Good Continuation in Layers: Shading flows, color flows, surfaces and shadows

Ohad Ben-Shahar
Computer Science
Ben Gurion University
Israel
ben-shahar@cs.bgu.ac.il

Andreas Glaser
Applied Mathematics
Yale University
New Haven, CT 06520-8285
andreas.glaser@yale.edu

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

Abstract

We extend the concept of good continuation in a uniform fashion from boundaries to shading, hue, and texture. Each has the property that local measurements yield an orientation, which we explicitly establish for hue using geometric harmonic techniques. Good continuation arises in a geometric sense, because these orientations all vary smoothly in an appropriate sense. Thus they correspond to flows. Taken together they define a layered set of flows, in the sense the "horizontal" computations within each flow provide global consistency while "vertical" computations across flows enable the identification of shading and shadowing and different types of edges. Evidence is reviewed that primate visual systems enjoy such an organization.[1]

> "...space and color are not distinct elements but, rather, are interdependent aspects of a unitary process of perceptual organization." Kanizsa [17]

1. Introduction

Image segmentation is normally taken to be that process of partitioning the image into a complete cover of non-overlapping regions, with the boundaries of these regions related to the (projected) boundaries of objects in the world. One source of complexity in this process is shadowing, by which image intensities vary both as a function of surface orientation (e.g., shading) and as a function of light sources (e.g., cast shadows). Land's *retinex theory* [19] suggested one way to manage this complexity, by ascribing abrupt image changes to material (or reflectance) discontinuities and smooth gradient changes to lighting. This developed into the intrinsic image concept [30], which emphasized that surface properties, geometry, and lighting all map into the image, and suggested representing them separately as images. Undoing this map clearly involves an inverse problem, which requires a model of some sort. One possibility is to try to learn the context of every possible measurement, a type of pseudoinverse [28]. Here we extend the notion of context in a different way, by considering natural images such as those in Fig. 1. Notice how space,reflectance, and lighting conspire together. We seek to find a representation rich enough to support unwinding this.

The first requirement for such a representation is that it be rich enough to capture the above phenomena. But unlike special purpose algorithms applicable in one situation (e.g., [16, 13]), our second requirement is that it be general purpose. That is, the information that it makes explicit must support computations for unraveling many such phenomena.

We do not yet have a formal solution to this problem that we can prove is complete. Instead, and consistent with the goals of this Workshop, we develop an argument based on a neurobiological analogy, several steps of which have been formalized and are complete. The demonstrations in the final section of this paper involve phenomena beyond the current capability of any single existing algorithm, and provide counterexamples to many. Constructively, however, we submit that any final solution will have an intermediate representation at least as rich as the one we describe. Thus we see the contribution of this Workshop submission as consisting of (i) an enlargement of the framework for perceptual organization informed by (ii) the rich foundation for perceptual organization in primate visual systems.

The core of our argument is that good continuation applies to several key domains: boundaries, intensity (shading); hue; texture; saturation, and so on, all of which enjoy a certain differential geometric structure. It is this structure that relates to the Gestalt notion of *good continuation*. Computationally we propose a layered representation—similar in spirit to intrinsic images [30]—but different in

[1] Acknowledgements: Research supported by AFOSR, DARPA, ONR and the Toman and Frankel Funds from Ben-Gurion University.

0-7695-2646-2/06 $20.00 (c) 2006 IEEE

Figure 1. The rich interaction between surfaces, lighting, pigmentation, and atmosphere work together to provide a diversity of appearance phenomena in natural images. To simply claim that "apples are red" or "bananas are yellow" or "the sky is blue" amounts to an assumption that physical processes in the world are constant in a way that only artifi cial examples can really achieve.

that all share the property that they are flows in a technical sense. This is what we meant by layered flows implied in the title, and computations across these flows then reflect subtle lighting, surface, and space interactions.

Fig. 1 illustrates this point in several different domains (see also [3]. Apples are not a single color; rather, fruits mature differentially and this is reflected in their pigmentation. Attempts to remove these slow variations as lighting are one reason why lightness and color constancy algorithms have problems. Atmospheric depth effects impose a blue tint with distance because of increased scattering and in spite of surface reflection effects. Mutual illumination and color bleeding mix everything.

