Statistics of natural scenes and cortical color processing

Guillermo A. Cecchi,
T. J. Watson IBM Research Center, Yorktown Heights, NY, USA

A. Ravishankar Rao,
T. J. Watson IBM Research Center, Yorktown Heights, NY, USA

Youping Xiao, and
Mt. Sinai School of Medicine, New York, NY, USA

Ehud Kaplan
Mt. Sinai School of Medicine, New York, NY, USA

Guillermo A. Cecchi: gcecchi@us.ibm.com; A. Ravishankar Rao: ravirao@us.ibm.com; Youping Xiao: youping.xiao@mssm.edu; Ehud Kaplan: ehud.kaplan@mssm.edu

Abstract

We investigate the spatial correlations of orientation and color information in natural images. We find that the correlation of orientation information falls off rapidly with increasing distance, while color information is more highly correlated over longer distances. We show that orientation and color information are statistically independent in natural images and that the spatial correlation of jointly encoded orientation and color information decays faster than that of color alone. Our findings suggest that: (a) orientation and color information should be processed in separate channels and (b) the organization of cortical color and orientation selectivity at low spatial frequencies is a reflection of the cortical adaptation to the statistical structure of the visual world. These findings are in agreement with biological observations, as form and color are thought to be represented by different classes of neurons in the primary visual cortex, and the receptive fields of color-selective neurons are larger than those of orientation-selective neurons. The agreement between our findings and biological observations supports the ecological theory of perception.

Keywords

structure of natural images; visual cortex; color vision

Introduction

The ecological theory of perception states that early sensory processing by the brain is adapted to the statistics of the natural environment. According to this view, the adaptation shaped by evolutionary pressure provides for an efficient transmission of information from the periphery to the higher order centers of the brain (Atick, 1992; Barlow, 1961; Field, 1987), given unavoidable constraints imposed on the transmission channel in terms of dynamic range, connectivity, and number of fibers and neurons. These ideas are best exemplified by the processing in the early visual pathway, where the expected redundancies

in the input are eliminated. This results in a drastic reduction of the high dimensionality of retinal information arising from sampling of the visual field (from the sheer number of rods and cones), thereby facilitating the transmission of this information by a limited number of retinal ganglion cell fibers. Perhaps the most striking use of this approach was shown by Atick and Redlich (1992), who derived the spatial receptive field properties of retinal ganglion cells as an optimal filter for the two-point correlations of contrast in natural images. These correlations, well characterized by a power-law spectrum $s(k) \sim k^{-1}$, imply that nearby photoreceptors are highly redundant (Ruderman, 1997; Ruderman & Bialek, 1994), and therefore, their information need not be transmitted in full. Similarly, the lagged-x cells of the geniculate are optimal filters for removing temporal redundancies (Dan, Atick, & Reid, 1996). Srinivasan, Laughlin, and Dubs (1982) discussed redundancy removal by predictive coding done by the center–surround antagonism of receptive fields. Similarly, Buchsbaum and Gottschalk (1983) derived the receptive field structure of retinal ganglion cells from optimality considerations.

There are few comparable results, however, for applications of the ecological theory to cortical processing. In particular, features that are statistically more common, such as vertical and horizontal lines or iso-orientated edges, tend to be over-represented (Betsch, Einhauser, Kording, & Konig, 2004; Li, Peterson, & Freeman, 2003; Sigman, Cecchi, Gilbert, & Magnasco, 2001). Here, we ask to what extent the organization of cortical response selectivity to color and orientation can be predicted by the statistics of natural images. While the physiology of orientation selectivity has been widely studied, color remains a more elusive feature. Some studies have suggested that color and orientation are represented in cytochrome oxidase (CO) blobs and interblobs, respectively (Lu & Roe, 2008; Ts'o & Gilbert, 1988). It was also found that color-selective cells in V1 and V2 that are unselective for orientation have larger receptive fields than orientation-selective ones (Johnson, Hawken, & Shapley, 2008; Roe & Ts'o, 1995; Solomon, Peirce, & Lennie, 2004). Because the blobs occupy less than half as much area as the interblobs do, these studies suggest that color information is processed in V1 by fewer cells with relatively larger receptive fields as compared to orientation information. One may wonder whether such differences are the result of the adaptation to the statistics of natural images. In the present study, we found evidence that supports the role of adaptation, by explicitly computing the spatial autocorrelation of color and orientation as vector fields in a large ensemble of natural images. Moreover, we collected evidence suggesting that color and orientation information are statistically independent.

The visual cortex is one of the most researched brain areas; in particular, its response to orientation has been very well characterized at the single-cell as well as the population levels, revealing a spatial structure of orientation columns, punctuated by point and line discontinuities (Blasdel & Salama, 1986; Bonhoeffer & Grinvald, 1991; Grinvald, Lieke, Frostig, Gilbert, & Wiesel, 1986). Other visual attributes such as visual field position (Tootell, Silverman, Switkes, & De Valois, 1982), spatial frequency (Everson, Prashanth, Knight, Sirovich, & Kaplan, 1998; Issa, Trepel, & Stryker, 2000; Xu et al., 2004), direction (Weliky, Bosking, & Fitzpatrick, 1996), and ocular dominance (Blasdel, 1992) are mapped in similar fashion. Hyvarinen and Hoyer (2001) have used a sparse coding framework to explain how cortical topographic orientation maps can arise from a learning process applied to natural image inputs.

In comparison with orientation, the cortical response to color attributes in the visual field has been more difficult to characterize. Functional brain imaging studies in humans have uncovered cortical areas that are preferentially activated by chromatic stimuli. Moreover, psychophysical and lesion reports suggest the existence of a specialized color processing system that is, to some extent, unconcerned with the processing of other visual features. The

pioneering work of Livingstone and Hubel (1984) revealed that the cytochrome oxidase dense blobs in the primary visual cortex (V1), and the regions between them (interblobs), contain populations of neurons that differ in their selectivity for color or orientation. More recently, Xiao, Casti, and Kaplan (2007), Xiao, Rao, Cecchi, and, Kaplan (2007, 2008), and Xiao, Wang, and Felleman (2003) demonstrated, using optical imaging, that color information is reliably encoded by spatial patterns of activity across V1 and the extra-striate area V2; moreover, color and orientation are mapped in segregated compartments and display different spatial properties. Using a support-vector machine (SVM) methodology to identify regions that contained the most information in discriminating color and orientation, they showed that areas that process color are relatively small, and far apart. On the other hand, areas that process orientation are comparatively larger and closer to each other. In the present study, we demonstrate that these differences can be related to the statistical properties of natural images.

