In the *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2003

Hue fields and Color Curvatures: A Perceptual Organization Approach to Color Image Denoising*

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 Electrical Engineering, Yale University
New Haven, CT 06520-8285
steven.zucker@yale.edu

Abstract

The denoising of color images is an increasingly studied problem whose state-of-the-art solutions employ a variety of diffusion schemes. Specifying the correct diffusion is difficult, however, in part because of the subtleties of color interactions. We address this difficulty by proposing a perceptual organization approach to color denoising based on the principle of good continuation. We exploit the periodic chromatic (hue) component of the color in its representation as a frame field. We derive two hue curvatures and use them to construct a local model for the behavior of the color, which in turn specifies consistency constraints between nearby color measurements. These constraints are then used to replace noisy pixels by examining their spatial context. Such a contextual analysis (combined with standard methods to handle the scalar channels, saturation and lightness), results in a robust noise removal process that preserves discontinuities, singularities, and fine chromatic structures, including those that diffusion processes are prone to distort. We demonstrate our approach on a variety of synthetic and natural images.

1. Introduction

While much of computer vision and image processing research focuses exclusively on the luminance domain, the analysis of color images has gained increasing interest in recent years. An important problem in the processing of color images is that of enhancement and denoising. We attempt a unification of these tasks by exploiting the geometry of color in a perceptual organization framework.

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 appropriately to the color *vectorial* data. While some studies explore vector median and directional filters [1, 26], most color image enhancement approaches are based on some form of anisotropic diffusion [20, 27, 25], either on an explicit vectorial representation of the color, or based on differential geometrical properties of a manifold representation in a higher dimensional space [24, 21, 22, 12].

Many of these efforts, as well as other related approaches [6, 5, 7], use the perceptually-motivated HSV color representation in which the achromatic (lightness, or value) and chromatic components of the color image are explicitly separated, and the latter is further described with two independent variables for hue and saturation. This representation has deep roots in the psychology of perception and the subjective experience of color, from Hering's opponent hues theory [9] and Munsell's book of colors [15], to accumulating psychophysical evidence on color sensitivity [8]. Combined with neurophysiological findings for color opponent cells, color specificity in V4, and perceptual impairments such as achromatopsia [14], it is now commonly speculated that neural structures that explicitly encode the hue, saturation, and lightness do exist and will eventually be found [17, p. 120].

Of the three perceptual dimensions of HSV, the hue is periodic and can be represented as an angle. Although it is not spatial in nature, this angle associates a *geometric* meaning to the hue component of the color (or alternatively, to the chromatic component as a whole) and thus suggests a geometrical treatment. Although this idea is not new and has been exploited in the color diffusion literature through the diffusion of orientations (e.g., [24, 7]), here we propose a different approach that treats the angular hue component in the context of perceptual organization and the principle of *good continuation*. 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. Most importantly, this is done while making explicit the observations that (1) the hue is typically piecewise smooth over the image, and (2) that it may vary greatly, albeit smoothly, even *within* perceptually coherent objects (Fig 1). In noisy

*This paper should be printed in color or viewed on a monitor. Please refer to the CDROM version of the proceedings. Research supported by AFOSR, DARPA, and ONR.

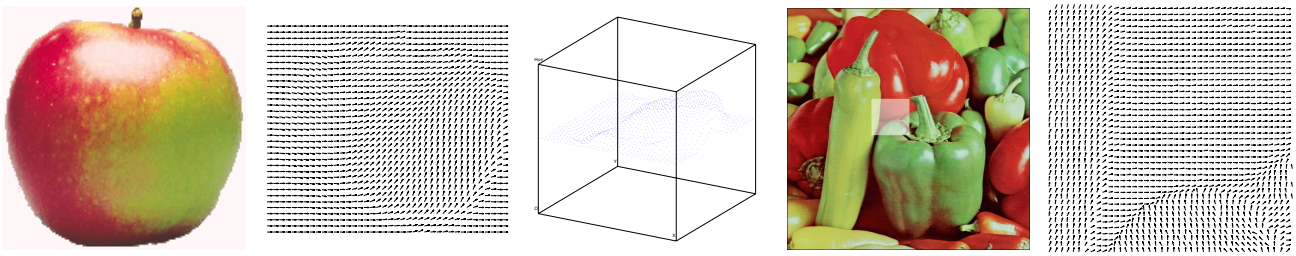

A B C D E

Figure 1: 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 that makes these variations explicit is necessary. **(A)** A natural image of an apple with varying hue. **(B)** The corresponding hue field (see Section 2). 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 its bottom one leads to the space $XY\mathcal{H} \stackrel{\triangle}{=} \mathbb{R}^2 \times \mathcal{S}^1$ in which the image's hue is a submanifold. **(D)** A natural image of peppers with a region of interest (ROI). **(E)** The hue field of the peppers image in the ROI 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 while noisy images are denoised.

color images, the noisy pixels do not posses the required good continuation with their context, thus can be detected, suppressed, and replaced with better ones.

