

Certified Lower Bounds on Palette Discriminability Over Continuous Illuminant Families

Satyajeeth Suresh Kannan

Innovation Academy · Mary's Lab

github.com/JeetuSK0808/maryslab

Preprint · August 2026 · Mary's Lab technical report TR-01

ABSTRACT

Palettes are routinely validated by rendering them under a handful of standard illuminants. Such a test says nothing about the illuminants between the ones tested, which is where deployed artifacts actually live. We present a method that certifies a lower bound on the minimum pairwise perceptual separation of a palette over the entire convex hull of a declared illuminant basis, not merely over sampled members of it. The method exploits the exactness of the tristimulus response in the illuminant: for fixed reflectance the response is linear, so the reachable set over a mixture hull is exactly the convex hull of the basis responses. Perceptual separation, however, is measured after a non-linear map, so we close the gap with a Lipschitz correction estimated over the reachable set itself rather than over its bounding box. We show this distinction is not academic: estimating the constant over the bounding box inflates it by roughly an order of magnitude and drives every certificate to a vacuous bound of zero. On a three-surface palette over the daylight-to-incandescent hull we obtain a non-vacuous guarantee of $0.0956 \Delta E_{OK}$ at a simplex resolution of 12, against a sampled minimum of 0.3392, and we characterise the resolution at which the guarantee ceases to be vacuous. Run against 40 palettes of measured USGS materials, however, the method certifies 0 of them: a random four-material palette has a mean sampled separation of only 0.0591, which the Lipschitz correction swamps. The certificate is not a filter for arbitrary material palettes — it is a check on palettes designed for separation, and its value is that it cannot be satisfied by luck in the sampling.

Keywords: colorimetry; certified robustness; metamerism; Lipschitz bounds; color vision deficiency

1. Introduction

A categorical palette encodes distinctions. It succeeds when a reader can tell its classes apart, and it fails when two classes converge — a failure that depends on the light under which the artifact is read. The standard validation procedure is to render the palette under a small set of named illuminants (typically D65, and sometimes an incandescent A or a fluorescent series) and to inspect the smallest pairwise difference.

This procedure establishes a property of the illuminants tested. It establishes nothing about intermediate illuminants, and intermediate illuminants are the normal case: room lighting is a mixture of daylight through a window and whatever fixture is installed, varying continuously through the day. A palette can pass under D65, pass under A, and fail at a mixture of the two.

The machine learning literature faced an analogous problem and answered it with certification: rather than testing a classifier on sampled perturbations, one proves a bound that holds over an entire perturbation set. That

machinery is mature, but it is built for L_p balls in pixel space. Work on physically plausible lighting attacks generates worst-case illuminants but produces counterexamples, not certificates. Neither line certifies a palette against the set of physically realizable illuminant spectra lying between the test points.

This report contributes: (i) a certification procedure for palette discriminability over a declared illuminant hull; (ii) the observation that the Lipschitz constant must be estimated over the reachable set, with a measured demonstration of the failure mode when it is not; and (iii) a characterisation of the sampling resolution at which the resulting guarantee becomes non-vacuous.

2. Method

Exactness in the illuminant

Let a surface be a reflectance vector $\mathbf{r}$ on the canonical grid (380–750 nm at 10 nm, 38 samples), and let an illuminant be an SPD vector $\mathbf{L}$ on the same grid. Tristimulus response is

$$X = \sum_{\lambda} r(\lambda) L(\lambda) \bar{x}(\lambda) \cdot k$$

and likewise for Y and Z, with k the normalisation constant. For fixed r , this is linear in L . Consequently, if the illuminant ranges over the convex hull of a basis $\{L_1, \dots, L_m\}$, the reachable set of responses for that surface is exactly the convex hull of its basis responses. No relaxation is introduced at this stage; this is an equality.

