



## Editorial

## Understanding the statistics of the natural environment and their implications for vision



The job of the visual system is to extract useful information about the 3D world from the 2D image projections formed by our eyes. How do the statistics of visual scenes—encompassing all relevant aspects of surface geometry, material properties and illumination—affect the architecture and function of the visual system and its performance on specified tasks?

Natural scenes are complex and varied, and since von Helmholtz (1867) at least, it has been understood that vision requires inference. The problems are ill-posed—multiple different scenes might account for any given set of image observations—although only one of these accounts is correct. To perform well in natural environments, the visual system must therefore somehow embody an understanding of the likelihood of these various accounts.

While Helmholtz expressed these ideas qualitatively, with the development of information theory in the mid-twentieth century, a view of visual perception as a fundamentally statistical task began to emerge (Attneave, 1954; Barlow, 1972; Brunswik & Kamiya, 1953): optimal visual coding and inference demands knowledge of the statistics of real-world environments and how projection onto the image affects these statistics. The vast cortical and subcortical resources devoted to vision attests to its adaptive value and predicts that neural coding will have evolved and developed to embody these statistics, in order to perform well in the environments we commonly encounter.

The importance of these statistics to understanding visual perception is by now generally agreed. But the problem of identifying these statistics and understanding how exactly they have shaped the visual system is daunting. The challenge springs from two fundamental difficulties. The first is the Curse of Dimensionality: the number of degrees of freedom needed to describe the possible covariation of visual scenes and images over space and time is so enormous as to defeat standard approaches for statistical estimation and inference. Hope lies in the evidence that the effective dimensionality of real scenes and images is much lower than the full dimensionality of the embedding space (Chandler & Field, 2007), underlining the importance of efficient coding concepts and dimensionality reduction to understanding visual perception. But this leads to the second fundamental difficulty: this lower-dimensional manifold appears to be highly curved, hence defeating standard linear systems approaches which dominated much of vision science in the 20th century.

Given these massive challenges, the road to understanding the statistical foundations of visual perception is clearly long and winding, and the journey has just begun. Nevertheless, the ten papers in this special issue collectively represent an important step

forward in understanding the statistics of natural scenes and how they have shaped visual coding, and in some cases at least, point to potential opportunities to overcome some of these challenges, a topic to which we will return at the end of this editorial.

### 1. Point statistics

Papers in the first half of this special issue address point statistics of the visual image, characterizing the local statistics of luminance and color contrasts. By marginalizing over (almost) all space and time dimensions, the analysis is restricted to one or two dimensions, and thus becomes tractable.

Galilei (1632) and von Helmholtz (1867) had both observed that spatial resolution is higher for dark stimuli than for light (Kremkow et al., 2014)—why is this? It has long been known that the distributions of luminance and contrast in natural scenes exhibits strong positive skewness (Laughlin, 1983; Richards, 1982). One consequence is that in natural images, negative contrasts are more numerous and carry more information than light contrasts. Given finite neural bandwidth, the principle of efficient coding (Barlow, 1961; Laughlin, 1981) predicts that this asymmetry would be matched by the brain, and recent evidence (Ratliff, Borghuis, Kao, Sterling, & Balasubramanian, 2010) suggests that it is: OFF retinal ganglion cells are more numerous and smaller than ON retinal ganglion cells.

Two papers in this issue show that these asymmetries are relevant to higher-level aspects of visual function as well. Sato et al. (2016) consider perceived blur: the tacit assumption has been that perceived blur is determined by the spatial frequency content of the image, treating darks and lights equally. However, the authors show that this is not the case, and that, along with the visual system's higher sensitivity and spatial resolution for darks, it gives darks higher weight in assessing blur.

While these results represent a concordance of visual tuning with natural scene statistics consistent with the efficient coding hypothesis, the results of Graham et al. (2016) do not. Graham et al. consider subjective preference for naturalistic images as a function of the skewness of the luminance distribution. Perhaps surprisingly, they find that observers prefer images when modified to have zero skewness (i.e., symmetric luminance distributions) rather than the positive skewness typically observed in natural scenes. What this seems to suggest is that aesthetic preference does not necessarily follow the principles of efficient coding.

Three papers address the point statistics of color. Provenzi et al. (2016) identify a simplifying feature of second-order statistics in the spatio-chromatic domain: they show that spatial and

chromatic covariance matrices of natural images commute with each other. This formalizes a sense in which spatial and chromatic statistics are separable, enabling a simplified description of the second-order spatio-chromatic statistics of natural images.

The other papers concerning color attempt to understand ecological aspects of the spatiochromatic variation of natural images. Color features in the image are co-determined by surface reflectance and illumination. It is generally agreed that the former is of greater interest to the observer, and so the latter appears largely as a nuisance variable. How then can the observer factor out the effects of varying illumination to obtain invariant estimates of surface color?

One possible assumption is that the illuminant can be approximated as roughly uniform within a scene, so that color variations in the image can mainly be attributed to variations in surface color. Examining the hyperspectral statistics of natural scenes over space and time, [Nascimento et al. \(2016\)](#) find that in fact the instantaneous variation in illumination color within a scene typically varies roughly as much as illumination varies over time or latitude. These findings argue against the use of overly simplistic models of illumination in natural scenes.

