Submitted to *Neural Networks*, Special issue on Vision and Brain, (D. Field and L. Finkel eds.), 2004.

This paper should be printed in color. Images are best viewed on a monitor.

Hue Geometry and Horizontal Connections

Ohad Ben-Shahar
Department of Computer Science,
Yale University
New Haven, CT 06520-8285
ben-shahar@cs.yale.edu

Steven W. Zucker
Departments of Computer Science
and Biomedical Engineering, Yale University
New Haven, CT 06520-8285
steven.zucker@yale.edu

Abstract

Primate visual systems support an elaborate specialization for processing color information. Concentrating on the hue component, we observe that, contrary to Mondrian-like assumptions, hue varies in a smooth manner for ecologically important natural imagery. To represent these smooth variations, and to support those information processing tasks that utilize hue, a piecewise smooth hue field is postulated. The geometry of hue-patch interactions is developed analogously to orientation-patch interactions in texture. The result is a model for long-range (horizontal) interactions in the color domain, the power of which is demonstrated on a number of examples. Implications for computer image processing, computer vision, visual neurophysiology and psychophysics are discussed.

1. Introduction

The visual world is a rich source of color-coded information. Ripe red fruit standing out from a background of green leaves immediately suggests one advantage of color selectivity to frugivore primates (Sumner and Mollon, 2000b,a), and the elaborate physiological specialization for processing color attests to its use. Photoreceptors evolved differential spectral responses, and cytochrome oxidase staining has revealed anatomical blobs in cortex. Neurons with opponent receptive field structure provide maps. But there is much more to the cortical processing of color: neural selectivities differ with context ((Wachtler et al., 2003)), and filling-in phenomena (Redies and Spillmann, 1981) suggest boundary and surface material interactions. The interaction between surface geometry, reflectivity, and light sources must matter (Beck, 1972; Horn, 1986); at some level, color-constant behaviors emerge (Zeki, 1993) in a manner such that, as Helmholtz famously observed, the illuminant is discounted. Artists have intuitively exploited this (Albers, 1987), and experience in computer graphics indicates how critical it can be. All of the above, taken together, suggest a much broader and deeper role for color processing in the inference of scene structure than the simple detection of ripe fruit, and our goal in this paper is to sketch one part of a this broader view of color processing.

We concentrate almost entirely on hue, and illustrate, in the next Section, different putative roles for it. Although hue played a fascinating role in the discovery of color opponent processing, it is surprisingly understudied in computer vision (Forsyth and Ponce, 2002). But the time is right to consider hue more directly, because neurophysiological evidence of hue maps is now emerging (Xiao et al., 2003; Wachtler et al., 2003; Hanazawa et al., 2000). In particular, it is unlikely that trichromacy would have re-evolved without providing basic information processing advantages. We argue, in particular, that when the full variability of fruit imagery is considered, the detection problem becomes rather more subtle. This serves, in turn, as a bridge to more abstract questions about the inferences of scene structure.

Our paper is mainly theoretical, and relies on an analogy between two physiological specializations in visual cortex. While differences between the cytochrome oxidase blobs and the interblobs in visual cortex are normally stressed, we shall take the opposite tack and stress a similarity. One primary biophysical difference–metabolic capability–is beyond the context of this paper, and has been considered elsewhere (Allman and Zucker, 1990).

The analogy that we shall exploit is based on long-range horizontal connections. In essence, the standard functional distinction is as follows: neurons in the blobs are not orientationally-selective while neurons in the interblobs are. Furthermore, the common view holds, neurons in the blobs are tuned for color (and contrast); while interblob neurons are not (Livingstone and Hubel, 1984a). While this is of course a crude oversimplification, conceptually it supported a "parallel pathways" view of early vision, with color processing proceeding separately from (and in parallel with) spatial processing (Ts'o and Gilbert, 1988; Zeki, 1993). The existence of separate systems of long-range horizontal connections to support sensory integration within–but not between–these systems has further supported the parallel pathways view.

Our goal is not to challenge the parallel pathways view at this time, although our study of hue does reveal limitations to it: some aspects of processing are most naturally expressed within the hue system and some between the hue and orientationally-selective systems. Nor are the implications of our research limited to V1; the specialization projects naturally into V2 and beyond ((Xiao et al., 2003; Roe and Ts'o, 1999; Kiper et al., 1997; Komatsu, 1998; Heywood et al., 1992; Livingstone and Hubel, 1984a; Zeki, 1993)). Rather, in more general terms, we seek first to understand the basic structure of hue and its uses as a way to inform the search for its natural neurophysiological implementation. Thus we exploit a generic interpretation of the long-range horizontal connections as the substrate for sensory integration. But we place a deeper technical interpretation on this than is normally done, in the following sense. One reads frequently that long-range horizontal connections allow "like" properties to be supported by "like" properties in their surround. For example, color connections support color connections, and shape connections support shape connections. More specifically, in V1 it is held that like (*e.g.,* vertical) orientations support like (*e.g.* vertical) orientations, because boundaries "typically" continue along the same direction (T'So et al., 1986; Livingstone and Hubel, 1984b; Field et al., 1993). (Just think of the edge of a Manhattan skyscraper.) Similarly, long-range horizontal connections preferentially couple red/green opponent cells in one cytochrome oxidase blob to those in another "red/green blob," and not to the "blue/yellow blobs." Long-range horizontal connections, in short, are taken to connect like properties with like properties. But there exists significant variation in the actual long-range connections from this presupposed ideal, and the question is whether this variation is noise or is functionally significant.

Our deeper analysis of the connections within the interblob system is based on a proof that a functionally-significant interpretation is sufficient to explain the orientation-based, interblob variations, predicting the variation through second-order statistics ((Ben-Shahar and Zucker, 2004a)). To continue the above parenthetic comment, our model exploits the observation that there are few man-made skyscrapers in the rain forest. Therefore (and this is the central point) one must study sensory integration in the context of general curves rather than straight lines. This requires a deeper theory of what "like properties" are, and our theory for the orientation systems (boundaries along one dimension; textures and shading in two) is briefly sketched in Sec. 3. In this paper we show formally that something analogous holds for the color system: While some fruits ripen to a uniform hue, some vary, and the intermediate case is clearly not constant. The analogy with hair patterns provides a conceptual dual: some are constant or slowly varying, and some vary more rapidly. All involve singularities, either point-wise, as around the crown of the head, or along a curves. Another example: embarrassment causes facial blushing but, again, not uniformly (Ekman, 2001) The result is that hue variation is like texture and shading variation: *the structure is in the flow*. We call the basis for this flow the *hue field*, and illustrate it in the next Section.

The technical contribution in this paper is a derivation of a differential-geometric model for hue interaction that makes this concept of hue flow concrete, and an illustration of how it can be applied in a range of scene applications. This provides a beginning for a more complete model of color processing, because it suggests how a color (specifically hue) Gestalt might emerge. To return to the opening example of finding ripe fruits in trees, current evidence suggests that there is much more to the feeding behavior of many primates than the detection of individual ripe fruits. In particular, differentiating young nutrient-rich leaves from more difficult-to-digest surrounds underscores the need to process extremely complex arrangements of color and shape simultaneously (Dominy and Lucas, 2001).

2 The Hue Field

Consider the HSV color space, in which a color image is a mapping $\mathcal{C} : \mathbb{R}^2 \to \mathcal{S}^1 \times [0,1]^2$, where $\mathcal{S}^1$ is the unit circle (see Fig.1). The hue component across the image is a mapping $\mathcal{H} : \mathbb{R}^2 \to \mathcal{S}^1$ and thus can be represented as a unit length vector field over the image plane , henceforth called the *hue field*. Displays of the hue field reveal that it may vary greatly, albeit smoothly, even *within* perceptually coherent objects (see Fig 2; many other examples will follow).

The hue field has deep roots in the psychology of perception and the subjective experience of color, from Hering's opponent hues theory (Hering, 1964) and Munsell's book of colors (Munsell, 1905), to accumulating psychophysical evidence on color sensitivity (De Valois and De Valois, 1990). Neurophysiological findings for color opponent cells, color specificity in V4, and perceptual impairments such as achromatopsia (Spillmann and Werner, 1990) indirectly imply the existence of neural structures that explicitly encode the hue; explicit evidence for them has now been found in V1 and V2 (Hanazawa et al., 2000; Xiao et al.,

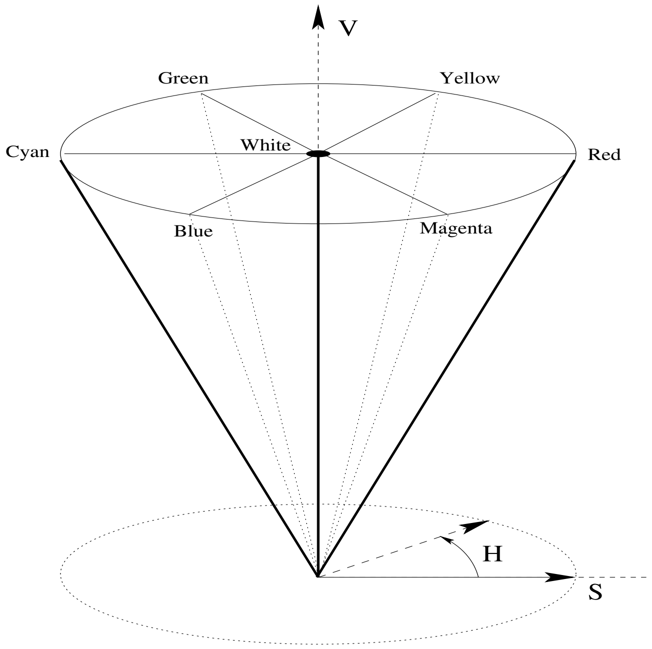

Figure 1: The HSV color representation in $\mathcal{S}^1 \times [0,1]^2$ and the color wheel.

2003; Wachtler et al., 2003).

Despite the prevalence of apples in our shops, color, even in modern analyses, is usually taken to be piecewise constant, as we now review. That it actually varies in many natural images, as shown above, gives rise to a number of problems in processing color imagery. The situation is strongly analogous to orientationally-defined textures, such as hair patterns, as we review in the final part of this Section. Curiously, most earlier analyses of these textures also assumed they were piecewise-constant.

2.1 Hue and Its Uses

Scene geometry, light sources, and surface properties conspire in the formation of images. Nevertheless, perceptually, we seem able to make relevant inferences about these different aspects of scene structure almost without effort; one normally needs an "artist's eye" to notice when our perceptual mechanisms fail.

There are two basic approaches to determining which aspects of color are used. Physiologists and psychophysicists seek to determine which aspects and representations of color are explicitly used by our visual systems. This *bottom-up* approach is complemented by a *top-down* approach of trying to understand which

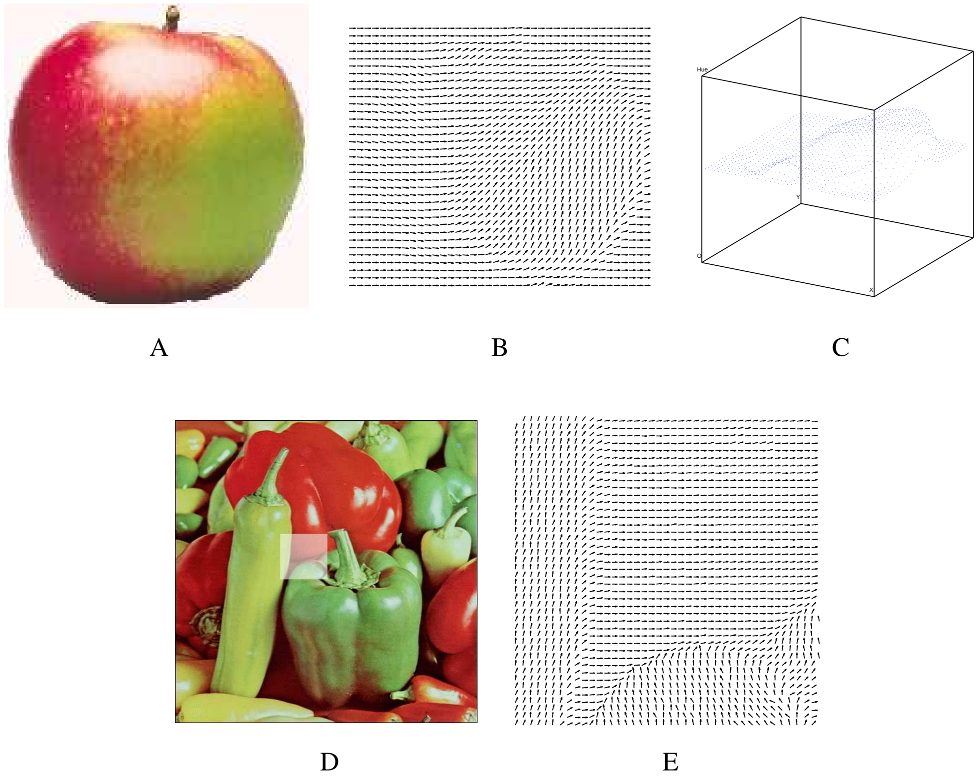

Figure 2: Color images, and their hue fields, are typically piecewise smooth. Most importantly, their hue can vary smoothly even within perceptually coherent objects. Thus, a representation sufficiently rich to support these variations is necessary. **(A)** A natural image of an apple with varying hue. **(B)** The corresponding hue field (see Section 3). Note how it changes smoothly across the apple's surface. **(C)** A 3D representation of the hue filed, where hue is represented as height. Identifying the top face with the bottom (since hue is a circle) leads to the space $XY\mathcal{H} \triangleq \mathbb{R}^2 \times \mathcal{S}^1$ in which the image's hue is a submanifold. **(D)** A natural image of peppers with a region of interest. **(E)** The hue field of the peppers image in the region of interest is piecewise smooth. In general, occlusion boundaries between objects in the world induce hue singularities (up to some blurring from the imaging process) that must be preserved in every processing of the color information. This, in effect defines the computational task that we shall be facing.

information processing tasks depend on color, and why.

