

Emergence of unique hues from sparse coding of color in natural scenes

ALEXANDER BELSTEN,^{1,2,*} E. PAXON FRADY,^{3,5} AND BRUNO A. OLSHAUSEN^{1,2,4,6}

¹Redwood Center for Theoretical Neuroscience, University of California, Berkeley, Berkeley, California 94720, USA

²Herbert Wertheim School of Optometry and Vision Science, University of California, Berkeley, Berkeley, California 94720, USA

³Neuromorphic Computing Laboratory, Intel, Santa Clara, California 95054, USA

⁴Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, California 94720, USA

⁵paxon.frady@gmail.com

⁶baolshausen@berkeley.edu

*belsten@berkeley.edu

Received 16 March 2026; revised 15 June 2026; accepted 16 June 2026; posted 16 June 2026; published 14 July 2026

Our subjective experience of color is typically described by abstract properties such as hue, saturation, and brightness that do not directly correspond to sensory signals arising from cones in the retina. Along the hue dimension, certain colors—red, green, blue, and yellow—appear unique in that they are not perceived as a combination of other colors, and the pairs red–green and blue–yellow appear opposites. However, the anatomical and physiological correlates of these “unique hues” within the brain and the reason for their existence remain a mystery. Here, we demonstrate a direct connection between these hues and the statistics of the natural visual environment. Analysis of simulated cone responses on a dataset of 503 calibrated natural images reveals a strongly non-Gaussian distribution in 3D color space, with heavy tails in distinct, asymmetrically arranged directions. A sparse coding model is then adapted to this data so as to minimize the total sum of coefficients on the basis vectors for representing the data. A six-basis-vector model converges to the four unique hues in addition to black and white. Moreover, we find that the nonlinear nature of inference in the sparse coding model yields both excitatory and inhibitory interactions among latent variables; the former facilitates combining adjacent pairs of unique hues to encode intermediate hues situated between them, while the latter enforces mutual exclusivity between opposite unique hues. Together, these findings shed new light on the distribution of color in the natural environment and provide a linking principle between this structure and the phenomenology of color appearance. © 2026 Optica Publishing Group. All rights, including for text and data mining (TDM), Artificial Intelligence (AI) training, and similar technologies, are reserved.

<https://doi.org/10.1364/JOSAA.598897>

1. INTRODUCTION

In the late 19th century, Ewald Hering and Hermann von Helmholtz engaged in a vigorous debate over their competing views regarding the mechanisms underlying color perception in the eye and brain. While Helmholtz advanced a trichromatic theory of vision [1], Hering, relying largely on introspection, proposed a theory based on four unique hues—red, green, blue, and yellow—which in his view constituted atomistic elements of color appearance [2]. Noting that certain pairs of these hues are never simultaneously perceived in any color (i.e., red–green and blue–yellow are mutually exclusive), Hering showed how they could be organized into opponent pairs and serve as a perceptual basis for describing all other colors. Ultimately, the Helmholtz–Young trichromatic theory was linked to the presence of three types of cones in the retina, and subsequently, a cone-opponent encoding strategy was identified in the retina and lateral geniculate nucleus (LGN) whereby neurons compute differences in activation between cone types [3,4]. At the same

time, Hering’s unique hues came to be adopted by psychology, vision science, and color technology as a canonical basis for describing color appearance [5–7]. However, the *cone-opponent* representation of the early visual system does not align with the *color-opponent* organization implied by the unique hues [8–11]. Attempts to find neurophysiological correlates of the unique hues have been largely inconclusive [12–14], with the exception of Forder *et al.* [15]. Psychophysical evidence reliably distinguishing their “uniqueness” as compared to other hues also remains inconclusive [14,16]. Thus, the alleged privileged status of the unique hues remains both enigmatic [17–19] and controversial [14,20], as they do not neatly map onto any well-understood property of light or the visual system.

It has long been suggested that color perception is shaped by the statistics of the natural visual environment [18,21–26]. Could this provide an account for the unique hues? Previous work by Webster and Mollon [22] studied the distribution of simulated cone activations to natural images and found that

a subset of images exhibited substantial variance along a blue-yellow axis. For such images, principal component analysis (PCA) yields a basis vector aligned with this chromatic axis. Ruderman *et al.* [27] applied PCA to simulated cone activations pooled across an ensemble of natural images and found principal components that correspond to physiological cone-opponent mechanisms found in the retina and LGN [4]. Models that characterize higher-order statistics have been adapted to spatiochromatic natural images and have been shown to produce orientation- and chromatically selective receptive fields [28–33], but these models have not resulted in bases aligned with the unique hues. One possible reason is that learning bases from spatiochromatic image patches mixes spatial and color statistics, obscuring higher-order structure that is specific to the color distribution itself. Yendrikhovskij [34] applied clustering algorithms to pixels of natural images in a perceptually uniform color space and found an alignment between the resulting clusters and color-naming lexicons, which points toward a relation to the unique hues. However, one drawback of these studies is that they were based on particular, limited datasets (e.g., approximately $10^4 - 10^5$ pixels). More importantly, these approaches did not directly examine the full structure of the 3D color distribution.

In this study, we analyze a large composite dataset comprising more than 10^8 simulated cone activations sampled from images collected by different investigators and spanning a diverse range of visual environments. We then derive encoding models adapted to the structure of these data, following the three-stage pipeline depicted in Fig. 1. Rather than examining the full spatiochromatic structure of image patches, we focus exclusively on the statistics of the three-dimensional color distribution defined by the joint activations of long-, medium-, and short-wavelength selective cones (henceforth referred to as LMS space). Thus, our analysis treats pixels as samples from a color distribution and does not model the spatial organization of chromatic structure, such as through spatial correlations, edges, or object-level color structure. Extending the present approach to the joint statistics of color and spatial structure is an important but more complex problem that we defer for future study.

2. RESULTS

The results below are based on a large composite dataset of simulated LMS cone activations in response to natural scenes assembled from five different, publicly available datasets. Details of each component dataset and the preprocessing procedures are provided in Section 4. Briefly, the data from each image are rescaled independently for each cone type and then passed through a compressive nonlinearity. Finally, the mean across all channels was subtracted so that the dataset was zero-centered before further processing.

The analysis proceeds through three stages. First, natural scenes are represented in the LMS cone activation space [Fig. 1(a)]. Second, these responses are linearly decorrelated and spherized, yielding a cone-opponent coordinate system in which second-order structure has been removed, and higher-order deviations from Gaussianity are made apparent [Fig. 1(b)]. Third, the spherized cone activations are represented using a non-negative sparse coding model, which learns a basis adapted to the asymmetric, heavy-tailed structure of the distribution of color in natural scenes. Nonlinear inference in this model transforms the spherized cone activations into a sparse, non-negative representation. In the six-basis-vector model, the learned basis vectors align with the unique hues, and the nonlinear inference dynamics produce mutual exclusivity between latent variables corresponding to opponent-color pairs [Fig. 1(c)].

A. Principal Components Align with Cone-Opponent Color Space

A natural first step in understanding the structure of multivariate data is to apply PCA, which linearly transforms the data into a new space where variables are pairwise decorrelated [35] (see Section 4). Applying PCA to the composite dataset of simulated cone activations yields the principal components and corresponding variances shown in Figs. 2(a) and 2(b) (see also Fig. S1 of Supplement 1). The first principal component (PC1) sums the activations of L, M, and S cones, corresponding to an achromatic channel. The remaining two principal components compare the activations of different cone types: the second component (PC2) contrasts the activation of S cones with the net activation of L and M cones, and the third (PC3) compares the activation of L and M cones. These components align nearly




Fig. 1. Three-stage model of chromatic processing in the visual system. (a) The pixels of a large dataset of natural images are represented in the LMS cone activation space (stage 1). (b) A linear sphering transform decorrelates the activations by first rotating to the principal components, yielding a set of *cone-opponent* bases, and then rescaling each axis to achieve unit variance in all directions (stage 2). (c) A set of bases that efficiently represent higher-order structure in the cone-opponent space is derived using non-negative sparse coding models (stage 3). A six-basis-vector model converges to directions that align with the unique hues, forming a new 6D, non-negative space whose axes correspond to perceptual color categories. The inference process of sparse coding leads to mutual exclusivity between opponent colors (i.e., red–green, blue–yellow, black–white) depicted here in a 3D space with opposing colors at opposite ends of each axis. Notably, these color-opponent axes, when projected back into the cone-opponent space, are non-orthogonal and not aligned with the cone-opponent axes. Circles in (b) and (c) indicate chromatic planes.