Statistics of natural images

The scaling analysis of contrast and luminance has been the focus of many studies. Scaling measurements involve studying how the probability of finding a co-occurring pair changes as a function of the relative distance. A classic result in the analysis of natural scenes is that the luminance of pairs of pixels is correlated and that this correlation is scale-invariant (Atick & Redlich, 1992; Field, 1987; Ruderman & Bialek, 1994). This invariance indicates that statistical dependencies between pairs of pixels do not depend on whether the observer zooms in on a small window or zooms out to a broad vista. The scale invariance results from stable physical properties such as a common source of illumination and the existence of objects of different sizes and similar reflectance properties (Ruderman, 1997).

Few studies have focused on the structure of long-range correlations of other visual attributes. A report by Sigman et al. (2001) analyzed the spatial distribution of orientation information in natural scenes. They showed that information about the presence of iso-oriented lines in natural images is correlated over relatively short distances, following specific power-law statistics for co-linearity; other pairwise arrangements display shorter correlations. This study also suggested a possible relationship between orientation statistics and the extra-classical receptive field properties of neurons in the visual cortex.

The long-range statistics of the color field, in particular, has been addressed only tangentially. A report by Johnson, Kingdom, and Baker (2005) studied the cross-correlation of the responses of band-pass filters applied to the luminance, R–G, and B–Y channels obtained from color images. Though the spatial structure of an image is taken into account at the local level, they did not compute correlations across different spatial locations. Similarly, Tailor, Finkel, and Buchsbaum (2000) derived the (local) independent components of natural images and showed that they contain oriented luminance edge filters and color-opponent red–green and blue–yellow filters; even though the authors did not explicitly compute the joint probability distribution between color and orientation, this result suggests that these fields are statistically independent and parallels the finding that luminance and contrast are statistically independent in natural scenes (Mante, Frazor, Bonin, Geisler, & Carandini, 2005).

Burton and Moorhead (1987) examined the spatial structure of color variation in natural scenes but restricted their analysis to vertical and horizontal axes and to individual colors (red, green, and blue receptor values). In contrast, our analysis is carried out over the full visual field and assumes a continuous 2D vector representation for color. A more detailed analysis of color space was performed by Ben-Shahar and Zucker (2004), who represented it by a 2D field formed of hue vectors, i.e., a hue field. Structural similarities between hue and

orientation fields were observed, in terms of the use of discontinuities in the fields to perform image segmentation; they did not, however, examine the spatial correlation properties of this hue field over a large set of natural images, as we do in our work.

Parraga, Brelstaff, Troscianko, and Moorehead (1998) investigated the spatial frequency content of color and luminance information in a set of 29 images of natural scenes. However, they did not compute orientation information in these images.

Johnson et al. (2005) analyzed color natural image scenes for their spatial frequency content in luminance, red–green, and blue–yellow channels. Their work differs from ours in that they computed correlation between spatial frequency bands, whereas we compute long-range spatial correlations for orientation and color representations.

Research on the relation between color image statistics and the neural representation of color has been focused on the spectral response properties of the LMS cones (Simoncelli & Olshausen, 2001). Several studies have examined the statistical relation between the components of color, as described by three channels comprising a luminance channel, an R–G channel, and a B–Y channel (Johnson et al., 2005; Ruderman, 1997). It is reasonable to expect that peripheral processing of color be highly determined by the statistics of the input; however, there is evidence of a link between the independent components of color natural images and the receptive field properties of cortical V1 simple neurons (Caywood, Willmore, & Tolhurst, 2004). The main motivation for our work was to provide a robust characterization of the joint color/orientation statistics, as we expect them to influence the organization of early visual cortical processing.

Natural image analysis

In order to compute the spatial autocorrelation of orientation and color in natural images, we computed their corresponding fields as described below. The orientation field was computed following similar approaches in the literature (Sigman et al., 2001). For color, we used a 2D hue field computed as described below. We first performed a conversion from RGB to CIE $L*a*b*$ coordinates using the ITU recommended D65 white point reference in the ITU BT. 709 standard, by using the equations given in Wyszecki and Stiles (1982, p. 166). For the sake of notational convenience, we use (L, a, b) to denote CIE $L*a*b*$. We also performed our analysis using a calibrated LMS space, as provided by the creators of the image database (Olmos & Kingdom, 2004). The results we obtained were virtually identical for the two color spaces.

We present a brief overview of the method for local orientation estimation in images, and details may be found in Rao and Shunck (1991).

There are five steps to estimating the local orientation in an image, i.e.,

1. Smooth the image with a Gaussian filter;

2. Compute the gradient of the smoothed image;

3. Find the local orientation angle;

4. Average the local orientation estimates over a small neighborhood

5. Compute a measure of the coherence (the degree of anisotropy) of the pattern.

The *coherence* is a measure for how strongly anisotropic the image edges are within a local neighborhood. The coherence can be interpreted as the magnitude of an orientation vector, and the direction of the vector is given by the angle of orientation.

Let the gradient vector at location $\vec{x}$ in the image have the polar representation $G_{\vec{x}} e^{i\theta \vec{x}}$. The estimate of the dominant orientation $\hat{\theta}$ at the center $\vec{x}$ of a neighborhood N of the image is given by

$$\widehat{\theta}_{\vec{x}} = \tan^{-1}\left(\frac{\left\langle G^2_{\vec{y}} \sin 2\theta_{\vec{y}} \right\rangle_{\vec{y} \in N}}{\left\langle G^2_{\vec{y}} \cos 2\theta_{\vec{y}} \right\rangle_{\vec{y} \in N}} \right) / 2,$$

(1)

where the angular brackets denote the average computed over the neighborhood. The estimated orientation angle at $\vec{x}$ is then $\hat{\theta}_{\vec{x}} + \pi/2$, since the gradient vector is perpendicular to the direction of anisotropy.

Let $\hat{\theta}_{\vec{x}}$ denote the estimated orientation angle at point $\vec{x}$, found in the earlier step. To find the coherence at point $\vec{x}$, consider point $\vec{y}$, which lies within a window W of prescribed size around point $\vec{x}$. The measure of coherence is defined by

$$\rho = G_{\vec{x}} \frac{\left\langle \| G_{\vec{y}} \cos(\widehat{\theta}_{\vec{y}} - \widehat{\theta}_{\vec{x}}) \| \right\rangle_{\vec{y} \in W}}{\left\langle G_{\vec{y}} \right\rangle_{\vec{y} \in W}}.$$

(2)

Thus, one can obtain a description of the orientation field by using Equations 1 and 2. The filter sizes used were g_1 for estimating the orientation angle in 1, and g_2 for estimating the coherence in 2. The values of g_1 and g_2 are described in Figure 1 and range from small to large spatial scales as indicated.

For calculating the color field, the following steps were used.

1. The images were converted into the LAB space, to yield for each pixel the vector (L, a, b).

2. The color field was computed as $\Psi_c(\vec{x}) = e^{i\theta_c}$, where $\theta_c = \arctan \frac{b}{a}$.

Note that both the color and orientation analyses produce a vector field over the spatial image coordinates. The vector fields are subject to the following analysis.