We believe these two complementary approaches are best combined into a hybrid, in which different aspects from each are combined. From neurophysiology we take the observation that hue is a viable representation, and the existence of long-range horizontal connections that could support hue good continuation. From psychophysics we take color constancy and filling-in phenomena. Hans Wallach observed that color constancy, like perceived lightness, is mainly the result of ratios computed at edges (Palmer, 1999), which implicates processes of edge detection and segmentation along surface borders. If edges are incorrect, then lightness and constancies are likely to be incorrect. But constancies are rather subtle, and more primitive information processing tasks relate to the role that color plays for primates in finding food. This, too, raises questions about segmentation, with detection as the goal.

Several of these different tasks are now reviewed, with the aim of demonstrating how local measurements are insufficient for these different tasks in general. Rather, the fact that hue can smoothly vary over an object leads us back to characterizing and seeking smooth hue representations. The foundation of the problem is to determine what is a coherent hue distribution and what are its singularities.

2.1.1 Image Segmentation and Color Borders

The first example in the Introduction—detecting ripe fruits against a background of green foliage—is the classical one. The computation is local, and hue is taken to be effectively constant; see Fig. 3. Statistical detection models classify pixels based on their spatiochromatic structure ((Fine and Boynton, 2003; Gegenfurtner and Rieger, 2000) possibly applied at different levels, from the photopigments ((Sumner and Mollon, 2000a)) to color opponent and double-opponent cortical receptive fields ((Shapley and Hawken, 2002; Wachtler et al., 2003)). Such receptive fields could reflect the statistics of natural scenes ((Bell and Sejnowski, 1997; Simoncelli and Olshausen, 2001; Tailor and Buchsbaum, 2000; Wachtler et al., 2001). As we show, however, much of the information about hue flows involves spatial organization as well, which suggests its statistical signature would be rather higher-order.

If arranged into proper circuitry, circular-surround receptive field models may underlie color edge detection (Shapley and Hawken, 2002) reflecting all of the questions regarding the circuitry underlying simple and complex cell receptive fields. Michaels (Michael, 1981) first identified such receptive fields, and similar color edge detectors have also been developed in computer vision; see, e.g., (Ohta and Sakai, 1980; Novak and S., 1987) to support color image segmentation.

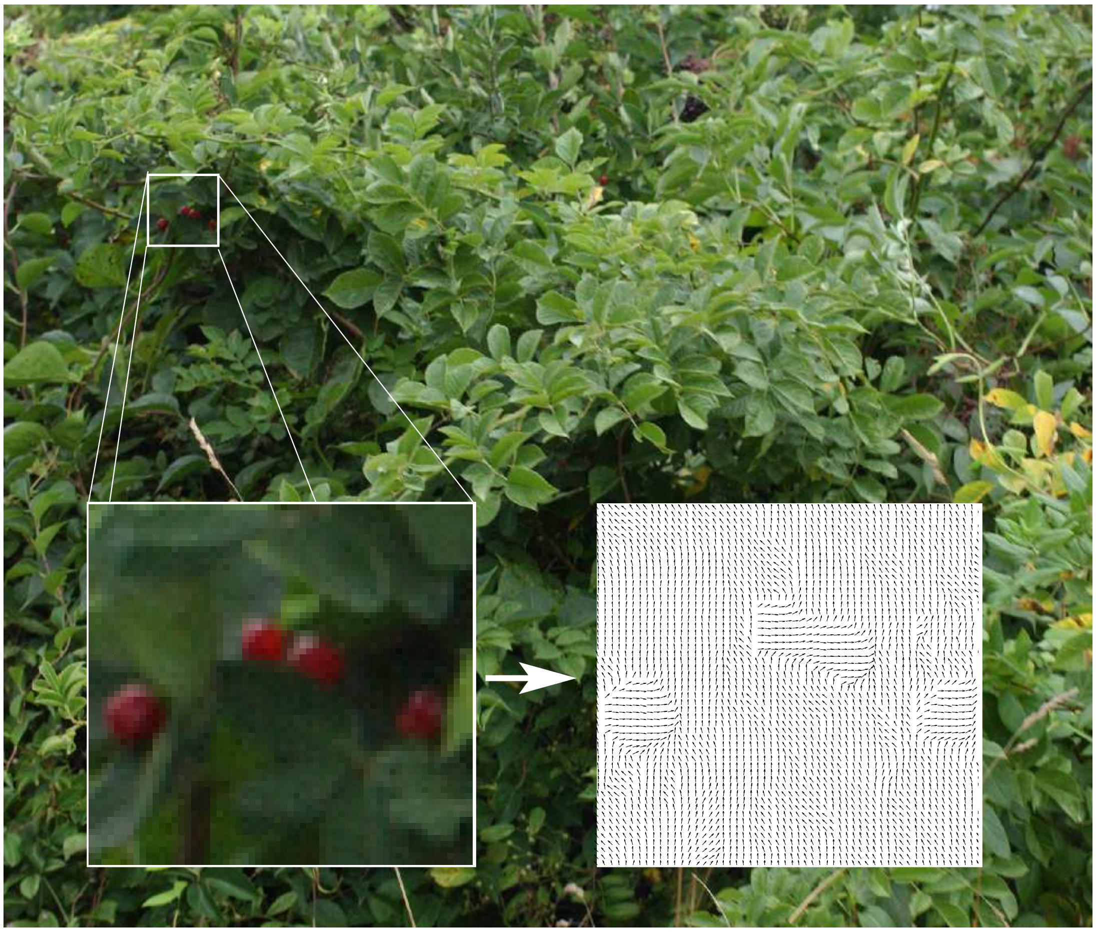

Figure 3: Illustration of the classical problem of recognizing red berries against a background of green foliage. Shown are the color image (left inset panel), with its hue field (right inset panel). Such tasks are naturally expressed in a local signal detection framework. While this example, which was photographed in bright sunlight, is relatively unchallenging, the task in general is one of detecting the berries when noise obscures the measurements.

The performance of segmentation algorithms is limited by the thresholding (or signal detection criterion), and typically a fixed value is sought (Ohta and Sakai, 1980; Novak and S., 1987; Healey, 1992). This approach, of course, encounters difficulties when color varies, as may be the case for single natural objects (recall the apple in Fig. 2). Attempts to use (possibly different) distributions at each local retinotopic point do not solve the problem, which suggests there does not exist a purely local solution (Ruzon and Tomasi, 1999).

As a second example examine Fig. 4, which shows a peach in different configurations. Peaches, like the apple shown previously, vary in color, and therefore proper thresholds based on local information may not exist.

Apples and peaches raise the issue of understanding the relationship between the evolution of color vision and diet, since it has been presumed for more than a century that the origin of color vision has been the detection of ripe fruits against a background of leaves (Allen, 1879). Comparing berries to peaches begins to reveal how complex this detection task might be. However, the relationship between color vision and diet is a topic that is currently undergoing a fascinating development (Dominy and Lucas, 2001; Lucas et al., 2003). Foraging advantages, in particular, are emerging as a link between the red-green subsystem and the detection of young, nutrient-rich leaves amid more mature foliage (Dominy and Lucas, 2001). However, the biology of leaves also leads directly to color-varying objects. The apparent color is due to the differential absorption of light by a leaf, e.g. for photosynthesis. Pigmentation also plays a definitive role in apparent color. There are two classes of carotenoid pigments, depending on whether they contain oxygen or not. Carotenes do not and make red and orange colors, while xanthophylls do have oxygen and are generally yellow. Beta-carotene gives carrots their color and zeanthin is the gold of corn. However, the fall foliage in New England illustrates clearly how variable this pigmentation can be–both in time and in space–thereby underlining the need to be able to process color-varying stimuli. Such advantages supplement the ones discussed above in detecting fruit against foliage (Sumner and Mollon, 2000a). Fig. 5 illustrates this variable green-red shift in a dramatic instance. Foliage, like fruit, gives rise to a rich, non-constant color Gestalt.

A second important aspect of considering hue is also apparent in Fig. 5. Since the hue is dictated largely by pigmentation, is remains somewhat invariant to light source effects such as shadowing and highlights. Re-examining the apple in Fig. 2 provides another example of this phenomenon. Taking this observation together with those from the leaf and the fruit examples above, it follows that mechanisms for determining smooth color flows could have a tremendous evolutionary advantage.

We show next that similar issues arise in lightness algorithms.

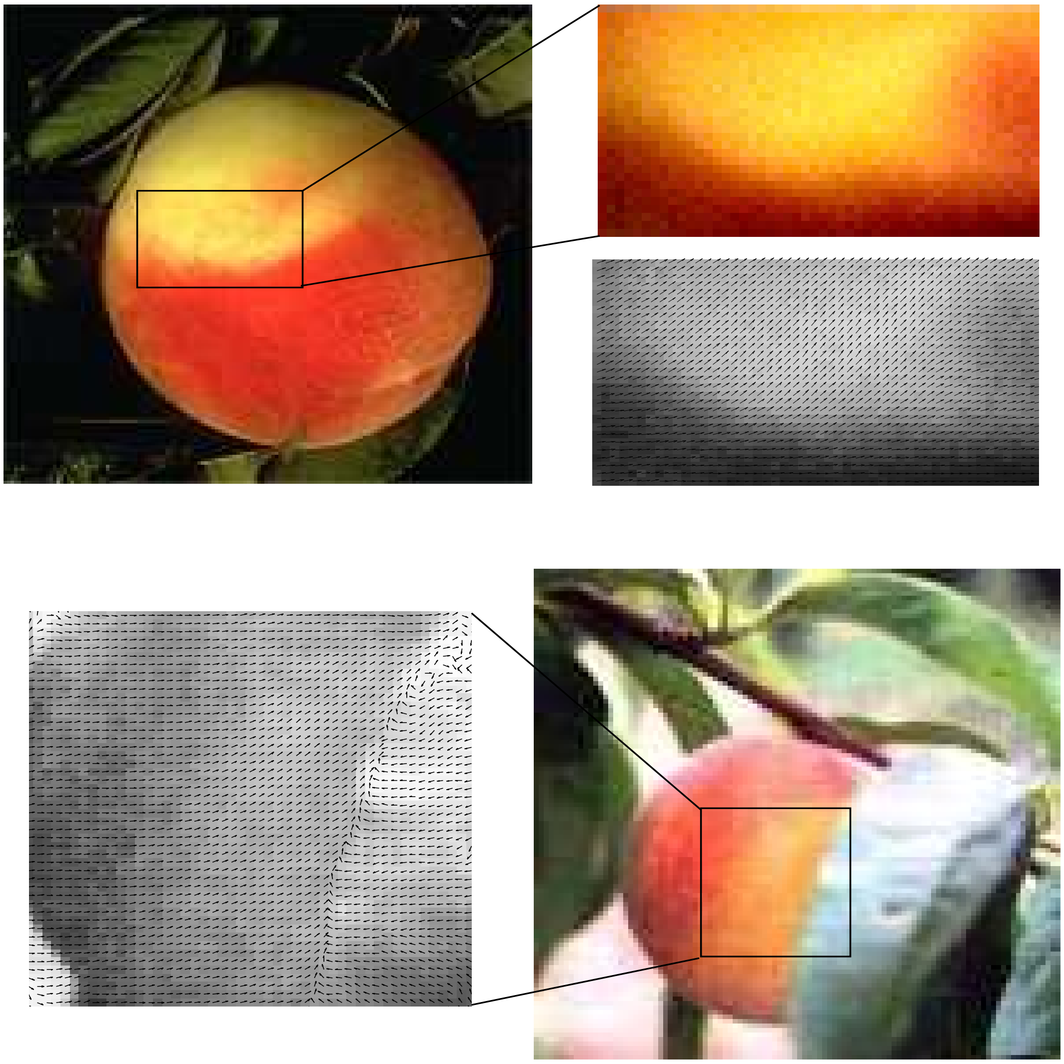

Figure 4: Peach images, and their hue fields, are piecewise smooth (but not necessarily constant), demonstrating that hue can vary smoothly within otherwise perceptually coherent objects. Proper representations of hue, that make these variations explicit, are therefore necessary, and segmentation algorithms must be developed that can tolerate such variation within the object but not between object and background. (*Top*) Image of a peach with a blowup of the color variation and its associated hue field. Such fields should remain coherent to support recognition of the peach. (*Bottom*) Illustration of a boundary between the peach (occluded) and a leaf (occluder). Such boundaries are more plausible for segmentation.

