Int J Comput Vis (2010) 86: 48–71
DOI 10.1007/s11263-009-0255-8

General Geometric Good Continuation: From Taylor to Laplace via Level Sets

Ohad Ben-Shahar · Steven Zucker

Received: 7 April 2008 / Accepted: 25 May 2009 / Published online: 17 June 2009

Abstract Good continuation is the Gestalt observation that parts often group in particular ways to form coherent wholes. Perceptual integration of edges, for example, involves *orientation* good continuation, a property which has been exploited computationally very extensively. But more general local-global relationships, such as for shading or color, have been elusive. While Taylor's Theorem suggests certain modeling and smoothness criteria, the consideration of level set geometry indicates a different approach. Using such first principles we derive, for the first time, a generalization of good continuation to all those visual structures that can be abstracted as scalar functions over the image plane. Based on second order differential constraints that reflect good continuation, our analysis leads to a unique class of harmonic models and a cooperative algorithm for structure inference. Among the different applications of good continuation, here we apply these results to the denoising of shading and intensity distributions and demonstrate how our approach eliminates spurious measurements while preserving both singularities *and* regular structure, a property that facilitates higher level processes which depend so critically on both of these classes of visual structures.

Keywords Perceptual organization · Good continuation · Grouping · Level set geometry · Singular and regular structure · Harmonic function · Anisotropic diffusion · Scale space · Relaxation labeling

O. Ben-Shahar (✉)
Department of Computer Science, Ben-Gurion University of the Negev, P.O. Box 653, Beer-Sheva, 84105, Israel
e-mail: ben-shahar@cs.bgu.ac.il

S. Zucker
Department of Computer Science, Yale university, P.O. Box 208285, New Haven, CT, 06520, USA

1 Introduction and background

Visual perception, by either man or machine, is a process of acquiring descriptions and information about the world from its images. Perceptual organization (Wertheimer 1955; Koffka 1935) facilitates this goal by providing a "*description that decomposes the image into constituents that capture regularity or coherence [and] therefore provides descriptive chunks that act as 'semantic precursors', in the sense that they deserve or demand explanation*" (Witkin and Tenenbaum 1983, p. 483). Such a goal, which presupposes only minimal a priori knowledge about the world, and employs only the most general models (Zucker et al. 1975), is typically embodied in a variety of tasks, from texture segmentation, through feature grouping, to figure-ground segregation and contour completion.

Of the different perceptual organization principles (Wertheimer 1955), one particularly prominent is *good continuation*. Described by Wertheimer as the "*inner coherence*" by which "*successive parts of a whole should follow one another*" (Wertheimer 1955, p. 83), good continuation had become a fundamental topic of computational investigation in a variety of ways and levels of abstraction. Notably, these efforts focused almost exclusively on the study of *curve-like* structures (e.g., Ullman 1976; Sha'ashua and Ullman 1988; Parent and Zucker 1989; Fischler 1994; Guy and Medioni 1996; Williams and Jacobs 1997; Kimia et al. 1999; Sharon et al. 2000; Cao 2004), perhaps because of the abundance of curves in visual stimuli, the importance of edge detection, and the fact that Wertheimer himself demonstrated this principle in terms of curve-like structures.

While the identification of good continuation with curves is historical, it is clear that Wertheimer meant more when he coined his general "*Factor of Direction*" (Wertheimer 1955).

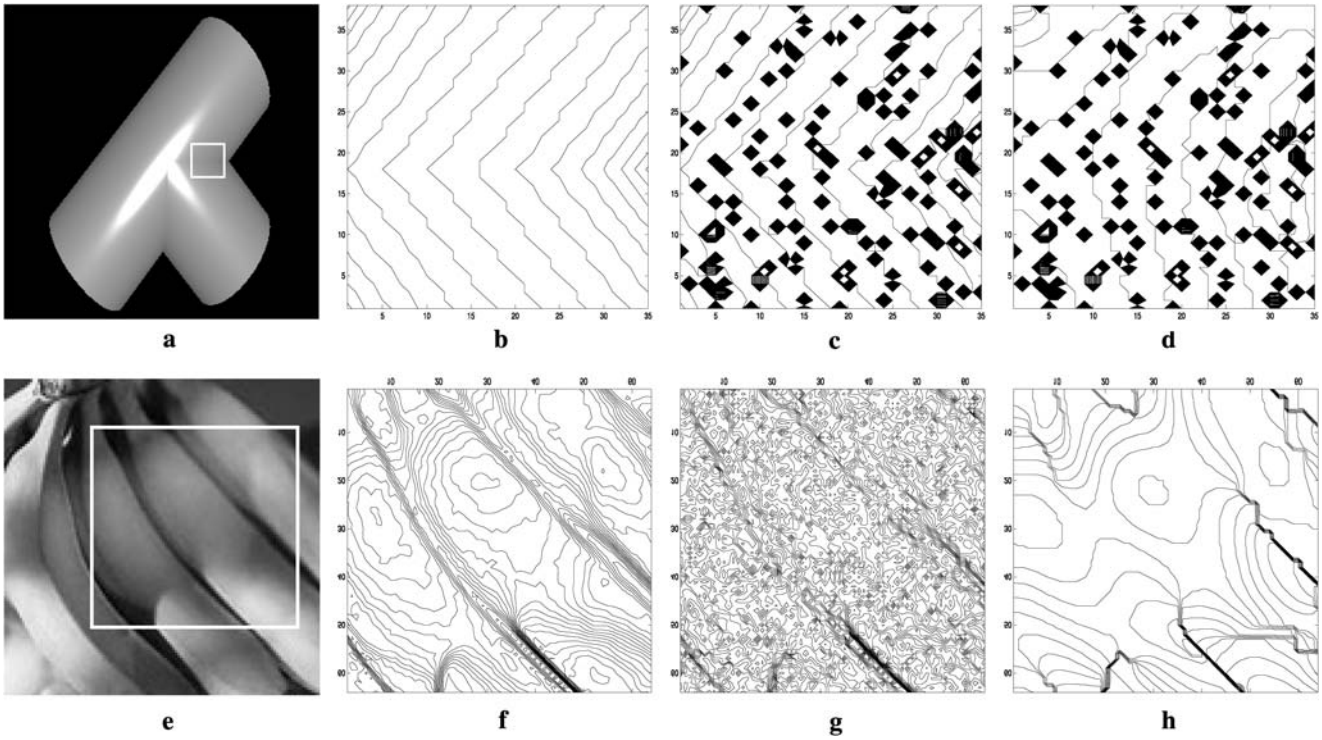

Fig. 1 The need for structure preservation in image operations for vision. (**a**) A two cylinder pipe-like object and region of interest (ROI). (**b**) The geometry of the shading level sets in the ROI. Note both the singularity and the regular structure. Both play a critical role in interpreting the shading correctly. (**c**) Noise (in this case—speckle) contaminates both the intensities and their level sets. For successful higher level visual processing, this noise must be eliminated with minimal qualitative change to the geometry of the shading. (**d**) A popular denoising procedure applied to the noisy image distorts both the regular structure and the singularity long before the noise is eliminated. Any subsequent shape inference procedure is therefore likely to fail (in the sense of reporting the correct physical structure). (**e**–**h**) A similar demonstration to a–d, this time using a natural image of bananas contaminated with additive Gaussian noise. Although the effect is less dramatic, it is evident that much of the geometrical structure is significantly distorted in the denoised image, putting at risk the success of subsequent shape inference attempts

Unfortunately, this general conceptual view of good continuation for arbitrary visual structures has barely penetrated into formal and computational investigations. Our goal in this paper is to develop such a computational generalization for a wide class of visual features—those that can be abstracted as scalar functions from the image plane to the real line **R** (or subintervals thereof). As we argue, developing this approach requires an explicit consideration of the geometry of the *level sets* of these functions.

The contribution of our paper is both theoretical and in addressing a critical computational necessity for computer vision. To motivate this argument, consider the image in Fig. 1a which is easily perceived as a fixture of two intersecting tubes. The intensity image has several cues that should enable a computer vision system make the correct 3D shape inference and part decomposition. For example, consider the (second-order) intensity *singularity* of co-dimension 1 in the marked region. Although it is not visually salient, this singularity is clearly visible in the representation of the shading as a set of iso-luminance level sets (Fig. 1b). Correspondingly, the *regular* geometrical structure away from this singularity is essential for correctly inferring the cylindrical shape of the two parts that intersect along the singularity. *Both* the regular and singular structures must survive any preprocessing before being handled by a higher level shape inference or part decomposition computation. But if the image were contaminated with noise, this task would be difficult to achieve. Fig. 1c shows the level set geometry for a noisy patch. Indeed, trying to remove the noise might distort both the singularity and the regular structure long before the noise itself is eliminated (Fig. 1d). If this happens, any subsequent shape inference process is thus likely to fail.

Noise elimination, as is required in Fig. 1, is of course nothing new to image analysis. However, typical denoising processes care explicitly about the noise, and frequently for the preservation of singularities, but rarely for the preservation of regular structures. Although the notion of good continuation is illusive and difficult to define, in this paper we argue that a contextual computation equipped with a formal model of good continuation can achieve all these tasks (i.e., noise reduction, preservation of singularities, and preservation of regular structure) simultaneously. Indeed,

various nonlinear diffusion techniques that smooth data sets (e.g., Alvarez and Mazorra 1994; Sochen et al. 1998; Tang et al. 2000) relate *indirectly* to good continuation by their virtue of minimizing certain variational functionals on a geometric representation of the image embedded in some higher dimensional space. However, these techniques typically do not account for the fact that certain transformations should be prohibited in this space, or that its dimensions are incommensurable (Koenderink and van Doorn 2002). Moreover, these techniques do not address good continuation *explicitly* or able to guarantee any kind of performance with regard to it. A more explicit consideration of good continuation of non curve-like features has been attempted within the tensor voting framework (Guy and Medioni 1997) for surface patches, and formal considerations based on a frame field representation were offered by Ben-Shahar and Zucker (2003) for locally parallel structures. We follow the latter in using a frame field representation, but note that it, too, is *incomplete* for general image features which are abstracted as 2D scalar functions. In technical terms we compensate for this incompleteness by extending our previously suggested helicoidal model (Ben-Shahar and Zucker 2003) to a much more elaborate model that provides not only good continuation for the level set geometry, but also for the structure of their values on top of that. Then, we incorporate this model in a relaxation labeling network that permits not only the organization of coherent regions, but also the preservation of the singularities between those regions.

As a final introductory comment we note that in this paper we apply our theory and computational algorithms directly to the image *shading* function so that the link to denoising can be easily demonstrated and illustrated visually. However, we stress that this work is theoretical and general, so our results are applicable not only to denoising, but to all cases where good continuation is a factor in the interpretation of visual signals. Among these additional applications are image restoration, segmentation, fragment and patch grouping, visual completion, image inpainting, and others. Similarly, our results are applicable not only to raw image data and intensity values, but also to all cases where a visual feature can be expressed in terms of a piecewise smooth 2D scalar function over the image plane. Hence, by proper representation one can also apply our results to the R, G, and B channels of color images; to color saturation and intensity in HSV representations; to disparity information along epipolar lines in stereo pairs; to texture energy channels (Malik and Perona 1990); to the magnitude of optical flow fields, and so forth.

2 The Geometry of 2D Scalar Features

Let $I : \mathbf{R}^2 \to \mathbf{R}$ be a generic image feature defined over the image plane. (We will frequently refer to I as the "shading" function, although this is a choice of convenience only; it may be one of the many image features described above.) For now, assume I is smooth.[1] Our goal is to group different measurements ("parts") of I into coherent "wholes". Following the Gestalt principle of good continuation, two nearby measurements should be considered part of the same coherent unit if and only if they are in good "succession" relative to each other (Wertheimer 1955). The question, of course, is how to define "good succession" formally.

One possibility is to approximate the "good" behavior of the function $I : \mathbf{R}^2 \to \mathbf{R}$ around a point $q = (a, b)$ based on the coefficients of its Taylor expansion. Depending on the order of approximation, this results in one ($I(q)$), three ($I(q)$ and the gradient $\nabla I(q)$), six ($I(q)$, $\nabla I(q)$, and the Hessian $H(q)$), or even higher number of parameters, plus the two parameters that describe the position. For example, a 2nd order approximation would be

$$\begin{aligned} I(x, y) \approx{}& I(q) + (x - a)I_x(q) + (y - b)I_y(q) \\ &+ \frac{1}{2!}\big[(x - a)^2 I_{xx}(q) + 2(x - a)(y - b)I_{xy}(q) \\ &+ (y - b)^2 I_{yy}(q)\big]. \end{aligned} \tag{1}$$

Another view—closer in spirit to good continuation—is that such approximations should be based on the level set geometry of these functions (Breton et al. 1992; Šára 1995; Koenderink and van Doorn 2002; Caselles et al. 2002). Because the level set geometry of smooth regions is generically *locally parallel* (a fact that follows from the classical existence theorem for ordinary differential equations), we can consider it to be a *visual flow* (Ben-Shahar 2003) and apply the construction from Ben-Shahar and Zucker (2003) developed in the context of texture. By that work, approximately locally parallel structures in the image plane can be characterized at point q by a triple $\{\theta(q), \kappa_T(q), \kappa_N(q)\}$. If we denote by $\hat{\mathbf{E}}_T(x, y)$ the unit length vector field tangent to I's level sets (see Fig. 2), then $\theta(x, y)$ is $\hat{\mathbf{E}}_T$'s orientation relative to a global coordinate frame

$$\hat{\mathbf{E}}_T = (\cos\theta, \sin\theta),$$

and $\kappa_T(x, y)$ and $\kappa_N(x, y)$ are the tangential and normal curvatures of this vector field, i.e., they are the initial rate of change of orientation in the tangential direction $\hat{\mathbf{E}}_T$ and the normal direction $\hat{\mathbf{E}}_N = (-\sin\theta, \cos\theta)$, respectively (Ben-Shahar and Zucker 2003). These two curvatures are derived from the covariant derivative of the field $\hat{\mathbf{E}}_T$ and they are

[1] Assuming that I is smooth seems to undermine one of our goals—that of preserving its singularities. However, at this point, and up to Sect. 5, we care to study the "good" parts of such function, not its singularities. The preservation of singularities will be handled later by the operation of the contextual computation in the relaxation labeling network.

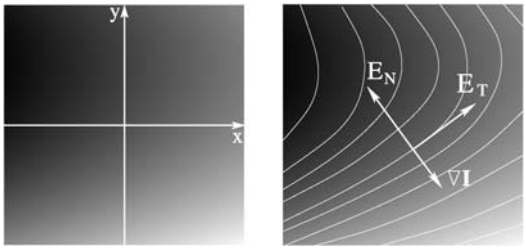

Fig. 2 Although any general 2D scalar function over the image plane can be viewed as a set of values over a global coordinate frame (*left*), we will stress the importance of its level set geometry (*right*) by making this aspect explicit in the representation that we use. Without loss of generality, and for convenience reasons only, this and the rest of the illustrations therein use a shading function as an instance of the abstraction

related to the gradient of the orientation function $\nabla\theta$ via orthogonal expansion (O'Neill 1966), i.e.,

$$\left.\begin{aligned}\kappa_T &= \nabla\theta\cdot(\cos\theta,\sin\theta)\\ \kappa_N &= \nabla\theta\cdot(-\sin\theta,\cos\theta)\end{aligned}\right\} \Leftrightarrow \begin{cases}\theta_x = \kappa_T\cos\theta - \kappa_N\sin\theta\\ \theta_y = \kappa_T\sin\theta + \kappa_N\cos\theta.\end{cases} \tag{2}$$

But $\{\theta(q),\kappa_T(q),\kappa_N(q)\}$ are only three parameters, capable of approximating the local level set geometry *but not* all of $I(x,y)$ in the neighborhood of q. To achieve that goal we must expand the $\{\theta(q),\kappa_T(q),\kappa_N(q)\}$ descriptor and add parameters that map *values* on top of the geometry. We therefore return to the Taylor descriptor, focus on 2nd order approximations (this is the lowest order that captures variations in the level set geometry), and seek an extension of the $\{\theta(q),\kappa_T(q),\kappa_N(q)\}$ descriptor that maps one-to-one with the set $\{I(q),I_x(q),I_y(q),I_{xx}(q),I_{xy}(q),I_{yy}(q)\}$. In other words, we need to expand the set $\{\theta(q),\kappa_T(q),\kappa_N(q)\}$ with other measurements at q such that all of the Taylor coefficients can be computed directly from the new set.

