

Use this checklist before any significant visual asset goes to market. **Score 1 point for each YES.** Photographic and AI-generated content now face the same credibility environment - audiences are skeptical of both until proven otherwise. This checklist helps you demonstrate provenance, not just assert it. From *Guardrails: How To Embrace AI in Your Business Without Damaging Your Brand* by Craig McDonogh.

SECTION 1 — ORIGIN: DO YOU KNOW WHAT THIS IMAGE ACTUALLY IS?

Apply to every visual asset before production sign-off

1

Can you state clearly whether this asset is: (a) a real photograph, (b) AI-assisted photography, (c) AI-generated imagery, or (d) a composite?

FOUNDATION CHECK

If you can't answer this without checking, the provenance trail has already broken down. The answer needs to be documented, not just known.

2

If the asset was photographed, do you have the original unedited file with intact EXIF metadata (camera model, timestamp, GPS if applicable)?

PROVENANCE

Metadata is the chain of custody for a real photograph. If it was stripped at any point in the production pipeline, the provenance claim is harder to defend. AI-generated images: N/A, score 1.

3

If AI tools were used at any stage - generation, editing, upscaling, background replacement, face smoothing - is that use documented with the tool name and the type of intervention?

AUDIT TRAIL

A photograph that had its background swapped by AI is not a pure photograph. The intervention needs to be logged, even if it doesn't require disclosure in this specific context.

SECTION 2 — PEOPLE: WHO IS IN THIS IMAGE, AND DO THEY KNOW?

Apply when the asset features real or AI-generated human likenesses

4

If the asset features real people, do you have signed model releases that explicitly cover AI-assisted production, digital alteration, and the specific channels where the asset will run?

LEGAL REQUIRED

Standard model releases predate AI-assisted production and may not cover it. Check the language - 'likeness' alone may not be sufficient for AI-processed images.

5

If the asset features AI-generated human likenesses, are they clearly non-identifiable as any real person, and is their synthetic nature disclosed in markets where this is required?

LEGAL REQUIRED

New York's Synthetic Performers Disclosure Law (in force June 9, 2026) requires conspicuous disclosure for AI-generated human performers in advertising served to New York audiences. Other states are following.

6

If the asset uses a real person's voice, likeness, or style to imply endorsement, has legal reviewed it for right of publicity and FTC endorsement compliance?

LEGAL REQUIRED

AI tools can now produce convincing voice and likeness approximations of public figures. Right of publicity law applies regardless of whether a human or an AI created the content.

SECTION 3 — DISCLOSURE: HAVE YOU TOLD AUDIENCES WHAT THEY NEED TO KNOW?

Apply based on the asset type and the channel where it will run

7	Does your channel strategy include appropriate AI disclosure language for assets where disclosure is required by law, platform policy, or your own governance standards?	COMPLIANCE
	Required by law: EU AI Act Article 50 (from 2 Aug 2026), New York Synthetic Performers Law, FTC endorsement guidelines. Required by platform: Cannes Lions entries, Meta political ads. Required by your standards: any AI content you've publicly committed to disclosing.	
8	For AI-generated or AI-assisted creative, does your disclosure language distinguish between 'AI-assisted' (human-led, AI-supported) and 'AI-generated' (AI-led output)?	CLARITY
	The distinction matters to audiences and regulators. 'AI-assisted' and 'AI-generated' describe meaningfully different production processes and carry different levels of audience scrutiny. Using them interchangeably erodes the credibility of both.	
9	If the asset will run in the EU or UK, does it comply with Article 50 of the EU AI Act - including machine-readable watermarking for AI-generated content that could mislead the public?	LEGAL REQUIRED
	Article 50 transparency obligations are in force from 2 August 2026. AI-generated image and video content on matters of public interest requires watermarking. Content created before 2 August has a grace period to 2 December 2026.	

SECTION 4 — RESILIENCE: WOULD THIS ASSET SURVIVE SCRUTINY?

Apply before any asset that could attract media or public attention

10	If this asset were questioned publicly - 'is this AI?' - do you have the documentation to answer clearly and quickly, without needing to investigate first?	CRISIS READY
	The McConnell photograph required a Drexel professor and a doctoral student to confirm it was real. You shouldn't need that for a product shot. If you'd have to check before answering, the documentation is inadequate.	
11	Does the asset carry authenticity markers that survive compression and platform processing - real location specificity, genuine human imperfections, verifiable context?	AUTHENTICITY
	AI detection tools are unreliable and inconsistent. The more durable protection is assets that feel visibly real: specific locations, non-idealized faces, contextual detail that a photographer captures naturally and an AI prompt struggles to specify.	
12	Has your agency or production partner confirmed, in writing, whether and how AI was used in producing this asset, and do you have that confirmation on file?	SUPPLY CHAIN
	Cannes Lions now requires this for award entries. Your legal team may need it if a right of publicity or copyright claim arises. Your client may ask. Get it in writing before the campaign runs, not after a question arrives.	

SECTION 5 — PAID MEDIA: DOES YOUR MEDIA PLAN REFLECT THESE REQUIREMENTS?

Apply at media planning stage, not after creative is locked

13	Have you confirmed the AI content disclosure requirements for each platform where this asset will run as paid media?	PLATFORM RULES
	Meta, Google, YouTube, TikTok, and X all have evolving AI disclosure requirements for paid advertising, particularly for political content and realistic synthetic media. Requirements change faster than media plans. Check at briefing, not at trafficking.	

14	If this campaign targets audiences in multiple jurisdictions, have you mapped the disclosure requirements by market rather than applying a single global standard?	MULTI-MARKET
	New York, the EU, and California have different rules that don't align neatly. A single disclosure label designed for one market may be inadequate or incorrect in another. Multi-market campaigns need a per-market disclosure matrix.	
15	Is there a named person accountable for monitoring AI disclosure compliance across this campaign's run, with authority to pull or amend assets mid-flight?	GOVERNANCE
	Regulations and platform policies change. A campaign that was compliant at launch may require amendment during its run. Someone needs to own that watch, with the budget authority to act without a three-week approval cycle.	

SCORING GUIDE

Score	What it means	Recommended action
12–15	Strong provenance posture	Maintain and document as standard practice
8–11	Moderate - gaps exist	Address missing disclosure and provenance steps before next campaign launch
4–7	Significant exposure	Audit current pipeline and implement disclosure policy before next major release
0–3	High risk	Stop and build a visual content governance policy before further AI-assisted production

QUICK REFERENCE — WHEN DISCLOSURE IS LEGALLY REQUIRED

Context	Requirement	In force
AI chatbots / virtual assistants (EU users)	Disclose AI nature to users	2 Aug 2026
AI-generated content (EU, public interest)	Machine-readable watermark / label	2 Aug / 2 Dec 2026
Synthetic performers in ads (New York)	Conspicuous on-screen disclosure	9 Jun 2026
Deceased performer digital replicas (NY)	Prior consent from estate required	In force
AI in political advertising (26 US states)	Disclosure or prohibition varies by state	Varies
Cannes Lions award entries	Mandatory AI use declaration	2026 festival
FTC endorsement / testimonial claims	Material connection disclosure if AI-generated	In force