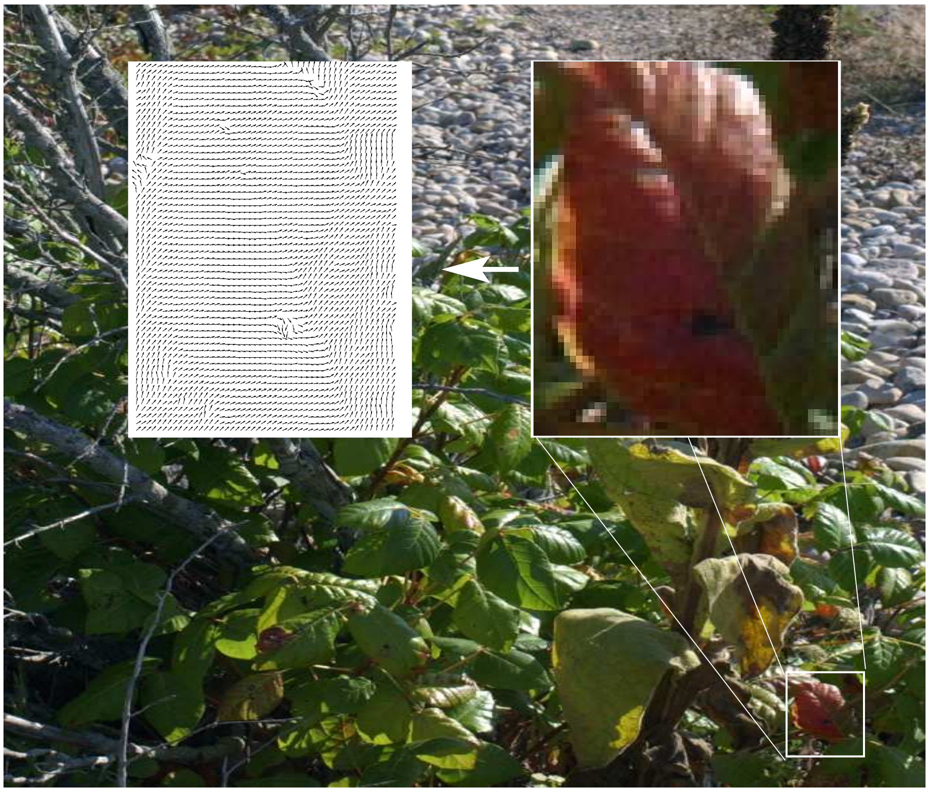

Figure 5: Illustration of color variation in leaves of poison ivy (Rhus radicans L.). The red coloration of the leaf is due to the accessory carotenoid pigments which broaden the range of wavelengths absorbed in photosynthesis and also serve to protect the sites of chlorophyll synthesis during leaf growth. Although this is a mature leaf, it illustrates similar biophysics to the young leaves important in some primate diets. Insets show a single leaf enlargement and its hue field. Notice that the hue field varies slowly with pigment variation and is largely invariant to lighting changes. (Color variations may not produce well in print. Please refer to the electronic version and view image on a monitor)

2.1.2 Lightness Computations and Thresholding

Helmholtz suggested discounting the illuminant by a process of unconscious inference, and lightness algorithms implement this with the assumption that color changes abruptly while lighting changes smoothly. Retinex theory (Land and McCann, 1971; Land, 1977) applied to Mondrian images leads to the standard implementation (Horn, 1986): since (by assumption) small, slow changes in illumination are due to light source effects such as shadowing, and large, abrupt changes to surface reflectance effects, such as edges between Mondrian patches, thresholding the logarithmic brightness image (or related edge maps) can isolate the edges as large jumps. While this works for some situations, revisions subsequently led to taking thresholded values for the gradient of the chromaticity field and comparing their location with luminance thresholds (Blake, 1985; Funt and Barnard, 1998). But problems remain around slow changes (e.g., shadows), which appear in this model as *steps* of edges (the number of steps depends on the threshold). Barnard *et al.* (Barnard et al., 1997) suggest reasoning about these sequences of step edges, a sort of heuristic return to Helmholtz, but this is, of course, just one aspect of the problem of slow changes discussed above. Our approach is to model these slow changes directly.

The confounding of information between shading and color confuses lightness algorithms, with the basic question remaining how to separate reflectance edges from illumination edges. The dominant approach–segregating different variations into two classes, abrupt and slow, by thresholding–was also used as the basis for texture segmentation, which we discuss below. This is actually a corollary to the analogy set up in the Introduction: in effect, Mondrian patterns are as informative to general lightness computations as piecewise-constant orientation fields are to general orientation-based texture segmentation.

2.1.3 Image Denoising and Diffusion

Our next example comes from image processing, and illustrates the practical need to process color imagery to remove noise.

Following early attempts to denoise color images through independent smoothing of the RGB channels, practically all contemporary approaches focus on a variety of filtering processes applied to the color data. While some studies explore vector median and directional filters (Astola et al., 1990; Trahanias and Venetsanopoulos, 1993), most color image enhancement algorithms are based on a form of anisotropic diffusion (Perona and Malik, 1990; Weickert, 1997; ter Haar Romeny, 1994), either on an explicit vectorial representation of the color, or based on differential geometrical properties of a manifold representation in

Figure 6: Color denoising of the Apple image (from Fig. 2) **(A)** Original image. **(B)** Noisy image. **(C)** Result of color diffusion in RGB space (Sochen et al., 1998; Kimmel et al., 2000). Note that although the noise is not yet removed, the color structure across the apple is completely blurred and distorted, as is the case across the apple's boundaries.

a higher dimensional space (Tang et al., 2001; Sapiro and Ringach, 1996; Yezzi, 1998; Sochen et al., 1998; Kimmel et al., 2000; Kimmel and Sochen, 2000). While diffusion in color space can work within very smooth regions, it does have the tendency to blur inappropriately; see Figure 6. Thus the requirement is not just for techniques that yield smooth color variations, but rather for ones that smooth only when appropriate. Unconstrained blurring reveals, in a sense, problems which are the opposite of trying to find a single, global threshold.

Sometimes, however, hue does spread across otherwise natural boundaries. Next we consider possible physical causes to this phenomenon.

2.1.4 Color Bleeding

The detection view (Sec. 2.1.1) emphasizes point-wise information. Fine *et al.* (Fine and Boynton, 2003) connect this to surfaces: "Two pixels that fall on the same surface are likely to have the same luminance and color, and two pixels that fall on different surfaces are likely to differ in both luminance and color" (p. 1283). Thus, on average for some class of scenes, there is an alignment of chromatic and luminance variations (e.g, Fig. 2D). But hue may cross between surfaces, as the following phenomenon demonstrates.

Color bleeding is the result of one colored surface reflecting light onto another. This type of inter-reflection demonstrates an advantage to separating chromatic variation from luminance variation; see Figs. 7. Although we do not often notice color bleeding effects as salient, artists have been concerned with reproduc-

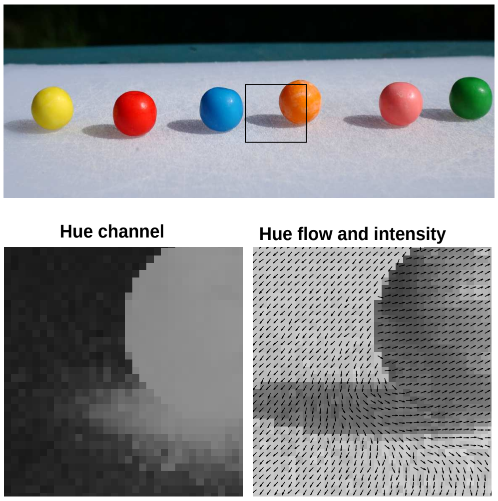

Figure 7: An example of color bleeding. It is most prominent in the deep concavity between the sphere and the tabletop. Notice how it is elongated in the direction of the light source and, in particular, how it differs from the cast shadow. The hue field shown here is based purely on local measurements of hue. While this local field is sufficient for us to view quickly, a more careful examination reveals noise in particular places and, more generally, the subtlety in localizing the hue borders between different regions.

ing it (Baxandall, 1995), and da Vinci wrote about it.[1] It follows, then, that color bleeding does carry some information. One clue is that Fine *et al.*'s conclusion above derives from a database of images with few deep shadows (Fine and Boynton, 2003). However, when shadows or shading or related types of inhomogeneous illumination are considered, pure or near-pure luminance variations can arise, and humans are sensitive to this psychophysically (Kingdom, 2003). Separating luminance and chrominance, and comparing and contrasting their relative structure, provides information about both surface color and light-source placement that differs from – and supplements – that of cast or attached shadows. We shall return to this topic.

2.2 The Analogy to Orientation Flows

While the apparent structure of hue may seem quite different from texture (Fig. 8), we claim that there is a sense in which hue fields are strongly analogous to oriented texture fields. To start with the obvious, both are

[1]Baxandall (Baxandall, 1995) claims that da Vinci intended to dedicate an entire section of his treatise on shadow to deal with it (p. 110).

related to descriptions of surface coverings, and therefore both can be modeled as (two-dimensional) fields. Both orientation and hue take values in the unit circle S^1. But the analogy becomes much more informative when one considers the segmentation task described earlier. We now review the segmentation of oriented textures, and show how it leads to the consideration of second-order (flow) models involving curvature. This suggests that some notion of curvature must underlie hue fields as well.

Orientation-based texture segmentation is the process of detecting texture boundaries, or abrupt changes in orientation that align to form a curve. This is a well-studied problem in visual psychophysics, and most formulations (Nothdurft, 1991; Landy and Bergen, 1991; Mussap and Levi, 1999) seek those points where the change in orientation *between* perceptually coherent regions (i.e., across perceptual boundaries) greatly exceeds that *within* coherent regions, quantities sometimes called orientation contrasts; see Fig. 9. These models predict that segmentation occurs reliably if and only if the ratio $\frac{\Delta\theta_{between}}{\Delta\theta_{within}}$ is significantly larger than 1. However, we have been able to show that this ratio is not a complete description of texture segmentation; see Fig. 10. In this example, $\nabla\theta$ (and thus, $\Delta\theta_{within}$) is constant within the figure and within the ground, and $\Delta\theta_{between} = 18^\circ$ is constant across the figure's edges (a square). Nevertheless, the saliency of the top edge is significantly higher than that of the bottom edge, which is hardly detectable without scrutiny.

We have conducted a large computational and psychophysical study of orientation-based texture segmentation in which we generalized the classical results using piecewise-constant patterns to smoothly varying ones. In brief, two curvatures are required to specify the local flow of orientation textures, one which describes it in the direction tangent to the flow at a point and another which describes it in the direction normal to the flow. The fingerprint image illustrates how richly flows can curve, and also illustrates the point singularities that emerge. These point singularities differ from the one-dimensional singularities along borders, but are also important. These texture curvatures are described in detail in (Ben-Shahar and Zucker, 2003) and the psychophysical results in (Ben-Shahar and Zucker, 2004b).

The second sense in which texture segmentation is like hue segmentation involves the refinement of local measurements. Since the overall flow structure holds over a neighborhood rather than just in a point, a technique must be available to refine local, noisy measurements into globally-coherent wholes. We have developed a relaxation procedure that takes initial noisy measurements and relaxes them to a smooth flow almost everywhere (Ben-Shahar and Zucker, 2003). It is necessary, of course, that the relaxation does not regularize true discontinuities or singularities away; see Fig. 11.