For each field, the autocorrelation function was computed over a representative ensemble of natural images, assuming translational and rotational invariance,

$$c(d) = \sigma^{-1} \left\langle \Re[(\Psi(\vec{x}) - \overline{\Psi})(\Psi^*(\vec{x} + \vec{x}') - \overline{\Psi}^*)] \right\rangle_{|x| = d},$$

(3)

where the brackets signify average over $\vec{x}$ and over the ensemble of images, $\overline{\Psi}$ is the average field, $*$ is the complex conjugate operator, $\Re$ is the real part operator, and σ is the variance, defined as

$$\sigma = \left\langle \Re[(\Psi(\vec{x}) - \overline{\Psi})(\Psi^*(\vec{x}) - \overline{\Psi}^*)] \right\rangle,$$

(4)

where, as in Equation 3, the average is over $\vec{x}$ and the ensemble of natural images. The invariance assumptions allow for a fast implementation via the Fourier transform; we have observed in previous work that rotational invariance might not be fully warranted (Sigman et

al., 2001), as vertical and horizontal orientations tend to be over-represented. For the sake of completeness, we will show in the Results section that these differences are present in our database; however, given their measured values, we do not expect them to affect the results on scaling behavior reported here.

We also computed the joint spatial statistics of color and orientation as follows. Consider the vector Ψ_{OC} formed by concatenating the vectors $\Psi_o(\vec{x})$ and $\Psi_c(\vec{x})$ as follows:

$$\Psi_{oc}(\vec{x}) = [\Psi_o(\vec{x}), \Psi_c(\vec{x})]. \tag{5}$$

We compute the correlation of the vector $\Psi_{OC}(\vec{x})$ over all translations of $\vec{x}$, in a manner analogous to Equation 4.

The database of natural images consists of 850 color calibrated pictures from the McGill Calibrated Colour Image Database and includes pictures of animals in their environment, foliage, land and water landscapes, shadows in natural and man-made settings, close-ups of natural and fabricated textures, close-ups of flowers, fruits, and vegetables, and city scenes (Olmos & Kingdom, 2004). We use this database because it contains color-calibrated pictures, as opposed to other databases in the literature that contain grayscale pictures.

Results

The results of the spatial autocorrelation analysis for orientation and color in natural images are presented in Figure 2. This plot shows that the orientation correlation decays with distance much faster than that of color. The orientation correlation approaches a power-law (i.e., algebraic) scaling, as demonstrated by the linear stretch in the double-log plot. The color correlation decays slowly as a power law until a distance of about 100 pixels, after which a more rapid decay is observed. We compute the correlations at three spatial scales: small, medium, and large. The size of the filters applied for orientation estimation in Equations 1 and 2 determines the scale, as shown in Figure 1.

The multi-scale variation in Figure 2a shows that the large-scale orientation vectors have distinctly higher spatial correlation than medium- and small-scale orientation vectors. However, the difference between the spatial correlations for large-scale and small-scale color vectors is not as high as that of the orientation vectors.

We used a two-sample *t*-test of the distributions for color and orientation that are summarized in Figure 2. For each displacement, the autocorrelation value for color over all the samples produces a first distribution, and that for orientation produces a second distribution. We used the function *ttest*2 in MATLAB for this computation using data generated from the small-scale filters. The result rejected the null hypothesis that the data in these two distributions are independent random samples from normal distributions with equal means but unknown variances. The *p* value associated with the test was less than 10^{-10} for each positive pixel distance.

In Figure 2, we computed the autocorrelations averaged over all directions, as this provides a compact way of viewing the differences in these functions. The orientation and color vectors were scaled such that the maximum magnitude of the orientation and color vectors in each image was set to 1. This facilitates a comparison with the joint orientation–color statistics, to be shown later in Figure 6.

We also calculated the autocorrelations across specific directions, and the results are presented in Figure 3. Note that the same overall trend is observed, in that the correlation

function for orientation decays faster than that of color. Moreover, the correlation function for color is anisotropic and falls off more slowly in the horizontal direction as compared to the other directions. The implication of this finding is examined in the Discussion section. Finally, we observe that the orientation field is more correlated along the horizontal and vertical directions than others. This finding is consistent with the so-called oblique effect, as reported by Betsch et al. (2004).

We also performed a randomization in order to provide a reference for the comparison of the spatial structures. We first computed the orientation and color vectors at each pixel, and then randomized the locations of these vectors. We then computed the spatial correlations over these randomized orientation and color fields. The spatial correlation decays significantly faster in the randomized images compared to the correlation in the original images. We do not show this explicitly in a 2D plot, as the correlation function decays too fast to be meaningfully displayed. Rather, the 1D correlation function is displayed in Figures 2c and 2d.

In order to compare the results summarized by Figures 2a–2d to a different, more structured null hypothesis, we compared the correlation structure of natural images with a specific type of man-made imagery, which consisted of five paintings of the American painter Jackson Pollock. We chose these abstract art paintings because of the apparent lack of correlation structure in them. We applied the same methodology as before to compute the correlation statistics for orientation and color. The result is shown in Figure 2e. Next, we randomized the locations of the computed orientation and color vectors in these paintings and computed the correlation statistics, as shown in Figure 2f. The results in Figure 2e show that the correlation values for color and orientation in Pollock's paintings fall more rapidly than corresponding values in natural images.

We applied our algorithm to compute image statistics on a second image database, known as the Berkeley Segmentation Dataset and Benchmark (Martin, Fowlkes, Tal, & Malik, 2001). This database consists of 300 color images. Though the database contains segmentation labels, these were ignored for the purpose of computing orientation and color statistics. The results are shown in Figure 4 and demonstrate the same relationship between orientation and color correlations that exist in Figure 2 for a different data set.

We also computed the joint statistics of orientation and color as follows. We estimated the pairwise joint probability distributions of orientation angle θ_e with each of the color (L, a, b) values, giving rise to three distributions, (θ_e, L), (θ_e, a), and (θ_e, b). The plots in Figure 5 clearly indicate that the pairwise product of the marginal probability density functions is equal to the joint probability distribution functions $P(\theta_e, L)$, $P(\theta_e, a)$, and $P(\theta_e, b)$. Thus, $P(\theta_e, L) = P(\theta_e)P(L)$, where $P(\theta_e)$ is the probability density function of θ_e and $P(L)$ is the probability density function of luminance, L.

In order to make the measurement of independence more rigorous, we computed the pairwise mutual information between (θ_e, L), (θ_e, a), and (θ_e, b), respectively. The mutual information between two variables A and B is defined in terms of their entropies, $H(A)$ and $H(B)$ as follows. Let A possess N finite states, $\{a_1, a_2, \ldots a_N\}$. The entropy $H(A)$ is given by

$$H(A) = \sum_{i=1}^{N} p(a_i) \log\, p(a_i).$$

(6)

The entropy $H(B)$ is similarly defined, where B possess M finite states. The joint entropy $H(A, B)$ is defined by

$$H(A, B)= -\sum_{i=1}^{N}\sum_{j=1}^{M}p(a_i, b_j)\log p(a_i, b_j).$$

(7)

The mutual information, $MI(A, B)$, is defined by

$$MI(A,B)=H(A)+H(B) - H(A,B) \geq 0.$$

(8)

The mutual information is zero when A and B are statistically independent.

