Image AI Detection Report

AI likely

TrustOriginality.ai — Forensic authenticity analysis

Run: 3f2a9c1e-8b4d-4e2a-9f12-6c7d8e9f0a1b

2026-06-16 14:18:33 UTC · Ref: Campaign asset — hero banner

product-hero-banner.png (2400×1350)

87% AI-generated likelihood score

0% = human-origin signal · 100% = AI-origin signal (probabilistic, not legal proof)

Frequency-domain artefacts and local noise residuals consistent with diffusion upscaling. No valid C2PA Content Credentials signature on file.

Forensic evidence

Detector	Confidence	Summary
forensic_score	87%	Weighted ensemble across CNN + FFT cues
frequency_grid	82%	Periodic high-frequency grid in luminance
noise_residual	76%	Non-camera sensor noise profile
c2pa_provenance	absent	No signed manifest attached to file
consistency_audit	warning	Declared human photo; forensics disagree

Blockchain anchor — immutability record

```
Network: Polygon PoS
Tx: 0x8f3c2d1e0a9b8c7d6e5f4a3b2c1d0e9f8a7b6c5d4e3f2a1b0c9d8e7f6a5b4c3d2
Anchor date: 2026-06-16 (daily Merkle root)
Content hash: sha256:7c4e2f8a1b9d3e6c0f5a8b2e1d4c7f9a3b6e0d2c5f8a1b4e7d0c3f6a9b2e5d8
IPFS CID (full report): bafybeigdyrzt5v7w9x1y3z5a7c9e1g3i5k7m9o1q3s5u7w9y1a3c5e7g9i1k3
```

Analysis summary is anchored on-chain; detailed PDF is stored on IPFS. Third parties can verify integrity without accessing raw content.

Certificate key: TO-CERT-20260616-3F2A9C1E8B4D-6C7D8E9F0A1B

TrustOriginality

Verify this report

Scan the QR code or open the link below to confirm this report was issued by TrustOriginality.ai and has not been altered. The link is cryptographically signed.

https://panel.trustoriginality.ai/verify/r/3f2a9c1e8b4d4e2a9f126c7d8e9f0a1b?e=1781612313&s=2f9a8c1e4d7b0e3f6a9c2d5b8e1f4a7

Attestation: https://panel.trustoriginality.ai/api/attestation/analysis/3f2a9c1e-8b4d-4e2a-9f12-6c7d8e9f0a1b.json

SAMPLE DOCUMENT — Illustrative data for customer preview. Production reports are generated per analysis with live signed verification links. See page 2 for the technical appendix.

Certificate key: TO-CERT-20260616-3F2A9C1E8B4D-6C7D8E9F0A1B

Technical Appendix — Image Analysis

Run 3f2a9c1e-8b4d-4e2a-9f12-6c7d8e9f0a1b · TrustOriginality.ai detection engine v1

1. Analysis pipeline

The submitted asset is analysed by an independent detector ensemble. Each detector emits an Evidence record (type, confidence 0–1, weight, metrics JSON, optional samples). Detectors run sequentially on the same byte stream; no single detector can override the ensemble alone. After forensic detectors complete, the Verify Suite cross-check layer evaluates contradictions between declared provenance (C2PA, watermarks, registry) and forensic signals.

2. Ensemble scoring

The headline AI-likelihood score is a weighted mean of detector confidences:

$$\text{score} = \Sigma(\text{confidence}_i \times \text{weight}_i) \,/\, \Sigma(\text{weight}_i)$$

For this sample run the aggregate is 0.87 (87%). The default decision threshold is 0.60: scores ≥ 0.60 are labelled “AI likely”. Weights are higher for cryptographically verifiable signals (C2PA manifest declaring generative origin: weight 2.0) and the optional ONNX SigLIP classifier (weight 1.3). Heuristic forensic detectors use weights 0.7–0.8.

3. Detector reference

Detector	Wt.	Method & output
classifier (ONNX)	1.3	SigLIP-based binary model (224×224, optional). Outputs ai_probability in metrics.
c2pa_provenance	2.0	JUMBF/COSE manifest parse: claim generator, digital source type, COSE signature validity.
watermark_signature	1.0	Metadata scan for generator strings (Stable Diffusion, Midjourney, DALL·E, SynthID, IPTC trainedAlgorithmicMedia).
trustoriginality_watermark	1.2	Invisible watermark decode + provenance registry lookup by content hash.
metadata	1.0	EXIF/IPTC/XMP tag extraction; flags AI-related keywords in embedded metadata.
ela	0.8	Error-level analysis proxy: JPEG recompress at Q=90, mean/max per-channel residual.
compression	0.7	8×8 block-boundary gradient — blockiness score for recompression artefacts.
consistency_audit	—	Verify Suite rule engine: manifest/watermark declarations vs forensic score.

4. Sample run — per-detector notes

- forensic_score (0.87): ensemble aggregate; not a separate model.
- frequency_grid (0.82): 2-D FFT magnitude peaks at diffusion upscaler lattice frequencies.
- noise_residual (0.76): PRNU-style residual deviates from camera sensor models.
- c2pa_provenance (absent): no JUMBF manifest store; signature not evaluated.
- consistency_audit (warning): declaration implied human capture; forensic 0.87 contradicts — FORENSIC_VS_DECLARATION.

5. Integrity, attestation & verification

Certificate key: TO-CERT-20260616-3F2A9C1E8B4D-6C7D8E9F0A1B

Content hash: SHA-256 over raw file bytes. The daily Merkle root (Polygon PoS) commits hash, score, verdict and run ID. Full PDF and evidence JSON pinned to IPFS; CID in anchor payload. Report URL: HMAC-signed, 1-year TTL. Attestation: trustoriginality-attestation/1.0 JSON, Ed25519-signed (key ID at /.well-known/trustoriginality-attestation.json).

6. Interpretation & limitations

Probabilistic forensic opinion, not legal proof. Adversarial post-processing, compression or novel generators may reduce sensitivity. C2PA absence does not prove AI origin; valid generative manifest is a strong positive signal. Human review recommended for high-stakes decisions.

SAMPLE DOCUMENT — Technical appendix aligned with production EvidenceBase / PdfBuilders schema.

Certificate key: **TO-CERT-20260616-3F2A9C1E8B4D-6C7D8E9F0A1B**