

The Width of the Hidden Crowd: Measuring the Metamer Set of a Color Wavelength by Wavelength

Satyajeeth Suresh Kannan

Innovation Academy · Mary's Lab

github.com/JeetuSK0808/maryslab

Preprint · August 2026 · Mary's Lab technical report TR-07

ABSTRACT

A color measurement does not determine a spectrum. The set of reflectances producing a given tristimulus response under a given light is an infinite family — the metamer set — and it is a convex polytope: the intersection of a 35-dimensional affine subspace with the physical box $[0,1]$. Metamer sets have been characterised in the literature, but their practical consequence is usually stated qualitatively. We measure it. By sampling the polytope boundary and recording, at each of the 38 grid wavelengths, the extremal values any member attains, we obtain an envelope whose width is the amount of spectral freedom a color measurement leaves unconstrained. For a mid-gray the mean envelope width is 0.9581, and the width is strongly wavelength-dependent: narrowest at 570 nm (0.8833) and widest at 380 nm (1). The envelope is a direct picture of which parts of a spectrum a color measurement actually pins down, and we argue it should accompany any spectral reconstruction. On 120 measured USGS material colors the picture is tighter than the synthetic analysis implies: 43 of them have metamer sets from which no boundary member can be sampled at all, because real materials sit near the edge of the object-color solid where the freedom disappears. Whether a given color is nearly determined or badly underdetermined is not guessable from the color, which is the argument for computing the envelope.

Keywords: metamerism; convex polytopes; spectral reconstruction; null space; color measurement

1. Introduction

Software that converts a color into a spectrum is common: rendering pipelines need spectral inputs, pigment matching needs reflectances, and "spectral upsampling" is a recognised operation. Every such tool faces the same underdetermination. The map from reflectance to tristimulus is a linear functional applied three times, from a 38-dimensional space onto a 3-dimensional one. Its fibers are infinite.

Tools generally resolve this by returning a distinguished member — typically the smoothest — which is a defensible default and is not the truth. The user receives a single curve and no indication of how much of it was determined by the measurement and how much by the tie-breaking rule.

We propose that the resolution should be reported alongside the answer, and we compute it.

2. Method

Write a reflectance as a vector $\mathbf{r}$ with 38 components. Fix an illuminant and observer; the tristimulus response is $\mathbf{A}\mathbf{r}$ for a 3×38 matrix $\mathbf{A}$ whose rows are the color matching

functions weighted by the illuminant. The null space of $\mathbf{A}$ has dimension at least 35: any $\mathbf{n}$ with $\mathbf{A}\mathbf{n} = \mathbf{0}$ can be added to $\mathbf{r}$ without changing the color.

Physical realizability bounds this otherwise unbounded affine subspace: $0 \leq \mathbf{r}(\lambda) \leq 1$. The metamer set is therefore the intersection of an affine subspace with a box — a convex polytope. Its emptiness is decidable, and when it is empty the correct output is a structured failure rather than a nearest guess.

The envelope

We sample the polytope boundary in many directions and record, at each wavelength, the minimum and maximum value attained by any sampled member. The difference is the envelope width at that wavelength: the span of reflectance values consistent with the measured color. A narrow width means the measurement determines the spectrum there; a wide width means it does not.

3. Evaluation

Inputs are synthetic. Targets are computed from analytic reflectances so that a ground-truth member is known to exist. The envelope is a property of the target

color and the viewing condition, not of the particular surface used to generate it.

Table 1. Envelope width of the metamer set at four target colors under D65, from boundary sampling. Width is the span of reflectance values consistent with the measured color at that wavelength. "Pinned" counts how many of the 38 coordinates of the feasible center sit exactly on the [0,1] boundary.

target	samples	pinned	mean width	narrowest (nm)	widest (nm)
mid gray	48	0	0.9581	0.8833 @ 570	1 @ 380
olive	48	0	0.7845	0.4797 @ 460	0.9849 @ 380
muted blue	48	0	0.5575	0.3526 @ 580	0.7873 @ 720
saturated blue	0	5	0	—	—

The envelope is not uniform across the spectrum, and the pattern is interpretable. Width is narrowest inside the band where the observer is sensitive — 570 nm for the gray, 460 nm for the olive — because the measurement constrains the spectrum where the color matching functions carry weight. It opens out toward the ends of the range, where sensitivity approaches zero and large reflectance changes barely register. The envelope is thus a direct picture of what a colorimetric measurement can and cannot see.

Width also shrinks with saturation: the gray admits a mean span of 0.9581, the muted blue 0.5575. The last row is the limiting case and the most interesting one. For the saturated blue the sampler returns 0 boundary members, because 5 of the 38 coordinates of the feasible center are already pressed against the walls of the box. The metamer set has collapsed: there is essentially one physical reflectance producing that color under that light. We report this rather than substituting a target that behaves, because it is the practically important regime — for colors near the edge of what surfaces can do, the measurement nearly determines the spectrum, and reconstruction is nearly recovery.