We estimated the pairwise mutual information $MI(\theta_e, L) = 0.0022$ bits, $MI(\theta_e, a) = 0.0019$ bits, and $MI(\theta_e, b) = 0.0029$ bits. We also measured the mutual information between θ_e and the hue angle $\theta_c = \arctan\frac{b}{a}$ to be $MI(\theta_e, \theta_c) = 0.0028$ bits. Since these pairwise mutual information measures are close to zero, we conclude that orientation is independent of the (L, a, b) color components. This observation suggests that an efficient scheme of representing natural images is to represent color and orientation information in separate pathways.

Finally, we show the results of computing the joint spatial distribution of orientation and color information, using Equation 5. Figure 6 shows that the correlation of the joint orientation and color vector decays more rapidly than the color vectors alone. The plot in Figure 6d can be compared with the one in Figure 2a. Furthermore, Figure 7 explicitly compares the joint spatial distribution of orientation and color information with the individual distributions. The implication of this observation is that there is less redundancy that can be exploited if the visual system tries to encode color and orientation information jointly. Given the spatial statistics of natural images that we have obtained in this paper, and the independence of the color and luminance channels, the best coding strategy is one where color and orientation are processed through independent channels. Furthermore, color information needs to be sampled less densely due to the larger spatial scale that exists for color; this seems to be precisely the coding strategy employed by the visual cortex.

Discussion

We have presented novel results showing that orientation and color attributes in natural images display different spatial structure of autocorrelations: color is more significantly correlated over longer distances than orientation. Though orientation has been studied much more than color, the specific question that we have asked in the present manuscript, that of the relationship between the spatial properties of visual attributes and that of their cortical representations, also requires a new look at the well-established properties of orientation maps.

We also demonstrated the statistical independence of orientation and color information in natural images, similar to the reported independence of luminance and contrast (Mante et al., 2005), and in agreement with the finding of independent components for orientation and color (Tailor et al., 2000). Fine, MacLeod, and Boynton (2003) showed that the luminance and color information of surfaces in natural scenes are relatively independent. Hansen and Gegenfurtner (2009) showed that color and luminance edges in natural images are independent. This indicates that it would be desirable to process orientation and color information in separate channels, as is observed in the early stages of the human and primate visual systems. The spatial correlation structure of orientation and color information may

also govern the spatial organization of units in the orientation-processing and color-processing channels, respectively.

Our approach was explicitly motivated by the ecological theory of perception. This theory has been dominated by information-theoretic concepts, i.e., the idea that the function and architecture of the nervous system needs to maximize the efficiency of information transmission in the context of biological constraints. At present, experimental confirmation of this theory has come from local neural features, such as the response properties of individual neurons. Our results point, however, to a more general interpretation of the ecological theory, in which the input space shapes the organization of neural responses at a population level. This organization might be ultimately related to the need to satisfy information-theoretic constraints, but it seems more likely that it is the result of processes with explicit spatial constraints. In particular, models of cortical organization that emphasize the role of spatial structure in coding and processing can make testable predications about the relationship between the input ensemble and the neural map of responses.

The finding that color information in natural scenes is correlated over longer distances than orientation information implies that a coarser sampling of color information will suffice for cortical representation purposes. This prediction is consistent with physiological findings that V1 contains a larger proportion of orientation-selective cells than color-selective ones (Shapley & Hawken, 2002). Most studies of color-processing areas in the cortex (Livingstone & Hubel, 1984; Xiao, Rao et al., 2007, 2008) confirm that they are relatively small and far apart. However, the nature of cortical color processing is far from understood. For instance, it is still controversial which class(es) of V1 cells code information about color of object surfaces (for a review, see Lennie & Movshon, 2005). A likely candidate is the class of cells with strong color opponency, which are also called color-preferring cells (Johnson, Hawken, &, Shapley, 2001, 2004). Unlike other classes of color-sensitive cells, such as color-luminance cells (Johnson et al., 2001), strongly color-opponent cells have chromatic tuning that is invariant with stimulus orientation, size, and contrast (Johnson et al., 2001; Solomon & Lennie, 2005). Since the strongly color-opponent cells make up about 10% of V1 cells (Johnson et al., 2001; Lennie & Movshon, 2005), the argument that they are the color-coding cells in V1 is consistent with our result that color varies more gradually than orientation in natural scenes.

A class of V1 cells that may also code information about object color are double-opponent cells, although the spatial organization and cone inputs of their receptive fields are still under debate (Conway & Livingstone, 2006; Johnson et al., 2001, 2008; Livingstone & Hubel, 1984). Several studies also gave different estimates regarding the percentage of double-opponent cells among V1 cells. In one study that mapped receptive fields of V1 cells with cone-isolating stimuli, the estimate was about 6% (Conway & Livingstone, 2006); in the same study, double-opponent cells and other types of color-opponent cells were estimated to comprise about 10% of V1 cells. Given that the majority of V1 cells are selective for orientation (Hubel & Wiesel, 1968), these results are also consistent with our prediction.

Since a given visual field is covered by fewer color-coding cells than orientation-coding ones, the former needs to have larger receptive fields in order to fully tile the visual field. This prediction is supported by a recent finding that the average receptive field size of strongly color-opponent neurons in V1 is 1.5 times as large as that of orientation-selective ones (Solomon et al., 2004). Solomon and Lennie (2007) also observed that the receptive fields of color-preferring neurons lack spatial structure, which renders them unsuitable of encoding fine detail. This observation is consistent with the results of Johnson et al. (2004) that color-preferring cells in primate V1 have low-pass spatial transfer functions, compared

with luminance-preferring cells, which have band-pass spatial transfer functions. Double-opponent cells have band-pass spatial transfer functions, but their receptive fields are larger than those of neurons that subserve visual features such as form and motion (Conway, 2009). These observations are consistent with the prediction of our results.

Color-selective cells in V2 with no orientation selectivity tend to have larger receptive fields than similar cells with broad-band orientation selectivity (Roe & Ts'o, 1995), whereas the thin cytochrome oxidase stripes, which are selective for color, display greater receptive field size and scatter as compared to disparity stripes. While it is quite difficult to relate these findings to psychophysical measures of color, form, and disparity perception (Mullen, 1985), it has been argued that the color system would require a less precise mapping than either the form or disparity system (Roe & Ts'o, 1995). Our findings provide quantitative measurements for a different interpretation of these results: a sparse sampling of a redundant variable need not compromise resolution.

Likewise, in concordance with our results, an independent component analysis of the chromatic structure in 7×7 pixel patches of natural scenes showed that the achromatic basis functions are localized and oriented (Figure 5 in Wachtler, Lee, & Sejnowski, 2001). In contrast, the chromatic basis functions are relatively less localized, indicating that they sample from a wider image area.