We approach the *lift* of these images into layered flows in two stages, both of which are mathematical but motivated by biology. We concentrate on one flow (from the color pathway) because, as will become clear below, the others fit naturally into our framework and are more widely discussed in the literature. Specifically, we first consider the question of how to represent color information as a dimensionality-reduction problem, which leads formally to intensity-hue-saturation coordinates at each point. This is important for us, because it suggests that there is more to color processing than simple detection tasks (consider: locate a red fruit among green foliage [27]) for which the standard cone pigments are tuned. We next consider (hue) interactions between points and adopt a technique previously used to denoise color patterns to articulate the flow of hue across image coordinates. The resultant computations are then run on the examples in Fig. 1.

2. Representation of Color at a Point

Take as data the Munsell patches considered as points in wavelength space. While wavelength-space is rather high-dimensional, our strategy is motivated by the observation that colors are not randomly distributed thoughout wavelength space, but rather occupy only a small portion of it. One possibility, suggested by the visual photopigments in primates, is that this structured space of colors is 3-dimensional. While this is a classical view of color, many of the classical algorithms have been modified in an *ad hoc* fashion to take account of non-linearities among colors (e.g., Multi-Dimensional Scaling). For this reason we use a new algorithm ([10, 11]) derived from the *geometric harmonics* (reviewed below) that can handle inherently non-linear data. It is in the class of spectral methods, and is related to [4].

2.1. Geometric Harmonics

Let $X = \{x_1, x_2, ..., x_N\}$ be the set of data points, in this case Munsell patches, with each $x_i \in R^n$. We seek to find a projection of these data into much lower dimension, under the assumption that they are not randomly distributed thoughout R^n but rather that they lie on (or near) a lower-dimensional manifold embedded in R^n.

The structure of the data are revealed via a symmetric, positivity-preserving, and positive semi-definite *kernel* $k(x, y)$, which provides a measure of similarity between data points. The result is a graph, with edges between nearby (according to the similarity kernel) data points. (The similarity value can be truncated to 0 for all but very similiar points.)

From this we construct a diffusion kernel $a(x, y)$ on the data set using the weighted graph Laplacian normalized as follows:

$$a(x, y) = \frac{k(x, y)}{\nu(x)}, \tag{1}$$

where $\nu = \sum_{y \in X} k(x, y)$. Note that, although symmetry is lost, we do have $\sum_{y \in X} a(x, y) = 1$ so the kernel $a(x, y)$ can be interpreted as the transition matrix of a Markov chain on the data X. The kernel $a^{(m)}$ of the m^{th} power of this matrix then represents the probability of getting from x to y in m steps.

If we now define the averaging operator for a function f defined on the data:

$$Af(x) = \sum_{y \in X} a(x, y) f(y) \tag{2}$$

then A admits a spectral theory. To develop this we symmetrize a by:

$$\tilde{a}(x, y) = \frac{\sqrt{\nu(x)}}{\sqrt{\nu(y)}} a(x, y) \tag{3}$$

which makes $\tilde{a}$ symmetric and positive semi-definite (although no longer row-stochastic). The spectral decomposition is then given by $\tilde{a} = \sum_{i\geq 0} \lambda_i^2 \phi_j(x)\phi_j(y)$ with the important consequence

$$\tilde{a^{(m)}}(x,y) = \sum_{i\geq 0} \lambda_i^{2m} \phi_j(x)\phi_j(y) \tag{4}$$

where $\lambda_0 = 1$.

Increasing powers of the operator A can be obtained by running the chain through the spectral decomposition. This gives rise to the family of *diffusion maps* $\{\Phi_m\}_{m\in N}$ given by

$$\Phi_m(x) = \begin{pmatrix} \lambda_0^m \phi_0(x) \\ \lambda_1^{2m} \phi_1(x) \\ \vdots \end{pmatrix} \tag{5}$$

Diffusion distances $D_m^2(x,y) = \tilde{a}^{(m)}(x,x) + \tilde{a}^{(m)}(y,y) - 2\tilde{a}^{(m)}(x,y)$ within the high-dimensional measurement space then approximate Euclidean distance in the diffusion map space.

2.2. The Munsell Color Space

The Munsell [22] patches were chosen according to human psychophysics, with each step between patches perceptually equal, and they are now known to be physiologically relevant [31, 29, 15]. Thus they represent data spanning those portions of color space relevant to our interactions with the visible world. We now seek to understand whether these data lie on or near a well-defined structure in wavelength-space.