Fig. 2. (a) Principal components of the distribution of LMS activations. The weights' alignment with the horizontal dashed lines at $\pm\sqrt{1/6}$, $\pm\sqrt{1/3}$, $\pm\sqrt{1/2}$, and $\pm\sqrt{2/3}$ reveals the near-integer-multiple relationship in how they mix L, M, and S signals (as pointed out by [27]). (b) Variance of each principal component. (c)–(e) LMS activation distributions in the sphered color space exhibit anisotropic, non-Gaussian structure. Solid lines show isoprobability contours after projection onto three different 2D planes defined by each possible pairing of axes in the sphered color space, as indicated by the axis labels. Dashed contours indicate isoprobability contours for a multivariate isotropic Gaussian distribution with the same variance as the data distribution. The contours are equally spaced in log-probability, using the same color scale for both solid and dashed contours. The color circle inset in (e) indicates the hue of each angle within the cone-opponent chromatic plane.

identically with those obtained previously by Ruderman *et al.* [27]. Notably, PC2 and PC3 correspond to cone-opponent mechanisms that have been characterized both psychophysically [3] and physiologically [4].

The PCA transformation also reveals a striking anisotropy in the data distribution: the variance on PC1 is more than 10-fold greater than on PC2, and the ratio between the variance of PC2 and that of PC3 is even greater [Fig. 2(b)]. This strong anisotropy makes it difficult to visually inspect the full 3D structure of the joint distribution. Therefore, we equalize the variance in all directions of the space by rescaling each principal component axis by the inverse square root of its corresponding variance (see Section 4). Throughout the remaining sections of this paper, we will refer to the unit norm principal component vectors as PC1, PC2, and PC3, and their corresponding rescaled versions based on their predominant cone inputs: $L + M + S$, $2S - (L + M)$, and $L - M$, respectively. In addition, we refer to the 3D space constructed by the latter three axes as the sphered color space.

The motivation for rescaling is twofold. First, for clarity of visualization: when examining the 3D distribution of activations in the decorrelated space, having axes with equal variance makes any deviation from Gaussianity readily apparent. Second, we conjecture that the brain may also rescale sensory data in this way. As argued by various authors [36–38], neural systems adjust their gain to fill their output dynamic range. Thus, neurons tuned to low-variance directions in the data exhibit comparatively higher gain than those sensitive to high-variance directions in the data. Since the achromatic component has substantially higher variance than the chromatic components [Fig. 2(b)], amplification of chromatic relative to achromatic signals is necessary to achieve equal variance. This prediction aligns with psychophysical measurements of detection thresholds, which reveal substantially higher sensitivity to chromatic increments than to luminance increments [39,40].

B. Distribution of LMS Activations is Highly Non-Gaussian

Figures 2(c)–2(e) show the structure of the sphered distribution via three 2D projections of the full 3D distribution, using each possible pairing of L-M, 2S-(L + M), and L + M + S axes (see also Fig. S2 of Supplement 1). The resulting density of each projection is plotted in terms of its isoprobability contours. For comparison, contours from an isotropic Gaussian distribution with identical variance are overlaid on each plot. Comparing the isoprobability contours of the LMS activations to those of the Gaussian reveals pronounced asymmetries and substantial deviations from Gaussianity.

First, consider the achromatic dimension in Figs. 2(c) and 2(d). The low-probability contours of the LMS distribution are slightly more compact along the achromatic axis than the Gaussian contours, suggesting a mild sub-Gaussian shape. Quantitatively, this is confirmed by a negative excess kurtosis of -0.6 for 1D projection along the achromatic axis (compared to 0 for a Gaussian; see Table S1 of Supplement 1).

By contrast, the chromatically sensitive dimensions (L-M and 2S-(L + M)) exhibit strong heavy-tailed structure as the isoprobability contours in these directions extend farther from the origin than those of a Gaussian. The excess kurtosis for 1D projections along these directions is 4.3 for 2S-(L + M), and 27.0 for L-M. Additionally, the tails are asymmetric, with greater probability mass in the extremes of the positive L-M direction than that of the negative. Similarly, there is more probability mass in the extremes of the negative 2S-(L + M) direction than in the positive. The pattern of heavy-tailed structure within the chromatic plane and slightly sub-Gaussian structure along the achromatic axis is consistent across the different LMS datasets analyzed (Table S1 of Supplement 1).

Figure 2(e) shows the joint distribution within the plane defined by the cone-opponent axes. The heavy-tailed structure varies substantially with azimuth angle (the angle formed between a point's projection onto the plane and the positive L-M direction). Certain directions, such as 0° and 260° , display pronounced heavy tails, while others, such as 135° , exhibit more Gaussian or even sub-Gaussian tails. Notably, the structure is asymmetric between opposing directions (e.g., 0° versus 180° and 135° versus 315°).

Early observations by Hendley and Hecht [41] reported that natural scenes generally lack saturated colors. While most pixel activations indeed cluster near the achromatic axis, the large dataset of over 225 million pixels reveals a richer structure: saturated colors actually occur with greater frequency than expected from a random process with the same variance as the data (i.e., a Gaussian distribution). Detecting this heavy-tailed structure requires large sample sizes; smaller datasets (e.g., tens of thousands of pixels) would fail to capture these tails accurately.

What gives rise to the heavy-tailed structure observed in the chromatic plane? It is not likely to be an artifact of the preprocessing pipeline, because rescaling and histogram equalization would, if anything, reduce kurtosis rather than amplify it [42]. Moreover, the heavy-tailed directions do not show correspondence to the directions of the L, M, and S cones in the sphered space, which point in 330° , 210° , and 90° directions, respectively. The fact that the heavy-tailed structure survives after

averaging over datasets collected by different investigators using different cameras suggests that it is a persistent property of the natural environment as opposed to idiosyncratic properties of any given dataset. Previous work has shown that color in natural scenes tends to be spatially patchy, with large, relatively uniform regions compared to the more fine-grained variation of luminance [43]. This spatial structure may arise because surface color is a material property that tends to remain coherent within an object, whereas luminance varies both with surface reflectance and with other factors such as shading and illumination geometry. Additionally, Rosenthal *et al.* [24] found that objects tend to have higher chromatic saturation, suggesting that the heavy tails in the chromatic distribution [Fig. 2(e)] are driven by the presence of highly saturated, chromatically consistent objects (e.g., flora, foliage, sky, and fruit) within natural scenes. A hue-resolved analysis that links the heavy-tailed structure observed here to image content will be interesting to explore in future work.

C. Modeling Higher-Order Structure

While heavy-tailed structure is revealed by projection onto pairs of principal components, PCA per se cannot provide a model of this structure as it is computed solely from second-order statistics (the covariance matrix of the data), which are blind to such structure [35,44]. Thus, one must turn to a modeling approach capable of capturing higher-order statistics.

A linear generative modeling approach capable of capturing such higher-order statistics is independent component analysis (ICA), or its highly related cousin, sparse coding [45–47]. Intuitively, sparse coding seeks a set of basis vectors such that each data point can be reconstructed using only a small number of basis vectors out of a larger dictionary. This favors learning basis vectors that align with the directions in the data which occur rarely but with high magnitude, i.e., high-kurtosis or heavy-tailed directions [35,44,48]. Sparse coding applied to natural images provides a canonical example of this principle. When adapted to these data, the learned basis aligns itself with heavy-tailed directions in image space and exhibits receptive field properties similar to those of simple cells in primary visual cortex [49]. Theoretical and computational studies have demonstrated the utility of sparse representations: sparsity makes complex structure in sensory inputs explicit, improves storage capacity, enables easier readout by high-order neural populations, and produces an energy-efficient representation [50].

The sparse coding generative model assumes each 3D data vector $\mathbf{x}$ to be described as a linear combination of m unit-norm basis vectors $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m$ plus Gaussian noise $\boldsymbol{\epsilon}$:

$$\mathbf{x} = \sum_{i=1}^m \mathbf{a}_i s_i + \boldsymbol{\epsilon}. \quad (1)$$

The variance of the noise $\boldsymbol{\epsilon}$ is assumed to be small relative to that captured by the basis vectors $\mathbf{a}_i$. Importantly, there is no assumption of orthogonality on the $\mathbf{a}_i$, and they can even be overcomplete and hence linearly dependent, providing greater flexibility in capturing structure in the data. Sparsity in the

latent variables s_i is enforced via a factorial prior that encourages their values to be zero. Typically, the prior is chosen to be symmetric about zero (e.g., a Laplacian distribution), enabling the s_i to take on both positive and negative values. Given the asymmetric structure evident in the data shown in Fig. 2, we impose a nonnegativity constraint on their values (i.e., $s_i \geq 0$ for all i).

