In the *IEEE International Conference on Computer Vision (ICCV)*, Beijing, China, 2005

Good continuation of general 2D visual features: Dual harmonic models and computational inference

Ohad Ben-Shahar
Department of Computer Science
Ben Gurion University, Beer Sheva, Israel
ben-shahar@cs.bgu.ac.il

Steven W. Zucker
Department of Computer Science
Yale University, New Haven, CT, USA
zucker@cs.yale.edu

Abstract

Good continuation is a fundamental principle of perceptual organization that guides the grouping of parts based on how they should succeed one another within coherent wholes. Despite the general language that was used by the Gestalt psychologists in phrasing this principle, computational work has focused almost exclusively on the study of curve-like structures. Here we offer, for the first time, a rigorous generalization of good continuation to arbitrary visual structures that can be abstracted as scalar functions over the image plane. The differential geometry of these structures dictates that their good continuation should be based both on their value and on the geometry of their levelsets, which yield a coupled system of equations solvable for a formal model. We exhibit the resulting computation on shading and intensity functions, demonstrating how it eliminates spurious measurements while preserving both regular structure and singularities. Related implementations could be applied to color channels, motion magnitude, and disparity signals.

1. Introduction and background

Perceptual organization has its origins in Wertheimer's phenomenological observations about perception [33, 19] which identify the need for an inferential process that provides 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" ([35, p. 483]). *Good continuation*, perhaps the most prominent organizational principle, is described by Wertheimer as the "inner coherence" by which "successive parts of a whole should follow one another" [33, p. 83]. It has inspired investigation into curve-like structures from computational (e.g., [32, 29, 22, 8, 9, 34, 17]) psychophysical, physiological, and anatomical considerations (e.g., [11, 3, 28, 20, 2]).

While the identification of good continuation with curves is historical (see Fig. 1a,b), it is clear that Wertheimer meant more. Indeed, Wertheimer intentionally phrased his "Factor of Direction" in general terms, and our goal is to develop this computationally for arbitrary scalar functions from the image plane to the real line $\mathbf{R}$ (or subintervals thereof). The need is also illustrated in Fig 1, which shows how intensity distributions can conspire to obscure boundaries and frustrate image segmentation systems. Indeed, various nonlinear diffusion techniques for smoothing higher dimensional data sets (e.g., [30, 31]) relate indirectly to good continuation. These techniques, however, were not only criticized on a principled validity level (cf. [18]), but they do not address good continuation explicitly. We demonstrate this directly in our experimental section. A more explicit consideration of good continuation of non curve-like features has been attempted within the tensor voting framework [10] for surface patches, and formal considerations based on a frame field representation were offered by Ben-Shahar and Zucker [1] for locally parallel structures. We follow the latter in using a frame field representation, but note that it, too, is *incomplete* for general scalar functions.

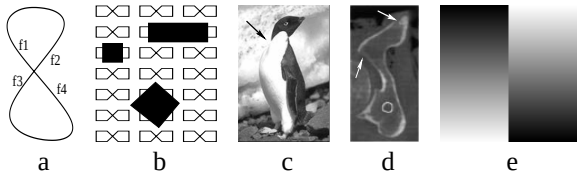

Figure 1. Good continuation in perceptual organization is classically phrased in terms of contours, e.g., to dictate the groups f1-f4 and f2-f3 in panel **a**, or to override context and visual experience in panel **b**. But issues of good continuation arise in more general cases also, for example in the context of segmentation of objects that reflect light comparable to the background, a problem that is exacerbated with noise. This happen frequently in natural images and in medical imagery, e.g., when there is a lack of bone density (note boundary regions pointed by arrows in panels **c** and **d**), and it is typified by the example of opposing intensity gradients (note the central region in panel**e**). How can such boundaries be structured without causing 'bleeding' between regions? To our knowledge no existing region or boundary detection system can solve this problem and here we show how it can be treated in the framework of good continuation.

We stress that this paper is theoretical and general, so our results are applicable not only to images and shading functions (presented therein), but also, for example, 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 a stereo pair; to texture energy channels;and to the magnitude of optical flow fields.