Obviously, $I(q)$ must be included explicitly in our expanded set. Since by definition $\hat{\mathbf{E}}_T$ is tangent to I's level sets, ∇I is parallel to $\hat{\mathbf{E}}_N$ and we arbitrarily select $\hat{\mathbf{E}}_N$ to point away from ∇I (Fig. 2), i.e.,

$$\left.\begin{aligned}\hat{\mathbf{E}}_T &= \tfrac{(I_y,-I_x)}{\|\nabla I\|},\\ \hat{\mathbf{E}}_N &= \tfrac{(-I_x,-I_y)}{\|\nabla I\|}\end{aligned}\right\} \Leftrightarrow \begin{cases}I_x = \|\nabla I\|\sin\theta,\\ I_y = -\|\nabla I\|\cos\theta,\end{cases} \tag{3}$$

which implies that $\|\nabla I\|(q)$ should be incorporated into our descriptor as well. To account for the second derivatives I_{xx},I_{xy},I_{yy} in the Taylor descriptor, we first differentiate I_x and I_y from (3)

$$\begin{aligned}I_{xx} &= \|\nabla I\|\theta_x\cos\theta \;+\sin\theta\tfrac{\partial}{\partial x}\|\nabla I\|,\\ I_{xy} &= \|\nabla I\|\theta_y\cos\theta \;+\sin\theta\tfrac{\partial}{\partial y}\|\nabla I\|,\\ I_{yx} &= \|\nabla I\|\theta_x\sin\theta \;-\cos\theta\tfrac{\partial}{\partial x}\|\nabla I\|,\\ I_{yy} &= \|\nabla I\|\theta_y\sin\theta \;-\cos\theta\tfrac{\partial}{\partial y}\|\nabla I\|.\end{aligned} \tag{4}$$

At first glance, since both $\frac{\partial}{\partial x}\|\nabla I\|$ and $\frac{\partial}{\partial y}\|\nabla I\|$ participate in the description of I's second derivatives, (4) suggests

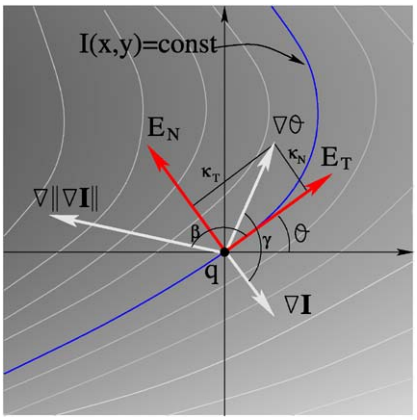

Fig. 3 A geometrical depiction of the parameters underlying the local description of a smooth 2D scalar function up to second order around a point of interest q. Each such point belongs to a certain level set (*in blue*) whose tangent at q defines the local frame $\{\hat{\mathbf{E}}_T,\hat{\mathbf{E}}_N\}$ (*in red*). The orientation gradient $\nabla\theta$ at q defines the direction and rate of the maximum change of level set orientation around q, and its projections on the frame $\{\hat{\mathbf{E}}_T,\hat{\mathbf{E}}_N\}$ defines the tangential and normal curvatures at q

that we need to add both of them as free parameters to our levelsets-centered descriptor, which will result in a total of seven parameters, one more than Taylor's. However, applying the integrability constraint $I_{xy}=I_{yx}$ to (4) translates to

$$\nabla I\cdot\nabla\theta = -\hat{\mathbf{E}}_T\cdot\nabla\|\nabla I\|$$

and removes one degree of freedom: either the direction of $\nabla\|\nabla I\|$ or its magnitude, but not both, should be used. Indeed, if γ is the angle between ∇I and $\nabla\theta$, and β is the angle between $\hat{\mathbf{E}}_T$ and $\nabla\|\nabla I\|$, then the cosine rule dictates that once we set the magnitude of $\nabla\|\nabla I\|$ the angle must satisfy

$$\cos\beta = \frac{\|\nabla I\|\cdot\|\nabla\theta\|\cdot\cos\gamma}{\|(\nabla\|\nabla I\|)\|}.$$

In conclusion, from this analysis we have:

Proposition 1 *With notation as above, the set* $\{I,\|\nabla I\|,\theta,\kappa_T,\kappa_N,\|(\nabla\|\nabla I\|)\|\}$, *evaluated* (*or measured*) *at point* q, *is a* 2*nd-order* levelset-centered *descriptor for* $I(x,y)$ *in the neighborhood of* q.

These parameters are illustrated in Fig. 3. We now use this descriptor to design a local model for $I(x,y)$ in the neighborhood of q such that good continuation of *both* level sets *and* values hold *simultaneously*. Moreover, by employing the specific criteria for "goodness" as discussed in the following sections, we are able to show that the higher order parameter of the intensity structure (i.e., $\|(\nabla\|\nabla I\|)\|$) is not needed, and hence to end up with a local descriptor that has only five parameters: $\{I,\|\nabla I\|,\theta,\kappa_T,\kappa_N\}$.

3 Level Sets Good Continuation Model

A primary goal of this work is to develop a model $\hat{I}(x, y)$ of good continuation based on the descriptor derived above. More specifically, we seek a family of functions $\hat{I}_{\{I,\|\nabla I\|,\theta,\kappa_T,\kappa_N\}}(x, y)$, parametrized by the different measurements at point q, which reflects "good continuation" in the neighborhood of q, both for its level set geometry and for the distribution of its values. Assuming that our model is constructed around an arbitrary point $q = (x_0, y_0)$ where the level sets are oriented with orientation $\theta(q) = \theta_0$, a model that reflects good continuation for the locally parallel structure of the level sets has already been proposed in the literature (see Ben-Shahar and Zucker 2003; Ben-Shahar 2003). According to this model, the local orientation of the level sets at any point (x, y) in the neighborhood of q takes the following general explicit form:

$$\theta_{x_0,y_0,\theta_0,K_T,K_N}(x, y) = \theta_0 + \tan^{-1}\left(\frac{(K_T(\Delta_x)+K_N(\Delta_y))\cos\theta_0+(-K_N(\Delta_x)+K_T(\Delta_y))\sin\theta_0}{1+(K_N(\Delta_x)-K_T(\Delta_y))\cos\theta_0+(K_T(\Delta_x)+K_N(\Delta_y))\sin\theta_0}\right) \tag{5}$$

where $\Delta_x = x - x_0$, $\Delta_y = y - y_0$ and K_T and K_N are the level set curvatures measured at q, i.e., $K_T = \kappa_T(q)$ and $K_N = \kappa_N(q)$. For simplicity, if we assume (without loss of generality) that the model is constructed around the origin $q = (0, 0)$ with the coordinate system aligned with the tangent of the level set at that point, i.e., with $\hat{\mathbf{E}}_T(q) = (1, 0)$ or $\theta(q) = 0$, then the general model in (5) reduces to

$$\theta(x, y) = \tan^{-1}\left(\frac{K_T x + K_N y}{1 + K_N x - K_T y}\right). \tag{6}$$

When viewed as a surface in the space $XY\theta \triangleq \mathbf{R}^2 \times \mathcal{S}^1$ (or, $\mathbf{R}^3$ whose Z axis represents orientation), this function (either the restricted or the general form) is known as a *right helicoid* (Ben-Shahar and Zucker 2003; Ben-Shahar 2003). Several instances of this model, depicted as texture patterns stripped of any value structure, are shown in Fig. 4.

The right helicoidal model possesses several geometrical properties that associate it with good continuation. For example, this model *uniquely* induces an *identical covariation* of the two curvature functions κ_T and κ_N and guarantees that their ratio remains invariant in a neighborhood $N(q)$ of q

$$\frac{\kappa_T(x, y)}{\kappa_N(x, y)} = \text{const} = \frac{K_T}{K_N} \quad \forall (x, y) \in N(q). \tag{7}$$

Therefore, unlike common ways to achieve "good behavior" of 2D (or higher dimensional) structures, typically via anisotropic diffusion and deformable models (e.g., Tang et al. 2000; Sochen et al. 1998), the helicoidal model emerges from explicit considerations of the behavior of curvatures in the image plane and thus it is closer in spirit to the methodology that was employed in the study of good continuation of curves (e.g., Ullman 1976; Parent and Zucker 1989; Mumford 1994; Kimia et al. 1999). As a consequence, this model guarantees that streamlines of the model's flow structure will have neither curvature extrema nor inflection points, both of which are considered significant geometrical events for segmentation and part decomposition (e.g., Richards et al. 1987; Hoffman 1998) and therefore are clearly undesired when perceptual coherence is needed. More importantly, the right helicoidal model is unique in guaranteeing a neighborhood behavior that is devoid of perceptual singularities (Ben-Shahar 2006) and non-smooth percepts that can trigger segregation in completely smooth visual signals (Fig. 5). Obviously, no good continuation model can afford to include such visual events in its pre-

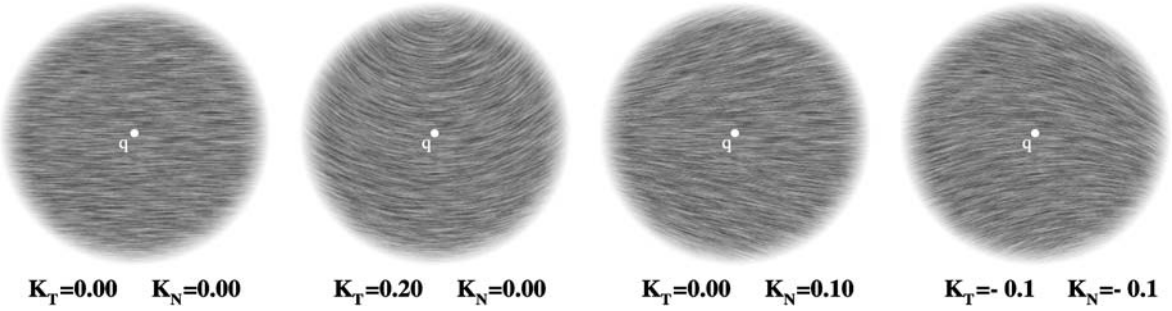

Fig. 4 Stripped of its value (say, *shading*) structure, underlying the model for good continuation $\hat{I}(x, y)$ is the geometry of its level sets, which, as we argue, should have good continuation qualities of its own. For this purpose we utilize an already proposed model for the good continuation of locally parallel structures—a right helicoid in $XY\theta$ (Ben-Shahar and Zucker 2003; Ben-Shahar 2003). Shown here are just several examples of helicoidal level set structures around a central point q having $\theta(q) = 0$ and different curvature values $\kappa_T(q)$ and $\kappa_N(q)$ as indicated. Geometry patterns were generated using the line integral convolution method (Carbal and Leedom 1993)

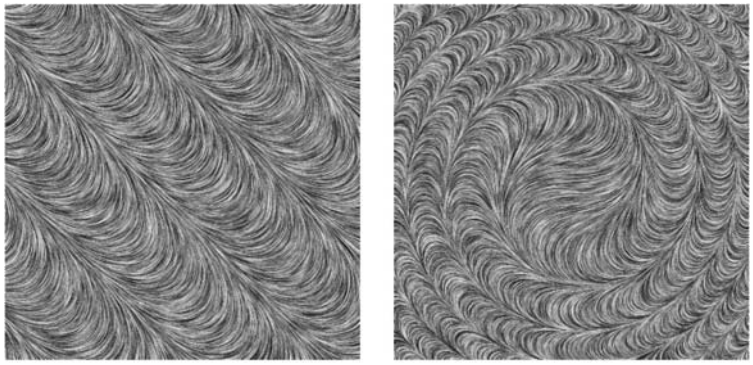

Fig. 5 Smoothness of a visual signal is by no means a sufficient condition for perceptual coherence. Here, for example, perceptual singularities (and therefore the emergence of global structure) occur in orientation defined textures whose dominant orientation function changes smoothly and slowly. (**a**) In this texture, the orientation function changes smoothly with constant orientation gradient and yet most observers report of salient diagonal perceptual singularities that give rise to the segregation of the pattern to diagonal bands, (**b**) Although this texture is completely smooth in terms of its orientation function, the emergence of certain perceptual singularities gives rise to a global structure of a double spiral. Obviously, such perceptual structure and singularities are undesired in a model that reflects good continuation

diction, and the right helicoid is unique[2] in possessing this property.

Finally, although it emerges from explicit good continuation considerations in the image plane, the helicoidal model does enjoy properties that link it to scale space. In particular, this model (as an orientation function) was proved to be both p-harmonic for all values of p and a minimal surface in $XY\theta$, thus exhibiting both properties that have been used extensively in the anisotropic diffusion literature for signal reconstruction that preserves discontinuities (e.g., Caselles et al. 1997b; Tang et al. 2000; Kimmel et al. 2000). For all these reasons we use this object as a starting point for the more general model that we seek for the good continuation of general 2D scalar functions over the image plane.

4 Extended Good Continuation Model for General 2D Functions

The helicoidal model defines only the level set structure of our sought-after model $\hat{I}_{\{I,\|\nabla I\|,\theta,\kappa_T,\kappa_N\}}(x,y)$. One way to proceed toward this general model is to develop a general function $I(x,y)$ whose level set orientation $\theta(x,y)$ obeys the helicoidal orientation function and then explore the constraints that this derivation entails. In principle, such an analysis is straight forward. Indeed, from the identity $\theta = \tan^{-1}\frac{-I_x}{I_y}$ we can derive $\nabla\theta$ in terms of I's derivatives

$$\begin{aligned}\theta_x &= \frac{I_x I_{xy} - I_y I_{xx}}{I_x^2 + I_y^2}\\ \theta_y &= \frac{I_x I_{yy} - I_y I_{xy}}{I_x^2 + I_y^2}\end{aligned} \tag{8}$$

and using (2) and (3) we can therefore express the two curvature functions in terms of I's derivatives

$$\begin{aligned}\kappa_T &= \frac{I_x^2 I_{yy} - 2I_x I_y I_{xy} + I_y^2 I_{xx}}{(I_x^2 + I_y^2)^{3/2}},\\ \kappa_N &= \frac{I_{xy}(I_x^2 + I_y^2) + I_x I_y (I_{xx} - I_{yy})}{(I_x^2 + I_y^2)^{3/2}}.\end{aligned} \tag{9}$$

Then, by applying the constraint in (7) we can therefore obtain a constraint that forces I to have helicoidal level sets

$$\begin{aligned}0 = {} & I_{xx}(K_N I_y^2 - K_T I_x I_y) + I_{xy}(K_T(I_x^2 - I_y^2) - 2I_x I_y)\\ & + I_{yy}(K_N I_x^2 + K_T I_x I_y).\end{aligned} \tag{10}$$

While this second order PDE defines a necessary condition which our sought-after model $\hat{I}$ must satisfy, its nonlinear nature suggests that a closed form solution may be difficult (or impossible) to obtain. A more productive way to generalize the level sets' helicoidal model to a general model

[2] Perceptual singularities in oriented patterns are modeled by the ridges of the scalar ratio PSM $= \frac{\kappa_N^2}{\kappa_T^2+\kappa_N^2}$ (see Ben-Shahar 2006). Since this ratio varies smoothly in the bounded domain [0, 1], the way to prevent it from having ridges is to keep it constant. However, doing so immediately entails the condition in (7), which is characteristic and unique to the helicoid. We do note that this is correct in the strict theoretical meaning, though not always in the practical perceptual sense. In particular, due to the finite sensitivity of the human visual system, it is possible to infinitesimally deform the right helicoidal model to a different model without any perceptual effect but with PSM-predictable perceptual singularities. This finite sensitivity is modelled in Ben-Shahar (2006) by a perceptual thresholding stage, which in practical terms permits some freedom in the selected model.