While the sparse coding model is a linear generative model, the latent variables $s_1, s_2, \dots, s_m$ are a nonlinear function of the data $\mathbf{x}$. The s_i are computed by minimizing an energy function E subject to a nonnegativity constraint:

$$\text{minimize } E = \frac{1}{2} \left\| \mathbf{x} - \sum_{i=1}^m \mathbf{a}_i s_i \right\|_2^2 + \lambda \sum_{i=1}^m s_i, \quad (2)$$

$$\text{subject to } s_i \geq 0 \quad \text{for } i = 1, 2, \dots, m, \quad (3)$$

where parameter λ controls the trade-off between reconstruction quality and sparsity. This minimization corresponds to performing maximum-a-posteriori (MAP) Bayesian inference, $\mathbf{s} = \arg \max_{\mathbf{s}'} p(\mathbf{s}' | \mathbf{x}, \mathbf{A})$, where vector $\mathbf{s}$ has elements s_i and matrix $\mathbf{A}$ has columns $\mathbf{a}_i$ [49].

The optimization problem of Eqs. (2) and (3) is solved using the locally competitive algorithm (LCA; see Section 4). The basis vectors $\mathbf{a}_i$ are then adapted to the structure of the data by gradient descent on E averaged over many data samples, using the MAP-inferred $\mathbf{s}$ for each data sample $\mathbf{x}$ (see Section 4). Note that the nonnegativity assumption limits the minimum number of basis vectors to $m = 4$, as at least four vectors are necessary to span three dimensions under non-negative weightings.

D. Progression of Basis Vector Solutions with Increasing Overcompleteness

Figure 3 shows the basis vectors $\mathbf{a}_i$ that emerge from non-negative sparse coding models trained on the data for $m = 4, 5, \dots, 8$. Each model's bases are plotted both within the 3D sphered color space [Fig. 3(a)] and as a projection onto the chromatic plane [Fig. 3(b)].

In the $m = 4$ model, two basis vectors lie mostly within the chromatic plane and are aligned with red and yellow-green, roughly corresponding to two dominant heavy-tailed directions seen in the data distribution [Fig. 2(e)]. The remaining two vectors are elevated along the achromatic axis toward black and white, but retain a chromatic component in the blue direction, as indicated by their slight tilt from vertical. The tilt toward blue is needed as this region of hue space would otherwise remain unrepresented by the other two vectors within the chromatic plane.

When the number of basis vectors increases to five, a third basis vector is placed within the chromatic plane oriented toward the blue direction. This placement is expected as the data exhibits a heavy-tailed extension in this direction, in addition to the red and yellow-green directions, which are represented by two other basis vectors. The remaining two basis vectors now align closely with the achromatic axis, providing an explicit description of luminance variation. Thus, the five-basis-vector model factorizes chromatic and achromatic structure into distinct groups of basis vectors: three basis vectors tile the chromatic plane, and two basis vectors encode positive and negative variations along the achromatic axis.

At $m = 6$, a fourth vector emerges within the chromatic plane, and the black and white vectors remain aligned to the achromatic axis. The new basis vector within the chromatic plane is aligned to green, nudging the previously yellow-green vector toward yellow, and the bluish vector toward the center of blue. This nudging of the other vectors arises from the dynamics of adaptation in the sparse coding model, whereby bases both




Fig. 3. (a) Adapted non-negative sparse coding basis vectors are shown in the sphered color space. The number of basis vectors is varied from left to right from four to eight. Each basis vector has unit length. The black ellipse depicts a unit-radius circle within the chromatic plane. (b) Basis vectors from (a) are shown projected into the chromatic plane. The black circle depicts the unit radius, and the outer ring shows the corresponding color for each hue angle. The gray curve reproduces the $\log_{10}$ probability = -4 contour from Fig. 2(e); it is uniformly rescaled for visual comparison, so shape is preserved but absolute scale is not. Vectors closely aligned with the achromatic axis project to (near) zero and may not be visible.

cooperate and compete to form an optimal tiling of the data distribution. The resulting arrangement—vectors aligned to red, green, blue, and yellow—forms a strong correspondence with the unique hues. Psychophysical characterizations of the unique hues indicate that they do not typically align with the cone-opponent axes, but instead (with the exception of red) often lie in the quadrants between them [9–11,51]. Moreover, color-opponent pairs (green/red and blue/yellow) are not necessarily collinear (e.g., the line between red and green does not pass through the origin). The adapted $m = 6$ model exhibits these properties: the hue-sensitive basis vectors do not always align with the cardinal cone-opponent directions, and the opponent pairs are not collinear. Note that subjective unique hue settings vary considerably across observers and datasets [51,52] and hue angles depend on the scaling of the cone-opponent axes. For this reason, we focus on the broader geometric correspondence rather than a precise match to any single set of unique-hue measurements. This general motif is maintained when adapting $m = 6$ models to each individual dataset (Fig. S3 of Supplement 1). It should be emphasized that, given the wide variety of ways six basis vectors could tile 3D space, it would be highly unlikely for this particular solution to emerge by chance. This result is examined more extensively below.

At $m = 7$, the basis vector aligned with black appears to split into two, each with a slight hue component, resulting in two vectors aligned to darker colors (i.e., brown and dark blue/purple). The other vectors remain largely unchanged from the $m = 6$ solution, albeit with the yellow and blue vectors nudged slightly toward white.

Finally, in the $m = 8$ model, an additional vector emerges within the chromatic plane aligned with violet/purple, enabling the dark blue/purple vector from $m = 7$ to revert to black. The directions of the other $m = 7$ vectors remain mostly unchanged.

Taken together, these results illustrate a progression in how sparse coding models with increasing degrees of overcompleteness adapt to the structure of the LMS distribution. All of the models form complete representations of the three-dimensional space, and all are adapted to the statistics of the data. They differ, however, in how parsimoniously they tile the heavy-tailed structure of the data. The four-vector model provides a minimal but awkward fit, requiring the achromatic vectors to participate in representing chromatic variation. The five-vector model gives a more natural separation between chromatic and achromatic structure, with three basis vectors spanning the chromatic plane and two aligned with the achromatic axis. The six-vector model is distinguished by the emergence of four chromatic vectors aligned with red, green, blue, and yellow, together with two achromatic vectors aligned with black and white. As shown below, this arrangement is unique among the other learned bases in producing three mutually exclusive opponent pairs that can be parsimoniously combined to form a three-dimensional opponent color space. Models with seven or eight vectors further subdivide this representation, suggesting finer-scale chromatic or luminance-related subcategories, but lacking a simple opponent organization.

While Fig. 3 reveals how the basis vectors $\mathbf{a}_i$ are arranged in the sphered color space, the manner in which color is encoded by these vectors requires examining how the corresponding




Fig. 4. Response maps of latent variable MAP activations s_i as a function of 2D location within the Munsell color chart (top). Each panel shows the iso-response contours for each s_i within a particular sparse coding model, with m increasing top to bottom. Solid contour lines indicate response levels from 55% to 95% of each latent variable's maximum, in 10% increments. Contour color indicates the direction of the corresponding basis vector $\mathbf{a}_i$ in sphered color space. The four-basis-vector model places strong weight on its two largely achromatic basis vectors to describe the blue/green directions. This is due to their opposing nature and the lack of other basis vectors that could be used to describe this region of color space. To show that these vectors are indeed sensitive to luminance variation, their iso-response contours are shown at 20% and 30% of each coefficient's maximum in dashed lines.

latent variables s_i activate in response to stimuli placed at different points in the space. These activations are computed via the MAP-inference of Eqs. (2) and (3), and hence are nonlinear functions of the stimulus vector $\mathbf{x}$. Visualizing these response functions, $s_i(\mathbf{x})$, within the 3D sphered color space is problematic, so we instead plot each latent variable's response as a function of its 2D location within the Munsell color chart, a color standard used in cross-cultural linguistic studies of color naming. These plots are shown in Fig. 4.