The rest of this paper addresses color denoising through geometrical good continuation of hue while treating the noise in the other two linear dimensions (saturation and lightness) through independent scalar anisotropic diffusion [20, 4]. Since diffusion processes are now well established, and for lack of space, we do not discuss it further in this paper. Rather, we first analyze the geometrical content of the hue and show how it necessarily results in a notion of hue curvatures. We then use these curvatures to devise a local, curvature-tuned model of coherent hue variations and study its properties. This model is subsequently used to devise consistency constraints between nearby pixels which we use in a relaxation labeling network to denoise color images. We conclude with experimental results on both synthetic and natural images and a short discussion of our future work in this area.

2. Color, hue, and hue curvatures

Consider the HSV color space, in which a color image is a mapping $\mathcal{C} : \mathbb{R}^2 \rightarrow \mathcal{S}^1 \times [0,1]^2$, where $\mathcal{S}^1$ is the unit circle. The hue component across the image is a mapping $\mathcal{H} : \mathbb{R}^2 \rightarrow \mathcal{S}^1$ and thus can be represented as a unit length vector field over the image plane, henceforth called the *hue field*. In many natural images, as well as visual artifacts, this hue field is typically piecewise smooth (Fig. 1) with orientation singularities that correspond to significant events in the scene (e.g., occlusion). Denoising color images requires suppressing spurious measurements *within* the smooth segments while preserving the boundaries *between* them. To achieve this goal, we first develop the notion of color coherence as expressed through the smoothness of the hue field.

An extension of the vector field representation that makes tools from differential geometry readily available is that of the *frame field* [16]. 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 provide a local coordinate system in which all other vectors can be represented in a natural, object centered view (Fig. 2). 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* [16]:

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

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 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:

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

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

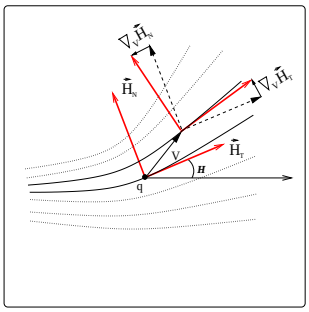

Figure 2: Any smooth hue field (depicted here with the set of locally parallel lines), can be represented as a differentiable frame field which is everywhere tangent and normal to the streamlines of the flow. An infinitesimal translation of the frame in a direction V rotates it by some angle determined 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 orthogonal expansion. The natural expansion based on 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* [3],

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 directly in terms of the hue field ($\mathcal{H}_T$) using the standard curl ($\nabla\times$) and divergence ($\nabla\cdot$) operators:

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

However, more useful is the expression of the hue curvatures in terms of the hue itself and its gradient (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):

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. More generally, this is characterized by the following constraint:

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

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

3. 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 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 flow [3]. 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. 1C), this model takes the form of a right helicoid[1]:

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

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 3. A detailed technical account on the development of the model in the texture flow domain can be found in [3].

[1] Unlike texture flows, the local model for the hue function is not a *double* helicoid since the hue function is 2π-periodic while texture flows are π-periodic. The basic proofs carry through, however.

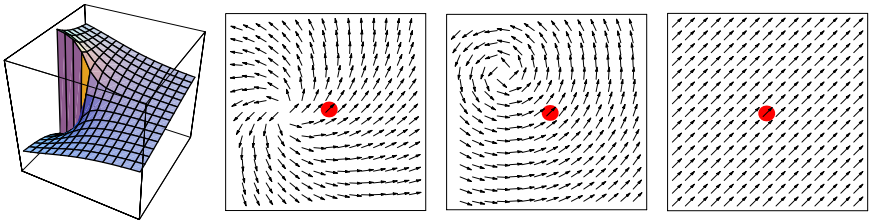

Figure 3: A local model for coherent local behavior of the hue can be depicted both as **(A)** a height function that wraps itself to $[-\pi, \pi)$ (hence the apparent discontinuity), and as **(B-D)** a unit length vector field in the image plane. Different orientation and curvature tunings at the central pixel (marked with red) yield different local behaviors around it, and the three variations shown correspond qualitatively to **(B)** $\kappa_T > 0$ and $\kappa_N > 0$, **(C)** $\kappa_T > 0$ and $\kappa_N = 0$, **(D)** $\kappa_T = 0$ and $\kappa_N = 0$. All three fields are tuned to the same hue value of $\mathcal{H} = 45\,^\circ$.

