PNAS COMMENTARY

Proximity matters

David H. Brainard[a,1]

About a century ago, color scientists learned how to predict when two lights with different spectral energy will look the same. About 50 y ago, practical standards were created to quantify the size of the perceptual difference between two lights. These developments are based on creating three-dimensional color spaces. One type of space, a color-matching space, captures equivalence. The second type of space builds on the color-matching representation to include a distance metric that represents perceptual differences. The mathematics and biology of the color matching space are well in hand; the specification of a space with a distance measure that represents perceptual color differences is an ongoing project.

In their paper "The non-Riemannian nature of perceptual color space," Bujack et al. (1) tackle the fundamental question of how to structure a color space that explains color differences as well as color matching. By deploying an elegant psychophysical procedure in which observers judged the relative magnitude of pairs of stimulus differences, they show that it is not possible to use either Euclidean or Riemannian geometry to describe both small and large perceptual color differences. Their results clarify the limits of current approaches that attempt to capture color perception using simple geometric representations and point the way toward future experimental and theoretical projects that might allow a more satisfactory account.

A long-standing goal of perceptual psychology is to develop principles and methods that provide accurate descriptions of our mental representations of physical phenomena (2). A celebrated example comes from color science. Maxwell (3) quantified the observation that humans can match the appearance of arbitrary light spectra with a weighted superposition of three primary lights, leading to the idea that the mental representation of light may be characterized as a point in a three-dimensional color mixture space, with the tristimulus coordinates of a particular spectrum taken as the weights on the primaries that provide a match to that spectrum (Fig. 1). At around the same time, Grassmann (4) formulated the key principles of linearity that describe color matching; the color coordinates of a match to the superposition of two spectra are the sum of the color coordinates of each alone. The quantitative description of color matching together with the machinery of linear algebra enables modern color reproduction technology, in which colors are reproduced as mixtures of three primaries on our televisions, phones, and computers. We now understand that the three-dimensional nature of color space as well as the linearity of matching derives because the initial visual encoding of spectra is the excitations of just three classes of light-sensitive photopigment contained in the cone photoreceptors of the retina (5).

What about representing not only color matching but also the perceptual differences between color stimuli? Specification of such differences is important at multiple scales. At the scale of small color differences, for example, tools are required to decide whether color reproduction has been accomplished to a satisfactory tolerance. Here, the goal is to define a metric over positions in color space so that the distance between two points allows accurate prediction of whether two color stimuli will appear sufficiently similar to satisfy a discerning customer. At the scale of large color differences, tools are required to determine whether the two colors intended to convey distinct meaning successfully do so. For example, are the green and red signals of a traffic light sufficiently different so as to have a low probability of being confused in an emergency? Here, a goal is to define a metric over positions so that the distance between two points allows for accurate prediction of the probability of confusion between two discriminable stimuli.

Color spaces that do a good job of both describing color matching and predicting the visibility of small color differences have been developed. A widely used example is the Commission Internationale de l'Elcairage (CIE) Lab uniform color space. In its original formulation, the Lab coordinates (referred to as L*, a*, and b*, respectively) were obtained as an invertible nonlinear transformation of a tristimulus space that describes human color matches (specifically, the CIE XYZ tristimulus space) (6). The transformation was designed so that the Euclidean distance between two closely located stimuli provided a reasonable prediction of their discriminability. This system is widely used in specifying industrial tolerances for color reproduction, with a more recent version applying a metric that varies systematically with the location of stimuli in the space (7). Although the CIE Lab system was designed to predict small color differences, it has also been widely adopted to equate larger color differences: for example, in studies of color in cognition (e.g., refs. 8 and 9 are just two of many examples). Although such use is reasonable given the current state of our knowledge, it is important to recognize that this generalization is not secure. It relies on the idea that large color differences may be predicted by integrating small color differences along a path between two well-separated points in color space. This idea has a long history, going back at least to Helmholtz (10). When the metric is the familiar L2 norm, the space is Euclidean. When the metric varies as one traverses the space, the

Author affiliations: [a]Department of Psychology, University of Pennsylvania, Philadelphia, PA 19104

Author contributions: D.H.B. wrote the paper.

The author declares no competing interest.