The progression of color tuning from broad to narrow as basis vectors are added to the model mirrors the increasingly refined tiling of the chromatic plane seen in Fig. 3. It is also reminiscent of the manner in which color terms are added to a color vocabulary. When additional vectors are added to the model, they subdivide a region that was previously represented by a single vector. For instance, the transition from five to six basis vectors introduces a new basis vector that separates blue and green, analogous to how languages with five basic color terms might introduce a distinct term for green or blue to resolve prior ambiguity. Moreover, earlier-acquired basis vectors tend to persist in models with more vectors. For example, vectors that

represent white, black, red, blue, yellow, and green are largely preserved in the seven- and eight-vector models.

E. Sparse Coding Account of the Unique Hues

Among the models discussed in the previous section, the six-basis-vector model holds particular interest due to its alignment with the unique hues. Given the extensive literature surrounding these hues and their foundational status in color science, we study this model in detail. In doing so, we demonstrate why this basis is useful for representing the statistics of color in natural scenes, and we draw similarities between the properties of this model and the phenomenology of the unique hues.

To understand why the six-basis-vector model arrives at its solution, it is important to consider how the sparse coding model represents data. In the model, any data point $\mathbf{x}$ is represented by scaling each unit-norm vector $\mathbf{a}_i$ according to its latent variable s_i and then adding them together. Graphically, this is equivalent to positioning the rescaled vectors end-to-end to reach the data point. Thus, the sum of latent variables [sparsity penalty in Eq. (2)] can be interpreted as the total path length traversed by basis vectors to describe a data point. The puzzle solved by the sparse coding model is how to arrange the basis vectors so that, when averaged over the distribution of the data, the total path length is minimized. In doing so, the model finds a configuration that efficiently spans the LMS activation distribution. The emergence of vectors resembling the unique hues can thus be understood as a consequence of the structure of the data and these model assumptions.

To demonstrate and verify this intuition quantitatively, we compare the efficiency and reconstruction quality of the unique-hue basis to those of two alternative six-vector bases [Fig. 5(a)]. One alternative is a canonical cone-opponent basis aligned with the axes of the spheroid color space (basis 1). The other is a basis roughly aligned with teal, lime, orange, and purple (basis 2), following [16], which used these hues to demonstrate that human subjects can use such a basis to describe color, and because they served as the subject of debate with respect to the unique hues [17,53]. Note that both alternative bases retain two achromatic vectors (black and white), ensuring a fair comparison in the full 3D space.

Each model's performance is evaluated by sweeping over a range of values for sparsity parameter λ and plotting the joint reconstruction error (mean-squared error; MSE) and average ℓ_1 norm of latent variables. As shown in Fig. 5(b), the adapted basis achieves consistently lower MSE and lower ℓ_1 norm across different settings of λ , indicating a more efficient representation under the model assumptions.

To further evaluate how this efficiency depends on the specific orientation of basis vectors within the chromatic plane, each basis is rotated in azimuth [Fig. 5(c)], and the total energy E [Eq. (2)] is computed as a function of rotation angle. As shown in Fig. 5(d), the adapted basis reaches a minimum at zero rotation, providing further confirmation that it is better matched to the structure of the data in comparison to a wide variety of alternatives.




Fig. 5. Comparison of the adapted six-basis-vector model with alternative models containing six basis vectors. (a) The three bases analyzed, projected onto the chromatic plane. All bases span the full 3D space and include two achromatic vectors (white and black). Left to right: the adapted basis, a cardinal cone-opponent basis, and a teal-lime-orange-purple basis. (b) Reconstruction quality (mean squared error) versus sparsity (average ℓ_1 norm) curves for each basis, sweeping across sparsity thresholds λ . Across thresholds, the adapted basis consistently achieves lower MSE and greater sparsity. (c) Each basis is rotated in azimuth (about the vertical axis), altering the hue tuning of the chromatic vectors. (d) Energy (the sum of MSE and ℓ_1 sparsity weighted by λ) as a function of azimuth rotation [see (c)]. The minimal-energy configuration occurs for the adapted basis at zero rotation.

F. Emergence of Color Opponency

Opponent pairs of unique hues are regarded as being perceptually opposite (e.g., there is no color that is simultaneously both red and green). However, opponent pairs of unique hues are not necessarily opposite to one another in the cone-opponent plane (i.e., collinear) [8–11]. This raises the question of why these directions in color space are perceived as opposites, and what mathematical principles could reproduce such an opponent organization.

We find that inference in the sparse coding model gives rise to mutual exclusivity between latent variables that represent perceptually opponent colors, as well as the achromatic (white and black) latent variables (Fig. 6). Figure 7(a) illustrates how inference represents stimuli within the chromatic plane by plotting




Fig. 6. Mutual exclusivity between opponent latent variables. Scatter plots show pairwise activations of inferred coefficients from the six-basis-vector sparse coding model for 10,000 randomly sampled LMS activations. Each axis corresponds to the activation of a latent variable associated with the indicated color direction. Opponent pairs (red–green, blue–yellow, and black–white) exhibit mutual exclusivity due to sparse inference, whereas non-opponent pairs can coactivate.

iso-response contours of the chromatic latent variable representation. The nonlinear interactions between latent variables s_i in the LCA inference dynamics [Eq. (8)] cause their iso-response contours to curve, producing a complete and evenly distributed encoding of the chromatic plane without any overlap between iso-response contours belonging to an opponent pair [Fig. 7(a)].

The interactions between latent variables during inference are mediated by the Gram matrix of inner-products between basis vectors $\mathbf{G}_{ij} = \mathbf{a}_i^T \mathbf{a}_j$ [Eq. (8)]. Positive entries in the Gram matrix lead to inhibitory interactions: when two basis vectors are similar, activity in one's latent variable suppresses the other. For example, the green and yellow basis vectors have strong positive overlap in the Gram matrix [Fig. 7(c)], leading to curved iso-response contours that bend away from each other [Fig. 7(a)]. This inhibitory interaction also ensures that pairs of basis vectors which represent opponent colors are mutually exclusive, as the contours of opponent pairs never cross despite a lack of collinearity or orthogonality.

Interestingly, negative entries of the Gram matrix (e.g., between blue and red) result in excitatory interactions between pairs of basis vectors. This interaction causes basis vectors to be recruited to represent regions of the chromatic plane despite the points within those regions having a negative projection on the basis vector. This prevents regions of input space from being represented by a single basis vector and thus enables a more even and complete tiling of the chromatic plane.

Somewhat paradoxically, opponent-pairs of basis vectors have large negative inner products, thus leading to an excitatory interaction. How is it that they remain mutually exclusive? The answer lies in the sparsity penalty. Inference selects the sparsest combination of active units to explain the data. Since simultaneous activation of opponent units would decrease sparsity, only one member of an opponent pair is used at a time.

However, the excitatory interactions between opponent pairs point to a broader issue: the model assumes a factorial (independent) prior over coefficients, but the empirical mutually exclusive distribution of inferred coefficients is not factorial (see Fig. 6). This emerges from the interaction between sparsity and the structure of the data. A direction for future work is to explicitly incorporate this mutual exclusivity into a non-factorial prior (see Section 3).

Importantly, mutual exclusivity among opponent units is not imposed by the model but instead emerges naturally from inference in the generative model [Eq. (1)]. Not all nonlinear transformations exhibit this property. Consider, for example, a representation generated by linear projection followed by rectification. Because responses in such a model are computed purely through linear projection, the iso-response contours in Fig. 7(b) form straight lines perpendicular to their corresponding basis vectors. The lack of orthogonality among the basis vectors leads to regions of the chromatic plane that are represented by combinations of hues that are considered perceptually opposed. For example, red and green units both respond in parts of the fourth quadrant, and blue and yellow responses overlap in the third quadrant. There are also significant regions of the chromatic plane encoded by only the blue vector or only the red vector, leading to metamerism: many distinct chromatic stimuli yielding identical responses. These problems cannot be solved by simply increasing the threshold (which would exacerbate metamerism) or introducing alternative pointwise nonlinearities.