Two experiments were performed. We used $N = 1269$ patches, each with $n = 421$ wavelengths (380nm - 800nm in 1nm steps). The kernel is $\exp(-d_{ij}^2/\sigma)$ where d_{ij} is the Euclidian distance between patch i and patch j. While the patch data are given in no particular order, the geometric harmonic map arranges them so that patches are close to one another provided the diffusion distance between them in wavelength space is small. The results are shown in Fig. 2. Note that the natural representation emerges—intensity, hue, saturation—even though the hue (color circle) is non-linear. The diffusion maps recover the Munsell representation, thus demonstrating that the structure is in the wavelength data. In the second experiment we first projected the wavelength data through the human cone photopigments; and again the color circle emerged (Fig. 2, bottom).

3. Spatio-spectral Interactions

Now that we know there is a preferred representation for color at a point, we next consider the question of how colors interact between nearby points. We first observe that the primate visual system is well organized to address this problem. While it is widely accepted that perceptual organization is first accomplished via the long-range horizontal connections in superficial V1, consideration of these connections has been limited to orientation good continuation for boundaries ([24, 1, 2]) and textures ([7]). However, there exists a specialized structure for color (and contrast) information in the cytochrome oxidase blobs, within which neurons also enjoy long-range horizontal interactions (Fig. 3 [32]). We submit that it is precisely these connections that implement a geometry for hue (and color) that is formally analagous to that for texture[7] and shading [9, 21] flows. A sketch of this geometry is developed next. The extension to include boundaries is in [5].

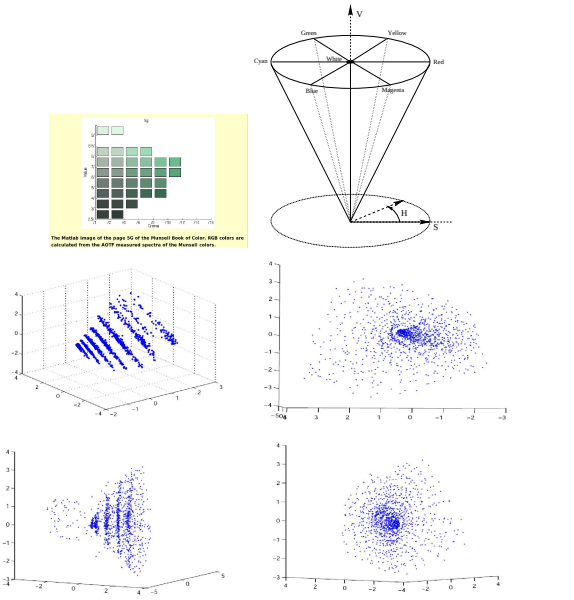

Figure 2. Geometric harmonics organize Munsell color patches. (**top row, left**) Typical "page" of the patch data used in the experiment. Data from `http://spectral.joensuu.fi-/databases/download/munsell_spec_matt.htm`. (**right**) Classical intensity, hue, saturation color space. Note that hue is organized around the circle. (**middle row**) The geometric harmonic organization of the Munsell data. Each point represents a single patch, and the scatterplots show the distribution of points in the subspace spanned by the first three non-trivial eigenfunctions. Two views are shown, with (**left**) illustrating different clusters according to the Munsell chromaticity parameters and (**right**) a view showing the hue circle. That this non-linear organization of the data is recovered by geometric harmonics is significant because it provides the foundation for the next, geometric stage of processing. (**bottom row**) Organization of the Munsell data first projected through the three human cone photopigments. Since the two views are essentially the same as (**middle**), the Munsell representation is largely invariant to the order of projection.

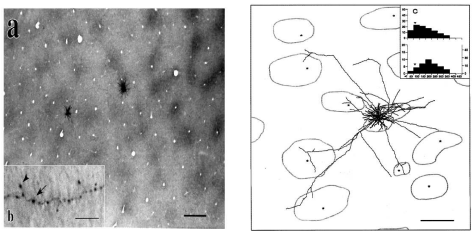

Figure 3. The cytochrome oxidase blobs in superfi cial primate visual cortex are specialized for the processing of color. The (**left**) fi gure shows the blobs selectively stained to highlight their locations regularly interspersed between orientation hypercolumns. (**right**) When single cells are fi lled with dye, their long-range connections become clear. Note how axons tend to terminate within (or near) other cytochrome oxidase blobs (drawn in outline). We submit that it is these long-range connections that enforce "good continuation" between hues at nearby positions. Images courtesy of E. Callaway, Salk Insitute.

