

Bracketing the Best Possible Palette: Covering Bounds on Categorical Discriminability Under Dichromatic Vision

Satyajeeth Suresh Kannan

Innovation Academy · Mary's Lab

github.com/JeetuSK0808/maryslabb

Preprint · August 2026 · Mary's Lab technical report TR-03

ABSTRACT

How many colors can a display show that a viewer can reliably tell apart? The folklore answer divides the volume of a color solid by the volume of a just-noticeable difference. That computation answers a question nobody asks: it assumes one observer, uniform packing, and no requirement that a specific number of classes be mutually separated. We compute the question practitioners actually face — given N categories and a set of observer models, what is the best minimum pairwise separation any palette could achieve — and we bracket it. An upper bound follows from a covering argument: if K balls of radius r cover the sampled gamut in the worst-case metric and $K < N$, then by pigeonhole no N -palette exceeds $2r$. A lower bound follows constructively from a maximin solver. For 6 classes under normal vision we obtain a bracket of $[0.2671, 0.6734] \Delta E_{OK}$; requiring the same palette to survive full protanopia and deuteranopia simultaneously moves it to $[0.1614, 0.553]$. The bracket, rather than a single number, is the contribution: it separates what is impossible from what a solver merely failed to find. We then ask what a real material library reaches: selecting from 1752 measured USGS reflectances attains between 0.317 and 0.409 of the certified ceiling, and that fraction is near-constant across class counts — working in materials rather than display colors costs about two thirds of the available separation, consistently.

Keywords: palette design; covering bounds; color vision deficiency; accessibility; sphere packing

1. Introduction

A designer choosing a categorical palette for twelve classes, required to remain legible for readers with common forms of color vision deficiency, has no principled way to know whether the task is achievable. If the resulting palette is disappointing, two explanations are available and they demand opposite responses: the design is bad and should be redone, or the requirement exceeds what the display gamut permits and should be renegotiated.

The literature offers a folklore estimate — on the order of ten million distinguishable colors — obtained by dividing the volume of a color solid by that of a just-noticeable difference. The estimate assumes uniform packability, a single observer, and one light. None of those hold for the design question, and the estimate cannot distinguish the two explanations above.

We compute a bracket instead. The upper end is a proof of impossibility; the lower end is a demonstration of possibility. Together they tell a designer whether to iterate or to restructure.

2. Method

The worst-case metric

Fix a set of observer models — normal trichromacy plus any declared dichromacies. For two colors a and b , define their separation as the *minimum* ΔE_{OK} over all declared observers. A palette is only as good as its weakest reader, so the minimum is the correct aggregator, and it is what makes the problem harder than the single-observer case: colors that separate cleanly for a trichromat frequently converge under simulation.

The ceiling, by covering

Sample the sRGB cube on a lattice with `gamutSteps` divisions per axis and map every point through every observer model. Greedily cover the sampled points with balls of radius r in the worst-case metric, and let $K(r)$ be the number of balls used. If $K(r) < N$, then any N colors must place two in a common ball by pigeonhole, and two points in a ball of radius r are at most $2r$ apart. Hence no N -palette achieves a minimum pairwise separation above $2r$. Binary-searching r for the smallest radius whose cover still uses fewer than N balls yields the tightest ceiling the sample supports.

This is a counting argument, not a search. It states that the palette a designer is hoping for does not exist, which is a different kind of claim from "we did not find one".

The floor, by witness

A ceiling alone is uninformative: a ceiling of 0.4 with a best-found palette of 0.05 tells you only that you are ignorant. We therefore run a seeded maximin solver and return the palette it found together with its achieved separation. The truth lies in `[floor, ceiling]`, and we report the gap.

3. Evaluation

Inputs are synthetic. The gamut is the sRGB cube sampled at 9 divisions per axis; no measured display characterisation is used. The observer models are Machado matrices at full severity. Results therefore describe the method's behaviour on the standard model, not a specific monitor.

Table 1. Bracket on minimum pairwise separation, normal vision only. The ceiling is a covering bound; the floor is a constructed palette. Gap is the residual uncertainty.

classes	floor (achievable)	ceiling (impossible above)	gap
2	0.548	2	1.452
3	0.4888	1.2098	0.721
4	0.331	0.9195	0.5885
5	0.331	0.8799	0.549
6	0.2671	0.6734	0.4064
8	0.244	0.6032	0.3592
10	0.2266	0.5364	0.3099
12	0.207	0.5015	0.2945

Table 2. The same bracket when the palette must additionally survive full protanopia and deuteranopia. Requiring three observers at once collapses the achievable separation.

classes	floor	ceiling	gap
2	0.3712	2	1.6288
4	0.2493	0.7537	0.5043
6	0.1614	0.553	0.3916
8	0.1614	0.4215	0.26
10	0.124	0.3586	0.2346
12	0.117	0.3002	0.1832

The achievable separation falls monotonically with class count, as expected. The design-relevant observation is the difference between the two tables: adding the dichromatic requirement reduces the achievable floor at 12 classes from 0.207 to 0.117. A designer told only the first number would plan a twelve-class encoding that fails for a substantial fraction of readers.