$\hat{I}_{\{I,\|\nabla I\|,\theta,\kappa_T,\kappa_N\}}(x,y)$ is to integrate the former (which is already parametrized by θ, K_T, and K_N) into a closed form which is parametrized by the newly introduced parameters I and $\|\nabla I\|$. Doing so requires several steps. First we seek a closed form expression for $\hat{I}$'s level sets by solving for the characteristics of a 1^{st} order differential equation derived from the helicoidal model. Second, we derive a non-characteristic Cauchy curve (John 1982) which is used to parametrize the entire solution as a dense collection of level sets. This parametrization is finally used to impose additional good continuation constraints for the derivation of the final desired model. All these steps are discussed next.

4.1 Level Sets Closed Form

Let $I(x,y)$ an arbitrary smooth scalar function over the image plane. The orientation of I's level sets satisfies

$$\theta = \tan^{-1}\left(\frac{-I_x}{I_y}\right) \tag{11}$$

and therefore, for I's level sets to have a certain given orientation function $\theta(x,y)$, it must satisfy the PDE

$$I_x + I_y \tan\theta(x,y) = 0. \tag{12}$$

Now, since $\theta(x,y)$ will be based on the helicoidal model for good continuation of locally parallel structure (Ben-Shahar and Zucker 2003), by plugging (6) into (12) we obtain the following PDE that our model should satisfy

$$(1 + K_N x - K_T y)I_x + (K_T x + K_N y)I_y = 0. \tag{13}$$

The characteristic curves (John 1982) of this PDE must satisfy

$$\begin{aligned}
\frac{dx}{dt} &= 1 + K_N x - K_T y,\\
\frac{dy}{dt} &= K_T x + K_N y,\\
\frac{dI}{dt} &= 0,
\end{aligned}$$

a system whose solution $\alpha(t) = (x(t), y(t), I(t))$ can be written in closed form as

$$\begin{aligned}
x(t) &= e^{tK_N}(c_1\cos(tK_T) - c_2\sin(tK_T)) - \frac{K_N}{\xi^2}\\
y(t) &= e^{tK_N}(c_1\sin(tK_T) + c_2\cos(tK_T)) + \frac{K_T}{\xi^2}\\
I(t) &= c_3,
\end{aligned} \tag{14}$$

where

$$\xi^2 = K_T^2 + K_N^2 \tag{15}$$

and the coefficients c_1, c_2, and c_3, are determined from the initial data

$$\left.\begin{aligned}
x(t=0) &= x_0,\\
y(t=0) &= y_0,\\
I(t=0) &= I_0
\end{aligned}\right\} \quad\Rightarrow\quad \begin{cases}
c_1 = x_0 + \frac{K_N}{\xi^2},\\
c_2 = y_0 - \frac{K_T}{\xi^2},\\
c_3 = I_0.
\end{cases} \tag{16}$$

Given any point in the image plane, we are now able to trace its level sets via (14) and (16). The fact that these level sets could be described in closed form should not be taken for granted. In fact, many level set orientation functions $\theta(x,y)$, including functions much simpler than the helicoidal one used here, would give rise to nonlinear characteristic PDEs with no closed form solution. The fact that we are able to do so with a function that was shown to have good continuation qualities is a property that will prove most valuable in the derivation of a closed form expression for $\hat{I}(x,y)$ as a whole.

4.2 Solution Parametrization via Level Sets

Since (14) and (16) transform point initial data to a curve, they will turn initial data along a curve into a surface patch in $\mathbf{R}^3$. Let $\Gamma(s) = (x_0(s), y_0(s), I_0(s))$ be such an initial Cauchy data parametrized by s and note that $\Gamma(s)$ can be interpreted as a parametrized family of initial data *points*, each having the form shown in the left part of (16). As long as Γ is not a characteristic curve of (14), the substitution of Γ into (14) and (16) yields the following solution surface which is parametrized by s and t:

$$\begin{aligned}
&x(s,t)\\
&\quad = \tfrac{-K_N + e^{tK_N}[(x_0(s)\xi^2 + K_N)\cos(tK_T) + (K_T - y_0(s)\xi^2)\sin(tK_T)]}{\xi^2},\\
&y(s,t)\\
&\quad = \tfrac{K_T + e^{tK_N}[(x_0(s)\xi^2 + K_N)\sin(tK_T) + (y_0(s)\xi^2 - K_T)\cos(tK_T)]}{\xi^2},\\
&I(s,t) = I_0(s).
\end{aligned} \tag{17}$$

Given the closed form of the solution surface in (17), it is left to find a "legal" (i.e., non characteristic) initial curve $\Gamma(s)$. To separate the "spatial" component $(x_0(s), y_0(s))$ of Γ from its "*value*" (e.g., shading) dimension $I_0(s)$, we first consider the projection of $\Gamma(s)$ on the image plane, i.e., the curve $\Gamma_{xy}(s) = (x_0(s), y_0(s))$. One safe course of action is to select a $\Gamma_{xy}(s)$ that is everywhere *orthogonal* to the characteristic curves of (14). If these latter characteristics are computed by tracing the direction tangential to the level sets, i.e., the direction of $\hat{\mathbf{E}}_T$, constructing a curve that is every-

where orthogonal will be based on integration along $\hat{\mathbf{E}}_N$. By repeating the process from Sect. 4.1 for the system

$$\frac{dx}{ds} = -K_T x - K_N y,$$
$$\frac{dy}{ds} = 1 + K_N x - K_T y$$

and the initial data $x_0 = 0$ and $y_0 = 0$, we yield the following solution for $\Gamma_{xy}(s)$

$$x_0(s) = \frac{-K_N + e^{-sK_T}(K_N \cos(sK_N) + K_T \sin(sK_N))}{\xi^2},$$
$$y_0(s) = \frac{K_T + e^{-sK_T}(-K_T \cos(sK_N) + K_N \sin(sK_N))}{\xi^2}. \tag{18}$$

Hence, substituting (18) into (17) results with the following expression for our sought-after model

$$\begin{aligned} x(s,t) &= \tfrac{-K_N + e^{(tK_N - sK_T)}[K_N \cos\tilde{K}(s,t) + K_T \sin\tilde{K}(s,t)]}{\xi^2}, \\ y(s,t) &= \tfrac{K_T + e^{(tK_N - sK_T)}[K_N \sin\tilde{K}(s,t) - K_T \cos\tilde{K}(s,t)]}{\xi^2}, \\ I(s,t) &= I_0(s) \end{aligned} \tag{19}$$

where $\tilde{K}(s,t) = sK_N + tK_T$.

4.3 Cartesian *Re*-parametrization

Equation (19) is a key result of special importance to this paper. It provides an explicit parametrization of the desired good continuation model based on its desired level set structure (i.e., the helicoidal good continuation model). It therefore parametrizes the sought-after model $\hat{I}_{\{I,\|\nabla I\|,\theta,\kappa_T,\kappa_N\}}(x,y)$ around the point q (which, without loss of generality, can be considered to be the origin) as a collection of integral curves that emerge from the generator curve $\Gamma_{xy}(s)$ (see Fig. 6). What is left to consider is the distribution of *values* $I_0(s)$ along $\Gamma_{xy}(s)$ from which we can complete the definition of $\Gamma(s)$ and the solution surface as a whole. Since nothing in the level set geometry can constrain this distribution of values, we will need to incorporate additional good continuation constraints, as discussed below. However, before doing so, one needs to realize that the parametrization provided by (19) is not Cartesian. This is a major limitation because eventually the good continuation model should be used with images provided in Cartesian coordinates. In this section we ask if this parametrization gap can be bridged. In particular, we ask if it is possible to re-parametrize (19) based on Cartesian coordinates, and whether or not that can be done in closed form. As we show next, all these questions can be answered in the affirmative despite the non nontrivial form of (19).

Formally, we seek a reparametrization $s = \tilde{s}(x,y)$ and $t = \tilde{t}(x,y)$ such that, when applied to s and t in (19) makes

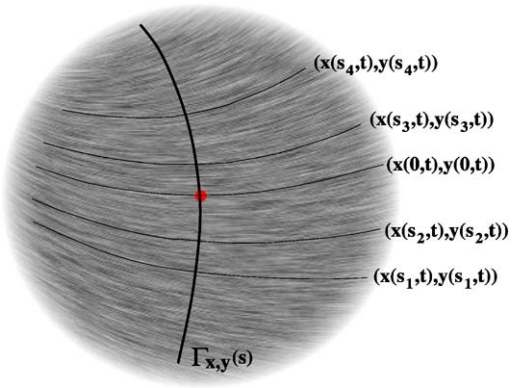

Fig. 6 Equation (19) parametrizes the plane as a collection of curves, each of which is a projection of a level set curve on the image plane. Shown here is the generator curve $\Gamma_{xy}(s)$ and several of the generated level set curves, superimposed on a helicoidal pattern of $K_T = 0.1$ and $K_N = 0.1$

both $x(\tilde{s}(x,y),\tilde{t}(x,y))$ and $y(\tilde{s}(x,y),\tilde{t}(x,y))$ the identity transformation

$$\left.\begin{aligned} s &= \tilde{s}(x,y), \\ t &= \tilde{t}(x,y) \end{aligned}\right\} \quad \Rightarrow \quad \left\{\begin{aligned} x(\tilde{s}(x,y),\tilde{t}(x,y)) &= x, \\ y(\tilde{s}(x,y),\tilde{t}(x,y)) &= y. \end{aligned}\right. \tag{20}$$

While here we skip most of the tedious algebraic manipulations, it can be verified that the transformation from (s,t) to (x,y) defined by (19) is reversible and that every coordinate (x,y) can be traced back to the pair (s,t) with the following transformation

$$\begin{aligned} \tilde{s}(x,y) &= \tfrac{\mathrm{sign}(K_T x + K_N y) K_N \cos^{-1}\omega(x,y) - K_T \log\mu(x,y)}{\xi^2} \\ \tilde{t}(x,y) &= \tfrac{\mathrm{sign}(K_T x + K_N y) K_T \cos^{-1}\omega(x,y) + K_N \log\mu(x,y)}{\xi^2}, \end{aligned} \tag{21}$$

where

$$\begin{aligned} \mu(x,y) &= \sqrt{1 + 2K_N x - 2K_T y + \xi^2(x^2+y^2)} \\ \omega(x,y) &= \frac{1 + K_N x - K_T y}{\mu(x,y)}. \end{aligned} \tag{22}$$

It can be shown that the expression under the square root is never negative, and therefore that $\mu(x,y)$ is always real valued and that the logarithm in (21) is always well defined. $\mu(x,y)$ vanishes at a single point $(\frac{-K_N}{\xi^2}, \frac{K_T}{\xi^2})$ which is the same singular point of the helicoidal model as a whole.

Having (21), and an explicit way to switch back and forth between the levelset-centered and the Cartesian parametrization, we are now able to plug $\tilde{s}(x,y)$ to the $I(s,t)$ component of (19) to obtain an explicit description of our desired model in terms of Cartesian coordinates

$$I(x,y) = I_0\big(\tfrac{\mathrm{sign}(K_T x + K_N y) K_N \cos^{-1}\omega(x,y) - K_T \log\mu(x,y)}{\xi^2}\big). \tag{23}$$

Equation (23) is one step short of a desired final model since it still includes the unspecified function $I_0(\cdot)$ to be

evaluated at the result of $\tilde{s}(x, y)$ for each coordinate (x, y). What is left to do is to devise the function $I_0(s)$ in a way that makes the *value* structure of the resultant $I(x, y)$ satisfy certain good continuation properties (i.e., in additional to the already guaranteed good continuation of its level sets).

4.4 Constraints Toward a Final Model

Equation (23), and the function $I_0(s)$ that defines it, represent the backbone on which any scalar function with helicoidal level sets can be constructed in Cartesian coordinates. But of all these functions, which one would best suit good continuation?

While this question was never addressed explicitly from a perceptual organization point of view, the notion of "nice" functions is of course ubiquitous in image analysis, especially for denoising and scale space analysis. Typically, the desired behavior in this context is achieved by low pass filtering (Jain 1989) or by minimization of some variation measure through anisotropic diffusion (Sochen et al. 1998; Tang et al. 2000). In this sense, the helicoidal model for the level sets already achieves that goal for the *geometry* underlying $I(x, y)$, as it was shown both to have a vanishing p-Laplacian (a measure used e.g., in Tang et al. 2000) for all values of p, and to be a minimal surface (used e.g., in Caselles et al. 1997a) in $XY\theta$ (Ben-Shahar and Zucker 2003). It is therefore natural to examine these measures also for the value structure of our sought-after model and attempt to derive a model $\hat{I}(x, y)$ that has the same good continuation properties for both its value and level set structures.

Fortunately, (23) allows to formally explore the last proposal in a straight forward way. Generally speaking, we seek to translate a desired differential constraint $\mathcal{F}$ on $I(x, y)$ to a differential constraint $\mathcal{G}$ on I_0,

$$\begin{aligned}&\mathcal{F}(I, I_x, I_y, I_{xx}, I_{xy}, I_{yy}; x, y) = 0\\ &\quad\Rightarrow\quad \mathcal{G}\left(I_0, \frac{\partial I_0}{\partial s}, \frac{\partial^2 I_0}{\partial s^2}; s\right) = 0,\end{aligned} \tag{24}$$

then solve $\mathcal{G}$ for $I_0(s)$ (if solvable) and plug back into (23) to obtain the final explicit solution.

Equation (23) can be differentiated for an arbitrary differentiable $I_0(s)$ using the chain rule

$$\begin{aligned}I_x &= \frac{\partial I_0}{\partial \tilde{s}}\frac{\partial \tilde{s}}{\partial x},\\ I_y &= \frac{\partial I_0}{\partial \tilde{s}}\frac{\partial \tilde{s}}{\partial y},\end{aligned}\qquad\begin{aligned}I_{xx} &= \frac{\partial^2 I_0}{\partial \tilde{s}^2}\left(\frac{\partial \tilde{s}}{\partial x}\right)^2 + \frac{\partial I_0}{\partial \tilde{s}}\frac{\partial^2 \tilde{s}}{\partial x^2},\\ I_{xy} &= \frac{\partial^2 I_0}{\partial \tilde{s}^2}\frac{\partial \tilde{s}}{\partial x}\frac{\partial \tilde{s}}{\partial y} + \frac{\partial I_0}{\partial \tilde{s}}\frac{\partial^2 \tilde{s}}{\partial x \partial y},\\ I_{yy} &= \frac{\partial^2 I_0}{\partial \tilde{s}^2}\left(\frac{\partial \tilde{s}}{\partial y}\right)^2 + \frac{\partial I_0}{\partial \tilde{s}}\frac{\partial^2 \tilde{s}}{\partial y^2}\end{aligned} \tag{25}$$

and the derivatives of $\tilde{s}(x, y)$ are readily computed from (21) (with some special care involving the sign function):

$$\begin{aligned}\frac{\partial \tilde{s}}{\partial x} &= -\frac{K_T x + K_N y}{\mu^2(x, y)},\\ \frac{\partial \tilde{s}}{\partial y} &= \frac{1 + K_N x - K_T y}{\mu^2(x, y)},\\ \frac{\partial^2 \tilde{s}}{\partial x^2} &= \frac{\mu^2(x, y)K_T - 2(\xi^2 x + K_N)(K_T x + K_N y)}{-\mu^4(x, y)},\\ \frac{\partial^2 \tilde{s}}{\partial x \partial y} &= \frac{\mu^2(x, y)K_N - 2(\xi^2 y - K_T)(K_T x + K_N y)}{-\mu^4(x, y)},\\ \frac{\partial^2 \tilde{s}}{\partial y^2} &= \frac{-\mu^2(x, y)K_T - 2(\xi^2 y - K_T)(1 + K_N x - K_T y)}{-\mu^4(x, y)}.\end{aligned} \tag{26}$$

Having these last two sets of equations we are able to impose any desired (second order) differential constraint $\mathcal{F}$ and examine the transformation to $\mathcal{G}$.