3.1. Geometry of Hue Fields

Within the (intensity, hue, saturation) color space, the hue component across the image is a mapping $\mathcal{H} : \mathbb{R}^2 \rightarrow \mathcal{S}^1$ and thus can be represented as a unit length vector field over the image. In many images this hue field is piecewise smooth (Fig. 4) with singularities corresponding to significant scene events (e.g., occlusion boundaries or material changes).

The frame field [23] obtained by attaching a (tangent, normal) frame $\{E_T, E_N\}$ to each point in the image domain is the representation suggested by modern differential geometry. This provides a local coordinate system in which the hue vector and related structures can be represented. Most importantly among these are the covariant derivatives of E_T and E_N, which represent the initial rate of change of the frame when it is moved in a direction v expressed by the connection equation [23]:

$$\begin{pmatrix} \nabla_V E_T \\ \nabla_V E_N \end{pmatrix} = \begin{bmatrix} 0 & w_{12}(V) \\ -w_{12}(V) & 0 \end{bmatrix} \begin{pmatrix} E_T \\ E_N \end{pmatrix} \quad (6)$$

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. $w_{12}(V)$ is a linear 1-form, so it can be represented with two scalars at each point:

$$\begin{aligned} \kappa_T &\triangleq w_{12}(E_T) \\ \kappa_N &\triangleq w_{12}(E_N) \quad . \end{aligned} \quad (7)$$

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 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 $\tilde{HUE}(x, y)$, in the neighborhood of q. Coherence of $HUE(q)$ to its spatial context $HUE(x, y)$ can then be determined by examining how well $HUE(x, y)$ fits $\tilde{HUE}(x, y)$ around q. Clearly, this 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

While many local coherence models $\tilde{HUE}(x, y)$ are possible, we exploit the fact that the hue field is a unit length vector field which suggests that it behaves similarly to oriented texture flows [6, 7] and adopt a similar curvature-tuned local model.

$$\tilde{HUE}(x, y) = tan^{-1}\left(\frac{\kappa_T(q)x + \kappa_N(q)y}{1 + \kappa_N(q)x - \kappa_T(q)y}\right) \quad . \quad (8)$$

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 $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right)$.

This local model possesses many properties that suit good continuation; in particular it is both a minimal surface in the $(x, y, \tilde{HUE}(x, y))$ 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 5. A detailed technical account of the model in the texture flow domain can be found in [7].

4. Examples of Flows

We now illustrate the above computations on several examples. We begin with artificial ones, to illustrate the points most clearly, then proceed to natural ones to illustrate the complexities that arise.

We stress that, for space reasons, some of these flows are not visible unless one zooms in to enlarge the manuscript.

In the first Fig. 6, we show one of the few examples from the psychophysical literature. In an important paper, Kingdom [18] created images consisting of superimposed sinusoids, one in brightness and the other in color. He demonstrated that it is the intensity component that drives the impression of shape-from-shading, while the color information appears "painted" onto the undulating surface. We reproduced this separation with our flows, from which it follows that the shading flow is sufficent (for these examples) to derive the shape.

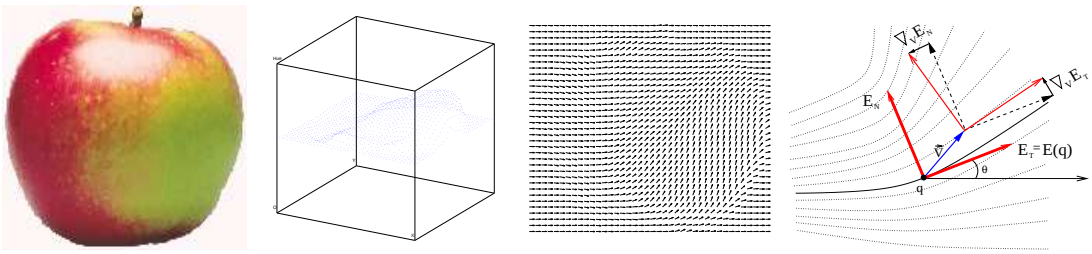

Figure 4. Color images of natural objects are piecewise smooth and the hue flow captures this. **(A)** An apple with varying hue. **(B)** A representation of hue as a scalar field, with value corresponding to height. **(C)** The hue field, with each value represented as a vector pointing to location on the hue circle. **(D)** The geometry of the hue flow, illustrating that nearby values can be represented as a differentiable frame field that is tangent (and normal) to the streamlines of the flow. Interations between nearby hue values then correspond to an (infinitesimal) transport of the frame in direction V, which rotates it according to the connection form of the frame field. Since E_T, E_N are unit length, their covariant derivative lies in a normal direction, regardless of V. This diagram also suggests a relationship between hue and texture and shading flows.