See companion article, "The non-Riemannian nature of perceptual color space," 10.1073/pnas.2119753119.

[1]Email: brainard@psych.upenn.edu.

Published June 23, 2022.

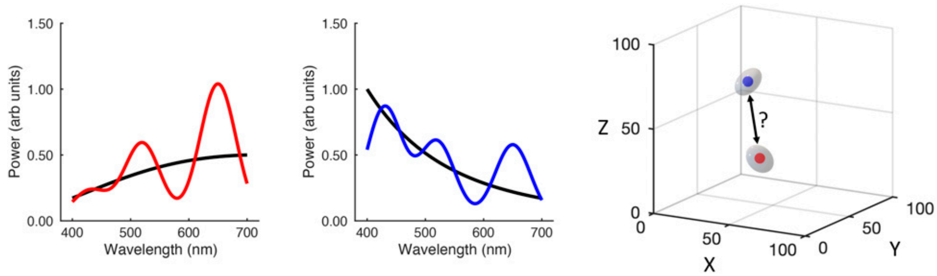

Fig. 1. (*Left*) The plot shows the spectral radiance of two different spectra that appear the same to a standard human observer. These two spectra map to the same point in the CIE XYZ tristimulus color space (red circle in *Right*). The three coordinates in this space give the intensities of three theoretical primaries required to match either of these spectra. (*Center*) Another pair of spectral radiances that match for a standard human observer. These two spectra are represented by the blue point in the tristimulus space in *Right*. (*Right*) Each point in the CIE tristimulus space represents a set of spectral radiances that appear the same to a human observer. The matches themselves, however, do not tell us about the perceptual distance between points, indicated by the question mark. A local metric around each point in the color space, referred to as a line element, can be created. For example, threshold measurements may be used to define a small ellipsoid around each point such that stimuli represented on the surface of the ellipsoid are equally discriminable from those represented by the point (16, 17). Two such ellipsoids are shown schematically. The question asked by Bujack et al. (1) is whether it is possible to construct a Riemannian perceptual color space such that integrating the line element along the geodesic between any two points predicts larger perceptual color differences.

underlying geometry is the more general Riemannian formulation. In the latter case, the distance between two points is taken to be that found by integrating along a geodesic—that is, along a path where the integrated distance is the least. In this tradition, the space-varying metric is called a line element (Fig. 1). Different realizations of the idea are based on different transformations of the underlying tristimulus color matching space and on different formulations of the line element (11).

Bujack et al. (1) step back from the details of the appropriate color space transformation and line element formulation, and ask more generally whether data collected simultaneously for small and large color differences are consistent with any Riemannian geometry. To do so, they focus on testing a key axiom of Riemannian geometry; if one considers three stimuli A, B, and C that lie along a geodesic, the distance between A and C should be the sum of the distances between A and B and between B and C. In the experiments reported, subjects make judgments on triads of stimuli placed along a single geodesic, indicating which of two comparison stimuli is most similar to a reference stimulus. In essence, the judgment on each trial is to determine which of two color differences is larger. Judgments are made for comparisons that are both close to and far from the reference and for multiple references. Thus, the data provide perceptual distance judgments that constrain the transformation of stimuli into an underlying perceptual space as well as the line element that defines distances between them.

To bring the data to bear on the question of color space geometry, the authors fit two scaling models (12) using maximum likelihood methods (13). One conforms to Riemannian structure; stimuli are placed along a line in perceptual space so that distances between them, computed with a line element, best predict the perceptual distance judgments. This model enforces the key axiom mentioned above; the distance between points A and B plus the distance between B and C is equal to the distance between A and C. In the second model, consistency with this key axiom is relaxed through the addition of a nonlinear function applied to the distances before judgments are predicted. This model is incompatible with an underlying Riemannian geometry. Model comparison, performed using cross-validation, shows that the non-Riemannian model provides a substantially better account of the data. The compressive nature of the nonlinearity obtained accounts for the empirical observation that subjects' responses show a shrinkage of large stimulus distances relative to those predicted by summing up small stimulus differences, a phenomenon the authors refer to as the law of diminishing returns. In essence, large color differences saturate relative to the Riemannian prediction. The fact that the model comparison rejects a general form of Riemannian representation makes the test applied by Bujack et al. (1) compelling.

