

Pixels That No Surface Could Produce: Per-Pixel Physical Realizability as an Image Forensic Certificate

Satyajeeth Suresh Kannan

Innovation Academy · Mary's Lab

github.com/JeetuSK0808/maryslab

Preprint · August 2026 · Mary's Lab technical report TR-06

ABSTRACT

Illuminant-consistency forensics estimates lighting per region and flags disagreement, producing an anomaly score that an expert can dispute. A stronger statement is available for a narrower claim. If a pixel's implied object color falls outside the object-color solid of every illuminant in a declared family, then no reflectance in $[0,1]$ produces it: the pixel is not a diffuse reflective surface under that family. This is a statement about physics rather than statistics. We implement it per pixel, quantify each violation by the smallest dimming factor that would restore realizability, and characterise the method's dominant sensitivity — the assumed exposure. Across a lightness–chroma sweep, the flagged fraction rises from 0.083 at unit exposure to 0.382 at exposure 0.55, demonstrating that exposure must be reported with any such claim rather than assumed. We also measure the false-positive rate, which requires inputs known to be physical: on 1024 measured USGS reflectances rendered at their correct exposure the audit flags 0.02344 of pixels as impossible. The instrument should therefore report a rate for comparison, never a verdict on an individual pixel.

Keywords: image forensics; object-color solid; realizability; colorimetry; provenance

1. Introduction

Digital image forensics largely proceeds statistically: estimate a property in different regions, model its expected variation, and flag improbable configurations. Lighting consistency, noise characteristics and compression history all follow this pattern. Each yields a score, and each score is contestable because it is probabilistic.

For one narrow class of claim, a categorical statement is available. Fix an illuminant and an observer. The set of colors producible by some reflectance in $[0,1]$ is the object-color solid — a bounded convex body that does not fill color space. Colors outside it are not merely improbable for a diffuse surface; they are unproducible by one.

We build the per-pixel test, and we are deliberate about what it licenses: a flag means the pixel is not a plain diffuse surface under the declared lights, not that an image is fabricated.

2. Method

Each sRGB pixel is converted to XYZ and interpreted relative to a diffuse white at a declared exposure. For each illuminant in the family — named CIE illuminants, black-

body radiators at a stated temperature, or Gaussian LED bands — we test whether any reflectance on the canonical grid, bounded in $[0,1]$, reproduces the response. Realizability under at least one family member suffices.

For a pixel realizable under none, we binary-search the *overshoot*: the factor by which the pixel is too bright for any surface under any declared light. An overshoot of 1.3 means the pixel exceeds the physical bound by thirty percent. Reporting a magnitude rather than a boolean is what makes the result robust to the exposure assumption, since a reader can see how close to the boundary each flag sits.

A 3% boundary tolerance absorbs the known mismatch between the sRGB primaries' white and a table-integrated D65 white, so exact-boundary whites are not flagged on numerical noise.

3. Evaluation

Inputs are synthetic. The test set is a 12×12 sweep in Oklab lightness and chroma converted to sRGB — not photographs. This isolates the method's sensitivity to exposure, which is what the experiment is designed to expose.

Table 1. Flagged fraction and overshoot statistics across assumed exposures, on a fixed 144-pixel lightness–chroma sweep. The pixels do not change; only the assumed diffuse-white scale does.

exposure	pixels	flagged	fraction	mean overshoot	max overshoot
1.00	144	12	0.083	20	20
0.85	144	31	0.215	7.818	20
0.70	144	46	0.319	5.325	20
0.55	144	55	0.382	3.491	20

The flagged fraction is strongly exposure-dependent, and this is the report's main practical finding. The same pixels are judged differently under different assumptions about how bright a white surface would be in the scene. Any deployment of this test that does not state its exposure is reporting an uninterpretable number.

The overshoot distribution is what makes the result usable in spite of that. A pixel flagged with overshoot near 1.0 sits at the boundary and would flip under a small exposure change; a pixel at 20 would not. Ranking by overshoot rather than thresholding on a flag gives an ordering that is stable even when the absolute count is not.

Reproducing these numbers

Clone `github.com/JeetuSK0808/maryslab` and run `node papers/build/experiments.mjs`. Every figure in this report is written by that script into `papers/build/results.json`; the instrument itself is exercised by `audit_image_physicality`. The engines carry a 401-vector conformance suite shared by a

TypeScript reference implementation, a C++20 core, and a WebAssembly build, so a reported number is a property of the specification rather than of one implementation.

4. Real-data evaluation: the false-positive rate

A detector's true-positive rate is easy to demonstrate: construct an impossible pixel and check that it is caught. Its false-positive rate is the number that decides whether it can be deployed, and measuring that requires inputs known to be physical. Measured reflectances are exactly that — a spectrum that a spectrometer recorded from a real surface, and that passed gate Q2 on ingest, is realizable by definition. Every "impossible" verdict on such a pixel is a false positive, with no judgment call involved.