4.5 Derivation of a "Dual Harmonic" Model

One sought-after property in the context of "good" or desired behavior of functions is the harmonic map constraint, a property which in its most basic form is related to heat diffusion and scale space analysis, and to applications ranging from denoising, through segmentation, to image inpainting and curve completion (to name but a few). Imposing the harmonic constraint on an image $I(x, y)$ means asking $I(x, y)$ to have a vanishing Laplacian, or requiring $I(x, y)$ to be a critical point of the harmonic energy $E = \int\int \|\nabla I\|^2 dx dy$ (which provides an intuitive appeal for its "goodness"). Remarkably, doing so in the context of (24) results in the following transformation

$$\mathcal{F} = \nabla \cdot \nabla I = I_{xx} + I_{yy} = 0 \quad\Rightarrow\quad \mathcal{G} = \frac{I_0''(s)}{\mu^2} = 0$$

and since μ never vanishes in the helicoidal domain of definition, it is safe to determine that under the harmonic constraint $I_0(s)$ must have a vanishing second derivative and therefore it must take the form

$$I_0(s) = I_0'(0)s + I_0(0).$$

Combining this form with (23) gives rise to the following result:

Proposition 2 (Dual harmonic good continuation model) *A general scalar good continuation model that is harmonic in both its value and level sets structures takes the form*

$$\hat{I}(x,y) = I_0(0) + \frac{I_0'(0)}{\xi^2}\left[\text{sign}(K_T x + K_N y) K_N \cos^{-1}\omega(x,y) - K_T \log\mu(x,y)\right] \tag{27}$$

where $\mu(x,y)$, $\omega(x,y)$, *and* ξ^2 *are all as defined in previous sections.*

It should be noted that although (27) appears ill-defined for $\xi^2 = 0$ (i.e., for $K_T = K_N = 0$), the limit actually converges to the following simple function

$$\lim_{\substack{K_T\to 0\\ K_N\to 0}} \hat{I}(x,y) = I_0(0) + I_0'(0)y.$$

The non-trivial expression in (27) represents a most unique object. It is a scalar function whose level set geometry (i.e., level set orientation function) *and* its value structure are both harmonic functions at the same time in their respective domains of definition! This "dual harmonic" model therefore complies both with the helicoidal good continuation model and the criterion commonly used to denote "good" scalar functions. Since $I_0(0) = I(0,0)$, and $I_0'(0) = \|\nabla I\|(0,0)$, this model is based on those levelset-centered local features which we discussed in Sect. 2, namely on the set $\{I, \|\nabla I\|, \theta, \kappa_T, \kappa_N\}$ as measured at the origin (or more generally, at a general point q). Note that this result also indicates that the higher order intensity parameter from our full 2nd-order, levelset-centered descriptor (i.e., the parameter $\|(\nabla\|\nabla I\|)\|$ from Proposition 1) is not needed for our good continuation model, hence reducing the descriptor's "complexity" to 5 parameters only. Figure 7 shows several instances of this model, for different values of curvature pairs, depicted as intensity patches with and without its level sets emphasized.

4.6 Derivation of a "Dual p-Harmonic" Model

Interestingly enough, the surprising result from Proposition 2 can be generalized even further. Indeed, consider the more general requirement which constrains $I(x,y)$ to have a vanishing *p-Laplacian* for some arbitrary value of p. This constraint (for values of p other than 2) has been used in the literature to obtain anisotropic diffusion processes with edge preserving behavior (e.g., Tang et al. 2000) by making $I(x,y)$ a critical point of the p-harmonic energy functional $E_p = \iint \|\nabla I\|^p dxdy$. In our case, representing this constraint in terms of (24) via the definition of the p-Laplacian (or the Euler-Lagrange equation of E_p) yields

$$\mathcal{F} = \Delta_p I = \nabla\cdot(\|\nabla I\|^{p-2}\nabla I) = 0, \tag{28}$$

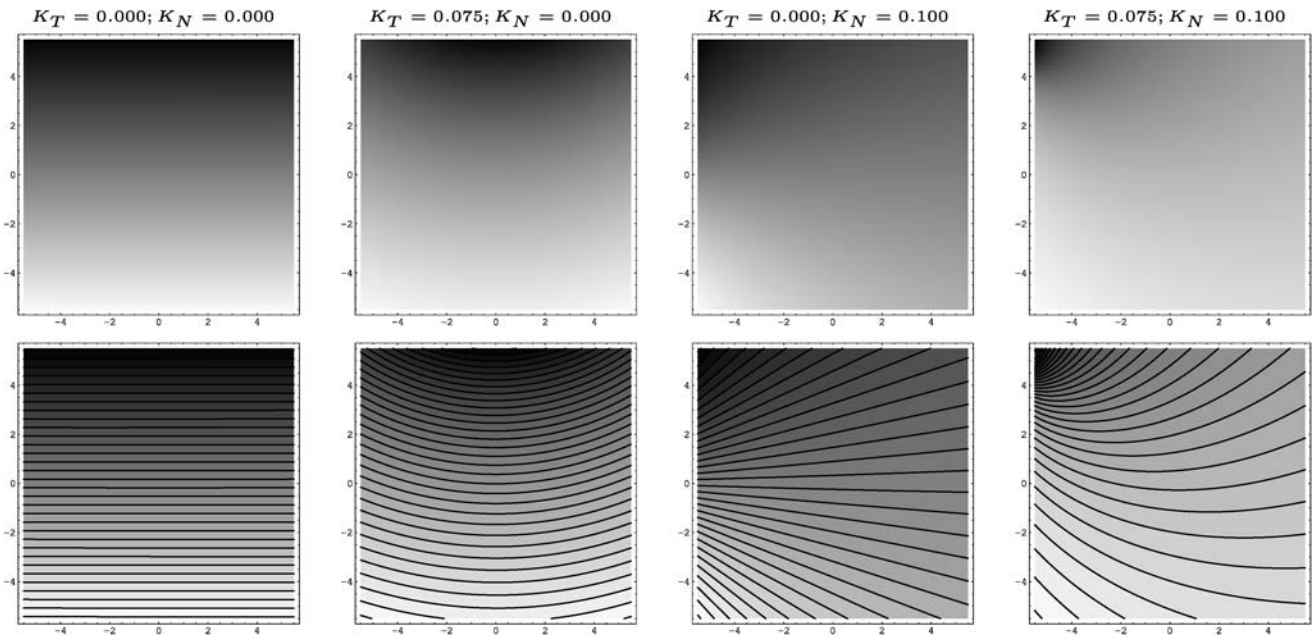

Fig. 7 Several instances of the dual harmonic good continuation model of (27), both as a shading patch and as a set of level sets (coarsely quantized for clarity of display). These instances correspond to the descriptor set defined by $I(0,0) = 0.5$, $\|\nabla I\|(0,0) = 0.25$, $\theta(0,0) = 0$, and the curvature pairs specified over the patches. (The origin $q = (0,0)$ is at the center of the plots)

or more explicitly,

$$\mathcal{F} = \|\nabla I\|^{p-4}\big[(p-2)(I_x^2 I_{xx} + 2I_x I_y + I_y^2 I_{yy}) + (I_x^2 + I_y^2)(I_{xx} + I_{yy})\big] = 0. \quad (29)$$

With the transformation from $\mathcal{F} = 0$ in (29) to $\mathcal{G} = 0$ as described in (24)–(26), we obtain

$$\mathcal{G} = I_0'(s)^{p-2}\,\frac{(p-2)K_T I_0'(s) + (p-1)I_0''(s)}{\mu^p(x,y)} = 0. \quad (30)$$

Since μ is always positive in the helicoidal domain of definition (Sect. 4.3), and since the singular condition of $I_0'(s) = 0$ implies an accidental case of constant intensity structure (where no level set geometry exist in the first place), (30) is satisfied only when

$$(p-2)K_T I_0'(s) + (p-1)I_0''(s) = 0. \quad (31)$$

It is easy to see that when $p = 2$ this constraint reduces back to $I_0''(s) = 0$ and therefore to the solution obtained in the previous section. When $p = 1$ we again obtain the singular case of a vanishing $I_0'(s)$ and constant $I_0(s)$. In all other cases, however, i.e., when $p \neq 2$ and $p \neq 1$, (31) is a second order linear ODE whose solution can be expressed in the following closed analytical form

$$I_0(s) = I_0(0) + \frac{I_0'(0)}{\eta_p}\left[1 - e^{-\eta_p s}\right], \quad (32)$$

where η_p is a constant that depends on p and defined as follows

$$\eta_p \triangleq K_T \cdot \frac{p-2}{p-1},$$

and $I_0(0)$ and $I_0'(0)$ are the integration constants which represent the initial conditions at the origin.

Combining the solution in (32) with the form of the final model in (23), and remembering that the level set geometry of our model is already p-harmonic for any value of p (a property of the helicoidal model, Ben-Shahar and Zucker 2003), we get the following result:

Proposition 3 (Dual p-harmonic good continuation model) *A general scalar good continuation model that is p-harmonic in both its value and level sets structures takes the form*

$$\hat{I}_p(x,y) = I_0(0) + \frac{I_0'(0)}{\eta_p}\left[1 - e^{-\frac{\eta_p}{2\xi^2}\left(2\cdot\mathrm{sign}(K_T x + K_N y)K_N \cos^{-1}\omega(x,y) + K_T \log\mu(x,y)\right)}\right] \quad (33)$$

where $\mu(x,y)$, $\omega(x,y)$, ξ^2, and η_p are all as defined previously, and $p \notin \{1,2\}$.

Again, it can be shown that although (33) appears ill-defined for $\xi^2 = 0$ (i.e., for $K_T = K_N = 0$), the limit converges again to the same function derived in the harmonic case, i.e.,

$$\lim_{\substack{K_T \to 0 \\ K_N \to 0}} \hat{I}_p(x,y) = I_0(0) + I_0'(0)y.$$

Not unlike the expression in Proposition 2, (33) represents an object of remarkable property—it is a scalar function whose level set geometry (i.e., level set orientation function) *and* its value structure are simultaneous p-harmonic functions in their respective domains of definition! Since $I_0(0) = I(0,0)$, and $I_0'(0) = \|\nabla I\|(0,0)$, this model, by construction, is based on those levelset-centered local features which we discussed in Sect. 2, namely on the descriptor set $\{I, \|\nabla I\|, \theta, \kappa_T, \kappa_N\}$ as measured at the origin (or more generally, at a point of interest q). Figure 8 shows one instance of this model for the same descriptor set and several selected values of p.

4.7 Can a "Dual-Minimal" Model Exist?

In addition to being p-harmonic, another property that has been used frequently in the literature in the context of "good" or "desired" behavior is *minimality of the surface representation* of the features of interest (e.g., Caselles et al. 1997b; Kimmel et al. 2000; Ben-Shahar and Zucker 2003; Savadjiev et al. 2007). Indeed, minimal surfaces, the mathematical objects which minimize area (or tension) under certain boundary conditions, have numerous geometrical properties (Dierkes et al. 1992) that motivated their use as a model for many physical phenomena. Unique among the minimal surfaces is the helicoid, which is the only non-planar minimal surface that is also harmonic at the same time (Hamel 1923; Graustein 1940; Nitsche 1989). Since we already use the helicoid as a good continuation model for the geometry of the level sets (after Ben-Shahar and Zucker 2003), it is tempting to consider the possibility of a general model which is "dual minimal", i.e., a model which is a helicoid (and hence a minimal surface) in its level set geometry *and* at the same time is also a minimal surface in its value structure. Unfortunately, the answer with regard to the existence of such an object is negative:

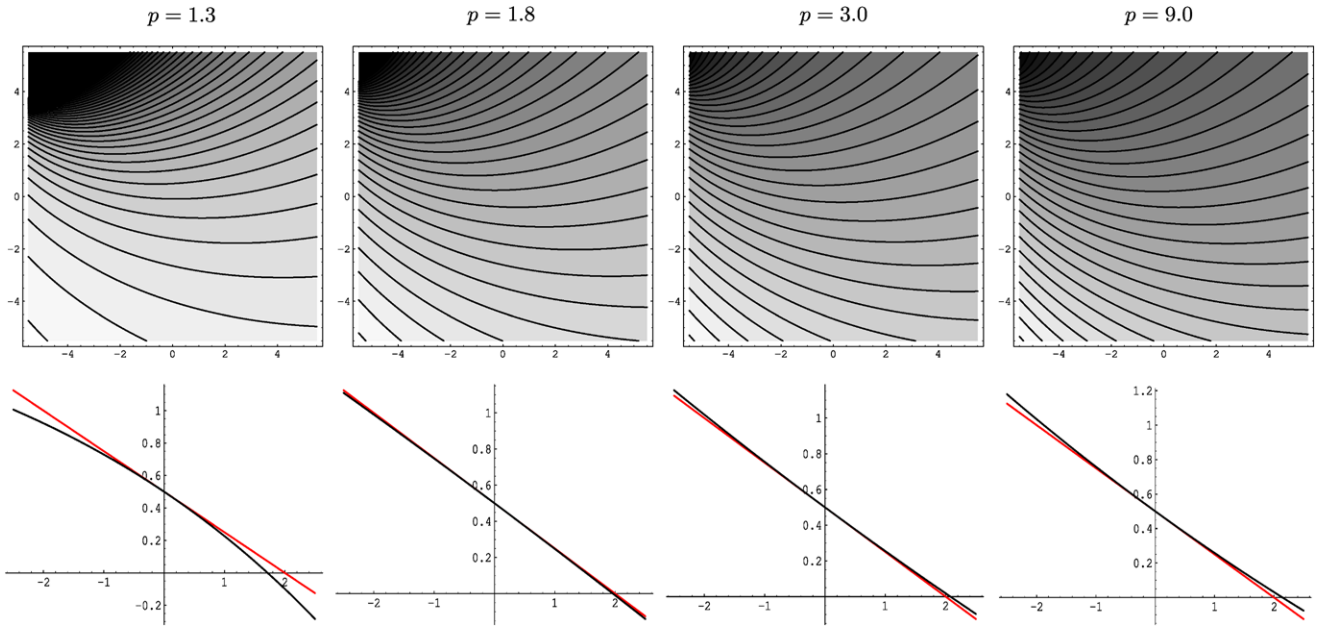

Fig. 8 One example of the dual p-harmonic good continuation model of Proposition 3, for several different values of p but the same descriptor set $\{I, \|\nabla I\|, \theta, \kappa_T, \kappa_N\}$ at $q = (0, 0)$. In particular, these instances correspond to $I(0,0) = 0.5$, $\|\nabla I\|(0,0) = 0.25$, $\theta(0,0) = 0$, $K_T = 0.075$, and $K_N = 0.100$. Several level sets are quantized and emphasized for clarity and to demonstrate how changing p affects the intensity structure only rather than the underlying level set geometry. Compare to the qualitative changes obtained by changing the curvatures (e.g., in Fig. 8). The *bottom graphs* show in black the intensity profile $I_0(s)$ as a function of arc length along the generator curve $\Gamma_{xy}(s)$ (see Fig. 6), compared to the intensity profile in the harmonic case ($p = 2$, shown in *red*)

Proposition 4 *No general good continuation model exists such that its level set geometry is helicoidal* (*and therefore a minimal surface in* $XY\theta$) *while its value structure is also a minimal surface in* its *respective domain*.