4. 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* [11, 13], *recurrent neural networks* [10], and *belief propagation networks* [18]. Here we present results using a relaxation labeling network whose nodes $i = (x, y)$ 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, is 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. 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. 4. 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 [11, 19] ensures that such a rule will converge to a consistent labeling while extremizing the average local consistency over the entire image.

5. Experimental results

We applied the proposed model for hue good continuation and the corresponding relaxation labeling network for denoising images 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. 4). Step size was set to $\delta = 0.5$.

Fig. 5 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 5E-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. 5H). 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

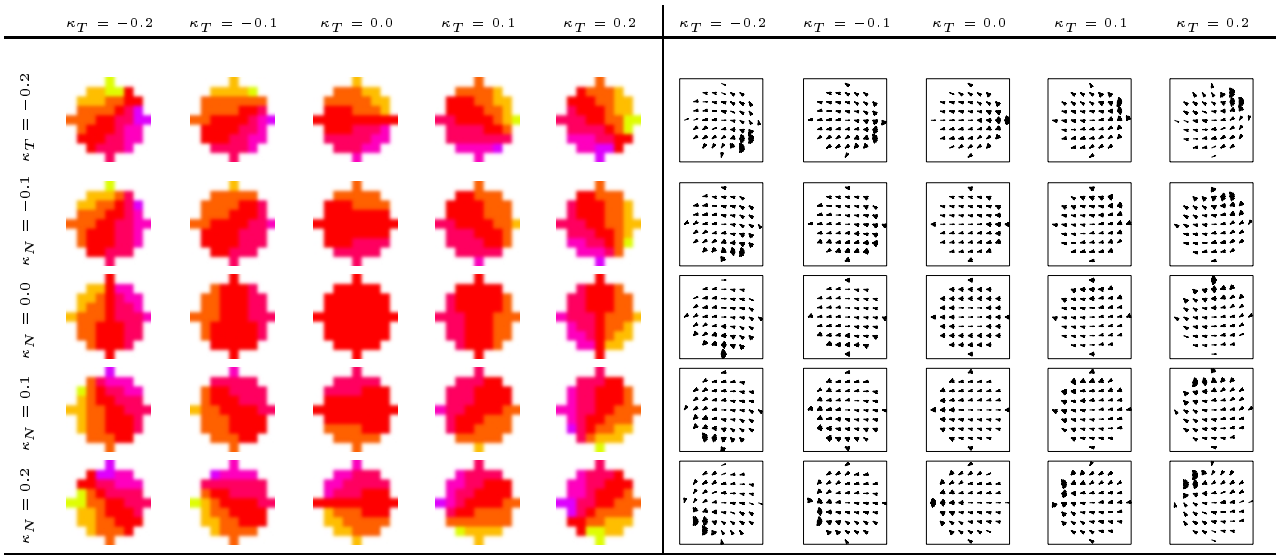

Figure 4: A collection of 9×9 hue compatibility fields for $\mathcal{H} = 0^\circ$ (the red color), where both κ_T and κ_N are selected from the set $\{-0.2, -0.1, 0.0, 0.1, 0.2\}$. Each field represents all compatible hue values in the neighborhood of the central label; on the left these compatible 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" amounts to a rotational transformation of the field.

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. 6 illustrates the results of our denoising approach on the Apple image from Fig. 1. 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 (c.f. Fig. 5D). The stability of the process once convergence has been achieved is demonstrated in Fig. 7.

Finally, Fig. 8 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. Summary

We have presented an alternative approach for the denoising of color images. Our approach is based on principles of perceptual organization, and in particular the application of good continuation to the two dimensional orientation field of the hue channel. Based on a notion of hue curvature, we derived a model for the local behavior of coherent hue. Subsequently, we developed a contextual method in which every pixel reinforces, or suppresses, the confidence in its color through examination of its spatial context. Implemented as a relaxation labeling network, this allows for a robust noise removal while maintaining the underlying structures in the color image. As we have argued, our approach is able to preserve not only step like edges in color space but more delicate structures as well, including those which orientation diffusion is prone to distort.