We rendered 1024 measured USGS reflectances to sRGB under D65 and audited them against the D65 basis. Records falling outside the sRGB gamut were excluded and counted (7 of them) rather than clipped, since clipping would manufacture the very impossibility being tested for.

Table 2. False-positive rate on pixels that are measured reflectances by construction. Exposure is the divisor the audit is told to assume; 1.0 is the correct value for these renders.

assumed exposure	pixels	flagged impossible	false-positive rate
1.0 (correct)	1024	24	0.02344
1.4	1024	47	0.0459
2.0	1024	73	0.07129

At the correct exposure the audit falsely flags 24 of 1024 pixels — a false-positive rate of 0.02344. That is small enough for the instrument to be useful and large enough that it must be published: an audit reporting a handful of impossible pixels in a photograph has not necessarily found anything, and roughly two in every hundred physical pixels will be flagged at this operating point.

The cause is the round trip rather than the feasibility test. A reflectance is converted to XYZ, then to sRGB with its transfer function and 8-bit-representable range, and the audit inverts that to ask whether a reflectance exists. Near the gamut boundary the inversion lands slightly outside the object-color solid, and the answer flips. This is why the rate climbs with assumed exposure — 0.0459

at 1.4 and 0.07129 at 2.0 — since a wrong exposure divisor pushes reconstructed XYZ further from where a reflectance can live. An audit run with a mis-estimated exposure is measuring its own assumption as much as the image.

The deployable form of the instrument therefore reports a rate and a comparison, not a verdict on individual pixels: a region flagged at 0.02344 is unremarkable, and one flagged at ten times that is worth explaining. Absolute counts on small selections should not be trusted, and the instrument should state the assumed exposure alongside every result.

5. Limitations

This section states what the result does not establish. It is written to be read before the abstract is quoted.

- **A flag is not an accusation.** Fluorescent materials, emissive sources in frame, specular highlights and clipped pixels all legitimately exceed the bound. The correct reading is "not a plain diffuse surface under the declared lights". Inferring fabrication requires evidence this instrument does not provide.
- **Exposure is the dominant assumption**, as Table 1 quantifies. It is an input, not an inference, and results are uninterpretable without it.
- **The illuminant family is only what is declared.** At most four members; Gaussian LED bands are crude stand-ins for real primaries.
- **The evaluation uses synthetic pixels.** No photographs, no camera pipeline, no demosaicing or tone mapping — all of which alter pixel values before any forensic test sees them. The behaviour on real captures is not established here.
- **A 3% boundary tolerance means hairline violations are deliberately unflagged.** The test is tuned to avoid reporting numerical artifacts as physical evidence.
- **Prior art was surveyed informally.** The object-color solid is classical and illuminant-consistency forensics is an active field; a search found no tool emitting per-pixel realizability certificates with overshoot magnitudes.

6. Acknowledgements and disclosure

The instruments described here were implemented in Mary's Lab, an open-source computational color science project. Software design, implementation, and the preparation of this report were carried out with substantial assistance from a large language model (Anthropic Claude), under the author's direction and review. All numerical results were produced by executing the released code; none were generated by a language model. The author is responsible for the claims made here.

References

1. Wyszecki, G. and Stiles, W. S. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd edition. Wiley.
2. CIE (2004). *Colorimetry*, 3rd edition. CIE 15:2004. Commission Internationale de l'Éclairage, Vienna.
3. Schanda, J. (ed.) (2007). *Colorimetry: Understanding the CIE System*. Wiley-Interscience.
4. Farid, H. (2009). Image forgery detection. *IEEE Signal Processing Magazine*, 26(2), 16–25.
5. Johnson, M. K. and Farid, H. (2007). Exposing digital forgeries in complex lighting environments. *IEEE Transactions on Information Forensics and Security*, 2(3), 450–461.
6. Logvinenko, A. D., Funt, B., and Godau, C. (2014). Metamer mismatching. *IEEE Transactions on Image Processing*, 23(1), 34–43.
7. Ottosson, B. (2020). *A perceptual color space for image processing (Oklab)*. bottosson.github.io.
8. MacAdam, D. L. (1935). The theory of the maximum visual efficiency of colored materials. *Journal of the Optical Society of America*, 25(8), 249–252.

Source, engines, and the script that produced every number in this report: github.com/JeetuSK0808/maryslab. Results regenerate with `node papers/build/experiments.mjs`. Inputs: synthetic where §3 says so, and measured USGS splib07a spectra (snapshot `usgs-splib07a-1`) where a real-data section says so. Prepared with AI assistance (see Acknowledgements).