The third analogy between texture and hue was presented in the introductory argument: just as the long-

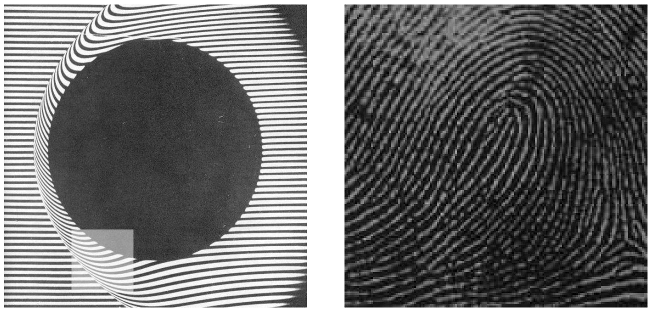

A

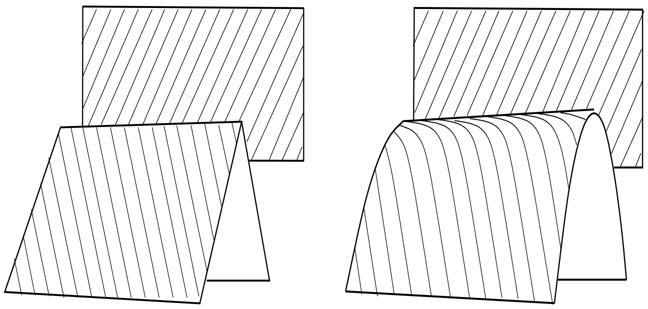

B

Figure 8: Examples of orientation flows. (**A**) Fluid flows and fingerprints are just two examples of orientation flows, whose line and point singularities are significant precursors to their interpretation and understanding. Other natural examples of include surface markings, such as the hair on a primate, the grass on a field, the pattern of stripes on a zebra. (**B**) Line drawings are another example of oriented texture patterns. These examples illustrate how the texture lines on a surface smoothly curving away from the observer tangentially approach the boundary (so-called ”folds”, *right*), while those on surfaces abruptly ”cut” off from the viewer approach the boundary transversely (*left*).

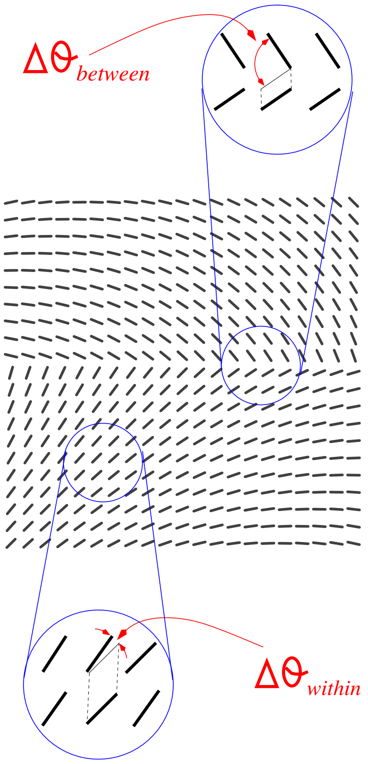

Figure 9: Standard models for orientation-based texture segmentation predict that it depends on the relationship between two orientation gradients, one within and the other between perceptually coherent regions. After (Ben-Shahar and Zucker, 2004b).

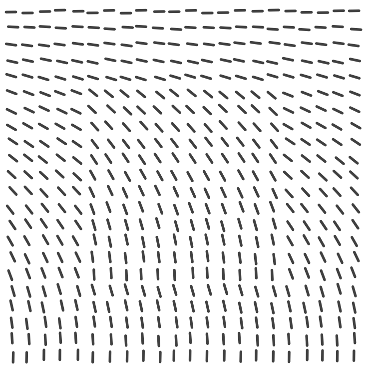

Figure 10: A demonstration that $\Delta\theta_{between}$ and $\Delta\theta_{within}$ are insufficient to determine the perceptual outcome of orientation-based texture segmentation. $\Delta\theta_{within}$ is constant within the figure and within the ground, and $\Delta\theta_{between}$ is constant across the figure's edges (a square). Nevertheless, the saliency of the top edge is significantly higher than that of the bottom edge. This is because of jumps in the curvature, which are larger at the top than at the bottom. (After (Ben-Shahar and Zucker, 2004b).)

range horizontal connections within the interblob system are a sufficient substrate for boundary, texture, and shading computations, we submit that the substrate for processing hue interactions will be similar.

While discussing the analogies between orientation and hue, we must caution that there are differences as well. One is suggested in Fig. 8, (*bottom*). Forshortening generically bends apparent surface textures, or shading (Breton and Zucker, 1996), to meet boundaries tangentially, for smooth surfaces folding away from the viewer (Huggins and Zucker, 2001; Huggins et al., 2001; Ben-Shahar et al., 2002), but we have seen no generic counterpart to this forshortening effect for hue.

2.2.1 Replacing the Constant Color Assumption

The brief discussion of lightness and segmentation algorithms above illustrates how common a "piecewise-constant color" assumption is in computational vision. The brief discussion of orientation-based texture segmentation suggested how common a "piecewise-constant orientation" assumption was in visual psychophysics (for more detailed review, see (Ben-Shahar and Zucker, 2004b)). And a consequence of this constancy assumption is normally made regarding the long-range horizontal connections between interblob neurons: in the strictest form, only like (i.e., constant) orientations are supported. Although the long-range

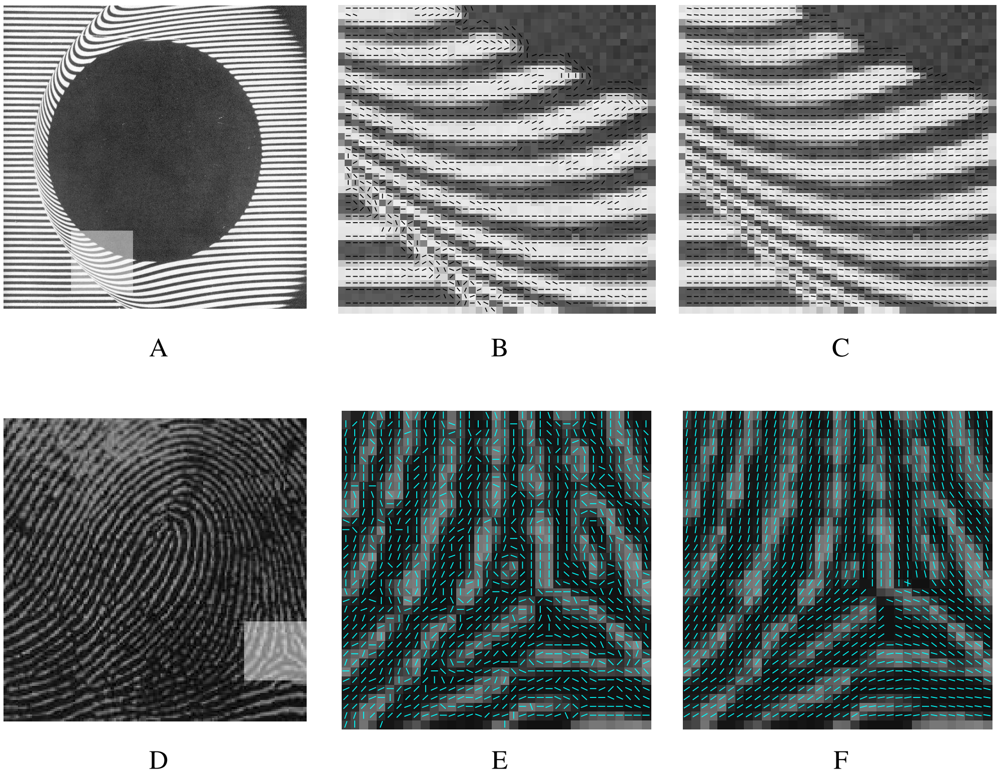

Figure 11: A relaxation procedure that enforces good continuation of orientation through curvature results in a coherent orientation flow. (**A**) A supersonic fluid flow and ROI. (**B**) Initial measurements are both noisy and non dense. (**C**) Convergence state of the relaxation process. Notice the well-defined line singularity, as well as the preservation of the flow boundary. (**D**) Fingerprint image and ROI around one of the point singularities. (**E**) Noisy initial measurements. (**F**) Convergence state of the relaxation process. Note the well-defined singularity. After (Ben-Shahar and Zucker, 2003)

horizontal connections between neurons in the cytochrome oxidase blobs are not nearly as well studied as those between the interblobs, the same view is being adapted. Quoting T'so and Gilbert: "Like the specificity of the interblob connections for columns of matched orientation preference, among the blob connections we found a specificity of blob connections for receptive field type and color selectivity."(Ts'o and Gilbert, 1988, p. 1726). However, our examination of natural images shows how rich and varied colors can be, just as orientation flows are essentially never constant.[2]

As we shall now show, this variation is naturally modeled in terms of hue flows. The basic idea is the following: if hue behaves like orientation, two nearby color patches should be mutually coherent if their hue orientations posses mutual *geometrical* good continuation. This amounts to assessing the degree to which each hue measurement is *geometrically* compatible with the context in which it is embedded, and whether or not that context is part of a single whole.

3. Hue Geometry and Hue Curvatures

Recall our earlier definition of the hue field. In HSV color space, a color image is a mapping $\mathcal{C} : \mathbb{R}^2 \to \mathcal{S}^1 \times [0,1]^2$, where $\mathcal{S}^1$ is the unit circle. The hue component across the image is a mapping $\mathcal{H} : \mathbb{R}^2 \to \mathcal{S}^1$ and thus gives rise to a unit length vector (hue) field over the image plane.

An extension of the vector field representation that makes tools from differential geometry readily available is that of the *frame field* (O'Neill, 1966). More specifically, by attaching a frame field $\{\mathcal{H}_T, \mathcal{H}_N\}$ to each point in the image domain, we now not only represent the hue vector itself, but also a local coordinate system in which all other vectors can be represented in a natural, object centered view (Fig. 12). Perhaps the most important vectors (other than the frame vectors themselves) are the covariant derivatives of $\mathcal{H}_T$ and $\mathcal{H}_N$. These covariant derivatives represent the initial rate of change of the frame in any given direction $\vec{\mathbf{V}}$, a quantity which in the $\{\mathcal{H}_T, \mathcal{H}_N\}$ coordinates is captured by Cartan's *connection equation* (O'Neill, 1966):

$$\begin{pmatrix} \nabla_V \mathcal{H}_T \\ \nabla_V \mathcal{H}_N \end{pmatrix} = \begin{bmatrix} 0 & w_{12}(V) \\ -w_{12}(V) & 0 \end{bmatrix} \begin{pmatrix} \mathcal{H}_T \\ \mathcal{H}_N \end{pmatrix} \quad (1)$$

The coefficient $w_{12}(V)$ is a function of the tangent vector V, which represents the fact that the local behavior of the flow depends on the direction along which it is measured. Fortunately, $w_{12}(V)$ is a *1-form* and thus

[2]The smooth surface variation around "fold" type edges reveals how rare and non-generic constant orientation flows are; for an orientation flow to be constant, the flow, the surface, and the viewer must all be in a special arrangement. Just imagine a pin-striped shirt; while the stripes are parallel and straight on the cloth, they rarely are in images because the surface slants and curves.

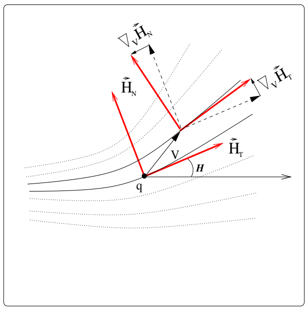

Figure 12: Any smooth hue field (depicted here with the set of locally parallel lines), can be represented as a differentiable frame field which is everywhere tangent (and normal) to the streamlines of the flow. An infinitesimal translation of the frame in a direction V rotates it by some angle determined by the connection form of the frame field. Since $\mathcal{H}_T, \mathcal{H}_N$ are unit length, their covariant derivative lies in a normal direction, regardless of V. Since the connection form is a linear operator, it is fully characterized by two numbers obtained by projection onto two independent directions. The natural choice to use the directions defined by the frame itself yields the two hue curvatures κ_T and κ_N. This diagram also suggests a relationship between hue fields and texture flows (Ben-Shahar and Zucker, 2003),

linear. This allows us to fully represent it with two scalars at each point since

$$w_{12}(V) = w_{12}(a\,\mathcal{H}_1 + b\,\mathcal{H}_2) = a\,w_{12}(\mathcal{H}_1) + b\,w_{12}(\mathcal{H}_2).$$

The freedom in selecting a basis $\{\mathcal{H}_1, \mathcal{H}_2\}$ for the representation of the tangent vectors V is naturally resolved by making, once again, the choice of $\mathcal{H}_1 = \mathcal{H}_T$ and $\mathcal{H}_2 = \mathcal{H}_N$. This yields the following two scalars:

$$\begin{aligned} \kappa_T &\triangleq w_{12}(\mathcal{H}_T) \\ \kappa_N &\triangleq w_{12}(\mathcal{H}_N) \quad . \end{aligned} \tag{2}$$