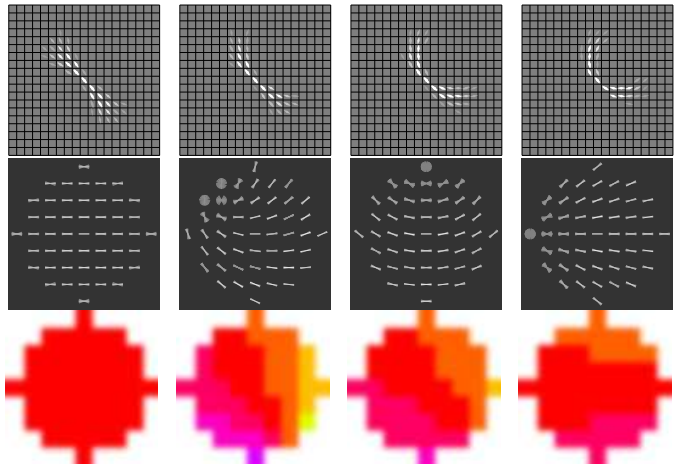

Figure 5. Illustration of the different types of compatibility fields that can be used for early forms of good continuation. In each case the central unit is supported by the contextual arrangement of surrounding units, and can be used as the constraints within quadratic programming, relaxation labeling, and belief propagation engines. **(top)** For boundary continuation, the orientation at a position is enhanced by consistent tangential (co-circular) boundary measurements at nearby positions [24, 14] **(middle)** For oriented texture measurements, both tangential and normal curvatures arise. Similar models can be used for shading flows, which are the tangent fields to the intensity level sets [8]. **(bottom)** For hue flows the orientations are replaced by colors. In the first column zero curvature continuations are shown. In the last column, a single large curvature is shown. For the texture and hue compatibilities, the tangential curvature is zero and the normal curvature is not. Note the emergence of singularities.

The shading flow is estimated by evaluating a gradient operator (an orientationally-selective receptive field tuned to low spatial frequency) over the image. It demonstrates one role for the long-range interactions: correcting local artifacts in shading flow estimation.

Our next examples (Fig. 7) on artificial images confirm the classical view that color remains invariant across shadows while shading effects surface percepts [25]. This is most clear in the plastic sphere, and the same effect is reproduced in the Google logo, which appears both 3-dimensional and colored. However, unlike the plastic sphere, there are no mutual illumination effects.

The next examples show how hue can vary over a natural object. Fig. 4 shows the hue flow for an apple, and Fig. 8 is a close-up of a woman's face in which a blush has been introduced. Note in particular how variant the "color" is, a point of some relevance to both face identification and emotional estimation. Hue can also vary systematically over a scene. Atmospheric depth scattering is shown in Fig. 9.

Our next two examples illustrate the beautiful complexity of shading, hue, and boundary interactions. The first shows an apple photographed on a highly reflective surface in bright sunlight (Fig. 10). The flows are varied with respect to one another and with respect to the boundaries (of both the apple and the shadow). In particular, the mutual illumination modulating the shadow [20] introduces a smooth shading flow not unlike the one for the plastic sphere or the Kingdom examples but this time due to a lighting effect and

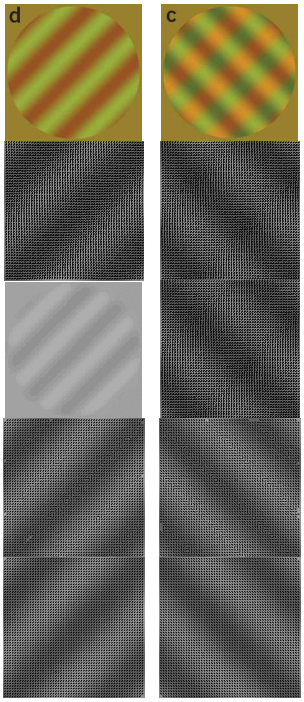

Figure 6. Results on the test Kingdom images. Note how both provide the impression of an undulating surface with color on it. The left column is Kingdom Fig. 2d; the right column is Kingdom Fig. 2c. From top to bottom are original images; initial estimate of shading flow (tangents to intensity level sets); final estimate of shading flow; initial estimate of hue flow; final estimate of hue flow. The shading flow corresponds to the undulations; the hue flows are smooth and do not interfere with them.