If the color statistics are isotropic, we would expect color-sensitive cells to possess circularly symmetric receptive fields. However, the anisotropy of color statistics that we report in Figure 3 suggests that some group of color-sensitive cells should not display circular symmetry. This seems to be confirmed experimentally, although the findings are still tentative (Conway, 2009).

In summary, we have presented evidence suggesting that the selectivity and relative sparsity of color over orientation processing in cortical cells are linked to the statistical regularities of the corresponding visual attributes, namely the independence of color and orientation, and the relative longer range of their spatial correlations. More precise measurements of large-scale cortical orientation and color responses should provide for a firmer ground to further test this idea.

Our analysis technique can be applied to investigations of spatial correlations of image texture measures, and it is possible that there may be other visual features that show a gradual spatial variation of correlation, similar to that of color. For instance, the spatial autocorrelation of luminance information depicted in Figures 3e and 3f shows that luminance variations are similar to that of color.

Finally, we would like to emphasize that even though we focused on orientation and color, the ideas developed in the present manuscript might be valid for other significant visual attributes, as well as the other primary perceptual modalities.

Acknowledgments

References

Atick JJ. Could information theory provide an ecological theory of sensory processing? Network: Computation in Neural Systems. 1992; 3:213–251.

Atick JJ, Redlich AN. What does the retina know about natural scenes? Neural Computation. 1992; 4:196–210.

Barlow, HB. Possible principles underlying the transformation of sensory messages. In: Rosenblith, WA., editor. Sensory communications. Cambridge, MA: MIT Press; 1961. p. 217-234.

Ben-Shahar O, Zucker S. Hue geometry and horizontal connections. Neural Networks. 2004; 17:743–771.

Betsch B, Einhauser W, Kording K, Konig P. The world from a cats perspective—Statistics of natural videos. Biological Cybernetics. 2004; 90:41–50. [PubMed: 14762723]

Blasdel GG. Differential imaging of ocular dominance and orientation selectivity in monkey striate cortex. Journal of Neuroscience. 1992; 12:3115–3138. [PubMed: 1494950]

Blasdel GG, Salama G. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. Nature. 1986; 321:579–585. [PubMed: 3713842]

Bonhoeffer T, Grinvald A. Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. Nature. 1991; 353:429–431. [PubMed: 1896085]

Buchsbaum G, Gottschalk A. Trichromacy, opponent colours coding and optimum colour information transmission in the retina. Proceedings of the Royal Society of London B. 1983; 220:89–113.

Burton G, Moorhead I. Color and spatial structure in natural scenes. Applied Optics. 1987; 26:157–170. [PubMed: 20454092]

Caywood M, Willmore B, Tolhurst D. Independent components of color natural scenes resemble V1 neurons in their spatial and color tuning. Journal of Neurophysiology. 2004; 91:2859–2873. [PubMed: 14749316]

Conway B. Color vision, cones and color-coding in the cortex. The Neuroscientist. 2009; 15:274–290. [PubMed: 19436076]

Conway B, Livingstone M. Spatial and temporal properties of cone signals in alert macaque primary visual cortex. Journal of Neuroscience. 2006; 26:10826–10846. [PubMed: 17050721]

Dan Y, Atick JJ, Reid RC. Efficient coding of natural scenes in the lateral geniculate nucleus: Experimental test of a computational theory. Journal of Neuroscience. 1996; 16:3351–3362. [PubMed: 8627371]

Everson R, Prashanth A, Knight B, Sirovich L, Kaplan E. Representation of spatial frequency and orientation in the visual cortex. Proceedings of the National Academy of Sciences of the United States of America. 1998; 95:8334–8338. [PubMed: 9653187]

Field DJ. Relations between the statistics of natural images and the response profiles of cortical cells. Journal of the Optical Society of America A. 1987; 4:2379–2394.

Fine I, MacLeod DIA, Boynton GM. Surface segmentation based on the luminance and color statistics of natural scenes. Journal of the Optical Society of America A. 2003; 15:563–569.

Grinvald A, Lieke E, Frostig RD, Gilbert CD, Wiesel TN. Functional architecture of cortex revealed by optical imaging of intrinsic signals. Nature. 1986; 324:361–364. [PubMed: 3785405]

Hansen T, Gegenfurtner KR. Independence of color and luminance edges in natural scenes. Visual Neuroscience. 2009; 26:35–49. [PubMed: 19152717]

Hubel DH, Wiesel TN. Receptive fields and functional architecture of monkey striate cortex. The Journal of Physiology. 1968; 195:215–243. [PubMed: 4966457]

Hyvarinen A, Hoyer PO. A two-layer sparse coding model learns simple and complex cell receptive fields and topography from natural images. Vision Research. 2001; 41:2413–2423. [PubMed: 11459597]

Issa NP, Trepel C, Stryker MP. Spatial frequency maps in cat visual cortex. Journal of Neuroscience. 2000; 20:8504–8514. [PubMed: 11069958]

Johnson AP, Kingdom FA, Baker CL Jr. Spatiochromatic statistics of natural scenes: First- and second-order information and their correlational structure. Journal of the Optical Society of America. A. 2005; 22:2050–2059.

Johnson EN, Hawken MJ, Shapley R. The spatial transformation of color in the primary visual cortex of the macaque monkey. Nature Neuroscience. 2001; 4:409–416.

Johnson EN, Hawken MJ, Shapley R. Cone inputs in macaque primary visual cortex. Journal of Neurophysiology. 2004; 91:2501–2514. [PubMed: 14749310]

Johnson E, Hawken M, Shapley R. The orientation selectivity of color responsive neurons in macaque V1. Journal of Neuroscience. 2008; 28:8096–8106. [PubMed: 18685034]

Lennie P, Movshon J. Coding of color and form in the geniculostriate visual pathway (invited review). Journal of the Optical Society of America A. 2005; 22:2013–2033.

Li B, Peterson M, Freeman R. Oblique effect: A neural basis in the visual cortex. Journal of Neurophysiology. 2003; 90:204–217. [PubMed: 12611956]

Livingstone MS, Hubel DH. Anatomy and physiology of a color system in the primate visual cortex. Journal of Neuroscience. 1984; 4:309–356. [PubMed: 6198495]

Lu H, Roe AW. Functional organization of color domains in V1 and V2 of macaque monkey revealed by optical imaging. Cerebral Cortex. 2008; 18:516–533. [PubMed: 17576751]

Mante V, Frazor RA, Bonin V, Geisler WS, Carandini M. Independence of luminance and contrast in natural scenes and in the early visual system. Nature Neuroscience. 2005; 8:1690–1697.

Martin, D.; Fowlkes, C.; Tal, D.; Malik, J. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics; Proceedings of the 8th International Conference of Computer Vision; 2001 July. p. 416-423.