We call κ_T the hue's *tangential curvature* and κ_N the hue's *normal curvature* - they represent the rate of change of the hue in the tangential and normal directions, respectively. Since $\mathcal{H}_T$ and $\mathcal{H}_N$ are rigidly coupled, we can rewrite the two curvatures in terms of $\mathcal{H}_T$ only using the standard curl ($\nabla\times$) and divergence ($\nabla\cdot$) operators:

$$\begin{aligned} \kappa_T &= ||\nabla \times \mathcal{H}_T|| \\ \kappa_N &= \quad \nabla \cdot \mathcal{H}_T \qquad . \end{aligned} \tag{3}$$

However, more useful is the expression of the hue curvatures in terms of the hue ($\mathcal{H}$) itself and its gradient

($\nabla \mathcal{H}$, measured relative to a fixed coordinate system):

$$\begin{aligned} \kappa_T &= \nabla \mathcal{H} \cdot (\cos \mathcal{H}, \sin \mathcal{H}) \\ \kappa_N &= \nabla \mathcal{H} \cdot (-\sin \mathcal{H}, \cos \mathcal{H}). \end{aligned} \tag{4}$$

Viewed this way, it is clear that if κ_T and κ_N were known functions of position q, Eq. 4 could be viewed as a PDE and be solved for $\mathcal{H}(q)$. This of course raises the question of the degree to which κ_T and κ_N are independent, which indeed leads to the following observation (proofs omitted for space considerations; they follow with minor alterations from those in (Ben-Shahar and Zucker, 2003)):

Proposition 1 *Unless κ_T and κ_N both equal zero, they cannot be simultaneously constant in a neighborhood around q, however small, or else the induced hue function will be nonintegrable.*

This observation has an important implication: Unless the hue function is constant, at least one of its curvatures must vary, or the two curvatures need to covary in any neighborhood of the color image. Formally, this is characterized by the following constraints:

Proposition 2 *Given any hue field $\{\mathcal{H}_T, \mathcal{H}_N\}$, its curvature functions κ_T and κ_N must satisfy the relationship*

$$\nabla \kappa_T \cdot \mathcal{H}_N - \nabla \kappa_N \cdot \mathcal{H}_T = \kappa_T^2 + \kappa_N^2 \quad .$$

4. A model for hue coherence

Since the local behavior of the hue is characterized (up to Euclidean transformation) by a pair of curvatures, it is natural to conclude that nearby measurements of hue should relate to each other based on these curvatures. Put differently, measuring a particular curvature pair $(\kappa_T(q), \kappa_N(q))$ at a point q should induce a field of coherent measurements, i.e., a hue function $\tilde{\mathcal{H}}(x, y)$, in the neighborhood of q. Coherence of $\mathcal{H}(q)$ to its spatial context $\mathcal{H}(x, y)$ can then be determined by examining how well $\mathcal{H}(x, y)$ fits $\tilde{\mathcal{H}}(x, y)$ around q.

Clearly, the local coherence model $\tilde{\mathcal{H}}(x, y)$ should be a function of the local hue curvatures $(\kappa_T(q), \kappa_N(q))$, it should agree with these curvatures at q, and it should extend around q according to some variation in both curvatures (as a consequence of the propositions above). While many such models are possible, the fact that the hue field is a unit length vector field over the image plane implies that it takes the form of a texture

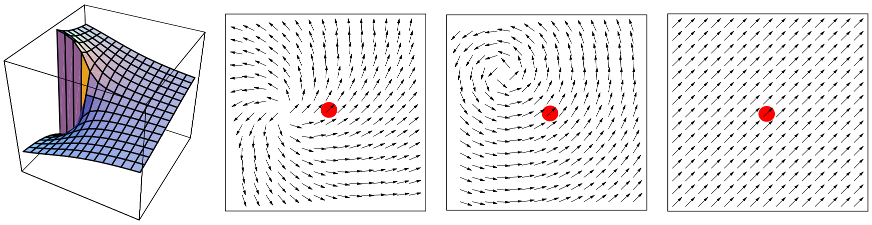

A B C D

Figure 13: A local model for coherent local behavior of the hue can be depicted both as (A) a height function that wraps itself to $[-\pi, \pi)$ (hence the apparent discontinuity), and as (B-D) a unit length vector field in the image plane. Different orientation and curvature tunings at the central pixel (marked with red) yield different local behaviors around it, and the three variations shown correspond qualitatively to **(A)** $\kappa_T > 0$ and $\kappa_N > 0$, **(B)** $\kappa_T > 0$ and $\kappa_N = 0$, **(C)** $\kappa_T = 0$ and $\kappa_N = 0$. All three fields are tuned to the same hue value of $\mathcal{H} = 45°$.

flow (Ben-Shahar and Zucker, 2003). Consequently, we adopt the same curvature-tuned local model developed recently for texture flows. Viewed as a surface in a three dimensional space whose Z axis represents the hue (as in Fig. 2C), this model takes the form of a right helicoid[3]:

$$\mathcal{H}(x,y) = tan^{-1}\left(\frac{\kappa_T(q)x + \kappa_N(q)y}{1 + \kappa_N(q)x - \kappa_T(q)y}\right) \quad . \tag{5}$$

This local model possesses many properties that suit good continuation, an in particular it is both a minimal surface in the $(x, y, \tilde{\mathcal{H}}(x,y))$ surface representation, and a critical point of the p-harmonic energy for *all* p. It is also the *only* local model that does not bias the changes in one hue curvature relative to the other, i.e., it satisfies

$$\frac{\kappa_T(x,y)}{\kappa_N(x,y)} = \text{const} = \frac{\kappa_T(q)}{\kappa_N(q)} \quad .$$

Examples of the model for different curvature tuning is illustrated in Fig 13. A detailed technical account on the development of the model in the texture flow domain can be found in (Ben-Shahar and Zucker, 2003).

[3] Unlike texture flows, however, the local model for the hue function is not a *double* helicoid since the hue function takes values in $[\pi, \pi)$ where texture flows are constrained to $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right)$. The basic properties and proofs (Ben-Shahar and Zucker, 2003) carry through, however.

5. A contextual approach to color denoising

The advantage of having a model for the local behavior of "good" hue flows lies in the ability to assess the degree to which a particular pixel is compatible, or consistent, with the context in which it is embedded. This, in turn, can be used to remove spurious measurements and replace them with consistent ones such that local ambiguity is reduced and global structures become coherent.

There are a few different frameworks in which one can pursue this task while maximizing some measure of global consistency or coherence over a domain of interest. Such frameworks include *relaxation labeling* (Hummel and Zucker, 1983; Kittler and Illingworh, 1985), *recurrent neural networks* (Hopfield and Tank, 1985), and *belief propagation networks* (Pearl, 1988). Here we present results using a relaxation labeling network, because it has been shown that such networks are (i) equivalent to polymatrix games (Miller and Zucker, 1992) and (ii) have a biophysically-plausible realization (Miller and Zucker, 1999). The nodes $i = (x, y)$ of our network are the image pixels, its labels at each node are drawn from the set $\Lambda = \{(\mathcal{H}, \kappa_T, \kappa_N) \mid \mathcal{H} \in [-\pi, \pi)\ ,\ \kappa_T, \kappa_N \in [-K, K]\}$ (after it has been quantized appropriately), and each label is assigned a confidence, or probability $p_i(\lambda)$ such that at each node $\sum_{\lambda\in\Lambda} p_i(\lambda) = 1$. The relaxation process itself drives an initial confidence distribution $p_i^0(\lambda)$ to a final (possibly ambiguous) distribution $p_i^\infty(\lambda)$. What governs the dynamics of this process, and ultimately its convergence state, are the compatibility relationships $r_{ij}(\lambda, \lambda')$ between different labels at different nodes. In our case, these compatibilities represent the degree to which two nearby pixels have consistent hue values. (It is these compatibilities that serve as a model for the long-range horizontal connections, as we discuss later.) Viewing the problem from a geometrical point of view we derive these compatibilities from the geometrical (helicoidal) model described above. Examples of a variety of compatibility fields for a single hue value and different combinations of hue curvatures are illustrated in Fig. 14. Note in particular the relationship between curvature tuning and the rotational organization in the hue domain.

With the network structure, labels, and compatibilities all designed, one can compute the support $s_i(\lambda)$ that label λ at node i gathers from its neighborhood. We define $s_i(\lambda)$ to be $s_i(\lambda) = \sum_j \sum_{\lambda'} r_{ij}(\lambda, \lambda') p_j(\lambda')$ and use it to update the confidence $p_i(\lambda)$ by gradient ascent followed by a confidences normalization Π

$$p_i^{t+1}(\lambda) \leftarrow \Pi\left[p_i^t(\lambda) + \delta s_i^t(\lambda)\right] \tag{6}$$

where δ is the gradient ascent step size. The relaxation labeling theory (Hummel and Zucker, 1983; Pelillo, 1997) ensures that such a rule will converge to a consistent labeling while extremizing the average local consistency over the entire image.

6. Experimental results

We applied the proposed model for hue good continuation and the corresponding relaxation labeling network to a number of the previous examples. We first consider the question of denoising images, because this is the easiest to demonstrate. We then show results from hue and shading flow relaxations, and conclude with implications to neurophysiology and psychophysics.

6.1 Denoising Images

Removing noise from the color channels of images is a challenging image processing problem, and in this section we illustrate our approach on a variety of synthetic and natural inputs. In all cases we quantized the hue uniformly to 32 equivalence classes and curvatures to 5 (as in Fig. 14). Step size was set to $\delta = 0.5$.

Fig. 15 illustrates the relaxation behavior around different kinds of synthetic color edges. Since our approach effectively considers only the coherent context (as defined through the geometrical model and the derived compatibilities) of each pixel, neither noisy pixels, nor information across edges, affect the support gathered by each label. This ensures not only the reliable elimination of noisy labels (or, as is necessarily true due to confidence normalization, their replacement with coherent ones), but also the robust preservation of edges. In this sense, the performance on the image in 15E-J is of particular interest because the input represents an edge configuration that a typical nonlinear diffusion is likely to distort. More specifically, note how the hue profiles along the two sides of the perceptual edge create a *cross*-like configuration in the hue domain (best seen in Fig. 15H). Since in the proximity of the cross point the hue gradient is very small the diffusion conduction increases and smoothing is encouraged. In practice this leads to the collapse of the edge from the inside out and to a distortion of the underlying structure. As is illustrated in the figure, this does not happen with our approach.

Fig. 16 illustrates the results of our denoising approach on the Apple image from Fig. 2. To emphasize the contribution of our method, which currently acts on the hue channel, here we show the result for a corruption along the hue dimension. In this, as well in all other images we tested, full reconstruction was achieved after 30 iterations or less. Unlike typical diffusion processes, which contain no natural indicators for stopping the progress along the scale space, in our approach convergence is readily signaled by convergence of the average local consistency (e.g., Fig. 15D). Fig. 17 demonstrates the stability of the process once convergence has been achieved, and how additional iterations do not distort essential structure.

Lastly, Fig. 18 demonstrates the result of using our approach on a variety of natural images. In all cases the noise was completely removed, and convergence achieved, after 30 iterations of relaxation labeling or less.

6.2 Fruit and Foliage Imagery

As described in Sec. 2.1.1, primates can utilize color image analysis for both finding fruits and edible leaves against a background of foliage. In both cases a complicated Gestalt is required.

We begin with the task of detecting berries in noisy images. Such noise might arise from reduced lighting, dirt, or other natural obscurations. Observe first that, since the berries can occupy only (spatially) small portions of the image, a linear smoothing process (to remove the noise) would tend to blur them away. Our relaxation system sharpens the hue differences from background; see Fig.19.

The second example illustrates the natural color variations indicative of edible leaves. Again, the relaxation process takes noisy local measurements into a piecewise smooth, coherent flow; Fig. 20. Although we might stress that this leaf is not edible, the chemistry underlying the pigmentation is consistent.

6.3 Hue and Shading Interactions

Our next example illustrates a few of the advantages of comparing hue flows to shading flows, concentrating on information relevant to inferring different aspects of scene structure. The concepts are illustrated in Fig. 21. We show both hue flows and shading flows[4], thereby illustrating hue and intensity differences and possible interactions.

First observe that in both cases the initial (local) measurements are noisy, and that the relaxation process maps them to a piecewise smooth result. For the hue flow note that it is most pronounced in the concavity between the ball and the tabletop and results from inter-reflection (Langer, 1999). It is elongated in the (projection onto the tabletop of the) direction of the source. While such hue bleeding is not necessarily salient perceptually, notice that it does capture some information about the scene configuration. Perhaps this is why artists have worked so diligently to represent it.

[4] Recall that the shading flow is the tangent field of the level curves for image intensity.

Second, notice how the apparent specularity on the ball has a roughly constant hue but a very different brightness.