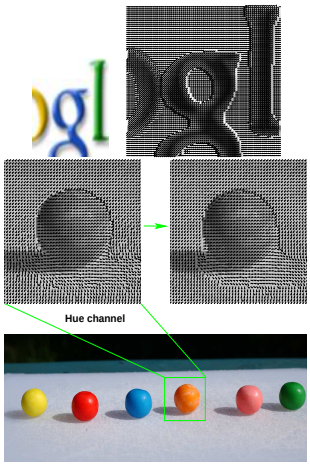

Figure 7. Shading and hue flows for artificial objects. Although the shading flow fields are not shown, notice how the hue flows (superimposed on the original image) are constant over the "plastic" objects. This is the way such materials were designed. The case of the sphere also introduces two more complex lighting effects. First, note how the hue flow remains constant through the shadow. This is a classical cue for separating shadow boundaries from surface boundaries. (Surface boundaries are taken to involve different materials, and therefore a hue discontinuity together with the intensity discontinuity.) Second, and less familiar, is the mutual illumination between the sphere and the tabletop, which is captured by the hue flow but not the shading flow. The left magnification shows the initial local measurements of hue; the right magnification shows the converged hue flow. A boundary has been introduced around the hue flow on the table top illustrating an elongation in the direction of the source.

not a surface normal effect. The mutual illumination effect is also strong on the bananas image (Fig. 11), which also illustrates a shading flow effect due to a highly diffuse cast shadow. In this case the cast shadow phenomenon is readily identified, because the hue flow is constant across it.

Our final example (Fig. 12) illustrates the complement to shading and hue; notice how the hue remains invariant through the highlight, even though it is a complex pattern for the pepper.

5. Summary and Conclusions

Perceptual organization was viewed within Gestalt psychology as pervasive in perception, but discussion of such issues in computer vision is significantly more limited. Our goal in this paper was to take a step back and raise the profile of questions for which P.O. is relevant. Following a biological analogy, we introduced the construct of multiple (spatially) aligned flows within which Gestalt good continuation can be enforced geometrically but between which information can be inferred about the many complexities of lighting, space, and geometry. The computation of each

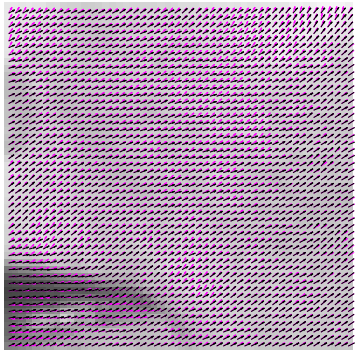

Figure 8. Hue flows vary for natural objects. This shows a portion of a woman's face (the lips are lower left) when she is blushing (blue vectors) and not blushing (black vectors). Note how hue varies both spatially and as a function of emotional and physical states.

flow was global, based on local measurements and differential (covariant derivative) constraints between them. At the same time the computation of each flow was local within an information (sometimes within a sensor) source, and logical relationships between flows provide a new foundation for

Figure 9. Hue fbws and atmospheric depth effects. The fbw is shown along a thin strip on the right side of the photograph. Note the dominant shift toward blue for the upper half.

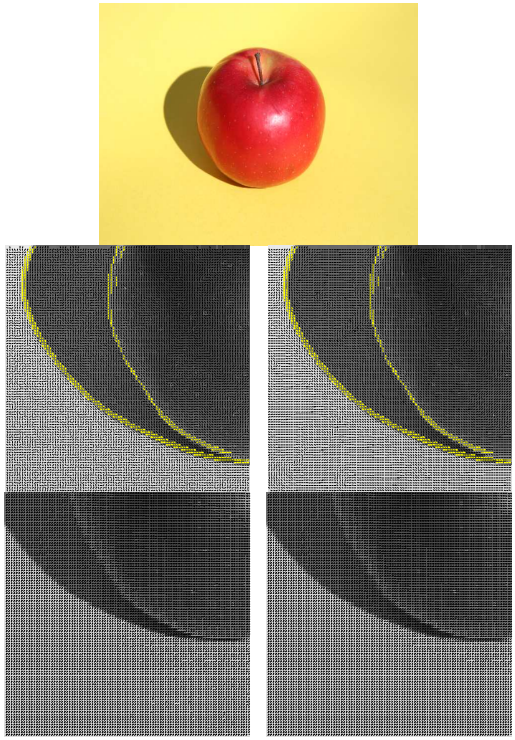

Figure 10. An image of an apple on colored cardboard in bright sunlight. It illustrates the complexities that can arise both for shading due to surface irregularities from packing and from mutual illumination. In particular, the shaded area now exhibits a shading fbw derived from mutual illumination, in which the gradient decreases in magnitude away from the concavity between the apple and the table. At the same time, there is strong mutual illumination between the apple and the cardboard and the cardboard and the apple. The result are smooth shading and hue fbws, with discontinuities at neither object nor shadow edges.

