

Colormap Topology: Counted Features a Rendering Displays Versus What the Field Contains, on Planted and Measured Fields

Satyajeeth Suresh Kannan

Innovation Academy · Mary's Lab

github.com/JeetuSK0808/maryslabb

Datasheet · August 2026 · Mary's Lab dataset DS-04 · CC-BY-4.0

ABSTRACT

A datasheet for the `colormap-topology` dataset: 2260 rows, each one scalar field rendered through one colormap and analysed under one observer model, with the number of displayed topological features counted by persistent homology on the rendered lightness channel. 89 fields are synthetic with an exactly known feature count; 24 are real reflectance matrices from the USGS library — the heatmap a spectroscopy paper actually prints. The headline is a count rather than an opinion: averaged over all fields and observers, rainbow displays 52.87 features per rendering that are not in the data, and on the measured fields alone that rises to 17.59. Faithfulness turns out to be almost entirely a property of the colormap rather than the observer, and the datasheet explains why that is a finding about lightness-monotonicity rather than a null result.

Keywords: dataset; data visualization; colormaps; persistent homology; accessibility

1. Motivation

That rainbow colormaps mislead is among the most repeated claims in scientific visualisation, and it is almost always argued perceptually and asserted qualitatively. This dataset turns it into a count with ground truth: on fields whose true feature count is known by construction, and on real measured fields where the reference is the

field's own persistence count, how many features does each rendering display, invent, and hide — and for which observers.

2. Composition

2260 rows = 113 fields × five built-in colormaps × four observer models. Train 1800, validation 320, test 140, grouped by field.

Table 1. Faithfulness by colormap — a row is faithful when the rendering neither invents nor hides a feature. Left: all fields. Right: the 24 measured USGS fields alone, the harder regime.

colormap	faithful (all)	mean spurious (all)	faithful (measured)	mean spurious (measured)
grayscale	95.6%	0.01	79.2%	0.04
viridis	94.9%	0.04	76%	0.19
coolwarm	2.2%	4.32	4.2%	8.34
jet	0%	19.66	0%	10.18
rainbow	0%	52.87	0%	17.59

On measured fields even the monotonic maps stop being perfectly faithful — grayscale holds 79.2% and loses features rather than inventing them (mean lost 0.29), which is the better failure but still a failure. The non-monotonic maps invent: jet averages 10.18 spurious features per measured rendering, rainbow 17.59.

3. Collection and labeling

Synthetic fields are sums of well-separated Gaussian bumps — separation enforced at 2.6σ , so the bump count

is the feature count. Measured fields are blocks of 48 real USGS reflectance spectra as material × wavelength matrices, min–max normalised; `true_features` is null there, and the reference is `data_features`, the persistence count of the un-rendered field. Displayed features are counted by union-find persistence with the elder rule on the rendered image's lightness channel, per observer model, at a fixed persistence threshold of 0.02.

Faithfulness is nearly flat across observers (spread under one percentage point) while spanning ~96 points across colormaps. That is a mechanism, not a null: features are counted on lightness, CVD simulation mostly rearranges chroma, so a non-monotonic lightness ramp fails for every observer alike. The narrow claim the data supports is that lightness-monotonicity is the property worth requiring — which is a cleaner design rule than "avoid rainbow".

4. Uses

Colormap evaluation by counting rather than argument: score any new colormap by rendering these fields and running the same audit. Learning tasks — predict spurious/lost from (colormap, observer, field statistics) — are well-posed with grouped splits. The per-row counts also make the dataset a citable source for the specific claim "this map invents N features on this class of field", which the visualisation literature usually states without a number.

5. Distribution, license, and maintenance

JSONL splits plus a dataset card; CC-BY-4.0. Measured fields derive from public-domain USGS data (cite Kokaly et al. 2017). Seeded and reproducible from `datasets/build/`; engine 0.4.1, snapshot `usgs-splib07a-1`. Corrections arrive as a new version, never a silent overwrite.

6. Limitations

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

- **The measured fields' reference is our own computation.** On synthetic fields the truth is planted; on measured fields spurious/lost compare the rendering against the field's persistence count, not against an independent census of the mineralogy.
- **Topology is counted on lightness only.** Structure carried purely by hue is not counted as displayed. That

choice is why observers barely separate here; a measure scoring hue discriminability would very likely separate them, and would be a different dataset.

- **One persistence threshold.** Feature counts at 0.02 are not feature counts at every scale; the threshold is stated so the audit is reproducible, not because it is canonical.
- **Five colormaps.** The built-in set spans the familiar failure modes but is not a survey of the ecosystem.

7. 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. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., and Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92.
2. Edelsbrunner, H., Letscher, D., and Zomorodian, A. (2002). Topological persistence and simplification. *Discrete & Computational Geometry*, 28, 511–533.
3. Borland, D. and Taylor, R. M. (2007). Rainbow color map (still) considered harmful. *IEEE Computer Graphics and Applications*, 27(2), 14–17.
4. Ottosson, B. (2020). *A perceptual color space for image processing (Oklab)*. `bottosson.github.io`.
5. Kokaly, R. F., et al. (2017). *USGS Spectral Library Version 7*. U.S. Geological Survey Data Series 1035. doi:10.3133/ds1035.
6. Kannan, S. S. (2026). *Counting the Features a Colormap Invents: A Persistent-Homology Audit of Scientific Color Maps Across Observer Models*. Mary's Lab technical report TR-02.

Generator, statistics, and this datasheet's renderer: `github.com/JeetuSK0808/maryslabb` under `datasets/build/`. The build is seeded — the same commit reproduces the same dataset byte for byte — and every number above resolves from the generator's own `stats.json`. Datasheet structure follows Gebru et al. (2021). Prepared with AI assistance (see Acknowledgements).