Thirdly, observe that the shading flow captures rather different information from the hue flow. We displayed the hue flow superimposed on the intensity image so that these differences would be pronounced (top panels in Fig.21). Notice in particular how the hue flows across the shadow away from the object, indicating the constant ground plane. The lower panels represent the shading flow information geometrically, composed with local edge measurements (in green, bottom panels). The right shading flow corresponds to a fold-type boundary. Note that the (green) boundary segments have not been relaxed, to illustrate the feedback possible to boundary re-enforcement from shading or hue fields. However, much remains to be developed regarding hue-shading-boundary interactions.

6.4 Neurophysiological and Psychophysical Implications

A number of different types of predictions arise from experiments suggested by the model presented in this paper. We begin with neurophysiological considerations. Perhaps the most basic of these derive from interpreting the hue compatibility fields (Fig. 14) directly as projection fields derived from the long-range horizontal connections. Just as surround effects have elaborated our understanding of long-range connections in the interblob regions of V1 for orientation, one would expect analogous effects in color. While chromatic contextual effects are classical, they are normally expressed by presenting a constant (in hue) annulus surrounding a (possibly different) central patch (Wachtler et al., 2003). This is analogous to pop-out configurations in orientation, when a single misaligned line segment is presented in an otherwise constant field of lines. But smooth orientation flows differ from these pop-out configurations and, returning to hue, the compatibility fields above suggest using surround annuli of *smoothly varying hue*. Since the spatially-varying fields that we show are supportive of the central hue, we predict in particular that excitatory as well as inhibitory effects will be seen.

Our next class of predictions relate to visual psychophysics. It clearly suggests an examination of color (and hue, in particular) segmentation from a geometrical point of view, similar to recent findings for orientation-based texture segmentation (Ben-Shahar and Zucker, 2004b). Versions of these experiments can be formulated in hue to test for sensitivity to both hue and hue-curvature changes. These can be compared with existing information about color-based contour integration (Mullen et al., 2000; Beaudot and Mullen, 2003), the latter of which includes evidence for a curvature dependency.

Finally, higher-level effects should be manifest with varying hue as well. We show one example of a filling-in phenomenon in Fig. 22 but many related versions of this exist.

Examining all these directions are part of our future research in this area.

7. Summary and Conclusions

We have presented a theoretical framework for hue interactions within color imagery. Our approach is based on principles of perceptual organization, and in particular on the application of good continuation to the two dimensional field of the hue channel represented as an orientation (or position within the ”color circle”) for each retinotopic position. Based on a notion of hue curvature, we derived a formal model for the local behavior of coherent hue. This extends the most studied class of color patterns from those in which the color is essentially constant and changes only at boundaries to those that allow smooth variations. We illustrate how natural images often enjoy smooth hue variations, and demonstrate, for a number of different examples, how this allows for rich hue-field descriptions and robust processing while preserving the underlying flow structures in the color image. The need to preserve a variety of singularities separates our approach from (related) orientation diffusion models.

The neurobiological substrate to support these computations is taken to be the long-range horizontal connections. One motivation for our work was that such connections exist between cells in the cytochrome oxidase blobs as they exist between cells in the interblobs, and we are seeking to develop theories and applications of color good continuation as rich as those for orientation good continuation. Our earlier theory for interblob orientation integration provided the mathematical foundations for this hue model.

The hue dimension is a natural candidate for application of the principle of good continuation, and different applications in finding hue boundaries, color bleeding, and denoising extend its use beyond straightforward image segmentation. The resultant hue boundaries could support rich lightnesss computations, and Interactions between the hue field, the shading flow, and boundaries provides information about shape and light source placement.

Our model for hue good continuation illustrates the contribution of *top-down* computational research: it creates theoretical connections between information processing tasks and the neural substrates that could carry them out. But the argument is only one of sufficiency, and the computational exploration of hue and its many uses, from the detection of nutrient foliage by primates to psychophysical demonstrations of hue-variant filling-in, is just beginning. Since the anatomical and physiological data for blob connections are not as complete as those for interblob connections, models such as the one presented here remain largely theoretical. While the geometry that we describe was analogous to V1 structures, it may well be realized in V2, or in V1/V2 interactions, or beyond. Thus it stands as a class of predictions for connections to be confirmed. And, we hope, it provides additional motivation to find them.

Acknowledgments

We thank John Allman, Nate Dominy and Alison Richard for discussions about primate ecology, Leo Hickey for plant biology, and Pamela Davis for the permission to use the Sidney image (third row in Fig. 18). Research supported by ONR, AFOSR, NIH, and DARPA.

References

Albers, J. (1987). *Interaction of Color*. Yale University Press.

Allen, G. (1879). *The colour-sense: its origin and development*. Trubner, London.

Allman, J. and Zucker, S. (1990). Cytochrome oxidase and functional coding in primate striate cortex: An hypothesis. *Cold Spring Harbor Symp. Quant. Biology*, 55:979–982.

Astola, J., Haavisto, P., and Nuevo, Y. (1990). Vector median filters. *Proc. IEEE*, 78(4):678–689.

Barnard, J., Finlayson, G., and Funt, B. (1997). Color constancy for scenes with varying illumination. *CVIU*, 65(2):311–321.

Baxandall, M. (1995). *Shadows and the Enlightenment*. Yale University Press.

Beaudot, W. and Mullen, K. (2003). How long range is contour integration in human color vision. *Visual Neurosci.*, 20:51–64.

Beck, J. (1972). *Surface Color Perception*. Cornell University Press.

Bell, A. and Sejnowski, T. (1997). The "independent components" of natural scenes are edge filters. *Vision Res.*, 37:3327–2228.

Ben-Shahar, O., Huggins, P., and Zucker, S. (2002). On computing visual flows with boundaries: The case of shading and edges. In *Workshop on Biologically Motivated Computer Vision*.

Ben-Shahar, O. and Zucker, S. (2003). The perceptual organization of texture flows: A contextual inference approach. *IEEE Trans. Pattern Anal. Machine Intell.*, 25(4):401–417.

Ben-Shahar, O. and Zucker, S. (2004a). Geometrical computations explain projection patterns of long range horizontal connections in visual cortex,. *Neural Comput.*, 16(3):445–476.

Ben-Shahar, O. and Zucker, S. (2004b). Sensitivity to curvatures in orientation-based texture segmentation,. *Vision Res.*, 44(3):257–277.

Blake, A. (1985). Boundary conditions for lightness computation in mondrian world. *CVGIP*, 32(3):314–327.

Breton, P. and Zucker, S. (1996). Shadows and shading flow fields. In *Proc. Computer Vision and Pattern Recognition*, pages 782–789.

De Valois, R. and De Valois, K. (1990). *Spatial Vision*. Oxford University Press.

Dominy, N. and Lucas, P. (2001). Ecological importance of trichromatic vision to primates. *Nature*, 410:363–366.

Ekman, P. (2001). *Telling lies: clues to deceit in the marketplace, politics and marriage*. W.W.Norton & Company.

Field, D., Hayes, A., and Hess, R. (1993). Contour integration by the human visual system: Evidence for a local "association field". *Vision Res.*, 33(2):173–193.

Fine, I. MacLeod, D. and Boynton, G. (2003). Surface segmentation based on the luminance and color statistics of natural scenes. *J. Opt. Soc. Am. A*, 20(7):1283–1291.

Forsyth, D. and Ponce, J. (2002). *Computer Vision: A Modern Approach*. Prentice Hall.

Funt, B. and Barnard, K. (1998). Is colour constancy good enough. In *ECCV*, pages 445–459.

Gegenfurtner, K. and Rieger, J. (2000). Sensory and cognitive contribution of color to the recognition of natural scenes. *Curr. Biol.*, 10(13):805–808.

Hanazawa, A., Komatsu, H., and Murakami, I. (2000). Neural selectivity for hue and saturation of colour in the primary visual cortex of the monkey. *Eur. J. Neurosci.*, 12:1753–1763.

Healey, G. (1992). Segmenting images using normalized color. *IEEE Trans. Sys. Man Cyber.*, 22:62–73.

Hering, E. (1878/1964). *Outline of the theory of the light sense*. Harvard University Press. (translated by L. M. Hurvich and D. Jameson).

Heywood, C., Gadoti, A., and Cowey, A. (1992). Cortical area v4 and its role in the perception of color. *J. Neurosci.*, 12(10):4056–4065.

Hopfield, J. and Tank, D. (1985). Neural computation of decisions in optimization problems. *Biological Cybernetics*, 52:141–152.

Horn, B., editor (1986). *Robot vision*. MIT Press, Cambridge, MA.

Huggins, P., Chen, H., Belhumeur, P., and Zucker, S. (2001). Finding folds: On the appearance and identification of occlusion. In *Proc. Computer Vision and Pattern Recognition*, pages 718–725.

Huggins, P. and Zucker, S. (2001). How folds cut a scene. In *Proc. of the* 4^{th} *Int. Workshop on Visual Form*.

Hummel, R. and Zucker, S. (1983). On the foundations of the relaxation labeling proceeses. *IEEE Trans. Pattern Anal. Machine Intell.*, 5:267–287.

Kimmel, R., Malladi, R., and Sochen, N. (2000). Images as embedded maps and minimal surfaces: Movies, color, texture, and volumetric medical images. *Int. J. Comput. Vision*, 39(2):111–129.

Kimmel, R. and Sochen, N. (2000). Orientation diffusion or how to comb a porcupine? CIS 2000-02, Technion - Israel Istitute of Technology.

Kingdom, F. (2003). Color brings relief to human vision. *Nature Neuroscience*, 6(6):641–644.

Kiper, D., Fenstemaker, S., and Gegenfurtner, K. (1997). Chromatic properties of neurons in macaque area v2. *Visual Neurosci.*, 14:1061–1072.

Kittler, J. and Illingworh, J. (1985). Relaxation labeling algorithms - a review. *Image and Vision Computing*, pages 206–216.

Komatsu, H. (1998). Mechanisms of central color vision. *Curr. Opin. Neurobiol.*, 8:503–508.

Land, E. (1977). The retinex theory of color vision. *Sci. Am.*, 37:108–129.

Land, E. and McCann, J. (1971). Lightness and retinex theory. *American Journal of Optical Society of America*, 61:1–11.

Landy, M. and Bergen, J. (1991). Texture segregation and orientation gradient. *Vision Res.*, 31(4):679–691.

Langer, M. (1999). When shadows become interreflections. *Int. J. Comput. Vision*, 34(2/3):193–204.

Livingstone, M. and Hubel, D. (1984a). Anatomy and physiology of a color system in the primate visual cortex. *J. Neurosci.*, 4(1):309–356.

Livingstone, M. and Hubel, D. (1984b). Specificity of intrinsic connections in primate primary visual cortex. *J. Neurosci.*, 4(11):2830–2835.

Lucas, P., Dominy, N., Riba-Hernandez, P., Stoner, K., Yamashita, N. Loria-Calderon, E., Petersen-Pereira, W., Rojas-Doran, Y., Salas-Pena, R., Solis-Madrigal, S., Osorio, D., and Darvell, B. (2003). Evolution and function of routine trichromatic vision in primates. *Evolution*, 57(11):2636–2643.

Michael, C. (1981). Columnar organization of color cells in monkey's striate cortex. *J. Neurophysiol.*, 46:587–604.

Miller, D. and Zucker, S. (1992). Efficient simplex-like methods for equilibria of nonsymmetric analog networks. *Neural Computations*, 2:167–190.

Miller, D. and Zucker, S. (1999). Computing with self-excitatory cliques: A model and and application to hyperacuity-sclae computation in visual cortex. *Neural Comput.*, 11:21–66.

Mullen, K., Beaudot, W., and McIlhagga, W. (2000). Contour integration in color vision: a common process for the blue-yellow, red-green and luminance mechanisms? *Vision Res.*, 40:639–655.

Munsell, A. (1905). *A Color Notation*. G.H.Ellis, Boston.

Mussap, A. and Levi, D. (1999). Orientation-based texture segmentation in strabismic amblyopia. *Vision Res.*, 39:411–418.

Nothdurft, H. (1991). Texture segmentation and pop-out from orientation contrast. *Vision Res.*, 31(6):1073–1078.

Novak, C. and S., S. (1987). Color edge detection. In *IUW*, volume 1, pages 35–37.

Ohta, Y.I. Kanade, T. and Sakai, T. (1980). Color information for region segmentation. *CGIP*, 13:222–241.

O'Neill, B. (1966). *Elementary Differential Geometry*. Academic Press.

Palmer, S. (1999). *Vision Science: Photons to Phenomenology*. The MIT Press.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann.

Pelillo, M. (1997). The dynamics of nonlinear relaxation labeling processes. *J. of Mathematical Imaging and Vision*, 7:309–323.

Perona, P. and Malik, J. (1990). Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Machine Intell.*, 12(7):629–639.