many computer vision computations. Hue flows smoothly through shadows, while intensity often jumps. Shading flows smoothly over many man-made objects, while hue is often constant. Natural objects often imply smooth shading and hue flows, although they are typically independent of one another. The involvement of boundaries is both necessary and complicated [12].

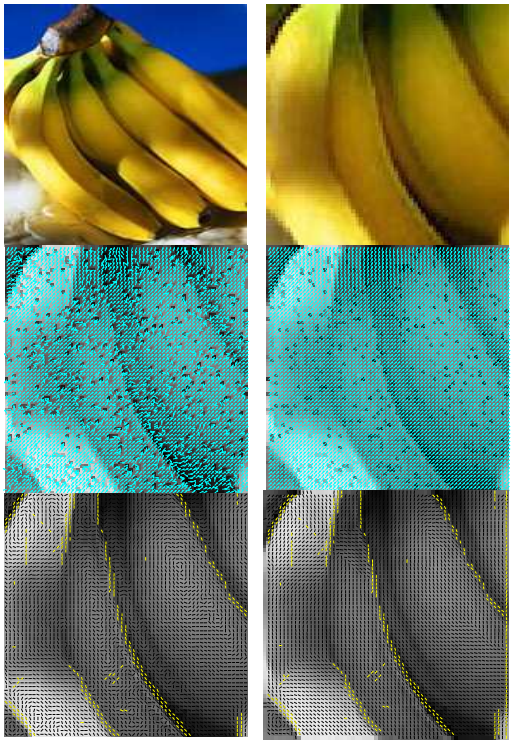

Figure 11. A photograph of bananas illustrates the richness of mutual illumination in a complex scene. The result is an essentially constant hue fbw (middle row, left: initial measurement; right: consistent fbw). The shading fbw (**bottom**) illustrates a special interaction between boundaries and shading fbws, in which multiple surface fold away from each other along them. Such situations are geometrically rare.

While the list of interactions must be extended (motion and stereo should at least be included), it is useful to conclude on an enlargement of the biological metaphor underlying this paper. The centrality of long-range horizontal connections as defining each flow suggests that the flows be layered on top of one another, enabling "vertical" connections for their interactions. Recent breakthoughs in color processing demonstrate that hue and orientation are not independent, as was once thought, and that such vertical connections exist [26]. Computationally it remains an open question whether only two interaction "dimensions" suffice.

References

[1] Y. Adini, D. Sagi, and M. Tsodyks. Excitatory-inhibitory network in the visual cortex: Psychophysical evidence. *Proc. Natl. Acad. Sci. U.S.A.*, 94:10426–10431, 1997. 3

[2] W. Beaudot and K. Mullen. How long range is contour integration in human color vision. *Visual Neurosci.*, 20:51–64, 2003. 3

[3] J. Beck. *Surface Color Perception*. Cornell University Press, 1972. 2

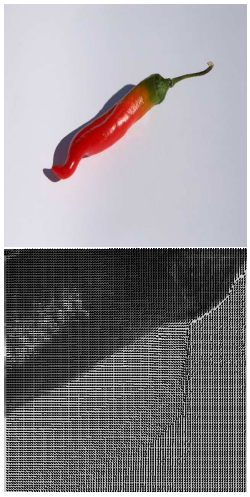

Figure 12. A photograph of a pepper illustrates how the hue fbw remains constant through highlights, even when the hue is varying. The mutual illumination on the background is very interesting as well. (**top**) Pepper image. (**bottom**) Hue fbw through a portion of the pepper image. The fbw is superimposed on the image for identifi cation purposes only.