Since it is represented as orientation, the hue dimension of the color is a natural candidate for an application the principle of good continuation, and the experimental results we obtained indicate the usefulness in such an approach. While thus far we have treated the saturation and lightness channels with scalar anisotropic diffusion, an integrated good-continuation approach is in place. In particular, the combination of hue and saturation (sometimes called the chroma) yield yet another orientation-like representation and while this representation resides on $\mathcal{S}^2$ (as opposed to $\mathcal{S}^1$), the same curvature-based principle of good continuations is applicable there as well. Although the lightness channel does not encode direction directly, the behavior of its levelsets does so indirectly and considering this information in color denoising may prove critical in preserving important visual

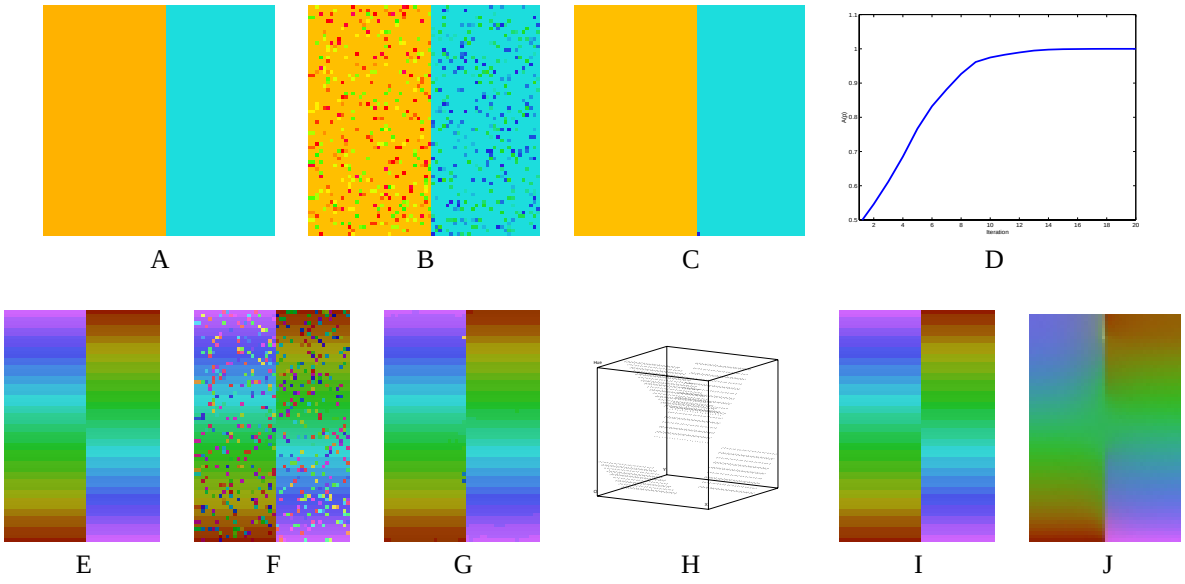

Figure 5: Color denoising of synthetic color images. **(A)** A color step edge. **(B)** A noisy step edge. **(C)** The result of 20 iterations of relaxation labeling. The noise is completely eliminated while the edge structure is preserved. **(D)** is the graph of the average local consistency as a function of iteration. it illustrates that convergence on this input was effectively achieved at the 12^{th} iteration, when the assignments $p_i(\lambda)$ reached a steady state. **(E-G)** 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 color images [22, 12], but similar results were obtained with other orientation diffusion schemes (e.g., [23, 21]).

structures (e.g., roof edges). Thus, color as a whole can be treated geometrically with a good continuation approach, a possibility which we study in our current research.

Finally, the approach proposed in this paper may have intriguing consequences for other vision sciences. Psychophysically, it suggests an examination of color segmentation and filling-in phenomena (Fig. 9) from a geometrical point of view, similar to recent findings for orientation-based texture segmentation [2]. Physiologically, it suggests to search for long range interactions between color sensitive cells (either in V1 cytochrome oxidase blobs or perhaps in V4) which facilitates coherent color perception. All these directions are part of our future research in this area.