Three practical consequences follow. A color match is not a material match: two manufacturers matching the same target may select different members, and the size of their possible divergence under a different light is

bounded by the crowd's width. Spectral reconstruction is selection, not recovery, and should report which member and on what principle — except near the gamut boundary, where the last row shows the distinction nearly vanishes. And metamer pairs make sharp test cases for any pipeline that treats color as intrinsic to a surface.

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 `sample_metamers`. 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: how much freedom real materials leave

The synthetic analysis chose three target colors and found envelope widths from 0.5575 to 0.9581, plus one degenerate case. Choosing targets is precisely what biases this measurement, since width depends strongly on where in the object-color solid the target sits. A measured corpus removes the choice: we sampled 120 reflectances from snapshot `usgs-splib07a-1` and computed the envelope for each.

Table 2. Metamer-set envelope width over measured material colors under D65, compared with the synthetic targets of §3.

quantity	measured corpus	synthetic mid gray
targets tested	120	1
with sampled members	75	1
collapsed (no members found)	43	0
solver infeasible	2	0
mean envelope width	0.539	0.9581
median envelope width	0.5645	—
p10 / p90 width	0.1584 / 0.9037	—
mean pinned coordinates (of 38)	3.14	0

The headline is the collapse count. 43 of 120 measured material colors — more than a third — have metamer sets from which the sampler recovers no boundary member at all. On average 3.14 of the 38 coordinates of the feasible center sit against the walls of the physical box, which is what blocks the null directions. The synthetic mid gray had none pinned and a mean width of 0.9581; the measured median is 0.5645, and the p10 is 0.1584.

Read plainly: for real materials, a color measurement constrains the spectrum far more tightly than the synthetic analysis suggested, and for a substantial minority it very nearly determines it. The reason is selection. A synthetic mid gray sits comfortably in the interior of the object-color solid, where there is room to move; real minerals are frequently saturated or dark, which puts them near the boundary, which is where the freedom disappears. The synthetic experiment was not wrong, it was sampling the easy part of the space.

This tightens the claim of §3 in one direction and weakens it in another. Spectral reconstruction from a color remains selection rather than recovery — the p90 width of 0.9037 shows there are real materials with large freedom. But the blanket statement that a color measurement leaves a spectrum badly underdetermined does not survive contact with a material library, where it is often nearly determined instead. Which regime a given target is in is not guessable from the color, and the envelope is how one finds out.

2 targets produced a feasibility failure. Since each target was computed from a reflectance that exists, every metamer set here is non-empty by construction and those failures are the solver's, not physics'. They are reported rather than dropped because a two-in-120 infeasibility rate on inputs known to be feasible is a fact about the implementation that a user of it should have.

5. Limitations

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

- **Boundary sampling is not exhaustive.** The envelope is computed from a finite set of sampled extreme members; the true polytope extent is at least as wide. The sample count is reported.
- **A zero-width row is a lower bound, not a proof of uniqueness.** The saturated-blue case returns no boundary members because random null directions are blocked by the pinned coordinates. That establishes that the set is small in the directions sampled; a di-

rected search could still find members this sampler missed.

- **Everything is conditional on the model.** The standard observer, the tabulated illuminant, the 38-sample grid, and the assumption of a passive non-fluorescent surface. Fluorescent materials are not described by a reflectance in $[0,1]$ at all and every bound here ceases to apply to them.
- **The standard observer is a population average.** Two spectra that match for it can visibly differ for an individual whose cone sensitivities sit away from the mean.
- **Grid resolution bounds the claim.** At 10 nm, narrowband structure is not represented, so the envelope cannot describe freedom that exists below that scale.
- **Prior art was surveyed informally.** Metamer set theory is well developed; a search found no tool reporting a per-wavelength envelope as a routine companion to reconstruction.

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. Logvinenko, A. D., Funt, B., and Godau, C. (2014). Metamer mismatching. *IEEE Transactions on Image Processing*, 23(1), 34–43.
2. Finlayson, G. D. and Morovic, P. (2005). Metamer sets. *Journal of the Optical Society of America A*, 22(5), 810–819.
3. Wyszecki, G. and Stiles, W. S. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd edition. Wiley.
4. CIE (2004). *Colorimetry*, 3rd edition. CIE 15:2004. Commission Internationale de l'Éclairage, Vienna.
5. Ottosson, B. (2020). *A perceptual color space for image processing (Oklab)*. bottosson.github.io.
6. Schanda, J. (ed.) (2007). *Colorimetry: Understanding the CIE System*. Wiley-Interscience.
7. Cohen, J. B. and Kappauf, W. E. (1982). Metameric color stimuli, fundamental metamers, and Wyszecki's metamer black. *American Journal of Psychology*, 95(4), 537–564.
8. Smits, B. (1999). An RGB-to-spectrum conversion for reflectances. *Journal of Graphics Tools*, 4(4), 11–22.

Source, engines, and the script that produced every number in this report: github.com/JeetuSK0808/maryslabb. 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).