[4] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 6(15):1373–1396, June 2003. 2

[5] O. Ben-Shahar, P. Huggins, and S. Zucker. On computing visual fbws with boundaries: The case of shading and edges. In *Workshop on Biologically Motivated Computer Vision*, 2002. 3

[6] O. Ben-Shahar and S. Zucker. On the perceptual organization of texture and shading fbws: From a geometrical model to coherence computation. In *Proc. Computer Vision and Pattern Recognition*, pages 1048–1055, 2001. 4

[7] O. Ben-Shahar and S. Zucker. The perceptual organization of texture fbws: A contextual inference approach. *IEEE Trans. Pattern Anal. Machine Intell.*, 25(4):401–417, 2003. 3, 4

[8] O. Ben-Shahar and S. Zucker. Geometrical computations explain projection patterns of long range horizontal connections in visual cortex,. *Neural Comput.*, 16(3):445–476, 2004. 5

[9] P. Breton and S. Zucker. Shadows and shading fbw fi elds. In *Proc. Computer Vision and Pattern Recognition*, pages 782–789, 1996. 3

[10] R. Coifman, S. Lafon, A. Lee, M. Maggioni, B. Nadler, F. Warner, and S. Zucker. Geometric diffusions as a tool for harmonic analysis and structure defi nition of data: Diffusion maps. *Proc. Nat. Acad. Sci. (USA)*, 102(21):7426 – 7431, 2005. 2

[11] R. Coifman, S. Lafon, A. Lee, M. Maggioni, B. Nadler, F. Warner, and S. Zucker. Geometric diffusions as a tool for harmonic analysis and structure defi nition of data: Multiscale methods,. *Proc. Nat. Acad. Sci. (USA)*, 102(21):7432 – 7437, 2005. 2

[12] J. Elder and S. Zucker. Evidence for boundary-specifi c grouping. *Vision Res.*, 38(1):143–152, 1998. 7

[13] K. Garg and S. Nayar. When does a camera see rain? *ICCV*, 2005. 1

[14] W. Geisler, J. Perry, B. Super, and D. Gallogly. Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res.*, 41(6):711–724, 2001. 5

[15] A. Hanazawa, H. Komatsu, and I. Murakami. Neural selectivity for hue and saturation of colour in the primary visual cortex of the monkey. *Eur. J. Neurosci.*, 12:1753–1763, 2000. 3

[16] B. Horn and M. Brooks, editors. *Shape from Shading*. MIT Press, Cambridge, MA, 1989. 1

[17] G. Kanizsa. *Organization in Vision: Essays on Gestalt Perception*. Praeger Publishers, 1979. 1

[18] F. Kingdom. Color brings relief to human vision. *Nature Neuroscience*, 6(6):641–644, 2003. 4

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

[20] M. Langer. When shadows become interreflections. *Int. J. Comput. Vision*, 34(2/3):193–204, 1999. 5

[21] S. Lehky and T. Sejnowski. Network model of shape-from-shading: neural function arises from both receptive and projective fi elds. *Nature*, 333:452–454, 1988. 3

[22] A. Munsell. *A Color Notation*. G.H.Ellis, Boston, 1905. 3

[23] B. O'Neill. *Elementary Differential Geometry*. Academic Press, 1966. 4

[24] P. Parent and S. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Trans. Pattern Anal. Machine Intell.*, 11(8):823–839, 1989. 3, 5

[25] M. Ruzon and C. Tomasi. Color edge detection with the compass operator. In *Proc. Computer Vision and Pattern Recognition*, pages 160–166, 1999. 5

[26] R. Shapley and M. Hawken. Neural mechanisms for color perception in the primary visual cortex. *Curr. Opin. Neurobiol.*, 12:426–432, 2002. 7

[27] P. Sumner and J. Mollon. Chromaticity as a signal of ripeness in fruits taken by primates. *Journal of Experimental Biology*, 203(13):1987–2000, 2000. 2

[28] M. Tappen, W. Freeman, and E. Adelson. Recovering intrinsic images from a single image. *IEEE Trans. Pattern Anal. Machine Intell.*, 27(9):1459–1472, 2005. 1

[29] T. Wachtler, T. Sejnowski, and T. Albright. Representation of color stimuli in awake macaque primary visual cortex. *Neuron*, 37:681–691, 2003. 3

[30] A. Witkin and J. Tenenbaum. On the role of structure in vision. In J. Beck, B. Hope, and A. Rosenfeld, editors, *Human and Machine Vision*, pages 481–542. Academic Press, 1983. 1

[31] Y. Xiao, Y. Wang, and D. Felleman. A spatially organized representation a colour in macaque cortical area v2. *Nature*, 421:535–539, 2003. 3

[32] N. Yabuta and E. Callaway. Cytochrome oxidase blobs and intrinsic horizontal connections of layer 2/3 pyramidal neurons in primate v1. *Visual Neurosci.*, 15:1007–1027, 1998. 3