Finally, because inference forms three mutually exclusive pairs of latent variables, they define a perceptually parsimonious $\mathbb{R}^3$ space in which each axis corresponds to a red–green, blue–yellow, or black–white opposition [Fig. 1(c)]. This structure reflects the phenomenological organization of color experience and is unique to the six-basis-vector model. The emergence of




Fig. 7. (a), (b) Iso-response contours (lines) colored by the corresponding basis vector. The hue saturation of the lines is set according to response magnitude (color bars at right). Contours are shown for both sparse inference (a) and linear projection (b). The linear projection representation is given by $s_i = g(\mathbf{a}_i^T \mathbf{x})$, where $\mathbf{x}$ is a point in the sphered color space, and $g(\cdot)$ is an element-wise rectifying nonlinearity [as in Eq. (9)]. Representations formed by sparse inference and linear projection are numbered $s_1 - s_4$ according to the azimuth angle of their basis vectors, from smallest to largest. (c) Gram matrix showing inner products between basis vectors, with positive (inhibition) and negative (excitation) values. (d) Tuning curves of the inferred representation (solid lines) in response to stimuli varying only in hue angle, compared to a linear projection with rectification (dashed lines).

these three opponent pairs is not a generic consequence of using six basis vectors, as this property depends on the specific orientations that the model learns from the data. For instance, a model with six basis vectors could instead represent three chromatic directions, one achromatic-increment direction (e.g., white), and two distinct decrement directions (e.g., black and brown); in such a case, pairwise mutual exclusivity would not occur, as white would be opponent to both black and brown. The other models also fail to construct a three-pair mutually exclusive opponent-color space: the four-vector model contains no opponent pairs, the five-vector model only one (black–white), and the seven-vector model introduces deviations from pairwise opponency, such as white being mutually exclusive with both black and brown. Thus, the six-basis-vector model uniquely combines numerical sufficiency with statistically driven directions to recover a perceptually meaningful, 3D opponent-color space.

G. Model Tuning Curves Resemble Perceptual Hue Scaling Functions

A well-established method for probing color appearance is hue scaling, in which participants rate the perceived proportions of a set of colors for a battery of colored stimuli that tile a color space [8,16]. To draw parallels with such psychophysical characterizations of the unique hues, we examined the responses of the six-basis-vector model to stimuli that vary only in azimuth angle within the chromatic plane of the sphered color space.

Figure 7(d) shows the tuning curves of the inferred sparse coding representation to unit-norm chromatic stimuli sweeping

through azimuth. For comparison, the tuning curves of the same basis under a linear projection [i.e., $g(\mathbf{A}^T \mathbf{x})$ as seen in Fig. 7(a)] are also plotted (dashed lines). Under linear projection, all curves have the same shape and differ only by phase shifts. By contrast, the inferred representation (solid lines) exhibits non-uniform tuning widths due to interactions between units during inference. For example, excitatory interactions broaden responses near 60° , while inhibitory suppression narrows responses around 240° . These interactions produce red and blue tuning curves that are broader than those of green and yellow. This resembles the asymmetry reported by De Valois *et al.* [8], in which blue and red were used to describe a broader range of azimuth angles than green and yellow in human hue-scaling judgments. Thus, the correspondence is not only in the learned basis directions, but also in how sparse coding inference combines those directions to describe intermediate hues.

3. DISCUSSION

The two main contributions of this paper are (1) the finding of strongly non-Gaussian structure of LMS activations in response to a large database of natural images and (2) showing that a linear generative model adapted to this data aligns with the unique hues. A key observation enabling the second contribution is the asymmetry in the data, particularly within the chromatic plane, which motivates the use of a non-negative code for its representation. Nonnegativity, when combined with sparse coding inference, introduces mutual exclusivity between coefficients and yields a color representation that mirrors human perceptual judgments. It is remarkable that not only unique hues but also

the structure of their interactions arise from fitting a single, linear generative model, with few other assumptions, to the data.

A. Implications and Model Predictions

When framed as a probabilistic generative model [49], the sparse coding model can be seen as constituting a prior over color appearance acquired from exposure to the natural environment. This raises the question of whether one can find evidence of the influence of this prior in judgments of color appearance, similar to how priors have been shown to influence judgments in other tasks [54]. For example, given a noisy or uncertain observation of a color, the model would predict an observer's estimate of the color to be biased toward regions of higher probability in the prior [55]. Alternatively, since the sparse coding model is an approximation, and the full color distribution is relatively low-dimensional, it may be feasible to use the empirical distribution (Fig. 2) directly to test whether it predicts color judgments under uncertainty.

Another question that arises from the sparse coding model is whether one should expect individual neurons in the visual system to exhibit tuning to the unique hues similar to that of the latent variables of the model (Fig. 4). This question is difficult to answer at present because we lack a theory for how space and color should be jointly represented. A widely held view is that the visual system aims to disentangle image data into independent factors of variation such as illumination, object shape, and reflectance properties [56]. Previous work applying sparse coding or ICA to color natural image patches recovers spatial Gabor-like basis functions with color-opponent structure [29,31,47]. However, information about color, shape, and illumination still remains entangled in such a representation since the latent variables co-vary with each of these properties. We conjecture that disentangling color and shape will require more complex, multistage models that have the ability to factorize components [57] as opposed to the simple type of linear generative models explored here and elsewhere. We speculate that the latent variables of such disentangled representations could be effectively represented and manipulated internally not by single neurons but via high-dimensional neural populations, as in vector symbolic architectures [58,59], which could explain the failure to find neurons tuned to the unique hues [14].

B. Relation to Prior Work

Other multistage models of chromatic information processing have also been proposed. Doi *et al.* [31] applied a similar three-stage model to spatiochromatic natural images, which were sampled by a realistic cone mosaic. The model consisted of a sphering step followed by ICA, which were both adapted to the structure in the spatiochromatic natural image patches. The ICA model derived cone-selective receptive fields; however, it is unclear to what extent these filters relate to color perception. De Valois and De Valois [60] demonstrated how the visual system could transform color representations in a multistage model proceeding from LMS to cone-opponent to red-green/blue-yellow color-opponent dimensions, given

biophysical constraints such as relative cone ratios and the spatially symmetric center-surround receptive field structure in the retina. Here we show how the specific computations in these stages can be derived within a first-principles framework applied to image statistics.

Yendrikhovskij [34] applied clustering algorithms to pixels from natural images in a perceptually uniform color space and found that the resulting clusters aligned with so-called “universally named” colors. However, later studies struggled to replicate this result [61,62]. Given the structure of the data shown in Figs. 2(c)–2(e), it is not clear that a clustering model per se is appropriate for capturing the structure of natural color statistics (see also below).

Philipona and O'Regan [63] provide an alternative account of the unique hues based on the observation that certain surface reflectance spectra, when filtered through the cone absorption functions, remain relatively invariant under changes in illumination. They show that these singularities in color space, which would be most informative of surface reflectance, correspond to the unique hues. This finding suggests that the asymmetric and heavy-tailed structure in the LMS distribution reported here may be related to the relative invariance of certain reflectances to illumination. Investigating this relationship would be an interesting direction for future work.

Webster and Mollon [22] found considerable variation in natural-image chromaticities along a blue–yellow direction, with greater variation along S-(L + M) than along L-M. This broad anisotropy has been replicated here and elsewhere [27]. Specifically, they found that the dominant axis of chromatic variation was not aligned with the cone-opponent axes. Direct comparison between their color distribution and the one reported here is limited by two differences. First, each image in our dataset was rescaled by its mean cone response (see Section 4), whereas their analysis did not include an analogous normalization. Second, their post-receptoral space differs slightly from the PCA-derived post-receptoral space used here and, more importantly, uses different scalings of the chromatic axes. They rescale the axes according to estimated detection thresholds, whereas we rescale the axes to have unit variance. Nonetheless, the sparse-coding models learn basis vectors aligned with the blue–yellow directions (Fig. 3). This off-axis bias could reflect scene content, as both datasets contain mixtures of foliage and sky. Understanding how the non-Gaussian structure of the LMS distribution arises from image content remains an important direction for future work.

C. Bases versus Categories

The unique hues are in some cases regarded as discrete color categories (as in color naming) and in other cases as a basis for representing the full continuum of colors. The sparse coding model is an example of the latter. Whereas categorization produces a discrete code that assigns many points in stimulus space to a single category, the sparse coding model produces a continuous code by linearly combining basis vectors to describe any point in color space, potentially without loss. However in the extreme case where sparsity is so heavily encouraged that only a single basis vector is used to describe a given point in color space, it becomes equivalent to a vector quantization model similar

to categorization or clustering (see [64] for a useful discussion of the relation between sparse coding and clustering). Previous attempts to measure categorical effects in color perception have produced mixed evidence for categorical boundaries consistent with the unique hues [65]. It is not clear that the sparse coding model would produce such categorical effects, especially if the prior were to be modified to account for facilitatory interactions between latent variables (see below).