To prove this proposition, one needs to consider the minimal surface constraint (i.e., a vanishing Mean curvature H) and apply to $I(x, y)$ as a constraint $\mathcal{F} = 0$ (cf., (24), i.e.,

$$\mathcal{F} = H = I_{xx}(I_y^2 + 1) - 2I_x I_y I_{xy} + I_{yy}(I_x^2 + 1) = 0.$$

In order to rephrase this constraint as $\mathcal{G} = 0$ as described in (24), one can first use (25) and (26) to obtain

$$\mathcal{G} = \mu^2(x, y) I_0''(s) - K_T I_0'(s)^3 = 0. \tag{34}$$

Since the function $\mu(x, y)$ itself can be re-parametrized in terms of the level set parameters s and t

$$\mu^2(x, y) = e^{2K_N t - 2K_T s}.$$

Equation (34) can be reduced to the following constraint on $I_0(s)$ and its derivatives

$$I_0''(s) \cdot e^{2K_N t} = I_0'(s)^3 \cdot K_T \cdot e^{2K_T s}. \tag{35}$$

Equation (35) is bad news since it depends on the level set parameter t (the one that changes *along* level sets) which does not constitute a parameter of $I_0(s)$. In other words, in order for the mean curvature constraint to hold, the solution must change along level sets, which is a contradiction to the very definition of the level set notion.

It should be emphasized that Proposition 4 must be considered with care. In particular, it should *not* be understood as if no "dual minimal" model can exist at all. The negative result in this proposition it true only when the minimal model used for the level set geometry is a right helicoid in $XY\theta$. However, it may be possible that a "dual-minimal" model does exist for level set geometry that defines a different minimal surface in $XY\theta$. While this open theoretical question may be interesting in other contexts, it is outside the scope of this paper because no other minimal surface satisfies the main good continuation constraints that are so needed in our case.

5 Contextual Inference of Coherent Structure

Having a model for the local behavior of "good" scalar features provides the ability to assess the degree to which a particular measurement at one point 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 several different frameworks in which one can pursue this

task while maximizing some measure of global coherency over a domain of interest, including, for example, *relaxation labeling* (Hummel and Zucker 1983; Pelillo 1997), *recurrent neural networks* (Hopfield and Tank 1985), and *belief propagation networks* (Pearl 1988). While the particular choice is independent from our main theoretical contribution, here we choose to use relaxation labeling while noting that other approaches could be used in a similar fashion. Relaxation labeling also offers an intrinsic ability to preserve singularities and boundaries in the signal (see Ben-Shahar and Zucker 2003), hence to facilitate the two goals discussed in the Introduction (i.e., preservation of both regular and singular structure).

5.1 Essential Elements of Relaxation Labeling

Following early studies on the interpretation of line drawings (Waltz 1975), it has become widely accepted that interpretation of sensory data is highly unreliable unless it involves contextual constraints. Relaxation labeling (Hummel and Zucker 1983; Kittler and Illingworh 1985; Pelillo 1997) is a formal approach for providing a contextual computational framework by representing the interpretation problem as the assignment of labels to nodes in a graph (a.k.a. network) whose edges represent the contextual structure.

Let $I = \{i \mid i = 1, \ldots, n\}$ be a set of nodes, each of which may take any label λ from the set Λ. Let $p_i(\lambda)$ denote the probability, or confidence, in the assignment of label λ to node i. In general, $p_i(\lambda)$ must sum to 1 at each node, i.e., $\sum_{\lambda\in\Lambda} p_i(\lambda) = 1\ \forall i$. and the space of all these assignments is denoted $\mathbb{K}$. The role of relaxation labeling is to start from a given (typically inconsistent) label assignments, and to employ contextual constraints in order to obtain optimally consistent one. The fundamental mechanism by which this is done in the relaxation labeling network is a *compatibility function* $r_{ij}(\lambda, \lambda')$ which quantifies the contextual information conveyed by label λ' at node j about label λ at node i. Having such a compatibility, one can defined the *contextual support* that label λ at node i receives from its neighborhood

$$s_i(\lambda) = \sum_{j=1}^{n} \sum_{\lambda'=1}^{m} r_{ij}(\lambda, \lambda') p_j(\lambda'), \tag{36}$$

which can be viewed as a sum of all neighborhood confidences weighted by the compatibilities. Using this machinery, relaxation labeling maps initial inconsistent labeling to consistent ones via an iterative process. Although different iterations have been proposed in the literature, here we follow the one by Hummel and Zucker (1983):

$$p_i^{t+1}(\lambda) \leftarrow \Pi_{\mathbb{K}}\left[p_i^t(\lambda) + \delta s_i^t(\lambda)\right], \tag{37}$$

where $\Pi_{\mathbb{K}}$ is an operator that projects its argument onto $\mathbb{K}$, and δ is a constant step size. A fundamental result from the theory of relaxation labeling states that this algorithm converges to a consistent labeling (Hummel and Zucker 1983) related to the Nash equilibria of the polymatrix game (Miller and Zucker 1999). Please refer to the aforementioned references for additional information.

5.2 A Relaxation Network for the Inference of Coherent 2D Scalar General Features

A direct abstraction of the relaxation process for the inference of coherent 2D image structure involves a 2D image-based network of nodes $i = (x, y)$ (i.e., pixels) whose labels are drawn from the set

$$\begin{aligned}\Lambda = \big\{&(I, \|\nabla I\|, \theta, K_T, K_N) \mid I \in [0, 1],\\ &\|\nabla I\| \in [0, \nabla_{\max}], \theta \in [-\pi, \pi),\\ &\kappa_T, \kappa_N \in [-K_{\max}, K_{\max}]\big\}.\end{aligned} \tag{38}$$

Ultimately, the desired relaxation labeling network would explore this entire *continuous*-labeling space. Although an ad hoc approach for continuous-labeling relaxation labeling has been proposed in Duncan and Birkholzer (1992), in this paper we follow the traditional discrete-labeling scheme for its rigorous theoretical background and convergence results (e.g., Hummel and Zucker 1983; Pelillo 1997; Torsello and Pelillo 2000). Hence, for the rest of this section we assume that each of the label coordinates in Λ has been quantized appropriately at the desired resolution.

As required by the relaxation process, a preliminary measurement procedure assigns an initial confidence value, or probability $p_i^0(\lambda)$, to each possible label such that at each node $\sum_{\lambda\in\Lambda} p_i^0(\lambda) = 1$. The relaxation process itself drives this initial confidence distribution $p_i^0(\lambda)$ to a final (more consistent but possibly ambiguous) distribution $p_i^\infty(\lambda)$, as discussed in Sect. 5.1. 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 labels from the set defined in (38). Following the bulk of this paper, we derive these compatibilities based on the model for consistency developed above, and in particular, for the rest of paper we will focus on the dual-harmonic model from Proposition 2.

Measurement quantization dictates that every possible node i represents an *equivalence class* C_i of measurements, each of which induces a dual-harmonic field of compatible labels in the neighborhood of i. In the continuum, the union of all these dual-harmonic fields that correspond to points in C_i forms a consistent seven dimensional hypersurface, or "7D volume" V_i. It is important to mention that since each node i corresponds to a different set of position, value, gradient magnitude, orientation, and curvatures, the shape of the

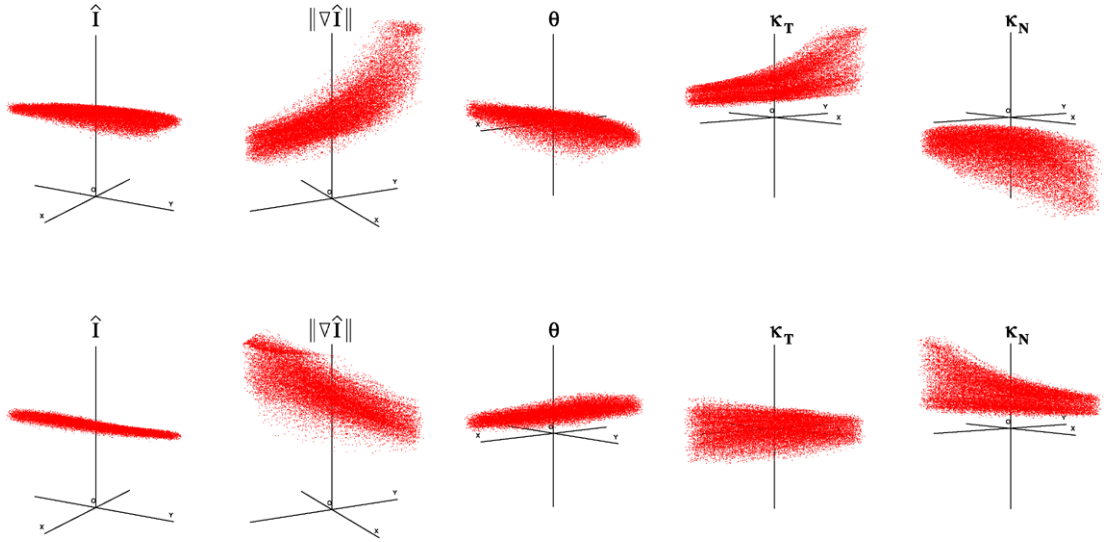

Fig. 9 A visualization of the 7D consistency volume corresponding to two particular values of the descriptor $\{I, \|\nabla I\|, \theta, \kappa_T, \kappa_N\}$ at the origin. Here we show this volume via its projections on the 3D subspaces XYI, $XY\|\nabla I\|$, $XY\theta$, $XY\kappa_T$, and $XY\kappa_N$. *Top*: $(I, \|\nabla I\|, \theta, \kappa_T, \kappa_N) = (0.5, 2.4, 0°, 0.1, -0.1)$. *Bottom*: $(I, \|\nabla I\|, \theta, \kappa_T, \kappa_N) = (0.5, 3.6, 45°, 0.0, 0.1)$

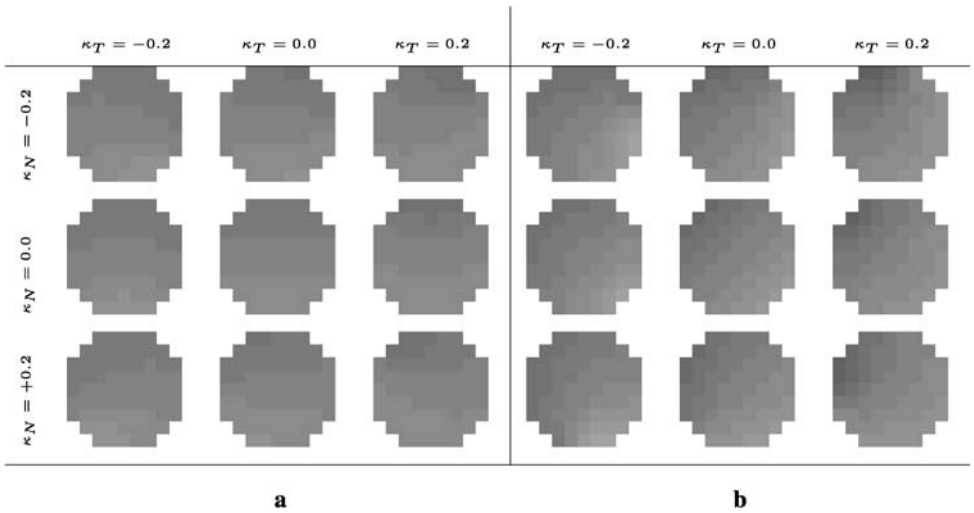

Fig. 10 A sample of 9×9 circular compatibility fields for $x, y \in [-4.5, 4.5]$, $I = 0.5$, two different combinations of $\|\nabla I\|$ and θ, and various combinations of the curvatures κ_T and κ_N. Each plot uses intensity to code the compatible values in the neighborhood of the central pixel. (Please use the electronic version of the paper if intensity variations are not reproduced well in print.) Due to quantization, especially that of curvatures, a given label at the center may be compatible with more than one label at a same nearby location in its neighborhood, an outcome that gives rise to the *volume* in Fig. 9. Since this multiplicity of compatible values cannot be depicted when coded with gray levels, the fields plotted here show only the *most compatible value* at each neighborhood position. Note how higher curvature values introduce more variations into the fields. (**a**) $I = 0.5$, $\|\nabla I\| = 1$, and $\theta = 0°$. (**b**) $I = 0.5$, $\|\nabla I\| = 3$, and $\theta = 45°$

7D consistency volume V_i around i will vary with i (Fig. 9). After quantization, this volume results in a set of consistent labels at nodes in the neighborhood of i. A sample of compatibility fields for two combinations of $\|\nabla I\|$ and θ and a variety of curvature pairs are illustrated in Fig. 10 using the most compatible value in the neighborhood of a given node. Note the influence of the different parameters on the pattern of compatible values.

While the consistency volume V_i determines *which* of the nodes are compatible with a given node i, it does not determine the *values* of r_{ij} which represent the likelihood that node j shares the same local (dual-harmonic) model with node i. Thus, we set r_{ij} to be the *probability* that C_j intersects the dual-harmonic function of a randomly selected point in C_i. In other words, if P_i is a random variable uniformly distributed in C_i, and $\hat{I}(P_i)$ is the dual-harmonic

function associated with it, we set

$$r_{ij} = \text{Prob}\left[\hat{I}(P_i) \cap C_j \neq \emptyset\right].$$

In practice, we estimate these values using Monte Carlo sampling.

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 based on (36). The support is then used to relax the label assignments based on the update rule (see (37)). The relaxation labeling theory (Hummel and Zucker 1983; Pelillo 1997) ensures that this relaxation process will converge to a consistent labeling while extremizing the average local consistency over the entire image. Recent studies also show that this process is stable at signal boundaries and singularities (Ben-Shahar and Zucker 2003), a capacity that facilitates the required preservation of singular structure.

Finally, it should be noted that the convergence state of the relaxation process may result in an unambiguous labeling assignment. However, such a labeling assignment can be naturally disambiguated at each node through a maximum confidence selection (i.e., a winner-takes-it-all principle, where only the label with maximum confidence at each node survives), a straight forward and parameter-free process that provides the final computational result.

6 Experimental Results

Among the different applications of good continuation in perceptual organization, we have examined our general theory in the context of image denoising and restoration. Toward that end we have tested the proposed dual-harmonic model and the performance of the described contextual inference computation on a variety of input images, both synthetic and natural. In all cases, we have examined the performance on clean versions of the input image, as well as on noisy versions. For our experiments we used two noise models—additive Gaussian noise and the non-additive Salt and Pepper noise. Part of our goal in using both noise models was to demonstrate how the good continuation approach can handle both equally well, something that to our best knowledge no denoising approach is able to achieve. Since, as we argue from the start, both values and level set geometry are critical aspects of the visual signal, we show both channels separately in all our result figures. For comparison, all our results are plotted against those obtained from popular anisotropic diffusion approaches for structure-preserving denoising—the Perona and Malik (1990) method and the Beltrami flow approach (Sochen et al. 1998). As we show, not only that our approach competes well against these methods in their ball park (i.e., smoothing additive noise), but at the same time it exhibits excellent denoising results for non-additive noise, unlike the very same diffusion techniques.