Redies, C. and Spillmann, L. (1981). The neon color effect in the ehrenstein illusion. *Perception*, 10:667–681.

Roe, A. and Ts'o, D. (1999). Specificity of color connectivity between primate v1 and v2. *J. Neurophysiol.*, 82:2719–2730.

Ruzon, M. and Tomasi, C. (1999). Color edge detection with the compass operator. In *Proc. Computer Vision and Pattern Recognition*, pages 160–166.

Sapiro, G. and Ringach, D. (1996). Anisotropic diffusion of multivalued images with applications to color filtering. *IEEE Trans. Image Processing*, 5(11):1582–1586.

Shapley, R. and Hawken, M. (2002). Neural mechanisms for color perception in the primary visual cortex. *Curr. Opin. Neurobiol.*, 12:426–432.

Simoncelli, E. and Olshausen, B. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24:1193–1216.

Sochen, N., Kimmel, R., and Malladi, R. (1998). A geometrical framework for low level vision. *IEEE Trans. Image Processing*, 7(3):310–318.

Spillmann, L. and Werner, J., editors (1990). *Visual Perception: The Neurophysiological Foundations*. Academic Press, Inc.

Sumner, P. and Mollon, J. (2000a). Catarrhine photopigments are optimized for detecting targets against a foliage background. *J. Exp. Biol.*, 203(13):1963–1986.

Sumner, P. and Mollon, J. (2000b). Chromaticity as a signal of ripeness in fruits taken by primates. *J. Exp. Biol.*, 203(13):1987–2000.

Tailor, D.R. Finkel, L. and Buchsbaum, G. (2000). Color-opponent receptive fields derived from independent component analysis of natural images. *Vision Res.*, 40(19):2671–2676.

Tang, B., Sapiro, G., and Caselles, V. (2000). Diffusion of general data on non-flat manifolds via harmonic maps theory: The direction diffusion case. *Int. J. Comput. Vision*, 36(2):149–161.

Tang, B., Sapiro, G., and Caselles, V. (2001). Color image enhancement via chromaticity diffusion. *IEEE Trans. Image Processing*, 10(5):701–707.

ter Haar Romeny, B., editor (1994). *Geometry-Driven Diffusion in Computer Vision*. Kluwer Academic Publishers.

Trahanias, P. and Venetsanopoulos, V. (1993). Vector directional filters - a new class of multichannel image processing filters. *IEEE Trans. Image Processing*, 2(4):528–534.

Ts'o, D. and Gilbert, C. (1988). The organization of chromatic and spatial interactions in the primate striate cortex. *J. Neurosci.*, 8(5):1712–1727.

T'So, D., Gilbert, C., and Wiesel, T. (1986). Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.*, 6:1160 – 1170.

Wachtler, T., Lee, T., and Sejnowski, T. (2001). Chromatic structure of natural scenes. *J. Opt. Soc. Am. A*, 18(1):65–77.

Wachtler, T., Sejnowski, T., and Albright, T. (2003). Representation of color stimuli in awake macaque primary visual cortex. *Neuron*, 37:681–691.

Weickert, J. (1997). A review of nonlinear diffusion filtering. In Romeny, B. t. H., Florack, L., Koenderink, J., and Viergever, M., editors, *Scale-Space Theory in Computer Vision*, volume 1252 of *Lecture Notes in Computer Science*, pages 3–28. Springer-Verlag.

Xiao, Y., Wang, Y., and Felleman, D. (2003). A spatially organized representation a colour in macaque cortical area v2. *Nature*, 421:535–539.

Yezzi, A. (1998). Modified curvature motion for image smoothing and enhancement. *IEEE Trans. Image Processing*, 7(3):345–352.

Zeki, S. (1993). *A Vision of the Brain*. Blackwell Scientific.

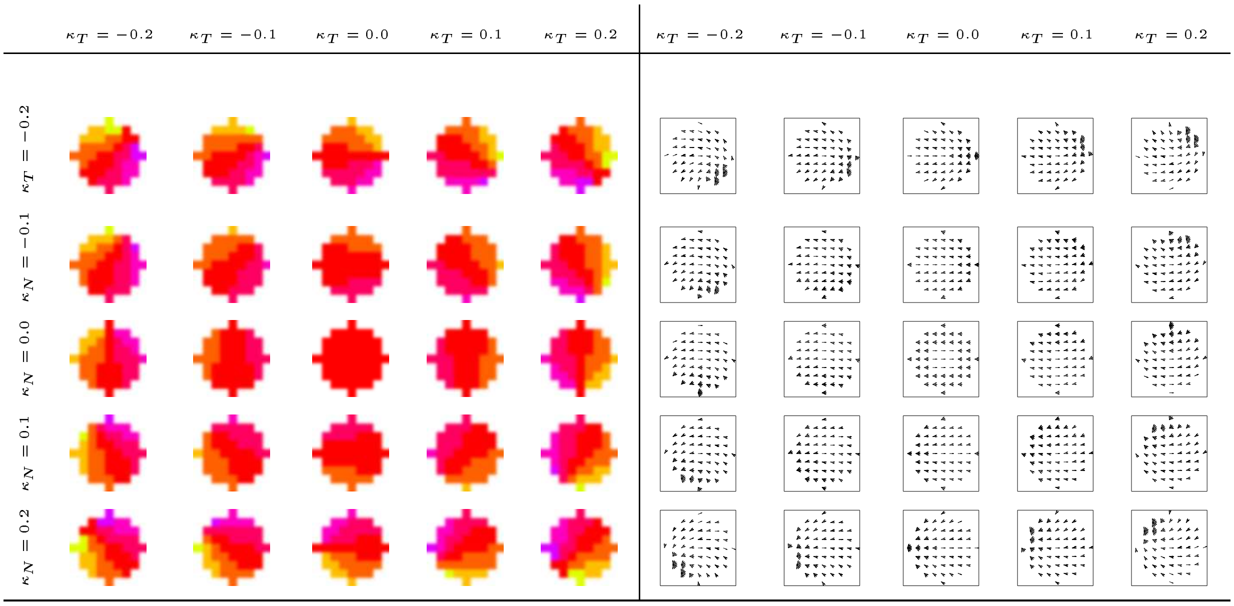

Figure 14: A collection of 9×9 hue compatibility fields for $\mathcal{H} = 0°$ (the red color), where both κ_T and κ_N are selected from the set $\{-0.2, -0.1, 0.0, 0.1, 0.2\}$. Each field represents all compatible hue values in the neighborhood of the central label; on the left these compatible values are depicted as color pixels while on the right they are depicted as hues fields. Due to quantization, especially that of curvatures, a given label at the center may be compatible with more than one label at the same nearby location in its neighborhood, an outcome depicted by multiplicity of vectors at certain positions in the fields. Since this aspect of the compatibilities cannot be depicted with the color representation, the fields on the left show only the most compatible hue value at each position. Note how higher curvature values introduce more variations into the fields, and how changing the curvature tuning while keeping constant the "total variation" measure of $\kappa_T^2 + \kappa_N^2$ amounts to a rotational transformation of the field. It is these hue compatibility fields that serve as an abstract model for hue-based long-range horizontal connections.

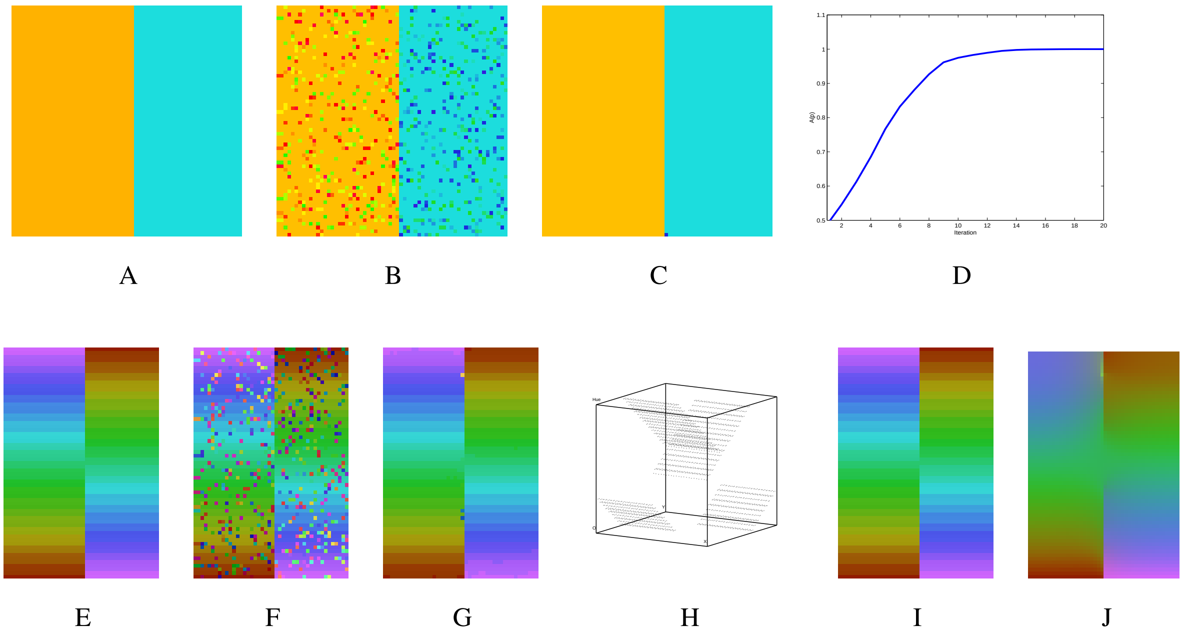

Figure 15: Color denoising on synthetic color images. **(A)** is a color step edge. **(B)** is a noisy version of it. **(C)** shows the result of 20 iterations of relaxation labeling. The noise is completely eliminated while the edge structure is preserved. **(D)** is the graph of the average local consistency as a function of iteration. it illustrates that convergence on this input was effectively achieved at the 12^{th} iteration, at which the distribution of assignments $p_i(\lambda)$ reached a steady state. **(E-G)** illustrates the denoising result on a different configuration of color edge. This time steady state (panel G) was achieved after 30 iterations. This example is particularly important because most anisotropic diffusion schemes are likely to fail on its cross-like edge configuration. This cross configuration is most apparent in panel H, where the hue values are represented as height (and the top face of the cube should be identified with its bottom). Since around the cross the hue gradient is very small, the diffusion is not suppressed and the edge collapses. Thus even without noise (shown again in panel I), diffusion of such a configuration yields an undesired distortion of the structure, as shown in panel J (compare to panel G with the result of our approach). Here we used the Beltrami flow for RGB images (Sochen et al., 1998; Kimmel et al., 2000), but similar results were obtained with other orientation diffusion schemes as well (e.g., (Tang et al., 2000; Sapiro and Ringach, 1996)).

A B C

Figure 16: Color denoising of the Apple image (from Fig. 2) **(A)** Original image. **(B)** Noisy image. **(C)** Result of 25 denoising iterations. Compare to Fig. 6C.

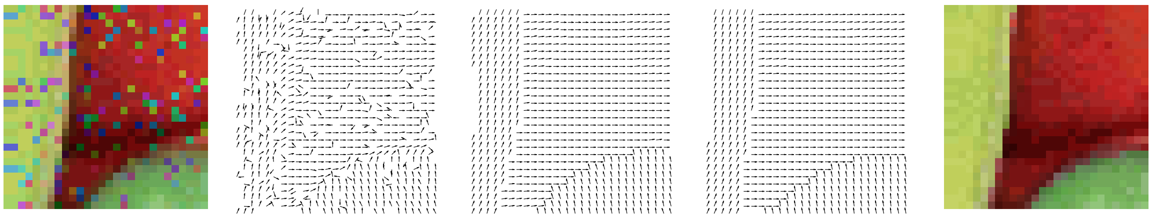

A B C D E

Figure 17: Stability of hue field after convergence of the local average consistency. **(A)** A detail from noisy Peppers image from the same ROI marked in Fig. 2A (but further zoomed in to allow better depiction of the hue field). **(B)** The corresponding hue field. **(C)** The result of 20 iterations. **(D)** The result of 50 iterations is virtually identical. **(E)** The denoised color image.