Mullen K. The contrast sensitivity of human colour vision to red–green and blue–yellow chromatic gratings. The Journal of Physiology. 1985; 359:381–400. [PubMed: 3999044]

Olmos, A.; Kingdom, FA. McGill calibrated colour image database. 2004. http://tabby.vision.mcgill.ca

Parraga CA, Brelstaff G, Troscianko T, Moorehead IR. Color and luminance information in natural scenes. Journal of the Optical Society of America A. 1998; 15:563–569.

Rao AR, Shunck BG. Computing oriented texture fields. Computer Vision, Graphics and Image Processing: Graphical Models and Image Processing. 1991; 53:157–185.

Roe AW, Ts'o DY. Visual topography in primate V2: Multiple representation across functional stripes. Journal of Neuroscience. 1995; 15:3689. [PubMed: 7751939]

Ruderman D. Origins of scaling in natural images. Vision Research. 1997; 37:3385–3398. [PubMed: 9425551]

Ruderman D, Bialek W. Statistics of natural images: Scaling in the woods. Physical Review Letters. 1994; 73:814–817. [PubMed: 10057546]

Shapley RM, Hawken M. Neural mechanisms for color perception in the primary visual cortex. Current Opinion in Neurobiology. 2002; 12:426–432. [PubMed: 12139991]

Sigman M, Cecchi GA, Gilbert CD, Magnasco MO. On a common circle: Natural scenes and Gestalt rules. Proceedings of the National Academy of Sciences of the United States of America. 2001; 98:1935–1940. [PubMed: 11172054]

Simoncelli EP, Olshausen BA. Natural image statistics and neural representation. Annual Review of Neuroscience. 2001; 24:1193–1216.

Solomon S, Lennie P. The machinery of colour vision. Nature Reviews Neuroscience. 2007; 8:276–286.

Solomon SG, Lennie P. Chromatic gain controls in visual cortical neurons. Journal of Neuroscience. 2005; 25:4779–4792. [PubMed: 15888653]

Solomon SG, Peirce JW, Lennie P. The impact of suppressive surrounds on chromatic properties of cortical neurons. Journal of Neuroscience. 2004; 24:148–160. [PubMed: 14715948]

Srinivasan MV, Laughlin SB, Dubs A. Predictive coding: A fresh view of inhibition in the retina. Proceedings of the Royal Society of London B. 1982; 216:427–459.

Tailor D, Finkel L, Buchsbaum G. Color-opponent receptive fields derived from independent component analysis of natural images. Vision Research. 2000; 40:2671–2676. [PubMed: 10958917]

Tootell RBH, Silverman MS, Switkes E, De Valois RL. Deoxyglucose analysis of retinotopic organization in primate striate cortex. Science. 1982; 218:902–904. [PubMed: 7134981]

Ts'o D, Gilbert C. The organization of chromatic and spatial interactions in the primate striate cortex. Journal of Neuroscience. 1988; 8:1712–1727. [PubMed: 3367218]

Wachtler T, Lee T, Sejnowski T. Chromatic structure of natural scenes. Journal of the Optical Society of America A. 2001; 18:65–77.

Weliky M, Bosking WH, Fitzpatrick D. A systematic map of direction preference in primary visual cortex. Nature. 1996; 379:725–728. [PubMed: 8602218]

Wyszecki, G.; Stiles, W. Color science: Concepts and methods, quantitative data and formulae. New York: Wiley; 1982.

Xiao Y, Casti A, Kaplan E. Hue maps in primate striate cortex. Neuroimage. 2007; 35:771–786. [PubMed: 17276087]

Xiao, Y.; Rao, AR.; Cecchi, GA.; Kaplan, E. Cortical representation of information about visual attributes: One network or many?; Proceedings of the International Joint Conference on Neural Networks; 2007. p. 1785-1789.

Xiao Y, Rao AR, Cecchi GA, Kaplan E. Improved mapping of information distribution across the cortical surface with the support vector machine. Neural Networks. 2008; 21:341–348. [PubMed: 18249089]

Xiao Y, Wang Y, Felleman DJ. A spatially organized representation of colour in macaque cortical area V2. Nature. 2003; 421:535–539. [PubMed: 12556893]

Xu X, Bosking W, Sáry G, Stefansic J, Shima D, Casagrande V. Functional organization of visual cortex in the owl monkey. Journal of Neuroscience. 2004; 24:6237–6247. [PubMed: 15254078]

|	Small scale	Medium scale	Large scale
Filter size, g_1	3	5	11
Filter size, g_2	5	11	21