6.1 Noise-Free Synthetic Images

Our first results are demonstrated on synthetic benchmark images, several of which are shown in Fig. 11. The use of synthetic images is particularly revealing since they provide both the "ground truth" which defines the ultimately desired

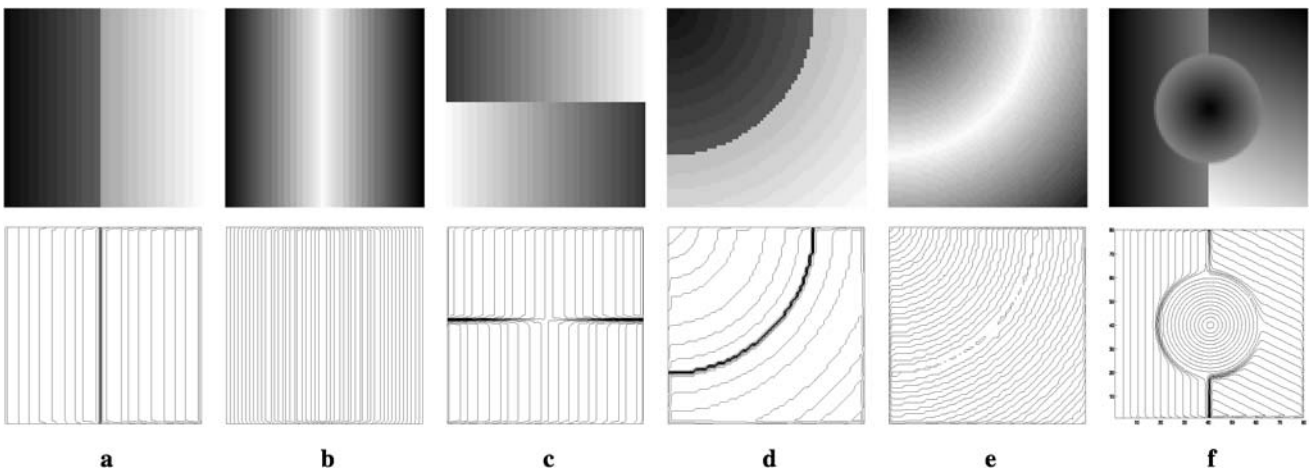

Fig. 11 A sample of synthetic images used for experimental evaluation, each represented as an intensity image (*top*) and via their level set geometry (*bottom*). Note that sometimes singular structure occurs in the shading pattern, sometimes in the level set geometry, and sometimes in both. (**a**) A vertical step edge. (**b**) A vertical positive roof edge (a.k.a. positive line). (**c**) A horizontal "cross"—the perceptual horizontal edge is obtained by two opposing gradient profiles which intersect at some point (hence the term "cross." See also the 3D representation in Fig. 12c). (**d**) A curved step edge. This and the following patterns could help evaluate the response of the various methods not only to the intensity profile across the edge, but also to the geometry of the edge itself. (**e**) A curved roof edge. (**f**) Several patches with different level set behavior

performance and they allow the exploration of various visual structures which otherwise may be difficult to isolate in natural images.

Furthermore, unlike natural images, which always are contaminated with some noise, synthetic images can also provide truly noise-free inputs (up to quantization) that allow to scrutinize the intrinsic qualitative effect that any particular method inflicts upon visual structures of interest. Indeed, our first test examined how our approach preserves singular and regular structure when no noise is present. To do so, we applied both the good continuation computation and the diffusion methods the same number of iterations that would typically be enough to handle noisy patterns (see Sect. 6.2), but we did so on noise-free inputs. Hence, the results of 5 iterations of the proposed good continuation approach (with step size $\delta = 0.25$) are compared to 200 iterations of the Perona-Malik diffusion (with step size $\lambda = 0.25$) and to 200 iterations of the Beltrami flow (with step size $dt = 0.002$). In all our experiments, intensity was quantized to 32 classes, orientation to 16 classes (in the range $(-\pi, \pi]$), intensity gradient to 3 classes (in the range $[0.0, 6.0]$), and curvatures to 3 classes each (in the range $[-0.2, 0.2]$). This selection reflects an empirical compromise between the *desired* highly-quantized (ideally, in the limit, continuous) labeling space, and the computational cost that grows rapidly (exponentially) with the quantization rate. As we show next, our results using these ad-hoc rates are already very good.

In order to emphasize even minute structure distortions, which sometimes are masked by perceptual limitations, the results of this initial test are depicted as *height functions* (see Fig. 12), a visual representation that can emphasize undesired effects especially around singular regions (like edges or peaks). We observe that all the test patterns remain virtually invariant to the application of the good continuation approach (up to the intrinsic quantization to 32 intensity values), as both singular structures and the regular regions between them exhibit almost no changes after the application of 5 relaxation iterations. The same, however, cannot be said about the diffusion methods which clearly do introduce undesired structural distortions in most cases examined.

6.2 Noisy Synthetic Images

While the good continuation approach seems to do excellent job in preserving noise-free structure, the motivation for its formulation was the need to behave similarly in the presence of noise and incoherent measurements. By design, this should be achieved by applying the developed dual harmonic model on the noisy signal and organizing geometrically coherent structure directly from it. In such cases, noisy measurements are iteratively replaced by more consistent values which correspond the best with their context in term of the defined 2D good continuation model. The relaxation process itself then drives these changes toward global consistency.

Figures 13 and 14 illustrate the results of such reconstruction of coherent structure from noisy versions of all images from Fig. 11. Following the signal aspects that motivated our theoretical analysis, in these and the rest of our result figures, the computational outcomes are depicted both as a shading structure and via the geometry of their level sets (drawn 10 gray levels apart). As before, the good continuation results are shown always after 5 relaxation iterations while diffusion results on *noise-free* images are shown after 200 iterations (as in Sect. 6.1). Diffusion results on *noisy* images are shown after 100 iterations on images with additive noise, and after 500 iteration on images with Salt and Pepper noise. Step sizes in all cases are as mentioned in Sect. 6.1.

It should be noted that, as is typically the case, the parameters of different numerical approaches cannot be made commensurate in a strict sense. However, what we set to examine in our comparison is a tradeoff between two conflicting goals—noise removal and preservation of (singular and regular) structure. Obviously, the more we diffuse (or relax), the more noise is removed. At the same time, the more we diffuse (or relax), the greater is the risk for structure distortion (especially, for regular structure). Hence, as long as we select the parameters of the different methods such that the process stops before (or just when) noise is removed, we are able to examine if and how much structure was distorted in the process. If structure was distorted before noise was removed, then clearly the two goals could not be met. Following this reasoning (and using the step sizes as in Sect. 6.1) 100 iterations were approximately needed in both diffusion methods to remove most or all of the additive noise, while 500 iterations were never enough to eliminate the Salt and Pepper noise, but more than enough to incur severe structural distortions.

In general, we observe that the good continuation process exhibits almost flawless behavior compared to these popular anisotropic diffusion schemes. These results apply not only in terms of noise removal and preservation of *singularities*, but also in terms of preserving and reconstructing the *regular* parts of the level set geometry. In comparison, anisotropic diffusion of the noisy images distorts these structures as, or before (and sometimes long before) noise is removed, hence putting in question the utility of these methods when used in conjunction with subsequent computer vision processes such as shape inference (see the Introduction).

6.3 Natural Images

In addition to the synthetic images used in the previous subsection, we have tested our approach on a variety of natural images as shown in Fig. 15. As before, each image is

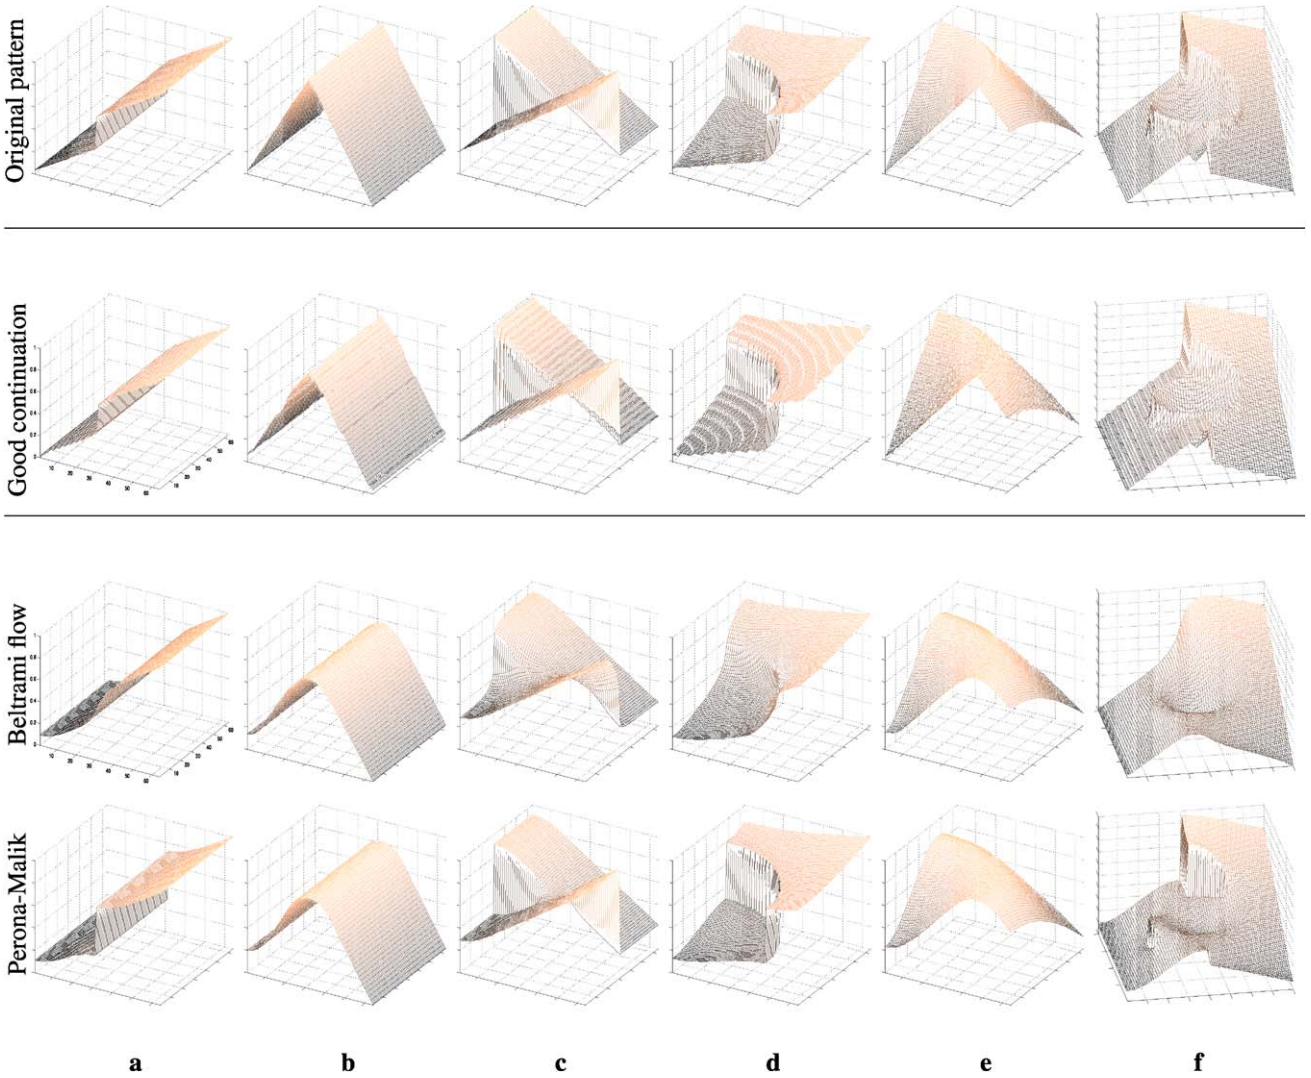

Fig. 12 The qualitative effect of the good continuation approach on all synthetic patterns from Fig. 11. All images are depicted as height functions. *Top row*: The original noise-free pattern (compare to Fig. 11). *Second row*: The result of 5 iterations of the good continuation approach with step size $\delta = 0.25$ (cf. (37)). Note that these results are virtually identical to the original pattern (up to the intrinsic quantization), which implies the excellent preservation of structure, both singular and regular. *Third row*: The results of 200 iterations of the Beltrami flow (Sochen et al. 1998) with step size $dt = 0.002$. *Bottom row*: The results of 200 iterations of the Perona-Malik anisotropic diffusion (Perona and Malik 1990) with step size $\lambda = 0.25$. Note that both methods inflict undesired distortions in almost all cases. Compare to the good continuation results in the second row

shown both as an intensity function and as a map of level sets (this time drawn 20 gray levels apart to reduce clutter). The results of applying our good continuation approach to these images is illustrated in Figs. 16 and 17 next to the results obtained from the selected anisotropic diffusion methods. We observe that in all cases (noise-free, additive noise, and Salt and Pepper noise), the good continuation approach typically exhibits better performance, both in terms of noise removal, and in particular, in terms of the preservation of level set geometry. It is interesting to note that this happens even when the perceptual results are comparable. Consider, for example, the result of the Perona-Malik diffusion vs. the good continuation result on the Lips image (second row in Fig. 17) and focus on the upper and low lip regions. Clearly, the result of the Perona-Malik diffusion in these regions suggests a "ramp" behavior—a roughly horizontal ramp in the upper lip and a roughly vertical ramp in the lower lip. These structures are qualitatively different than the original ones (from Fig. 15c) which have, in particular, local intensity maxima (two in the upper lip and one in the lower lip). In contrast, the result of the good continuation approach show that these fine structures were fully preserved. Following our motivation from the Introduction, if one were to use these two results in a subsequent shape inference procedure, the one based on the diffusion result would clearly fail to yield the correct qualitative shape of the lips.