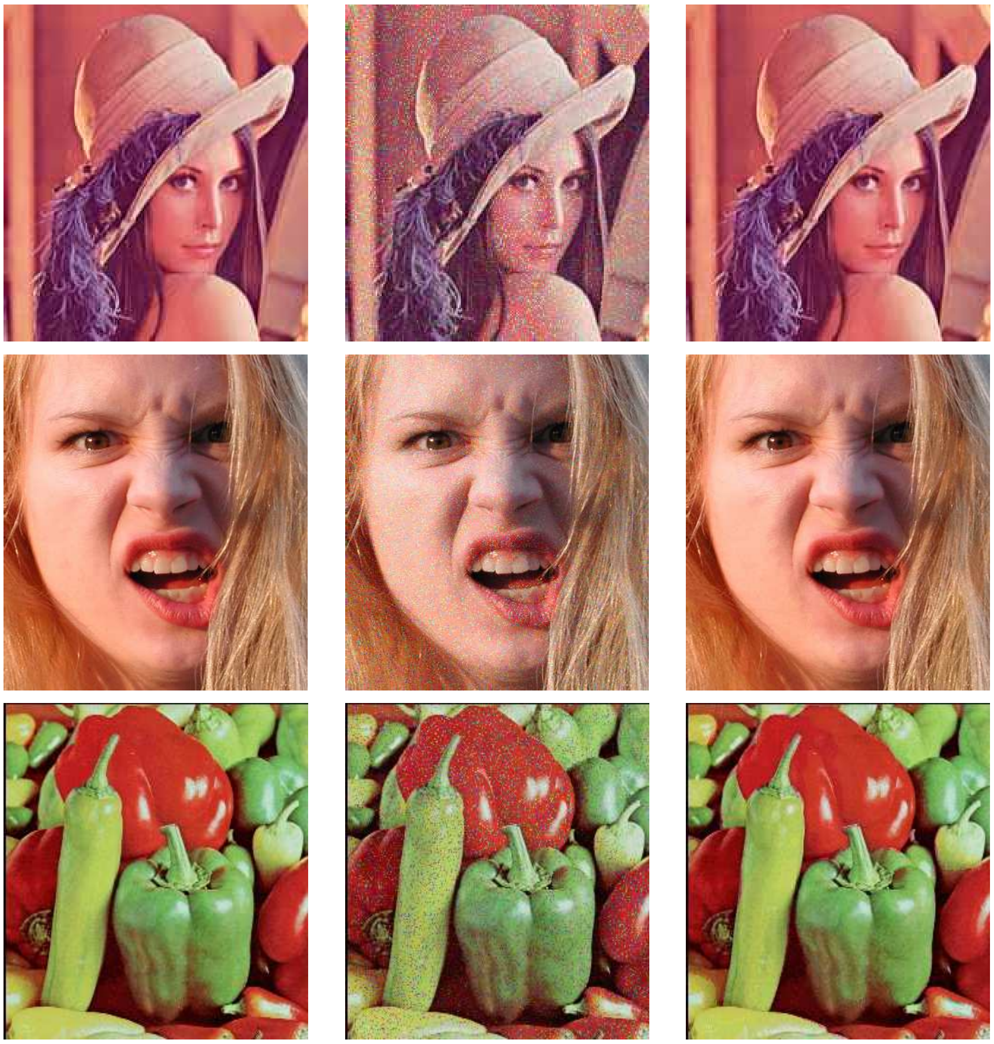

A B C

Figure 18: Color denoising of the Lena, Sidney, and Peppers images. **(A)** Original image. **(B)** Noisy image. **(C)** Convergence state of the relaxation process.

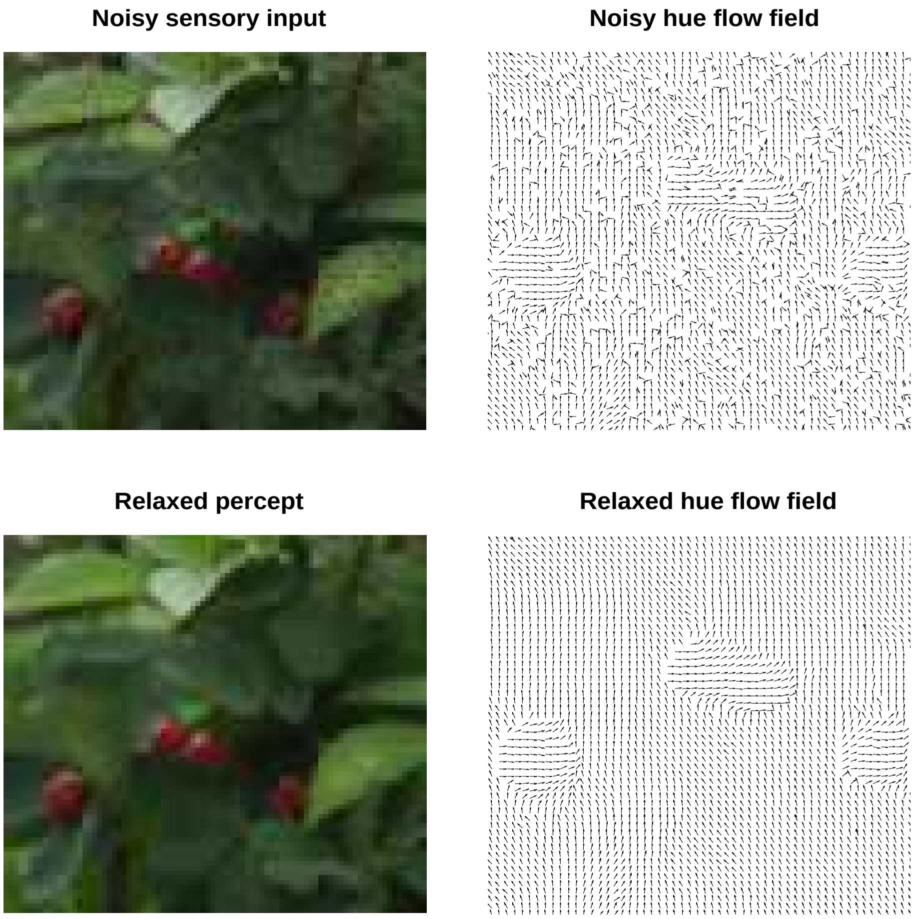

Figure 19: The detection of (noisy) berries from a (noisy) foliage background. Notice how the additive noise corrupts the local hue field measurements (**top**), as might occur in a low-light situation, and how the relaxation process (**bottom**) reduces the noise to make the detection essentially trivial.

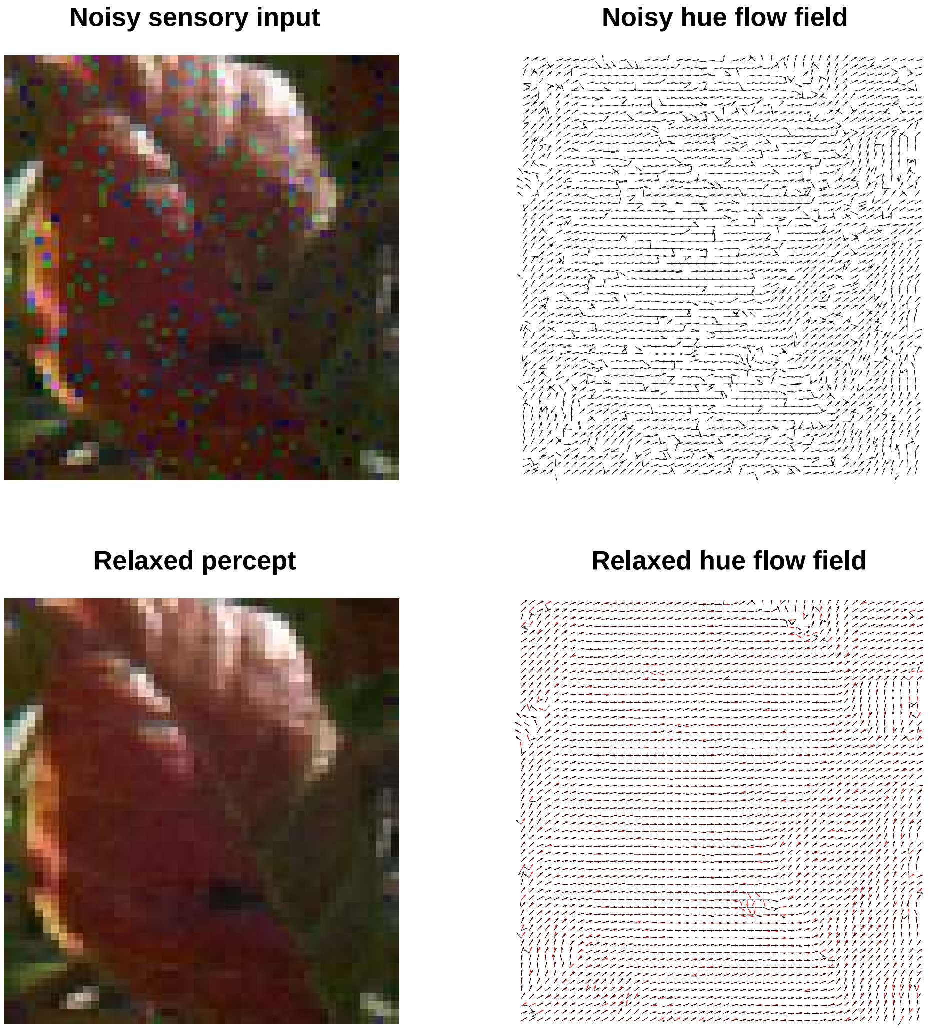

Figure 20: The relaxation process achieves a smooth coherent flow from noisy incoherent measurements of a leaf. Compare relaxed hue flow to the original one in Fig. 5 (also shown in red beneath the black arrows in the bottom right panel).

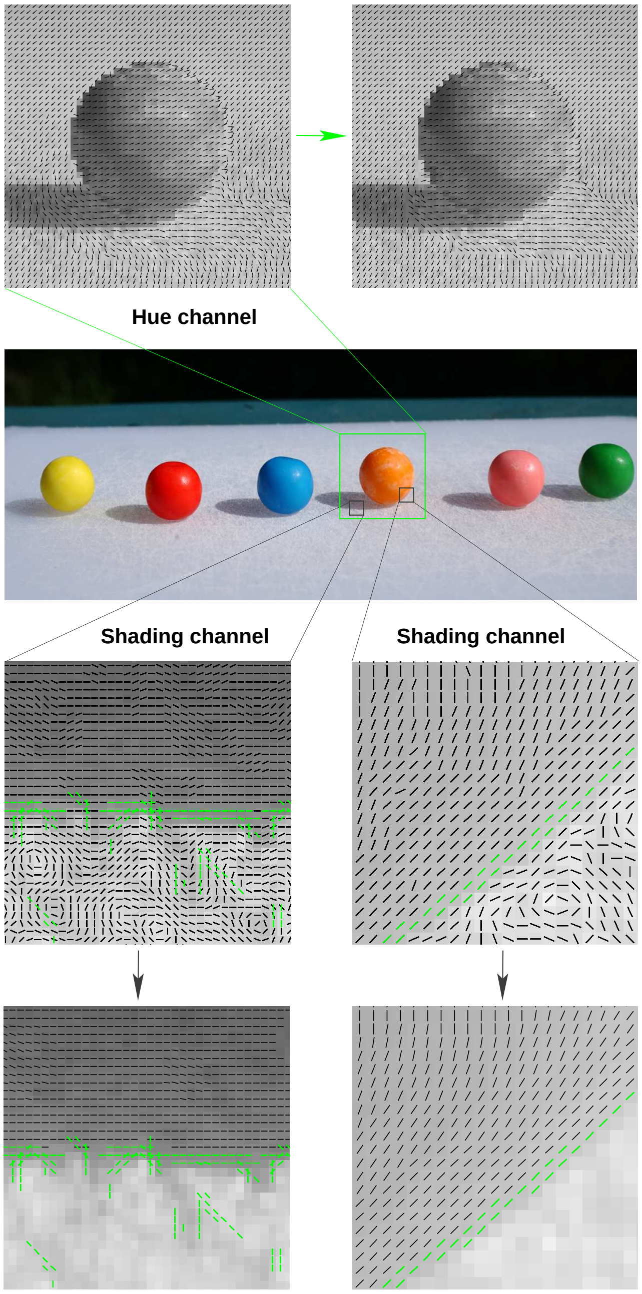

Figure 21: Illustration of hue and shading flows; see text for description.

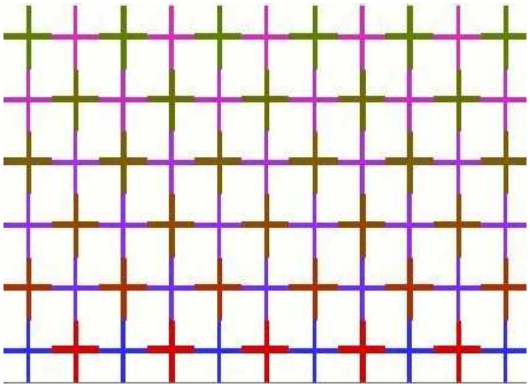

Figure 22: A modified Ehrenstein illusion shows how color filling in can occur between boundaries of different colors, thus suggesting a link between the approach proposed in this paper and the psychology of color perception. This illusion best viewed on a fluorescent (e.g., CRT) display.