Interestingly, the sparse coding model also captures the atomistic property of the unique hues originally noted by Hering, since a color aligned with any one basis vector would be described by that vector alone, whereas a color lying between basis vectors would be described as a linear combination of them.

D. Model Extensions and Future Work

There are natural extensions to the sparse coding model. First, while we assume a factorial prior over the latent variables, the inferred representations clearly deviate from this assumption. This is most notable through mutual exclusivity between opponent-color units. Future work could explore incorporating structured priors that explicitly model mutual exclusivity. Second, in this work, we compute the latent variables via MAP estimates rather than sampling from the full posterior distribution. A disadvantage of this approach is that these estimates are biased toward zero due to the sparse prior, requiring the basis function norms to be constrained to one (otherwise they would grow without bound). By instead sampling from the posterior (e.g., via Langevin sampling), the model could potentially learn the basis vector norms along with parameters of the prior, such as excitatory and inhibitory couplings between latent variables [66].

Lastly, cone spectral sensitivities vary greatly across the animal kingdom [67]. Prior work has asked whether primate cone sensitivities are optimized for sampling natural spectra [68]. The results presented here are all based on spectral sensitivities of the trichromatic primate visual system; however, this theoretical approach is general and can be applied to any species with spectrally distinct photoreceptors. In this sense, cone sensitivities define the photoreceptor basis through which spectra are measured, whereas the sparse-coding model addresses how the resulting cone responses may be organized into a post-receptoral color representation. Human spectral sensitivities vary across observers, yet unique-yellow settings can remain remarkably stable [69]. This raises the possibility that the higher-order structure captured here is robust to some variation in cone spectral sensitivities, a question that could be tested by repeating the present analysis under different photoreceptor sensitivities. Theoretically this could make sense in terms of a compressive sensing strategy, where the cones compress spectra into a low-dimensional representation from which higher-order (sparse) structure can still be recovered for a broad range of variations in the choice of cone absorption functions [70].

More broadly, applying the same framework across species could test whether similar coding principles operate in visual systems with very different photoreceptor bases. Recent findings on color vision in *Drosophila* have revealed a very similar coding

strategy to what we have presented here: photoreceptor activations are initially decorrelated via horizontal cells [71], and then a later stage is believed to use a recurrent non-negative encoding strategy, similar to LCA, to produce hue-selective neurons with curved iso-response contours similar to those in Fig. 7(b) [72]. To what extent these hue tuning properties can be explained by natural image statistics remains an open question. Nonetheless, these empirical findings bear a striking resemblance to the model proposed here, and their convergence offers the opportunity to test whether principles of efficient coding and perceptual inference can provide an account for color perception across the animal kingdom.

4. METHODS

A. Dataset of Simulated LMS Activations to Natural Images

We constructed a large dataset of simulated LMS activations in response to natural images by combining five calibrated LMS and hyperspectral datasets, the details of which are provided in Table 1. The collated dataset contains over 225 million LMS activations arising from 503 images. The motivation behind combining datasets is that it allows for a more diverse sampling of visual environments and reduces potential biases from any single camera, environmental setting, or selection bias of the investigator. The hyperspectral datasets were converted to cone responses using the ten-degree cone fundamentals from [77].

We eliminated from the dataset any images that contain man-made objects. We removed these images because the colors of man-made objects could inherit the very biases in human color perception that we seek to study, which would confound our results. That said, even scenes of only naturalistic objects could contain biases of what the photographer thought was worth photographing. It is hard to rule out this bias, and it could be present in any dataset. Additionally, only images in the “naturalistic” category from the Barcelona dataset [73] were used, and the gray sphere in the left portion of each image was removed by cropping out the leftmost 200 pixels. We did not include the large simulated LMS cone natural image dataset of [78] as its PCA-summarized statistics were substantially different from the other datasets. Understanding the reasons for this and its implications deserves additional study.

Table 1. Details of the Natural Image Datasets Used for Simulating LMS Cone Activations^a

Dataset	Type	N. Images	Location(s)
Kyoto natural images [31]	Calibrated LMS	62	Japan
Barcelona calibrated images [73]	Calibrated LMS	237	Spain
Ruderman <i>et al.</i> [27]	Hyperspectral	12	United States and Australia
2022 NTIRE spectral recovery [74]	Hyperspectral	165	Israel
Foster and Reeves [75,76]	Hyperspectral	27	Portugal

^aEach dataset is listed with its data type, number of images, and collection location(s).

With the exception of the Kyoto dataset, all of the calibrated LMS datasets are proportional to linear cone absorbances and can be thought of as cone photon isomerizations, which will be nonlinearly transformed to cone outputs. Following Doi *et al.* [31] (see also [33]), the transform from absorbance v to nonlinear response r is modeled using an exponential saturation function $r = 1 - e^{-kv}$ as proposed by Baylor *et al.* [79]. Interestingly, this function describes the cumulative density function of the exponential distribution $p(v) = ke^{-kv}$, which is the maximum entropy distribution for a non-negative variable v with mean $1/k$. We observed that cone absorbance distributions for a given image often resemble an exponential distribution, with k varying with cone type and illumination level. Therefore, by setting k equal to the inverse mean, $1/\langle v_i \rangle$, where $i \in \{L, M, S\}$ denotes a particular cone type, an approximately uniform distribution over $[0, 1]$ will be produced on the cone outputs. (A similar line of reasoning was used by Laughlin to derive the optimal nonlinear response function for neurons in the fly visual system [36]). Moreover, this choice of k effectively applies von Kries adaptation to the images, in which each cone type's response is rescaled independently according to its mean [80]. Mathematically, the transformation from linear cone inputs v_i to nonlinear outputs r_i is as follows:

$$r_i = 1 - \exp\left(-\frac{v_i}{\langle v_i \rangle}\right). \quad (4)$$

Note that the Kyoto dataset uses a similar function to convert to nonlinear outputs; however, they set k such that the median of r_i is equal to 0.5 (see [31] for details). As a final step of processing, the mean is subtracted from each image (computed over the spatial and channel dimensions), such that the mean over the entire dataset is zero prior to the redundancy reduction step, which is described in the subsequent section.

B. PCA for Redundancy Reduction

Following Ruderman *et al.* [27], we apply PCA to remove second-order statistical redundancies from simulated LMS activations.

Given an LMS cone activation vector $\mathbf{r} = (r_L \ r_M \ r_S)^T$, we seek a linear transformation $\mathbf{W} \in \mathbb{R}^{3 \times 3}$ such that

$$\mathbf{x} = \mathbf{W}\mathbf{r} \quad \text{with} \quad \langle \mathbf{x}\mathbf{x}^T \rangle = \mathbf{I}, \quad (5)$$

where $\langle \cdot \rangle$ denotes the expectation over the dataset, and $\mathbf{I}$ is the identity matrix.

Such a linear transformation can be computed using PCA. The covariance matrix of the LMS inputs is decomposed as

$$\langle \mathbf{r}\mathbf{r}^T \rangle = \mathbf{E}\mathbf{\Lambda}\mathbf{E}^T, \quad (6)$$

where $\mathbf{E}$ contains the eigenvectors (principal components) and $\mathbf{\Lambda} = \text{diag}(\sigma_1^2, \sigma_2^2, \sigma_3^2)$ is the diagonal matrix of their corresponding variances, sorted such that $\sigma_1^2 > \sigma_2^2 > \sigma_3^2$.

To decorrelate and rescale the data to have unit variance in all directions, the data are first rotated to the principal component basis using $\mathbf{E}^T$ and then rescaled by the inverse standard deviation using $\mathbf{\Lambda}^{-1/2} = \text{diag}(\sigma_1^{-1}, \sigma_2^{-1}, \sigma_3^{-1})$. Combined, this yields the linear transformation:

$$\mathbf{W} = \mathbf{\Lambda}^{-1/2}\mathbf{E}^T. \quad (7)$$

This maps LMS activations into the sphered color space.