Fig. 13 Results of applying the good continuation inference process on the first three synthetic patterns from Fig. 11 (unless printed on a quality printer, details in several images may be difficult to observe. In this case, please refer to the electronic version and zoom in on individual panels). Different rows relate to performance on different input images, where each input image is tested with no noise (*top row* of each triple), with Gaussian additive noise (*middle row*, $\sigma = 0.002$, Matlab's implementation), and with Salt and Pepper noise (*bottom row*, $p = 0.1$, Matlab's implementation), respectively. Note how the good continuation approach keeps both the shading patterns and their level set geometry almost intact while removing the noise. In contrast, note how both diffusion methods often induce undesired structural effects (i.e., the elimination or creation of singular structure and the distortion of regular structure, e.g., in the neighborhood of the cross's center) even before all noise is eliminated. Compare all results to the desired shading and level set patterns in Fig. 11

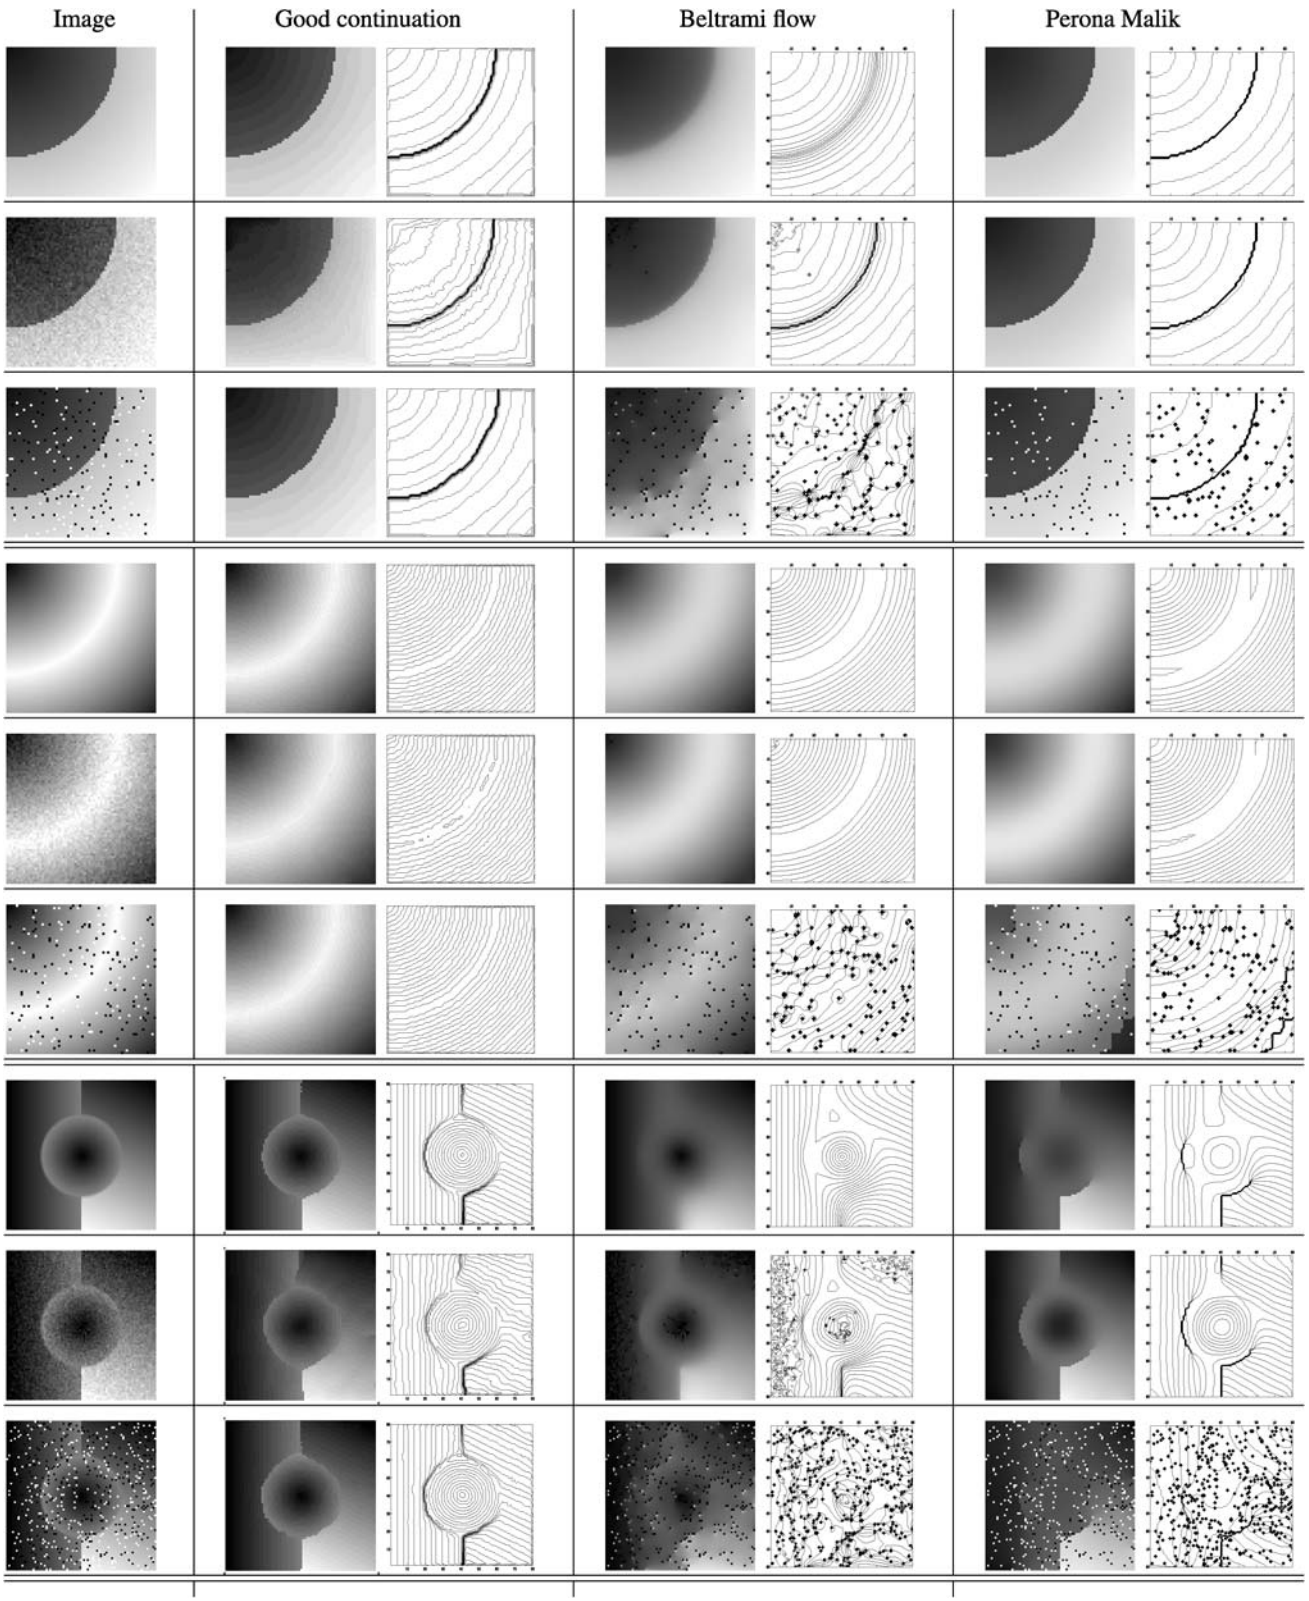

Fig. 14 Results of applying the good continuation inference process on the last three synthetic patterns from Fig. 11. Figure organization is as described in the caption of Fig. 13. As in the previous figure, compare results to desired shading and level set patterns in Fig. 11. Note, for example, how the good continuation approach recovers very accurately the structure in the patches image while both anisotropic diffusions introduce severe structural distortions to it

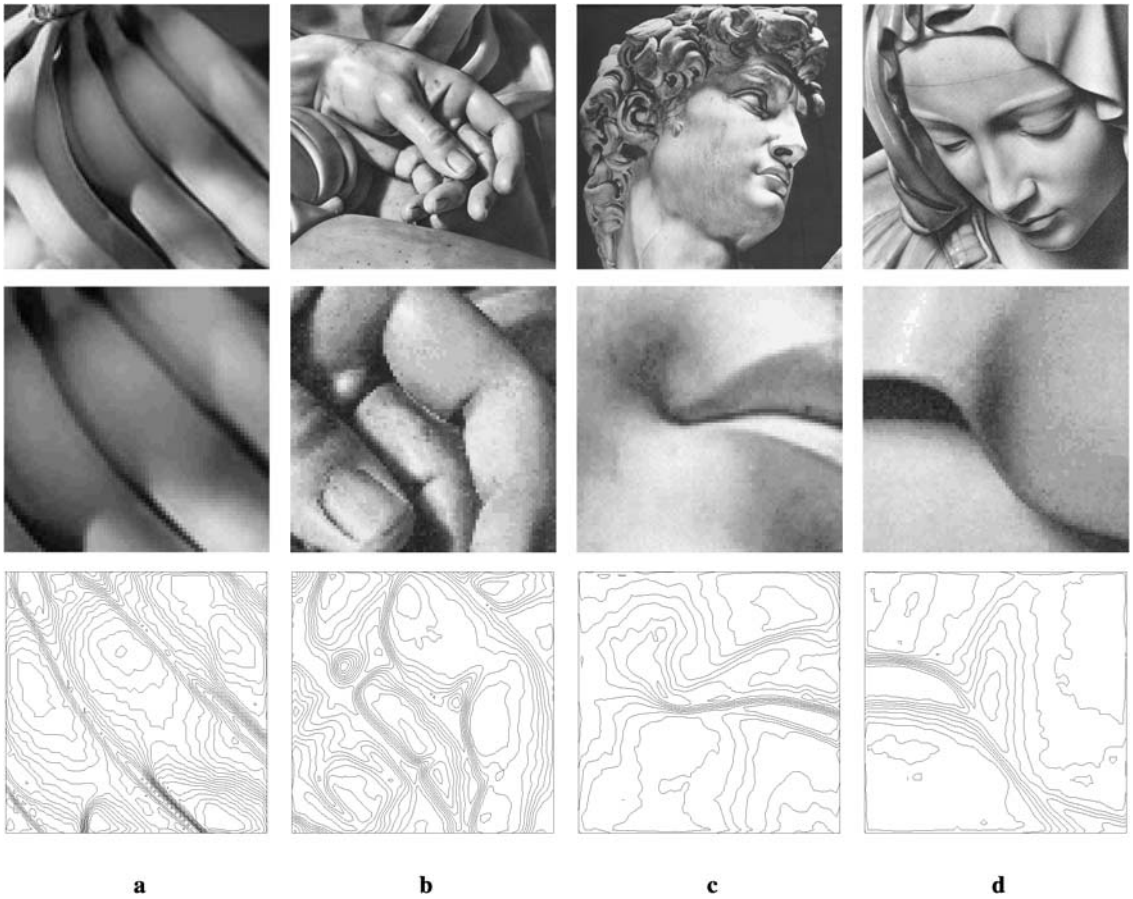

Fig. 15 A sample of natural images used for experimental evaluation (*top*). To allow clear depiction of the results, in particular those of the level set geometry, all computations are shown for selected closeup region of interest (*middle row*) which are represented both as an intensity image and via their level set geometry (*bottom*). (**a**) Bananas. (**b**) Fingers (from Michelangelo's *Pieta*). (**c**) Lips (from Michelangelo's *David*) (**d**) Hood (from Michelangelo's *Pieta*)

6.4 Other Domains of Application

It is important to reiterate that this paper is theoretical and general, and that our results are relevant to various perceptual organization applications, visual cues, and image features. For convenience reasons we choose to demonstrate applicative results from the domain of image denoising and the restoration of shading functions. However, as mentioned in the Introduction, the same geometrical theory and computational mechanisms could facilitate other applications which use good continuation, including (but not limited to) segmentation, grouping, visual completion, or image inpainting. Moreover, all these applications could be applied not only to the raw intensity function, but to any visual cue that can be abstracted as a piecewise smooth 2D scalar function over the image plane. To provide a glimpse to these other domains, Fig. 18 shows one example of *color* image restoration based on the inference of coherent structure in all of the R, G, and B color channels. Here, representing the image via these channels create three feature maps, each of which obeys the abstraction that underlies our theory. Hence, a corrupted color image could be restored by simultaneous application of our contextual inference on each of these channels, and stacking them back to a full restored color image. While we do no claim that this is the correct way to handle color or color images (e.g., see Ben-Shahar and Zucker 2004), Fig. 18 suggests that even this naive approach can achieve practical results.

7 Summary

This paper presents a general theory, a formal derivation, and a computational method for the organization of coherent visual structures based on a rigorous extension of the principle of good continuation. First, we have shown how geometrical good continuation may be applied to general functions by considering both their level sets and the distribution of values and have developed a unique formal model in which both aspects are *simultaneously* harmonic functions. We

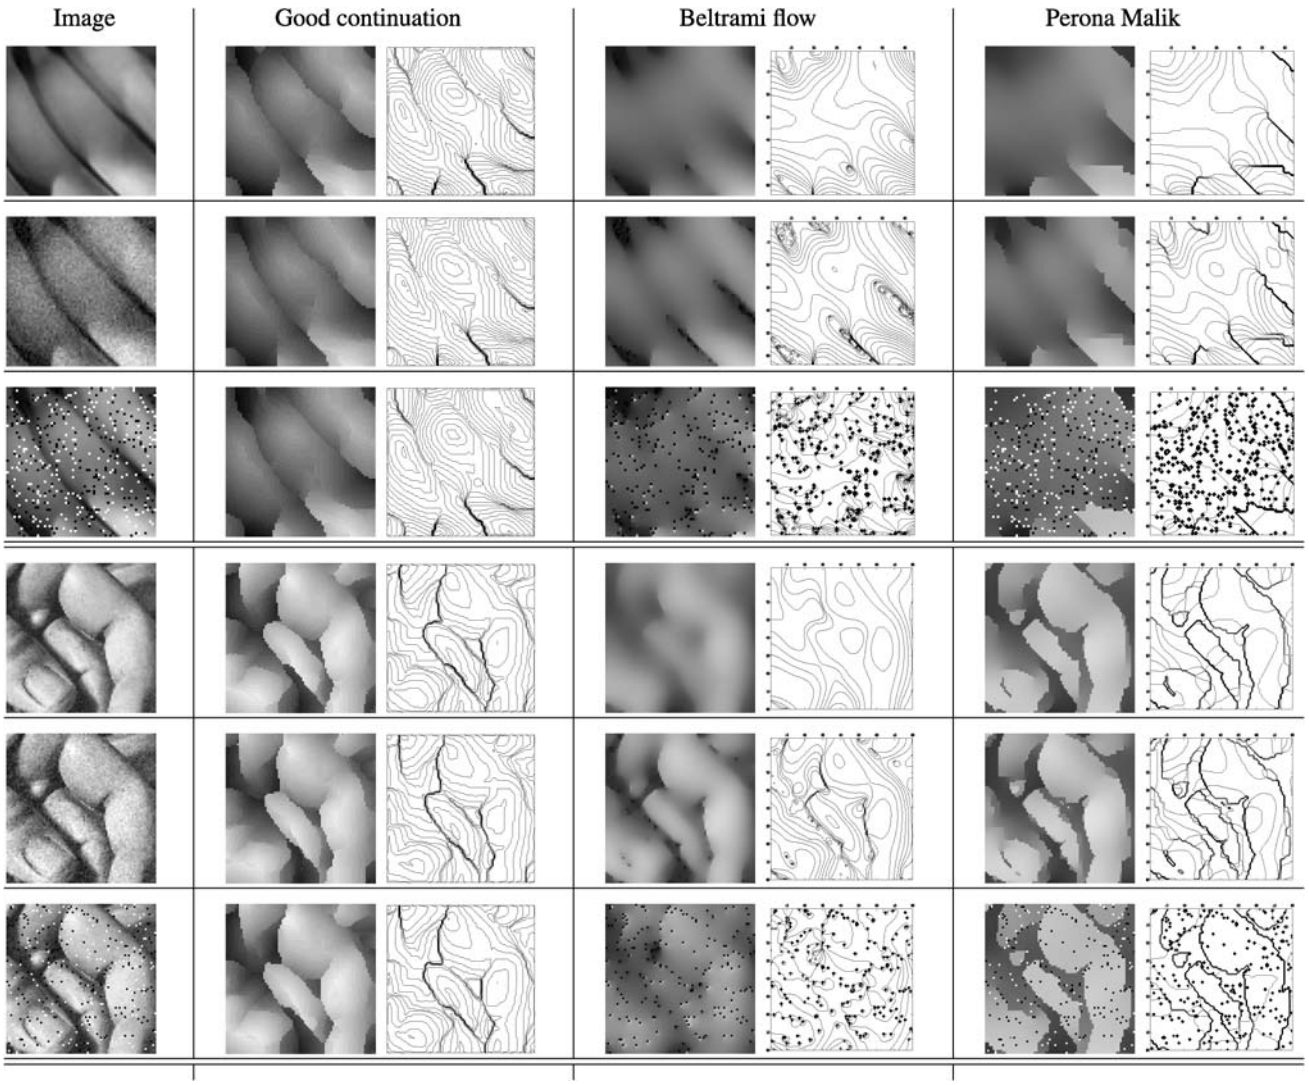

Fig. 16 Results of applying the good continuation inference process on the first two natural images from Fig. 15. Figure organization, as well as all execution parameters, are all as described in previous result figures. Compare all results to the desired original structures (in particular, the level set maps) from Fig. 15

have further extended this result to dual p-harmonic models to show that for any value of p there exist a unique good continuation model which is simultaneously p-harmonic in both its value structure and its level set geometry.