The gap between floor and ceiling is not small at higher class counts. We report it rather than splitting the difference: it is honest uncertainty about where the optimum sits, and it identifies which side needs work — a denser lattice tightens the ceiling, a longer solver budget raises the floor.

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 `certify_palette_ceiling`. 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: what a material library can actually reach

The ceiling above is computed over the sRGB gamut — every color a display can show. A designer working in pigment, ink, dye or mineral cannot reach every such color, so the ceiling is an upper bound on an upper bound. The question a materials designer actually asks is what a real library affords, and a measured corpus can answer it.

We took all 1752 measured reflectances in snapshot `usgs-splib07a-1`, converted each to Oklab under D65, and selected palettes by greedy farthest-point search — a constructive lower bound, not an optimum, restarted twelve times and reported at its best.

Table 3. Minimum pairwise separation achievable from a corpus of measured materials, against the certified ceiling over the full sRGB gamut. The last column is the fraction of the theoretical ceiling a real library reaches.

classes	corpus achieves	certified ceiling	gamut floor	fraction of ceiling
3	0.4341	1.2098	0.4888	0.359
4	0.354	0.9195	0.331	0.385
5	0.2791	0.8799	0.331	0.317
6	0.2756	0.6734	0.2671	0.409
8	0.2256	0.6032	0.244	0.374
10	0.1953	0.5364	0.2266	0.364
12	0.1774	0.5015	0.207	0.354

The fraction column is strikingly flat: a library of 1752 measured materials delivers between 0.317 and 0.409 of the certified ceiling at every class count tested. Roughly speaking, working in real materials rather than in display colors costs about two thirds of the available separation, and it costs that proportion consistently rather than biting only at high class counts.

The corpus column also tracks the gamut floor closely — for instance 0.2256 against 0.244 at eight classes. That is a coincidence of magnitude rather than a relationship: the floor is a witness palette the solver constructed inside the display gamut, while the corpus figure is the best selection from a fixed set of minerals. They are close because both are dominated by the same geometry — at eight or more classes, any set of realizable surface colors is crowded. Reporting them side by side is what makes the flatness legible.

This is the number a materials designer needs and it is not derivable from the synthetic analysis at all. The ceiling says how much separation physics permits; this says how much of it the available materials deliver. A brand specifying a 10-category palette in printed pigment should plan against roughly 0.1953, not against 0.5364.

Two caveats belong with the table rather than in the limitations section, because they bound the reading directly. The USGS library is a geological library — minerals, soils, vegetation, a handful of artificial materials — and it is not a pigment catalogue. A designer's real palette would come from an ink or dye library, which would be differently distributed. And greedy selection is a lower bound: a better search over the same corpus could only improve these numbers, so the fraction of the ceiling reported here is pessimistic.

5. Limitations

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

- **The ceiling depends on lattice density.** The cover is built on a finite sample; $2r$ is exact for sampled points, and off-lattice colors are covered only to within the lattice's slack. The density is part of the claim.
- **The floor is a found palette, not a proven optimum.** It is whatever the seeded solver reached within its budget. A wide gap may indicate a loose ceiling or a weak solver, and the experiment does not separate those.
- **Feasibility, not impossibility, at the threshold.** The threshold walk reports the largest class count at which the solver still met a requirement. It certifies achievability up to that N ; it is not a proof that $N+1$ is impossible.
- **Greedy covering is an upper bound on the optimal cover.** A greedy cover may use more balls than necessary, which makes the reported ceiling conservative — safe in direction, but not tight.
- **Vision models are models.** Machado matrices parameterise dichromacy; real color vision deficiency varies between individuals in ways the model does not capture.
- **Prior art was surveyed informally.** Sphere-packing estimates of gamut capacity are long established; a search found no tool producing a two-sided bracket for a specified class count under a declared observer set. That is a statement about a search.

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. 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.
3. CIE (2004). *Colorimetry, 3rd edition*. CIE 15:2004. Commission Internationale de l’Éclairage, Vienna.
4. Wyszecki, G. and Stiles, W. S. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd edition. Wiley.
5. Morovič, J. (2008). *Color Gamut Mapping*. Wiley.
6. 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.
7. Pointer, M. R. (1980). The gamut of real surface colours. *Color Research & Application*, 5(3), 145–155.
8. Conway, J. H. and Sloane, N. J. A. (1999). *Sphere Packings, Lattices and Groups*, 3rd edition. Springer.

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).