C. Locally Competitive Algorithm (LCA)

We solve Eqs. (2) and (3) using a non-negative variant of LCA [81]. LCA evolves subthreshold variables u_i whose thresholded outputs s_i are the sparse coefficients:

$$\tau \dot{u}_i + u_i = \mathbf{a}_i^T \mathbf{x} - \sum_{j \neq i} \mathbf{G}_{ij} s_j, \quad (8)$$

$$s_i = \max(0, u_i - \lambda) = g(u_i), \quad (9)$$

where $\mathbf{G}_{ij} = \mathbf{a}_i^T \mathbf{a}_j$. The rectifying nonlinearity $g(\cdot)$ enforces the ℓ_1 -sparsity penalty on the coefficients while imposing an infinite energy barrier at $s_i < 0$, satisfying the nonnegativity constraint. Letting the dynamics settle to a steady state yields the coefficients $\mathbf{s}$ that minimize E , corresponding to the MAP estimate under the model. The basis is learned by stochastic gradient descent on the energy E with respect to $\mathbf{A}$, using the inferred coefficients $\mathbf{s}$. This yields the following learning rule:

$$\Delta \mathbf{a}_i = \eta \left[\mathbf{x} - \sum_{j=1}^m \mathbf{a}_j s_j \right] s_i. \quad (10)$$

The update rate η is set small so that the energy decreases continuously until convergence. After each update, the basis vectors are normalized to have unit norm: $\|\mathbf{a}_i\|_2 = 1$ for all i .

The sparsity parameter λ was adjusted during learning such that a constant signal-to-noise ratio (SNR) was maintained in the reconstruction error. The SNR is a function of mean squared error (MSE) where $\text{MSE} = \|\mathbf{x} - \mathbf{A}\mathbf{s}\|_2^2$. SNR in decibels (dB) is computed as follows:

$$\text{SNR} = 10 \log_{10} \text{MSE}^{-1}. \quad (11)$$

The sparsity parameter λ was set such that 16 dB SNR was maintained for all the bases shown in this paper. We did not observe significant changes to the optimal bases when λ enforced higher or lower SNRs.

Funding. United States Air Force Office of Scientific Research (FA9550-21-1-0230).

Acknowledgment. A.B. and B.A.O. are supported by the Air Force Office of Scientific Research (AFOSR) MURI grant. The authors gratefully acknowledge the feedback and support of the MURI team. A.B., E.P.F., and B.A.O. designed research; A.B., B.A.O., and E.P.F. performed research; A.B., E.P.F., and B.A.O. analyzed data; and A.B. and B.A.O. wrote the paper.

Disclosures. The authors declare no competing interests.

Data availability. The natural image databases underlying the results are available for download online. Code is available upon request.

Supplemental document. See Supplement 1 for supporting content.

REFERENCES

1. H. Helmholtz, *Handbuch der Physiologischen Optik*, Vol. 2 of Die Lehre von den Tonempfindungen als physiologische Grundlage für die Theorie der Musik (Leopold Voss, 1860).
2. E. Hering, *Zur Lehre Vom Lichtsinne: Sechs Mittheilungen an die Kaiser* (Akademie der Wissenschaften, 1878).

3. J. Krauskopf, D. R. Williams, and D. W. Heeley, "Cardinal directions of color space," *Vision Res.* **22**, 1123–1131 (1982).
4. A. M. Derrington, J. Krauskopf, and P. Lennie, "Chromatic mechanisms in lateral geniculate nucleus of macaque," *J. Physiol.* **357**, 241–265 (1984).
5. K. K. De Valois and R. L. De Valois, "Chapter 4: Color vision," in *Seeing: Handbook of Perception and Cognition*, K. K. De Valois, ed. (Academic, 2000), pp. 129–175.
6. M. D. Fairchild, "Color models and systems," in *Handbook of Color Psychology*, A. J. Elliot, M. D. Fairchild, and A. Franklin, eds. (Cambridge University, 2015), pp. 9–26.
7. J. M. Wolfe, K. R. Kluender, and D. M. Levi, "The perception of colour," in *Sensation and Perception*, 6th International Edition (Sinauer Associates, 2022), pp. 131–165.
8. R. L. De Valois, K. K. De Valois, E. Switkes, *et al.*, "Hue scaling of iso-luminant and cone-specific lights," *Vision Res.* **37**, 885–897 (1997).
9. M. A. Webster, E. Miyahara, G. Malkoc, *et al.*, "Variations in normal color vision II Unique hues," *J. Opt. Soc. Am. A* **17**, 1545–1555 (2000).
10. S. M. Wuerger, P. Atkinson, and S. Cropper, "The cone inputs to the unique-hue mechanisms," *Vision Res.* **45**, 3210–3223 (2005).
11. L. E. Wool, S. J. Komban, J. Kremkow, *et al.*, "Salience of unique hues and implications for color theory," *J. Vis.* **15**(2):10, 10 (2015).
12. C. M. Stoughton and B. R. Conway, "Neural basis for unique hues," *Curr. Biol.* **18**, 698–699 (2008).
13. J. D. Mollon, "A neural basis for unique hues?" *Curr. Biol.* **19**, R441–R442 (2009).
14. B. R. Conway, S. Malik-Moraleda, and E. Gibson, "Color appearance and the end of Hering's Opponent-Colors Theory," *Trends Cognit. Sci.* **27**, 791–804 (2023).
15. L. Forder, J. Bosten, X. He, *et al.*, "A neural signature of the unique hues," *Sci. Rep.* **7**, 42364 (2017).
16. J. M. Bosten and A. E. Boehm, "Empirical evidence for unique hues?" *J. Opt. Soc. Am. A* **31**, 385–393 (2014).
17. B. A. C. Saunders and J. Brakel, "Are there nontrivial constraints on colour categorization?" *Behav. Brain Sci.* **20**, 167–179 (1997).
18. J. D. Mollon and G. Jordan, "On the nature of unique hues," in *John Dalton's Colour Vision Legacy*, C. Dickenson, I. J. Murray, and D. Carden, eds. (Taylor & Francis, 1997), pp. 381–392.
19. A. Valberg, "Unique hues: an old problem for a new generation," *Vision Res.* **41**, 1645–1657 (2001).
20. K. Jameson and R. G. D'Andrade, "It's not really red, green, yellow, blue: an inquiry into perceptual color space," in *Color Categories in Thought and Language*, C. L. Hardin and L. Maffi, eds. (Cambridge University, 1997), pp. 295–319.
21. R. N. Shepard, "The perceptual organization of colors: an adaptation to regularities of the world," in *Readings in Color*, A. Byrne and D. R. Hilbert, eds. (MIT, 1997), pp. 311–356.
22. M. A. Webster and J. D. Mollon, "Adaptation and the color statistics of natural images," *Vision Res.* **37**, 3283–3298 (1997).
23. J. D. Mollon, "Monge: the Verriest Lecture, Lyon, July 2005," *Visual Neurosci.* **23**, 297–309 (2005).
24. I. Rosenthal, S. Ratnasingham, T. Haile, *et al.*, "Color statistics of objects, and color tuning of object cortex in macaque monkey," *J. Vis.* **18**(11):1, 1–1 (2018).
25. A. E. Skelton, J. Maule, S. Floyd, *et al.*, "Effects of visual diet on colour discrimination and preference," *Proc. R. Soc. B* **291**, 20240909 (2024).
26. L. Hedjar, M. Toscani, and K. R. Gegenfurtner, "Importance of hue: color discrimination of three-dimensional objects and two-dimensional discs," *J. Opt. Soc. Am. A* **42**, B296–B304 (2025).
27. D. L. Ruderman, T. W. Cronin, and C.-C. Chiao, "Statistics of cone responses to natural images: implications for visual coding," *J. Opt. Soc. Am. A* **15**, 2036–2045 (1998).
28. P. O. Hoyer and A. Hyvärinen, "Independent component analysis applied to feature extraction from colour and stereo images," *Netw. Comput. Neural Syst.* **11**, 191–210 (2000).
29. T. Wachtler, T.-W. Lee, and T. J. Sejnowski, "Chromatic structure of natural scenes," *J. Opt. Soc. Am. A* **18**, 65–77 (2001).
30. T.-W. Lee, T. Wachtler, and T. J. Sejnowski, "Color opponency is an efficient representation of spectral properties in natural scenes," *Vision Res.* **42**, 2095–2103 (2002).
31. E. Doi, T. Inui, T.-W. Lee, *et al.*, "Spatiochromatic receptive field properties derived from information-theoretic analyses of cone mosaic responses to natural scenes," *Neural Comput.* **15**, 397–417 (2003).
32. M. S. Caywood, B. Willmore, and D. J. Tolhurst, "Independent components of color natural scenes resemble V1 neurons in their spatial and color tuning," *J. Neurophysiol.* **91**, 2859–2873 (2004).
33. T. Wachtler, E. Doi, T.-W. Lee, *et al.*, "Cone selectivity derived from the responses of the retinal cone mosaic to natural scenes," *J. Vis.* **7**(8):6, 1–14 (2007).
34. S. N. Yendrikhovskij, "Computing color categories from statistics of natural images," *J. Imaging Sci. Technol.* **45**, 409–417 (2001).
35. E. P. Simoncelli and B. A. Olshausen, "Natural image statistics and neural representation," *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
36. S. Laughlin, "A simple coding procedure enhances a neuron's information capacity," *Z. Naturforsch. C* **36**, 910–912 (1981).
37. D. J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," *J. Opt. Soc. Am. A* **4**, 2379–2394 (1987).
38. D. J. Graham, D. M. Chandler, and D. J. Field, "Can the theory of 'whitening' explain the center-surround properties of retinal ganglion cell receptive fields?" *Vision Res.* **46**, 2901–2913 (2006).
39. C. F. Stromeyer, III, G. R. Cole, and R. E. Kronauer, "Second-site adaptation in the red-green chromatic pathways," *Vision Res.* **25**, 219–237 (1985).
40. B. A. Wandell, "Color measurement and discrimination," *J. Opt. Soc. Am. A* **2**, 62–71 (1985).
41. C. D. Hendley and S. Hecht, "The colors of natural objects and terrains, and their relation to visual color deficiency," *J. Opt. Soc. Am.* **39**, 870–873 (1949).
42. R. Baddeley, "Searching for filters with 'interesting' output distributions: an uninteresting direction to explore?" *Network* **7**, 409–421 (1996).
43. A. Yoonessi, F. A. Kingdom, and S. Alqawlaq, "Is color patchy?" *J. Opt. Soc. Am. A* **25**, 1330–1338 (2008).
44. D. J. Field, "What is the goal of sensory coding?" *Neural Comput.* **6**, 559–601 (1994).
45. A. J. Bell and T. J. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," *Neural Comput.* **7**, 1129–1159 (1995).
46. B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature* **381**, 607–609 (1996).
47. A. Hyvärinen, J. Hurri, and P. O. Hoyer, *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision* (Springer, 2009), Vol. **39**.
48. S. Ganguli and H. Sompolinsky, "Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis," *Annu. Rev. Neurosci.* **35**, 485–508 (2012).
49. B. A. Olshausen and D. J. Field, "Sparse coding with an over-complete basis set: a strategy employed by V1?" *Vision Res.* **37**, 3311–3325 (1997).
50. B. A. Olshausen and D. J. Field, "Sparse coding of sensory inputs," *Curr. Opin. Neurobiol.* **14**, 481–487 (2004).
51. G. Malkoc, P. Kay, and M. A. Webster, "Variations in normal color vision. IV. Binary hues and hue scaling," *J. Opt. Soc. Am. A* **22**, 2154–2168 (2005).
52. R. G. Kuehni, "Variability in unique hue selection: a surprising phenomenon," *Color Res. Appl.* **29**, 158–162 (2004).
53. J. Broackes, "Could we take lime, purple, orange, and teal as unique hues?" *Behav. Brain Sci.* **20**, 183–184 (1997).
54. T. L. Griffiths and J. B. Tenenbaum, "Optimal predictions in everyday cognition," *Psychol. Sci.* **17**, 767–773 (2006).
55. J. R. Macyszko, O. B. Kavcar, M. A. Crognale, *et al.*, "Asymmetries in hue percepts and early cortical color coding: evidence from chromatic visual evoked potentials," *J. Vision* **26**, 3 (2026).
56. E. H. Adelson and A. P. Pentland, "The perception of shading and reflectance," in *Perception as Bayesian Inference*, D. C. Knill and W. Richards, eds. (Cambridge University, 1996), pp. 409–423.
57. E. P. Frady, S. J. Kent, B. A. Olshausen, *et al.*, "Resonator networks, 1: an efficient solution for factoring high-dimensional, distributed representations of data structures," *Neural Comput.* **32**, 2311–2331 (2020).