Figure 1.
The sizes of the filters used for extracting small-, medium-, and large-scale orientation and color fields.
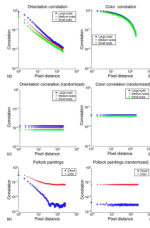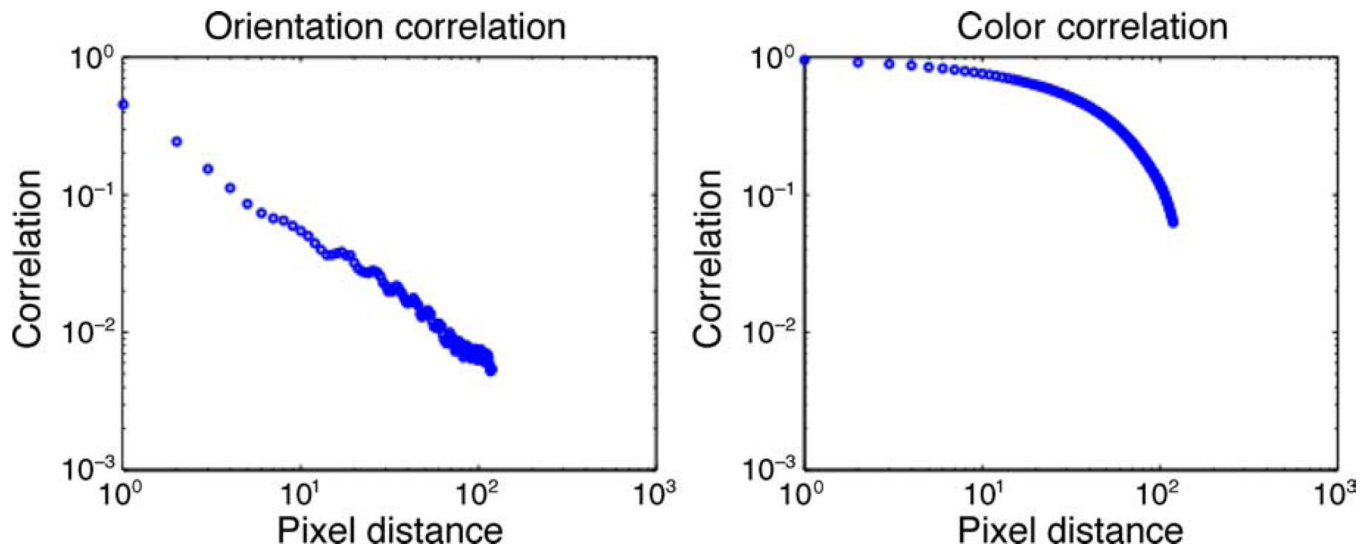
Figure 2.
(a) Spatial autocorrelation of the orientation field in natural images. The orientation is computed at three different spatial scales, ranging from large to small scale. The standard error for the measurements shown is too small to be meaningfully depicted in this figure. For instance, the standard error for the orientation correlation at small scale is 0.0028796 for a pixel distance of 1 and 0.00037322 for a pixel distance of 100. (b) Spatial autocorrelation of the color field in natural images. The color is computed at three different spatial scales, ranging from large to small scale. The original images were smoothed with a Gaussian filter of varying size, as described in Figure 1. The standard error is 0.0022356 for a pixel distance of 1 and 0.0039698 for a pixel distance of 100. (c) The spatial correlation statistics of orientation when the location of orientation vectors is randomized. (d) The spatial correlation statistics of color when the location of color vectors is randomized. In all these plots, the correlations for a given distance have been averaged over all directions. (e) The correlation statistics gathered over five Jackson Pollock paintings. (f) The correlation statistics computed over randomized versions of Jackson Pollock paintings. The locations of existing orientation and color vectors were randomized.
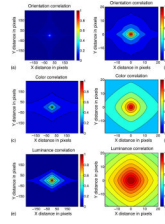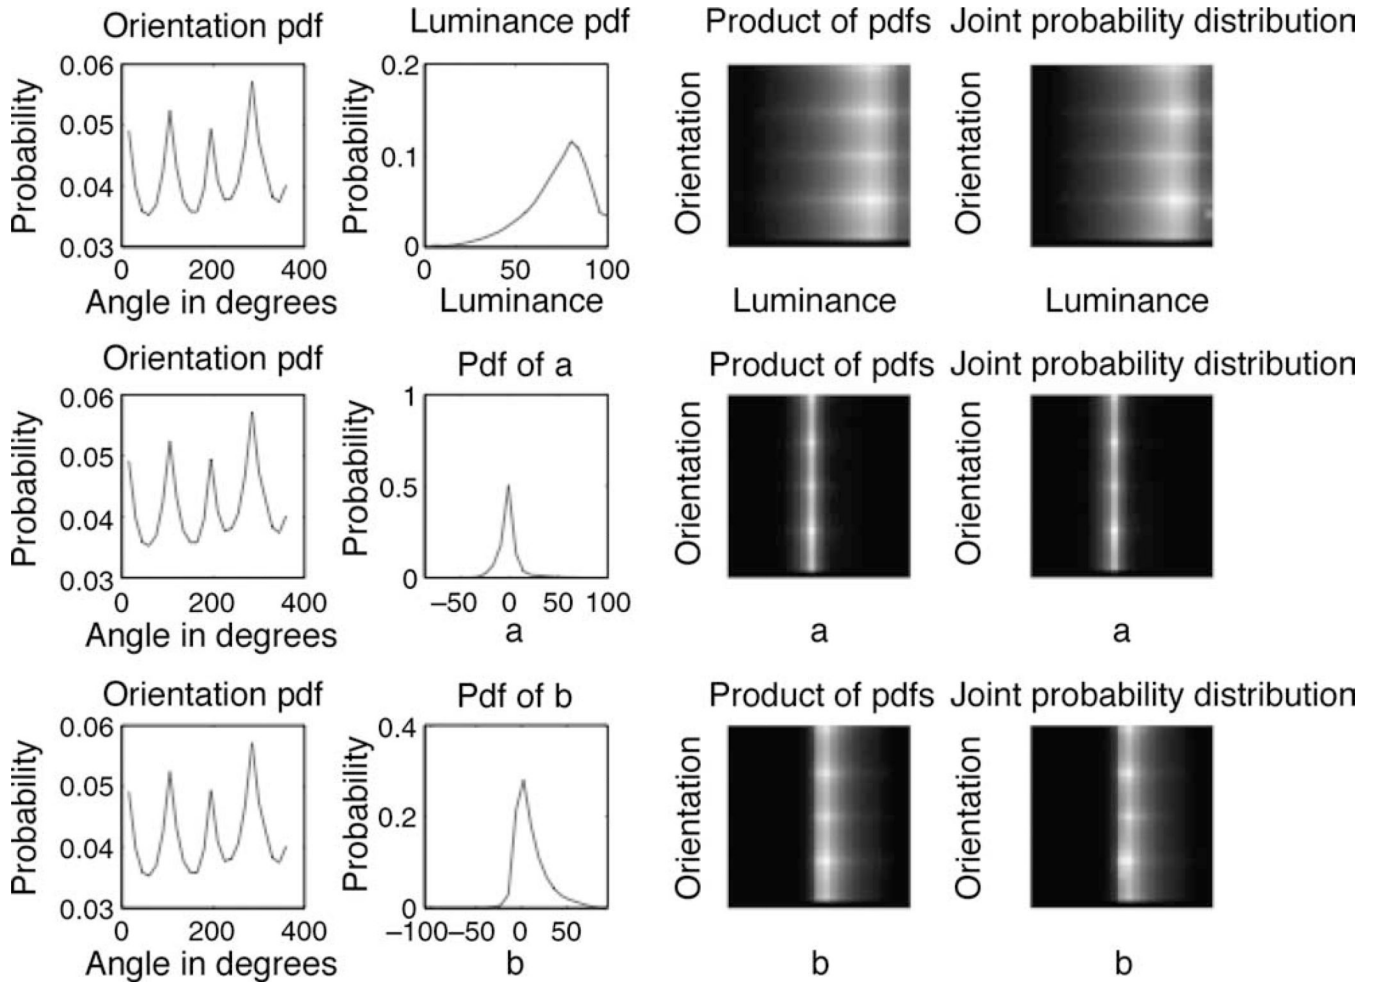
Figure 3.
(a) Spatial autocorrelation of the orientation field in natural images, shown over specific directions. The large-scale filters were used for smoothing the image before orientation and color computations. (b) An enlarged version of the autocorrelation for orientation. Contour plots are used to depict these functions, using the MATLAB command *contourf*. (c) Spatial autocorrelation of the color field. (d) An enlarged view of the autocorrelation for color. (e) Spatial autocorrelation of luminance. (f) An enlarged view of the autocorrelation for luminance.