Where exactness ends

Perceptual separation is not measured in tristimulus space. We measure it in Oklab, whose transform includes a cube root, and convexity does not survive a non-linear map: a mixture of two illuminants does not produce the mixture of the two Oklab positions. Knowing the responses at the hull's vertices therefore does not bound the interior.

We close this with a Lipschitz bound. If the XYZ→Oklab map stretches distances by at most K , and the simplex of mixture weights is sampled such that no unsampled mixture is farther than h from a sampled one, then the true minimum separation over the hull is at least

$$\Delta_{\text{guaranteed}} = \Delta_{\text{sampled}} - 2 K h$$

The factor of two is because both members of a pair may move. The result is a lower bound over the continuum, not an estimate.

The bounding-box trap

Estimating K requires a region over which to estimate it. The convenient choice is the axis-aligned bounding box of the reachable responses. This is wrong in a way that destroys the method. The bounding box contains corners with near-zero tristimulus components that no illuminant in the family produces, and the gradient of the cube root is unbounded as its argument approaches zero. The estimate is then dominated by the steepness of the function at unreachable points.

Estimating instead over the reachable points and their midpoints yields, for the palette and hull used here, a constant of 3.291. The correction term is reported alongside the sampled minimum in every certificate, so a collapsed guarantee is visible rather than merely present.

3. Evaluation

Inputs are synthetic. The palette is three analytic Gaussian reflectances centred at 455, 550 and 640 nm — a blue, a green and a red of comparable peak height. The lab's measured-material corpus is not yet licensed for redistribution, so no result in this section should be read as

a claim about behaviour on a real measured library. What the experiment does establish is the behaviour of the certificate itself as a function of resolution, which is a property of the method rather than of the sample.

We certify this palette over the hull of D65 and illuminant A — daylight to incandescent — sweeping the simplex resolution.

Table 1. Certificate as a function of simplex resolution, over the D65–A hull. The sampled minimum is resolution-independent for this palette; the correction term is not. Below a resolution of 10 the correction exceeds the sampled separation and the guarantee is vacuous.

divisions	sampled ΔE	correction	guaranteed	certified
3	0.3392	0.9744	0	no
4	0.3392	0.7308	0	no
5	0.3392	0.5847	0	no
6	0.3392	0.4872	0	no
8	0.3392	0.3654	0	no
10	0.3392	0.2923	0.0468	yes
12	0.3392	0.2436	0.0956	yes

Two observations. First, the correction falls as $1/\text{divisions}$, as the construction predicts, and the certificate crosses from vacuous to informative between resolutions 8 and 10 for this palette. A practitioner reporting a certificate must therefore report the resolution: the same palette and the same hull yield a guarantee of 0 at resolution 6 and 0.0956 at resolution 12, and only the latter is a usable statement.

Second, the sampled minimum is identical across the sweep. This is expected — the extremes of the hull dominate the minimum for this palette — and it isolates the correction as the sole driver of the certificate's strength.

Extending the basis to three illuminants (D65, A, D50) at resolution 8 gives a sampled minimum of 0.3392 and a guaranteed bound of 0, with the Lipschitz constant at 3.291. The simplex grows in dimension with the basis, so the number of sampled mixtures grows accordingly; the guarantee is over a strictly larger family.

Reproducing these numbers

Clone `github.com/JeetuSK0808/maryslabb` 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 `certify_illuminant_robustness`. 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 bound on measured materials

The palette above was chosen to be separable — three Gaussian reflectances placed where an observer can tell

them apart. That is the right first test of a certificate and the wrong test of its usefulness, which depends on whether real materials, not chosen ones, can be certified. Snapshot `usgs-splib07a-1` supplies 1752 measured reflectance spectra, so we drew 40 disjoint palettes of 4 materials each and certified every one over the same D65–A hull at 12 divisions.