58. P. Kanerva, "Hyperdimensional computing: an introduction to computing in distributed representation with high-dimensional random vectors," *Cognitive Comput.* **1**, 139–159 (2009).
59. D. Kleyko, M. Davies, E. P. Frady, *et al.*, "Vector symbolic architectures as a computing framework for emerging hardware," *Proc. IEEE* **110**, 1538–1571 (2022).
60. R. L. De Valois and K. K. De Valois, "A multi-stage color model," *Vision Res.* **33**, 1053–1065 (1993).
61. L. Steels and T. Belpaeme, "Coordinating perceptually grounded categories through language: a case study for colour," *Behav. Brain Sci.* **28**, 469–488 (2005).
62. N. Zaslavsky, C. Kemp, N. Tishby, *et al.*, "Communicative need in colour naming," *Cogn. Neuropsychol.* **37**, 312–324 (2020).
63. D. L. Philipona and J. K. O'Regan, "Color naming, unique hues, and hue cancellation predicted from singularities in reflection properties," *Visual Neurosci.* **23**, 331–339 (2006).
64. G. E. Hinton and Z. Ghahramani, "Generative models for discovering sparse distributed representations," *Philos. Trans. R. Soc. London B* **352**, 1177–1190 (1997).
65. C. Witzel and K. R. Gegenfurtner, "Are red, yellow, green, and blue perceptual categories?" *Vision Res.* **151**, 152–163 (2018).
66. M. Y.-S. Fang, M. Mudigonda, R. Zarccone, *et al.*, "Learning and inference in sparse coding models with Langevin dynamics," *Neural Comput.* **34**, 1676–1700 (2022).
67. D. Osorio and M. Vorobyev, "A review of the evolution of animal colour vision and visual communication signals," *Vision Res.* **48**, 2042–2051 (2008).
68. A. Lewis and L. Zhaoping, "Are cone sensitivities determined by natural color statistics?" *J. Vis.* **6**(3):8, 285–302 (2006).
69. J. Neitz, J. Carroll, Y. Yamauchi, *et al.*, "Color perception is mediated by a plastic neural mechanism that is adjustable in adults," *Neuron* **35**, 783–792 (2002).
70. C. J. Hillar and F. T. Sommer, "When can dictionary learning uniquely recover sparse data from subsamples?" *IEEE Trans. Inf. Theory* **61**, 6290–6297 (2015).
71. S. L. Heath, M. P. Christenson, E. Oriol, *et al.*, "Circuit mechanisms underlying chromatic encoding in Drosophila photoreceptors," *Curr. Biol.* **30**, 264–275 (2020).
72. M. P. Christenson, A. Sanz Diez, S. L. Heath, *et al.*, "Hue selectivity from recurrent circuitry in Drosophila," *Nat. Neurosci.* **27**, 1137–1147 (2024).
73. C. A. Parraga, J. Vazquez-Corral, and M. Vanrell, "A new cone activation-based natural images dataset," *Perception* **36**, 180 (2009).
74. B. Arad, R. Timofte, R. Yahel, *et al.*, "NTIRE 2022 spectral recovery challenge and data set," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 863–881.
75. D. Foster, K. Amano, and S. Nascimento, "50 reduced hyperspectral reflectance images," figshare (2022) [accessed 13 August 2024], https://figshare.manchester.ac.uk/articles/dataset/Fifty_hyperspectral_reflectance_images_of_outdoor_scenes/14877285.
76. D. H. Foster and A. Reeves, "Colour constancy failures expected in colourful environments," *Proc. R. Soc. B* **289**, 20212483 (2022).
77. A. Stockman and L. T. Sharpe, "The spectral sensitivities of the middle-and long-wavelength-sensitive cones derived from measurements in observers of known genotype," *Vision Res.* **40**, 1711–1737 (2000).
78. G. Tkačik, P. Garrigan, C. Ratliff, *et al.*, "Natural images from the birthplace of the human eye," *PLoS One* **6**, e20409 (2011).
79. D. A. Baylor, B. J. Nunn, and J. L. Schnapf, "Spectral sensitivity of cones of the monkey *Macaca fascicularis*," *J. Physiol.* **390**, 145–160 (1987).
80. G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae* (Wiley, 2000).
81. C. J. Rozell, D. H. Johnson, R. G. Baraniuk, *et al.*, "Sparse coding via thresholding and local competition in neural circuits," *Neural Comput.* **20**, 2526–2563 (2008).