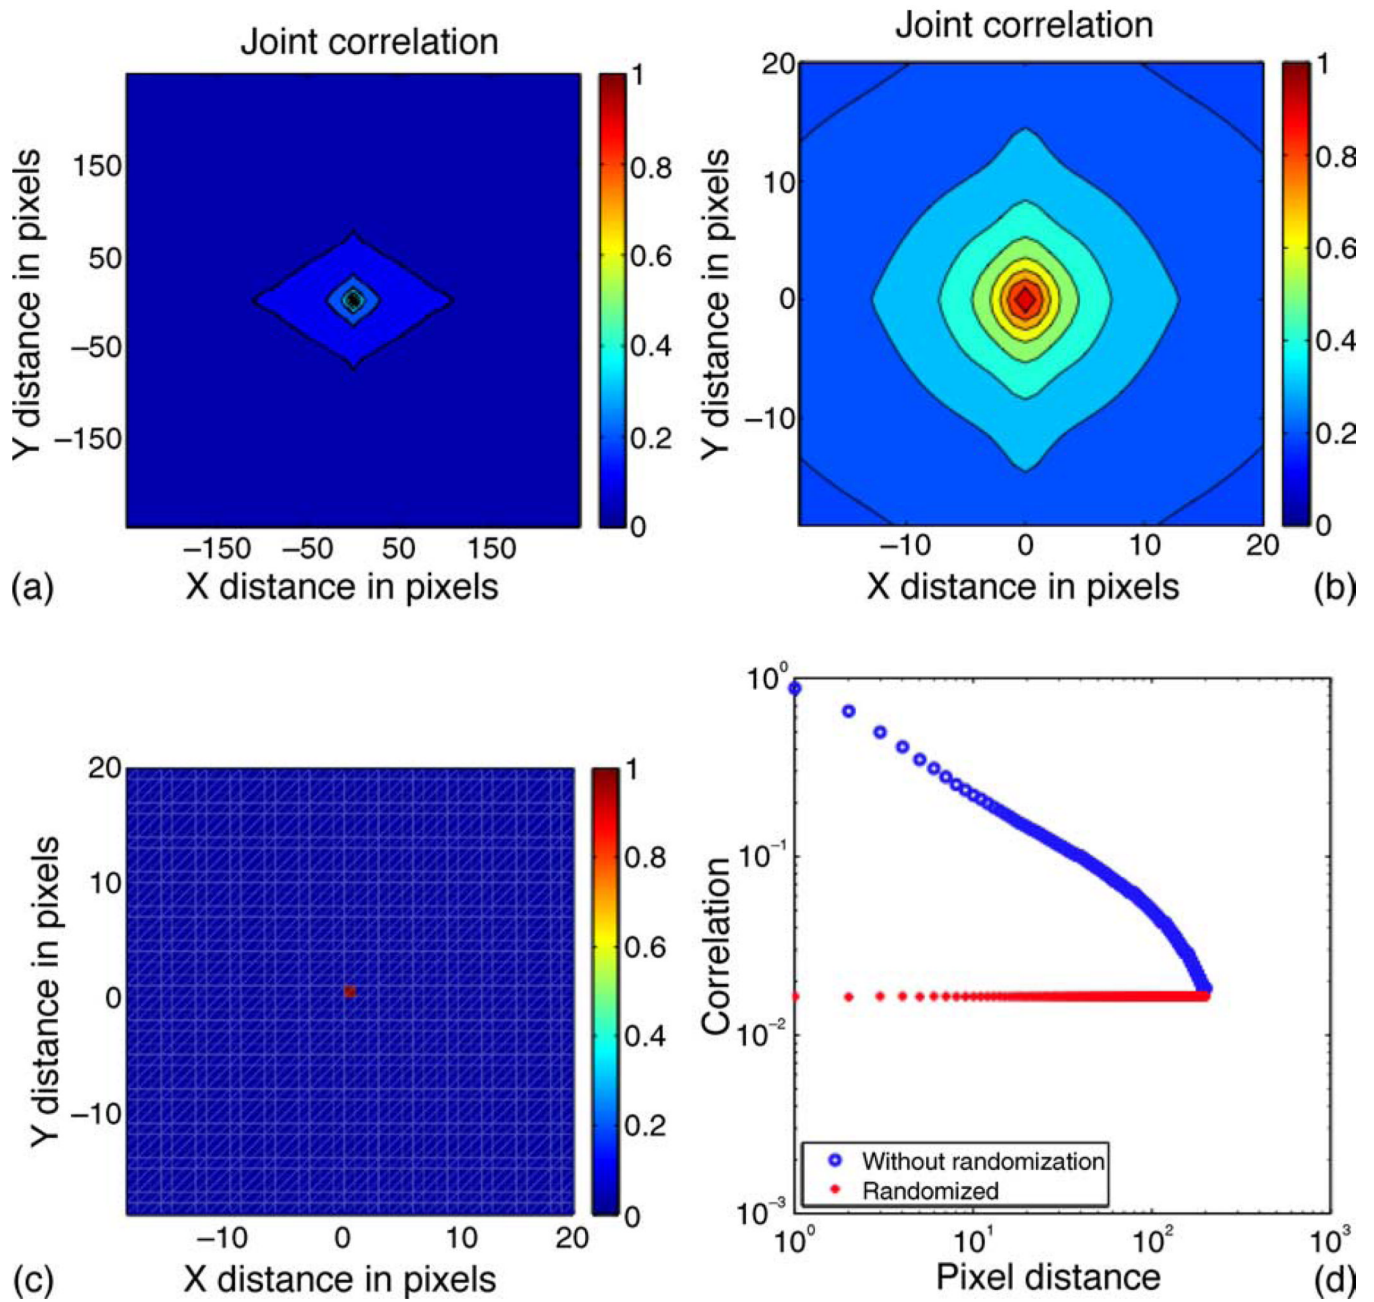
Figure 4.
The correlation statistics for orientation and color computed over the images in the Berkeley Segmentation Dataset and Benchmark (Martin et al., 2001). Small-scale filter sizes were used in this computation, as shown in Figure 1. (a) The correlation statistics for orientation. (b) The correlation statistics for color.
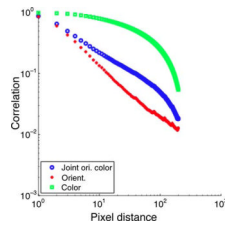
Figure 5.
The marginal probability density functions are displayed in the first two columns. The probability density function for orientation $P(\theta_e)$ is shown in the first column. The second column contains probability density functions for *L, a*, and *b*. The third column shows the pairwise product of the two marginal probability density functions in the first two columns. The fourth column shows the joint probability distribution for the variables in the first two columns. These plots show that the joint probability distributions appear similar to the marginal probability density functions, suggestive of statistical independence of the variables represented in the first two columns.

Figure 6.
(a) The 2D correlation statistics of the joint orientation and color vector as described in Equation 5. (b) An enlarged view of the correlation function around the origin. (c) For this correlation plot, the locations of the 4D vectors were randomized. We show an enlarged view around the origin. (d) A 1D plot for the correlation, generated by summing the correlations within an annulus at a given radius of the function shown in (a).

Figure 7.
The 2D correlation statistics of the joint orientation and color vector as described in Equation 5 are compared with the statistics of the individual color and orientation fields. A 1D plot for each case is generated by summing the correlations within an annulus at a given radius of the 2D correlation function.