Second, we have used the dual harmonic model in a contextual computation and demonstrated how this perceptual organization approach provides superior performance in eliminating spurious measurements while preserving singularities and the underlying level set geometry. Unlike existing methods for denoising visual signals, ours can handle both additive and non-linear noise while preserving both the singular and much of the regular structure, hence offering much greater chance for success to computer vision algorithms that follow the denoising step.

The approach developed in the present paper not only exhibits state-of-the-art computational results but it extends the possible contribution of perceptual organization beyond the natural boundaries of this field. While good continuation has been considered mostly in the context of curve-like structures, here we show that it can be defined rigorously and formally for general 2D functions as well. We assert that this extension is significant not only from a theoretical and computational perspectives, but also for biological and human vision. Indeed, the analysis and algorithms discussed here are all biologically-plausible in a sense that they fit naturally to a massively distributed (i.e., brain-like) computational network in which each node is tuned to some properties of the signal and its interactions to its neighboring nodes facilitate the inference of coherent global structure. As it has been suggested previously for the organization of coherent oriented structure (e.g., Ben-Shahar and Zucker 2004), here our work suggests both new tuning properties for early vi-

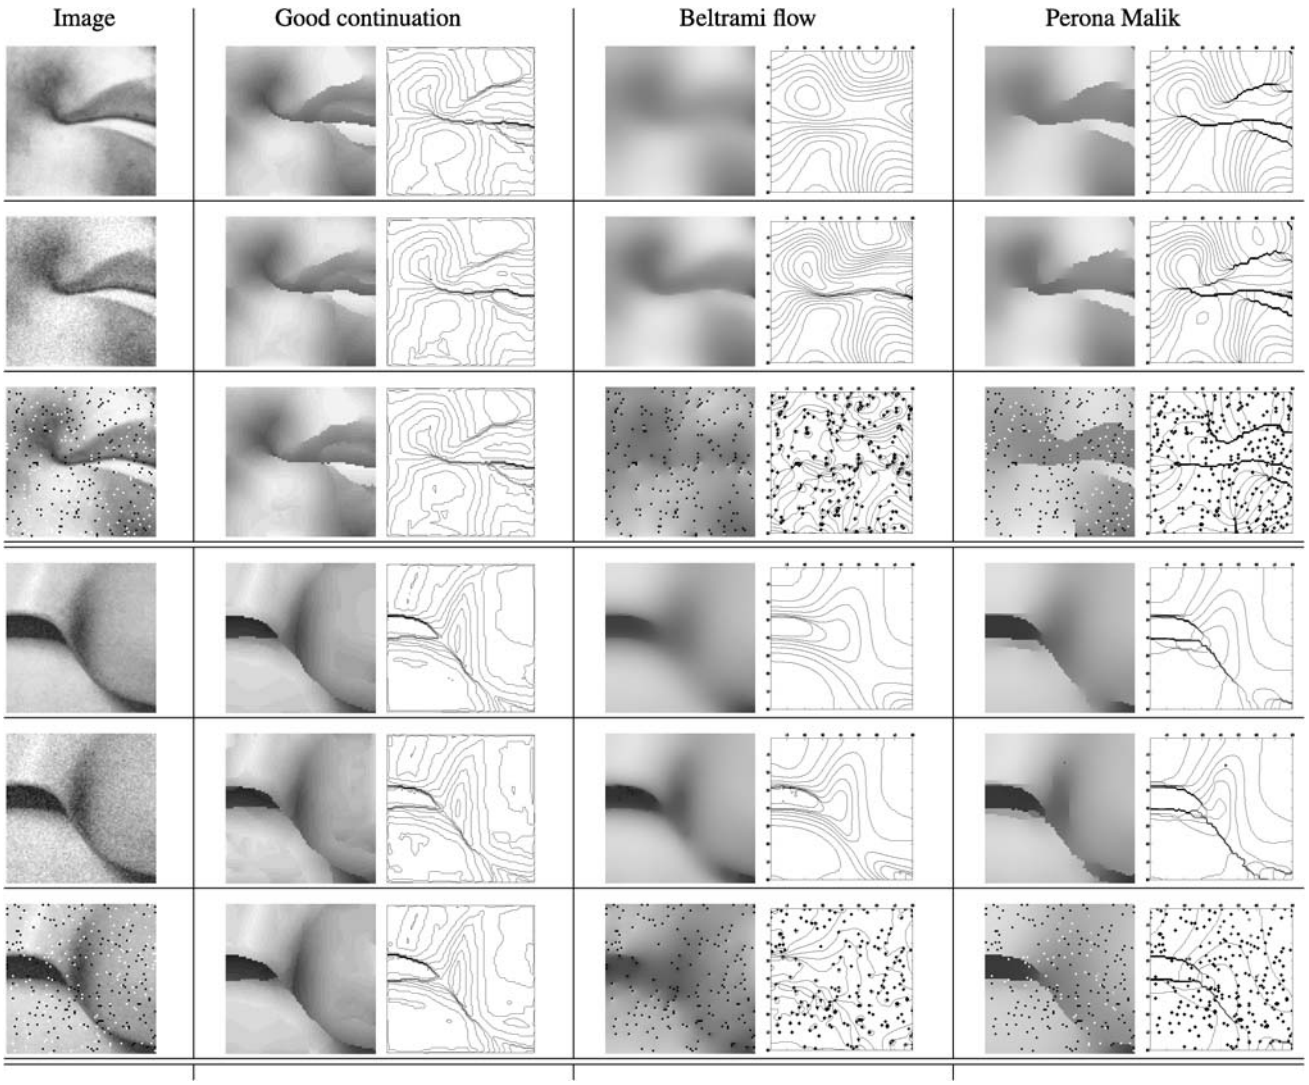

Fig. 17 Results of applying the good continuation inference process on the last two natural images from Fig. 15. Compare all results to the desired original structures (in particular, the level set maps) from Fig. 15

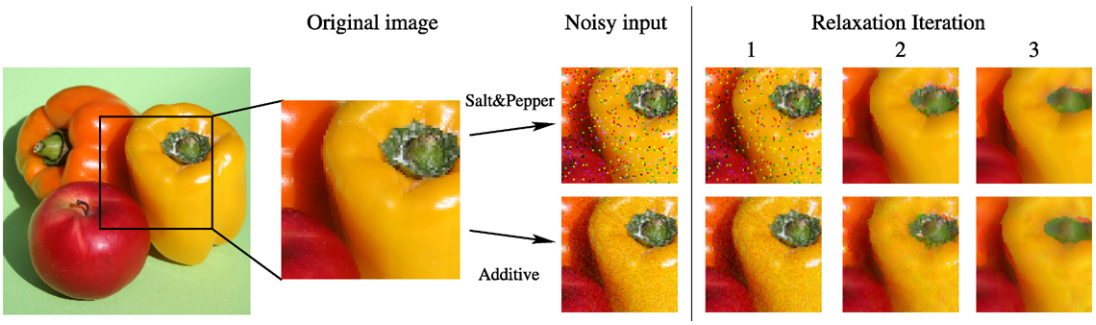

Fig. 18 The restoration of a color image corrupted with either additive Gaussian noise (*top*) or Salt and Pepper noise (*bottom*) by the simultaneous application of the good continuation contextual inference to the R, G, and B channels

sion receptive fields and a particular way by which neurons in the visual cortex should interact and connect in order to facilitate the inference of structure for which the network is designed. Both implications constitute quantitative predictions which could (and hopefully will) inform biological vision research to new directions.

Interestingly enough, one of the new tuning properties that is suggested by our analysis is sensitivity to the *gradient* of visual signals and cues. Such tuning properties have been supported *indirectly* through the classical observation of scale tuning and spatial frequency tuning in visual cortical cells (Hubel and Wiesel 1977; De Valois et al. 1982). Furthermore, recent new studies have employed other computational reasons to argue for *explicit* neural representation of gradients and their processing in the early visual system (e.g., Keil et al. 2006). Once such tuning properties are confirmed experimentally, the well defined set of connections predictable by our dual-harmonic model could also be tested directly using available neuroanatomical and neurophysiological techniques.

Acknowledgements This work was funded in part by the Psychobiology Young Investigator grant 207-07-08, and the Israel Science Foundation (ISF) grant No. 1245/08. Additional support was provided by the ARO, AFOSR, NGA, ONR and NSF. O.B.-S. also thanks the generous support of the Frankel fund, the Paul Ivanier center for Robotics Research and the Zlotowski Center for Neuroscience at Ben-Gurion University.

References

Alvarez, L., & Mazorra, L. (1994). Signal and image restoration using shock filters and anisotropic diffusion. *SIAM Journal on Numerical Analysis*, *31*(2), 590–605.

Ben-Shahar, O. (2003). *The perceptual organization of visual flows*. Ph.D. thesis, Yale university.

Ben-Shahar, O. (2006). Visual saliency and texture segregation *without* feature gradient. *Proceedings of the National Academy of Sciences of the USA*, *103*, 15704–15709.

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

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

Ben-Shahar, O., & Zucker, S. (2004). Hue geometry and horizontal connections. *Neural Networks*, *17*, 753–771.

Breton, P., Iverson, L., Langer, M., & Zucker, S. (1992). A new approach to shape from shading. In G. Carpenter & S. Grossberg (Eds.), *Neural networks for vision and image processing*, Chap. 5 (pp. 111–131). Cambridge: Cambridge MIT Press.

Cao, F. (2004). Application of the gestalt principle to the detection of good continuation and corners in image level lines. *Comput. Vis. Sci.*, *7*, 3–13.

Carbal, B., & Leedom, L. (1993). Imaging vector fields using line integral convolution. In *Proceedings of SIGGRAPH* (pp. 263–270).

Caselles, V., Kimmel, R., Sapiro, G., & Sbert, C. (1997a). Minimal surfaces: a three dimensional segmentation approach. *Numerische Mathematik*, *77*(4), 423–451.

Caselles, V., Kimmel, R., Sapiro, G., & Sbert, C. (1997b). Minimal surfaces based object segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *19*(4), 294–298.

Caselles, V., Coll, B., & Morel, J.-M. (2002). Geometry and color in natural images. *Journal of Mathematical Imaging and Vision*, *16*, 89–105.

De Valois, R., Albrecht, D., & Thorell, L. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research*, *22*, 545–559.

Dierkes, U., Hildebrandt, S., Küster, A., & Wohrlab, O. (1992). *Minimal surfaces*. Berlin: Springer.

Duncan, J., & Birkholzer, T. (1992). Reinforcement of linear structure using parametrized relaxation labeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *14*(5), 502–515.

Fischler, M. (1994). The perception of linear structure: a generic linker. In *Proceedings of the image understanding workshop* (Vol. 2, pp. 1565–1579).

Graustein, W. (1940). Harmonic minimal surfaces. *Transactions of the American Mathematical Society*, *47*, 173–206.

Guy, G., & Medioni, G. (1996). Inferring global perceptual contours from local features. *International Journal of Computer Vision*, *20*(1/2), 113–133.

Guy, G., & Medioni, G. (1997). Inference of surfaces, 3d curves, and junctions from sparse, noisy 3d data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *19*(11), 1265–1277.

Hamel, G. (1923). Zum Gedachtnis an Hermann Amandus Schwarz. *Jahresberichte D.M.V.*, *32*, 6–13.

Hoffman, D. (1998). *Visual Intelligence*. Norton.

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

Hubel, D., & Wiesel, T. (1977). Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society London Series B*, *198*, 1–59.

Hummel, R., & Zucker, S. (1983). On the foundations of the relaxation labeling processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *5*, 267–287.

Jain, A. (1989). *Fundamentals of digital image processing*. Prentice-Hall: Englewood Cliffs.

John, F. (1982). *Partial differential equations*. New York: Springer.

Keil, M., Cristobal, G., & Neumann, H. (2006). Gradient representation and perception in the early visual system—a novel account of Mach band formation. *Vision Research*, *46*, 2659–2674.

Kimia, B., Frankel, L., & Popescu, A. (1999). Euler spiral for shape completion. In *Proceedings of the IEEE Computer Society workshop on perceptual organization in computer vision*.

Kimmel, R., Malladi, R., & Sochen, N. (2000). Images as embedded maps and minimal surfaces: Movies, color, texture, and volumetric medical images. *International Journal of Computer Vision*, *39*(2), 111–129.

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

Koenderink, J., & van Doorn, A. (2002). Image processing done right. In *Proceedings of the European conference on computer vision* (pp. 158–172).

Koffka, K. (1935). *Principles of Gestalt Psychology*. London: Routledge & Kegan Paul.

Malik, J., & Perona, P. (1990). Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America*, *7*(5), 923–932.

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

Mumford, D. (1994). Elastica in computer vision. In B. Chandrajit (Ed.), *Algebraic geometry and its applications*. Berlin: Springer.

Nitsche, J. (1989). *Lectures on Minimal Surfaces*. Cambridge: Cambridge University Press.

O'Neill, B. (1966). *Elementary differential geometry*. New York: Academic Press.

Parent, P., & Zucker, S. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *11*(8), 823–839.

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

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

Perona, P., & Malik, J. (1990). Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *12*(7), 629–639.

Richards, W., Keonderink, J., & Hoffman, D. (1987). Inferring 3d shapes from 2d silhouettes. *Journal of the Optical Society of America*, *4*, 1168–1175.

Šára, R. (1995). Isophotes: the key to tractable local shading analysis. In *Proceedings of the international conference on computer analysis of images and patterns* (pp. 416–423).

Savadjiev, P., Zucker, S., & Siddiqi, K. (2007). On the differential geometry of 3d flow patterns: generalized helicoids and diffusion MRI analysis. In *Proceedings of the IEEE international conference on computer vision* (pp. 1–8).

Sha'ashua, A., & Ullman, S. (1988). Structural saliency: the detection of globally salient structures using a locally connected network. In *Proceedings of the IEEE international conference on computer vision* (pp. 321–327).

Sharon, E., Brandt, A., & Basri, R. (2000). Completion energies and scale. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *22*(10), 1117–1131.

Sochen, N., Kimmel, R., & Malladi, R. (1998). A general framework for low level vision. *IEEE Transaction on Image Processing*, *7*(3), 310–318.

Tang, B., Sapiro, G., & Caselles, V. (2000). Diffusion of general data on non-flat manifolds via harmonic maps theory: the direction diffusion case. *International Journal of Computer Vision*, *36*(2), 149–161.

Torsello, A., & Pelillo, M. (2000). Continuous-time relaxation labeling processes. *Pattern Recognition*, *33*, 1897–1908.

Ullman, S. (1976). Filling in the gaps: the shape of subjective contours and a model for their creation. *Biological Cybernetics*, *25*, 1–6.

Waltz, D. (1975). Understanding line drawings of scenes with shadows. In P. Winston (Ed.), *The psychology of computer vision* (pp. 19–91). New York: McGraw-Hill.

Wertheimer, M. (1955). Laws of organization in perceptual forms. In W. Ellis (Ed.), *A source book of Gestalt Psych.* (pp. 71–88). London: Routledge & Kegan Paul.

Williams, L., & Jacobs, D. (1997). Stochastic completion fields: a neural model of illusory contour shape and salience. *Neural Computing*, *9*(4), 837–858.

Witkin, A., & Tenenbaum, J. (1983). On the role of structure in vision. In J. Beck, B. Hope, & A. Rosenfeld (Eds.), *Human and machine vision* (pp. 481–542). New York: Academic Press.

Zucker, S., Rosenfeld, A., & Davis, L. (1975). General purpose models: expectations about the unexpected. In *Proceedings of the international joint conference on artificial intelligence* (pp. 716–721).