References

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

[2] O. Ben-Shahar and S. Zucker. Curvature and the perceptual organization of texture flows. In *VSS*, 2002.

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

[4] M. Black, G. Sapiro, D. Marimont, and D. Heeger. Robust anisotropic diffusion. *IEEE Trans. Image Processing*, 7(3):421–432, 1998.

[5] P. Blomgren and T. Chan. Color TV: Total variation methods for restoration of vector-valued images. *IEEE Trans. Image Processing*, 7(3):304–309, 1998.

[6] G. Boccignone, M. Ferraro, and T. Caelli. Generalized spatio-chromatic diffusion. *IEEE Trans. Pattern Anal. Machine Intell.*, 24(10):1298–1309, 2002.

[7] T. Chan and J. Shen. Variational restoration of non-flat image features: Models and algorithms. CAM-TR 99-20, UCLA, 1999.

[8] R. De Valois and K. De Valois. *Spatial Vision*. Oxford University Press, 1990.

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

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

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

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

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

A B C D

Figure 6: Color denoising of the Apple image (from Fig. 1) **(A)** Original image. **(B)** Noisy image. **(C)** Result of 25 denoising iterations. **(D)** Result of color diffusion [22, 12]. Note that although the noise is not yet removed, the color structure across the apple is completely blurred and distorted, as it is the case across the apple's boundaries. Compare to C.

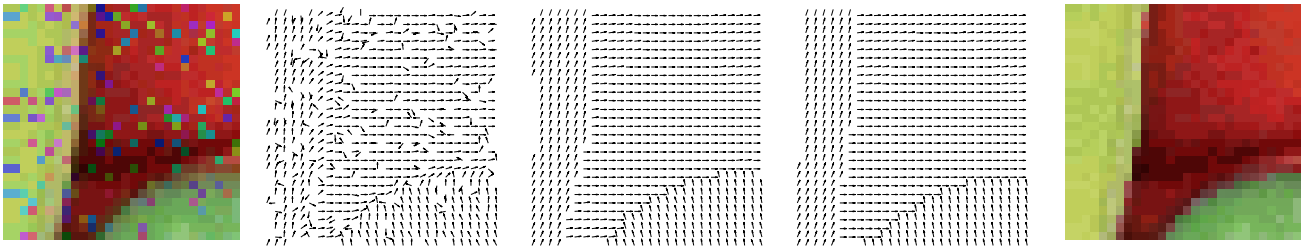

A B C D E

Figure 7: 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 noisy hue field. **(C)** The result of 20 iterations. **(D)** The result of 50 iterations. **(E)** The denoised color image.

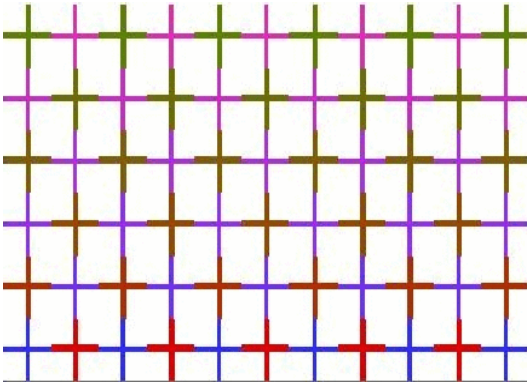

Figure 9: A modified Ehrenstein illusion shows how color filling in can occur between boundaries of different colors, thus suggesting a link between the approach proposed in this paper and the psychology of color perception. This illusion best viewed on a fluorescent (e.g., CRT) display.

[14] S. L. and J. Werner, editors. *Visual Perception. The Neurophysiological Foundations.* Academic Press, Inc., 1990.

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

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

[17] S. Palmer. *Vision Science: Photons to Phenomenology*. The MIT Press, 1999.

[18] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.

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

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

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

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

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

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

[25] B. ter Haar Romeny, editor. *Geometry-Driven Diffusion in Computer Vision*. Kluwer Academic Publishers, 1994.