There are some caveats. First, the authors make a first principles assumption that the achromatic locus is a geodesic and use this in their choice of stimuli. This assumption is intuitively appealing in that it would be surprising that the shortest path in color space between two achromatic stimuli would involve a detour through a chromatic stimulus and back. However, the achromatic locus as a geodesic was not empirically established, and more work could be considered to shore up this aspect of the argument. Second, the data were collected using online methods and combined across subjects prior to the analysis. This raises the question of whether the aggregate performance analyzed could be non-Riemannian even when the performance of each individual subject was itself Riemannian. Although it is not immediately obvious whether this could occur, it might be further considered as a possibility.

From the work of Bujack et al. (1), we learn that attempts to place color stimuli in a space described by Euclidean or Riemannian geometry so as to simultaneously account for small and large color differences cannot fully succeed; such attempts need to be regarded as approximations. As noted above, this caution is particularly relevant because a wide range of studies use color stimuli as a model system, with conclusions sometimes based on the use of a uniform color space to equate perceptual distance between stimuli. The current work emphasizes the importance of making sure conclusions are robust to violations of the assumption that perceptual distance has been equated. Indeed, the fact that intuitions about angle and distance gleaned from the familiar geometries do not hold even for a perceptual system as simple as color sounds a more general note of caution about applying metric distance concepts to other perceptual and neural representations. An available tool in this regard is nonmetric multidimensional scaling (14, 15).

More generally, Bujack et al. (1) motivate the development of representations that more accurately capture color matching, small color differences, and large color differences within a single framework. The authors' modeling in the current paper suggests one empirically driven approach. The generality of that approach, both for stimuli off the achromatic locus and for additional judgements that depend on color difference, remains to be assessed.

1. R. Bujack, E. Teti, J. Miller, E. Caffrey, T. L. Turton, The non-Riemannian nature of perceptual color space. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2119753119 (2022).
2. R. N. Shepard, Toward a universal law of generalization for psychological science. *Science* **237**, 1317–1323 (1987).
3. J. C. Maxwell, On the theory of compound colours and the relations of the colours of the spectrum. *Philos. Trans. R. Soc. Lond.* **150**, 57–84 (1860).
4. H. Grassmann. Zur theorie der farbenmischung. *Annalen der Physik und Chemie* **165**, 69–84 (1853).
5. D. H. Brainard, A. Stockman, "Colorimetry" in *The Optical Society of America Handbook of Optics, 3rd Edition, Volume III: Vision and Vision Optics*, M. Bass *et al.*, Eds. (McGraw Hill, New York, NY, 2010), pp. 10.11–10.56.
6. CIE, CIE recommendations on uniform color spaces, color-difference equations, and metric color terms. *Color Res. Appl.* **2**, 5–6 (1977).
7. CIE, *Improvement to Industrial Colour-Difference Evaluation* (Bureau Central de la CIE, Vienna, Austria, 2001).
8. T. Regier, P. Kay, N. Khetarpal, Color naming reflects optimal partitions of color space. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 1436–1441 (2007).
9. M. Olkkonen, S. R. Allred, Short-term memory affects color perception in context. *PLoS One* **9**, e86488 (2014).
10. H. Helmholtz, *Handbuch der Physiologischen Optik* (Thoemmes Press, 1896).
11. G. Wyszecki, W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae* (John Wiley & Sons, New York, NY, ed. 2, 1982).
12. L. L. Thurstone, A law of comparative judgment. *Psychol. Rev.* **34**, 273–286 (1927).
13. L. T. Maloney, J. N. Yang, Maximum likelihood difference scaling. *J. Vis.* **3**, 573–585 (2003).
14. J. B. Kruskal, Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika* **29**, 1–27 (1964).
15. T. F. Cox, M. A. A. Cox, *Multidimensional Scaling* (Chapman & Hall/CRC, Boca Raton, FL, ed. 2, 2001).
16. C. Noorlander, J. J. Koenderink, Spatial and temporal discrimination ellipsoids in color space. *J. Opt. Soc. Am.* **73**, 1533–1543 (1983).
17. A. B. Poirson, B. A. Wandell, The ellipsoidal representation of spectral sensitivity. *Vision Res.* **30**, 647–652 (1990).