[Foster et al. \(2016\)](#) explore the consequences of these illumination variations on color measurements made in the image. A potential way for the visual system to attain a degree of invariance with respect to these illumination variations is to rely not upon absolute but upon relative color measurements. The latter are computed as ratios of cone responses measured at different points in the image, which are invariant to purely temporal variations of the global illumination. The authors show that in fact, due to the complex geometry of natural scenes, the local illuminant and hence these ratios vary considerably. However, placing either spatial or temporal limits on these ratios can reduce variation below detectability, suggesting that cone ratios may serve as useful pseudo-invariants for more localized spatiotemporal neighborhoods within a scene.

## 2. Oriented filter statistics

These first five studies focused on first- and second-order point statistics of the luminance and color distributions of natural images. However, important features in the scene, such as the boundaries of objects, result in locally oriented structure in the image, and since Hubel & Wiesel's discovery that most neurons in primate area V1 are tuned for orientation ([Hubel & Wiesel, 1959, 1968](#)), the local processing of oriented information in the image has been seen as a primary goal of early visual cortex.

Hubel & Wiesel sometimes referred to cells in V1 as edge- and line-detectors, which suggests that their receptive fields should not only be oriented, but should have specific selectivity for the detection of edges and lines versus other forms of oriented structure. In tension with this proposal was the spatial frequency account of early visual coding ([Campbell & Robson, 1968](#)), which envisioned these neurons as uniformly tiling a two-dimensional Fourier representation of the image. In such a representation, tuning for edges and lines should manifest as a bias toward particular bandwidths and phases (odd and even). While early physiological studies in cat (e.g., [Field & Tolhurst, 1986](#)) did not confirm this prediction, later studies in primate (e.g., [Ringach, 2002](#)) did.

[MaBouDi et al. \(2016\)](#) address whether the statistics of natural images may account for these physiological findings. In particular, using a standard quadrature filter model of an oriented V1 complex cell, they analyze the phase distribution of local image patches as a function of the frequency tuning, bandwidth and aspect ratio parameters of these filters. While for most parameter values they find a relatively uniform distribution of phases, for a specific range of filters they find that phase response is bimodal, consistent with

tuning for specific local features such as edges and lines. They speculate that the visual system makes use of these inhomogeneities to code visual inputs efficiently: phase-independent complex cells for configurations that correspond to uniform local phase distributions, and phase-tuned complex cells whose configurations correspond to bimodal phase distributions.

The nonlinearity involved in the quadrature filter model of complex cells is just one example of the many nonlinear properties of early visual cortical neurons that complicate statistical analysis. [Golden et al. \(2016\)](#) seek to relate these nonlinearities to principles of sparse coding ([Bell & Sejnowski, 1997](#); [Olshausen & Field, 1996](#)). Employing a representation of neural tuning in curved coordinate frames introduced by Zetzsche and colleagues ([Zetzsche, Krieger, and Wegmann \(1999\)](#), [Zetzsche and Nuding \(2005\)](#) and [Zetzsche and Rohrbein \(2001\)](#)), they show that the sparse coding network of Olshausen and Field produces a form of curvature that accounts for nonlinearities that give rise to key selectivity and invariance properties of oriented V1 neurons. Since the objective function used to learn these sparse networks does not explicitly pursue this curvature, these findings provide insight into the connections between sparsity, response nonlinearity and curved representational manifolds for neural coding.

The prominence of orientation tuning in early visual cortex is often attributed to the importance of identifying the boundaries of objects in the scene. However these boundaries are signaled not just by gradients in luminance and color but also in motion and binocular disparity. How important is each of these modalities to the detection of an object boundary? [Mély et al. \(2016\)](#) address this question by training classifiers that use individual modalities and ensembles. While motion and binocular disparity are often cited as most diagnostic for the discrimination of object boundaries, the authors find that in fact color and luminance play dominant roles. However, they also find that more accurate boundary extraction can be achieved by combining all cues, and, importantly, motion and stereoscopic disparity play a much larger role when combined with the other cues than would be expected from their performance in isolation. This suggests a tighter coupling of these modalities than might be suggested by a strict modular view of visual cortex ([Livingstone & Hubel, 1988](#)).

## 3. Disparity statistics

The detection of object boundaries is only one expression of the general role stereoscopic disparity plays in revealing the 3D structure of visible surfaces in the scene. [Hibbard et al. \(2016\)](#) consider the statistical cues in the natural visual environment that are available to compute disparity. While early "first-order" models relied upon luminance correlations between corresponding image patches in the two eyes, psychophysical ([Hess & Wilcox, 1994](#)) and physiological ([Tanaka & Ohzawa, 2006](#)) work has revealed mechanisms that allow stereoscopic disparity to be computed from stimuli with no first-order inter-ocular correlations. Such mechanisms must rely upon nonlinear processing (e.g., rectification or squaring) to reveal "second-order" disparity cues. Hibbard et al. show that such cues are indeed present in natural images and, while correlated with first-order cues, can improve the accuracy of disparity estimation. These findings suggest that stereoscopic mechanisms have evolved to take advantage of the diverse cues to stereoscopic matching available in complex natural scenes.

Despite this richness of cues, recovering an accurate global surface estimate from stereo is one of the classic ill-posed vision problems; there are many ways to interpolate a surface through noisy local disparity measurements, particularly when the measurements are sparse. To perform well, the visual system must therefore incorporate a statistical prior over these possible surfaces. The question of how these scene priors can be learned

Download English Version:

<https://daneshyari.com/en/article/6202956>

Download Persian Version:

<https://daneshyari.com/article/6202956>

[Daneshyari.com](https://daneshyari.com)