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

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

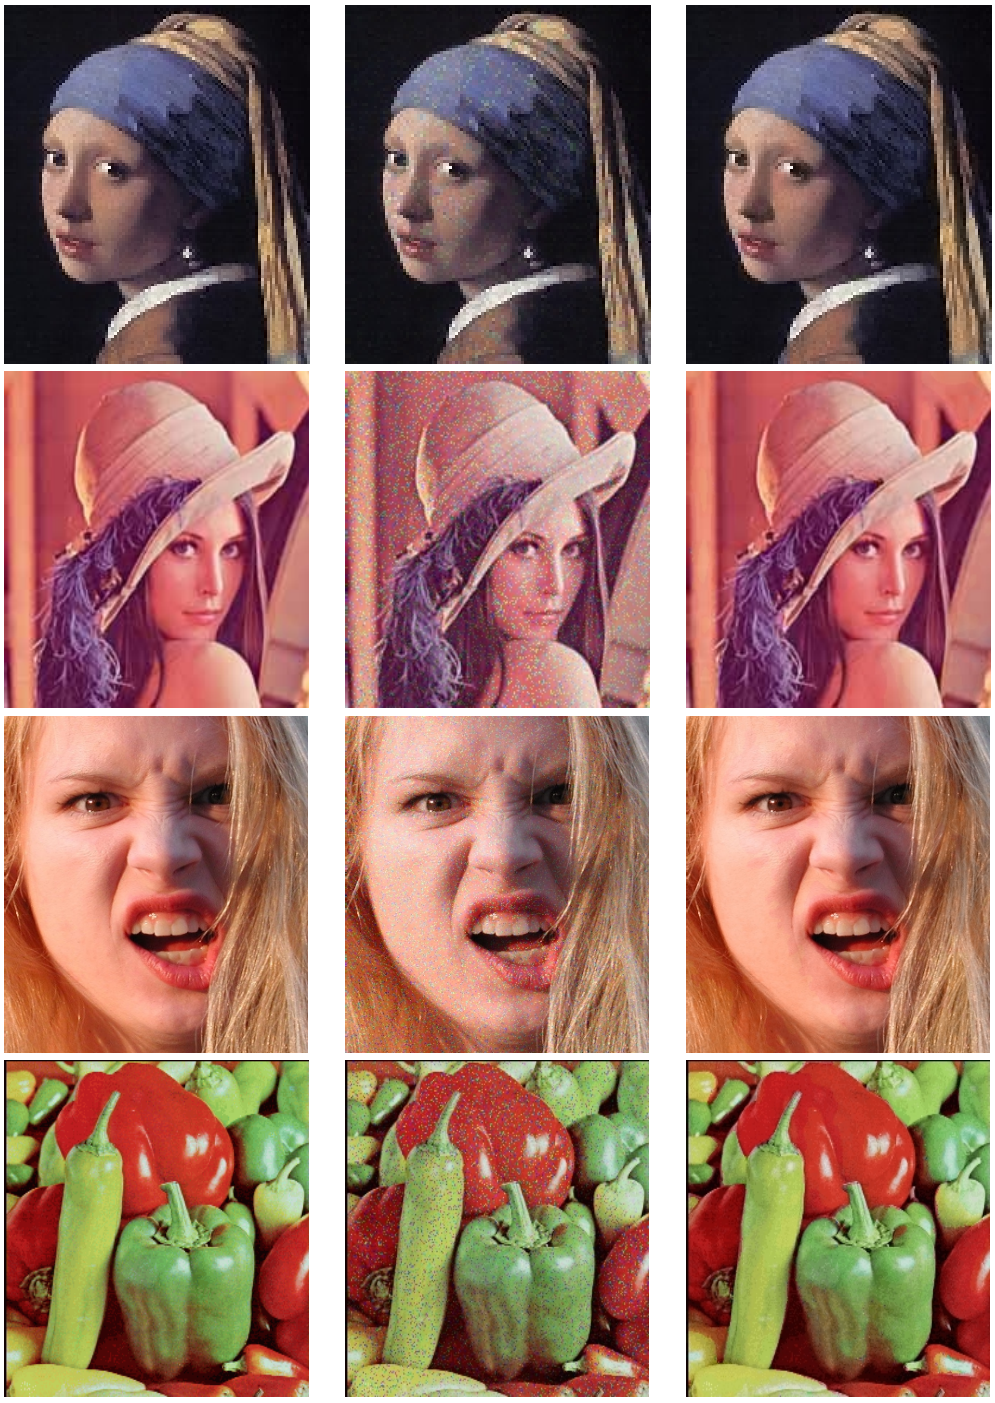

A B C

Figure 8: Color denoising of the Vermeer, Lena, Sidney, and Peppers images. **(A)** Original image. **(B)** Noisy image. **(C)** Convergence state of the relaxation process. We are grateful to Pamela Davis for the permission to use the Sidney image.