This article has endeavored to provide a foundation for that knowledge. We've examined what deepfakes are and how they work. We've explored the metadata that accompanies genuine digital evidence and developed discovery strategies for obtaining it. We've catalogued the visual and auditory tells that betray synthetic media and reviewed the authentication framework under both Texas and federal law. And we've developed practical approaches to building and presenting deep fake challenges.

The knowledge in these pages will require updating as technology evolves. But the fundamental principles will endure. Authentic digital evidence usually has a provenance; fabricated evidence usually doesn't—or has a fabricated provenance riddled with inconsistencies. Authentic evidence can ordinarily be traced to its source; fabricated evidence typically can't withstand that scrutiny. Metadata isn't infallible—it can be stripped, altered, or spoofed—but the effort required to create a convincing false provenance is substantial, and most fabricators fail the test.

In a world where seeing is no longer believing, metadata and forensic rigor are our anchors to truth. Learn to read them, demand them in discovery, and deploy them in advocacy. Your clients—and the integrity of our system of justice—are counting on you.

Craig Ball is a Texas attorney, law professor, certified computer forensic examiner, and electronic evidence expert based in Austin and New Orleans. He has served as the court-appointed special

master in electronic discovery disputes and has published extensively on forensic examination, electronic discovery, and digital evidence authentication. He can be reached at craig@ball.net.