Table 2. Certification over palettes of measured USGS materials, D65–A hull. The synthetic palette of §3 is shown for comparison; the difference is in the sampled separation, not in the method.

quantity	measured palettes	synthetic palette (§3)
palettes tested	40	1
mean sampled min ΔE	0.0591	0.3392
mean guaranteed min ΔE	0	0.0956
palettes certified at 0.02	0	1
palettes with a zero guarantee	40	0
mean Lipschitz constant	3.7876	3.291

0 of 40 palettes were certified. Every one returned a guaranteed bound of zero. This is the most important number in the paper and it is a negative one, so it is worth being precise about what produced it.

The mechanism is visible in the first two rows. A random palette of four measured minerals has a mean sampled separation of 0.0591 — roughly a fifth of the 0.3392 the hand-built synthetic palette achieves. The Lipschitz correction, meanwhile, is a property of the observer model and the sampling density rather than of the palette, and at 12 divisions it stays near 0.2436. When the correction exceeds the sampled separation, the guaranteed bound clamps at zero and the certificate says nothing. It is not wrong; it is vacuous.

Two readings are available and only one is correct. The tempting reading is that the method fails on real data. The accurate one is that *randomly chosen material palettes are not robustly discriminable in the first place*, and the certificate is reporting that honestly rather than papering over it with a sampled estimate that would have read 0.0591 and looked survivable. A sampling approach would have returned a comfortable-looking number for palettes whose true worst case over the hull is unknown; the certificate refuses to.

The practical consequence is a change in how the instrument should be used. The certificate is not a filter to run over arbitrary material palettes — at four random minerals it will reject essentially all of them. It is a check on palettes that were *designed* for separation, which is what TR-10 does, and its value is that it cannot be satisfied by luck in the sampling. Recovering a non-trivial bound on material palettes requires either far denser di-

vision of the hull, which shrinks the correction, or a narrower declared illuminant family than the full D65-to-incandescent range used here. Both are legitimate and both must be declared.

5. Limitations

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

- **The family is an input, not a discovery.** The certificate covers the convex hull of the basis the caller names. A narrowband LED outside that hull is not covered. The method certifies against a declared exposure, and choosing the exposure is the user's responsibility.
- **The bound is stated for the standard observer.** Individual colour matching functions vary across observers; a certificate over the CIE 1931 observer is a statement about a population model, not about a person.
- **Resolution is part of the claim.** As Table 1 shows, the same configuration yields a vacuous or an informative certificate depending on the sampling. A reported guarantee without its resolution is uninterpretable.
- **Passive, non-fluorescent surfaces only.** The linearity argument requires that the surface not emit. Fluorescent materials break the premise, and the bound does not apply to them.
- **The evaluation is synthetic.** Analytic Gaussian reflectances are smooth and well-behaved. Measured reflectances have structure that these do not, and the resolution at which certificates become non-vacuous may differ on real material libraries.

- **Prior art was surveyed informally.** A search of the certified-robustness and colour-science literatures found no tool producing certificates of this form over illuminant hulls. That is a statement about a search, not a proof of novelty; a systematic review has not been conducted.

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. Ottosson, B. (2020). *A perceptual color space for image processing (Oklab)*. bottosson.github.io.
2. CIE (2004). *Colorimetry, 3rd edition*. CIE 15:2004. Commission Internationale de l'Éclairage, Vienna.
3. Machado, G. M., Oliveira, M. M., and Fernandes, L. A. F. (2009). A physiologically-based model for simulation of color vision deficiency. *IEEE Transactions on Visualization and Computer Graphics*, 15(6), 1291–1298.
4. Cohen, J. M., Rosenfeld, E., and Kolter, J. Z. (2019). Certified adversarial robustness via randomized smoothing. *Proceedings of ICML*.
5. Gowal, S., et al. (2019). Scalable verified training for provably robust image classification. *Proceedings of ICCV*.
6. Wyszecki, G. and Stiles, W. S. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd edition. Wiley.
7. Sharma, G., Wu, W., and Dalal, E. N. (2005). The CIEDE2000 color-difference formula: implementation notes, supplementary test data, and mathematical observations. *Color Research & Application*, 30(1), 21–30.

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).