

Error diffusion shaped by palette geometry

Michail Semoglou; College of Design and Innovation, Tongji University; Shanghai, China

Abstract

Designers often dither an image against a small, hand-chosen or image-derived palette, and the visible quality of the result depends on how the quantization error is distributed, not only on which colors are chosen. This paper presents a dithering method that reads the palette’s own geometry and steers error away from the directions the palette cannot represent, at a cost of one small matrix multiplication per pixel. On 124 images, the method reduces perceptual color error by up to a fifth at small palette sizes and matches or outperforms a costlier matrix filter from the prior art. A companion measurement study shows that two standard quality metrics disagree on which dithering variant is best, and establishes how such output should be judged.

Proposal Motivation

Designers encounter constrained-palette halftoning wherever the palette is fixed in advance: brand and editorial palettes, textile and craft work, and limited-ink printing. Palettes are also found rather than devised: a photograph can be reduced to a small set of colors that keeps its mood, ready to dither new work against. In either case, the quality of the result is governed substantially by how the quantization error is distributed; error spread without regard to palette structure appears as visible banding, worm-like texture, and tone shift. Yet how that error is distributed has received less attention than palette selection itself, and the small, irregular palettes of design practice are precisely where ordinary diffusion struggles most. A method that reads the palette’s own geometry, and sound measurement practice for judging it, are both of direct interest to the imaging community.

Problem

The work addresses two questions. The first is steering: whether the palette’s geometry can direct where the quantization error goes, keeping error out of the directions the palette cannot absorb. The second is measurement: how such output should be judged, given that the field’s standard ways of scoring dithered output need not agree. The research question is therefore twofold: whether palette-structure-aware error diffusion improves perceived quality (and in which palette-size range, against which baseline), and how the resulting output should be judged.

Approach

The method, palette-structure-aware error diffusion, derives a shaping matrix from the covariance eigenstructure of the target palette in Oklab [1], a perceptually uniform color space, and applies it to the error signal during Floyd–Steinberg error diffusion [2]. The matrix damps error along the palette’s degenerate, near-flat directions and leaves it unaltered along directions the palette can represent. The matrix is closed-form and computed once per palette; dithering then costs one 3×3 matrix multiplication per pixel beyond ordinary error diffusion. Global and per-cell variants are developed; the

global variant outperforms the per-cell variant under S-CIELAB at all but the smallest palette size tested, where the two are comparable. Experiments use a 124-image corpus (24 Kodak [3] and 100 CQ100 [4] images), with the palette extracted from each image at six sizes $K \in \{8, 16, 32, 64, 128, 256\}$, and three diffusion filters: Floyd–Steinberg, the Kolpatzik–Bouman scalar HVS filter [5], and an Oklab-domain reimplementation of the matrix-valued Damera–Venkata filter [6]. The primary endpoint is S-CIELAB ΔE [7]; hypotheses, subgroups, and the analysis plan were preregistered before the experiments were run.

Results

Shaping improves S-CIELAB ΔE across the small-to-moderate range, peaking near a 21% median reduction at $K = 16$ (Figure 1, top), with an 18% reduction at $K = 8$ and a monotonic decline to no detectable effect at $K = 256$ ($p = 0.53$). The two corpora diverge at large K : from $K = 64$ the pooled effect is carried by CQ100, while the Kodak subgroup is harmed at $K \geq 128$ (median -18% at $K = 256$). The measurement study shows that two standard scoring methods disagree systematically: an unweighted lightness-channel error measure and S-CIELAB agree on overall magnitude (Spearman $\rho = 0.92$) yet choose the same winning variant in only 13.2% of the 744 image–palette-size cells, below the 33% chance level. A second, exploratory finding concerns the Kolpatzik–Bouman filter itself: it behaves in two regimes separated by the image’s baseline low-frequency error, with a low-frequency error floor where the baseline low-frequency energy fraction falls below about 0.02, and a low-frequency benefit above it; the boundary marks where DC conservation begins to matter more than HVS weighting. The Damera–Venkata filter, the exact prior art for this problem, is the strongest unshaped baseline tested (a median 17 to 22% below Floyd–Steinberg at $K \leq 32$, and 17 to 24% below Kolpatzik–Bouman across the range), and the overlap between palette-geometry shaping and HVS filter tuning is nearly complete: shaping on top of that filter adds a pooled median of at most 2.4%. The practical consequence favors the cheaper method (Figure 1, bottom): shaped Floyd–Steinberg matches Damera–Venkata at $K = 8$, significantly beats it at $K = 16$ – 64 ($p \leq 0.028$), and is comparable thereafter.


Conclusion

Constrained-palette halftoning should be measured with a spatially weighted perceptual metric, against the correct baseline, and judged in the setting in which design practice actually operates: a few chosen colors, not 256. In that setting, the cheapest method is also the best one. An open-source release of the method, together with the experiment harness and corpus manifests, accompanies the paper.

References

- [1] B. Ottosson, “A perceptual color space for image processing,” 2020. [Online]. Available: <https://bottosson.github.io/posts/oklab/>

a — shaping over Floyd–Steinberg: median S-CIELAB ΔE reduction



b — against the Damera–Venkata baseline: median ΔE difference, %




Figure 1. Top: median S-CIELAB ΔE reduction from palette-geometry shaping over Floyd–Steinberg diffusion (124 images, interquartile band). Bottom: median ΔE against the Damera–Venkata (DV) baseline; shaped Floyd–Steinberg (accent) matches or beats it throughout the design range, while shaping applied to DV itself (black dots) adds little.

- [2] R. W. Floyd and L. Steinberg, “An adaptive algorithm for spatial grey scale,” *Proc. Soc. Inf. Display*, vol. 17, no. 2, pp. 75–77, 1976.
- [3] Eastman Kodak Company, “Kodak Lossless True Color Image Suite (PhotoCD PCD0992),” 1993. [Online]. Available: <http://r0k.us/graphics/kodak/>
- [4] M. E. Celebi and M.-L. Pérez-Delgado, “CQ100: a high-quality image dataset for color quantization research,” *J. Electron. Imaging*, vol. 32, no. 3, 033019, 2023.
- [5] B. W. Kolpatzik and C. A. Bouman, “Optimized error diffusion for image display,” *J. Electron. Imaging*, vol. 1, no. 3, pp. 277–292, 1992.
- [6] N. Damera-Venkata and B. L. Evans, “Design and analysis of vector color error diffusion halftoning systems,” *IEEE Trans. Image Process.*, vol. 10, no. 10, pp. 1552–1565, 2001.
- [7] X. Zhang and B. A. Wandell, “A spatial extension of CIELAB for digital color-image reproduction,” *J. Soc. Inf. Display*, vol. 5, no. 1, pp. 61–63, 1997.