2. The geometry of 2D scalar features

Let $I : \mathbf{R}^2 \rightarrow \mathbf{R}$ be a generic image feature defined over the image plane. For now, assume I is smooth. (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.) 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"

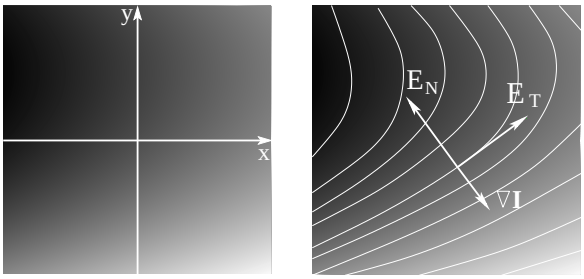

Figure 2. Although any general 2D scalar function over the image plane can be viewed as a set of values over a global coordinate frame (left), we will stress the importance of its levelset geometry (right) by making this aspect explicit in the representation that we use. Without loss of generality, and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

relative to each other [33]. The question is how to define "succession".

One possibility is to approximate the "good' behavior of the function $I : \mathbf{R}^2 \rightarrow \mathbf{R}$ around a point $q = (a, b)$ based on the coefficients of its Taylor expansion. Depending on the level 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 q itself), for example by

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

Another view – closer in spirit to good continuation – is that such approximations should be based on the levelset geometry of these functions [4, 27, 18, 6]. Because the levelset geometry of smooth regions is generically *locally parallel* (Fig. 2), as follows from the classical existence theorem for ordinary differential equations, we can apply a construction from Ben-Shahar and Zucker [1] developed for texture. By their 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 levelsets, 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. These two curvatures are derived from the covariant derivative of the field $\hat{\mathbf{E}}_T$ and they are related to the gradient of the orientation function $\nabla\theta$ via the formulas:

$$\begin{array}{ll} \kappa_T = \nabla\theta \cdot (\cos\theta, \sin\theta) & \theta_x = \kappa_T \cos\theta - \kappa_N \sin\theta \\ \kappa_N = \nabla\theta \cdot (-\sin\theta, \cos\theta & \theta_y = \kappa_T \sin\theta + \kappa_N \cos\theta \ . \end{array} \quad (2)$$

But $\{\theta(q), \kappa_T(q), \kappa_N(q)\}$ are only three parameters, capable of approximating the local levelset 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 2^{nd} order approximations (this is the lowest order that captures variations in the levelset 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 levelsets, ∇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.,

$$\begin{array}{ll} \hat{\mathbf{E}}_T = \frac{(I_y, -I_x)}{||\nabla I||} & I_x = ||\nabla I|| \sin\theta \\ \hat{\mathbf{E}}_N = \frac{(-I_x, -I_y)}{||\nabla I||} & I_y = -||\nabla I|| \cos\theta \quad , \end{array} \quad (3)$$

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

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

At first glance, these equations suggest that we need to add both $\frac{\partial}{\partial x}||\nabla I||$ and $\frac{\partial}{\partial y}||\nabla I||$ as free parameters in our descriptor, which will result in a total of seven parameters, one more than Taylor. However, the integrability constraint $I_{xy} = I_{yx}$ translates to

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

and removes one degree of freedom. Thus either the direction of $\nabla||\nabla I||$ or its magnitude, but not both, should be used. 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

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

In conclusion, from this discussion 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 an appropriate* 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 levelsets *and* values hold *simultaneously*. Moreover, by extending the model for locally parallel structures [1] exactly to I's levelsets, we are able to remove one parameter and end up with a local descriptor that has the five natural parameters: $\{I, ||\nabla I||, \theta, \kappa_T, \kappa_N\}$.

3. Levelsets 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 de-

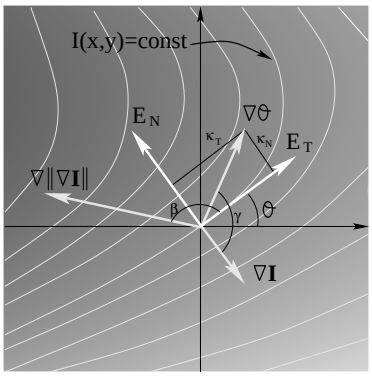

Figure 3. The geometrical parameters underlying the local description of general 2D scalar functions up to second order. Note that although the levelset curvatures are depicted as projections of $\nabla\theta$ on the frame $\{\hat{\mathbf{E}}_T, \hat{\mathbf{E}}_N\}$, their more intuitive depiction, avoided here to prevent clutter, is through two osculating circles, one tangent to $\hat{\mathbf{E}}_T$ and another tangent to $\hat{\mathbf{E}}_N$, having radii $1/\kappa_T$ and $1/\kappa_N$, respectively. Please refer to the text for additional explanations.

rived 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, that reflects good continuation in the neighborhood of q, both for its levelset geometry and for the distribution of its values. Assuming, without loss of generality, that our model is constructed around the origin $q=(0,0)$ with the coordinate system aligned with the tangent of the levelset at that point, i.e., with $\hat{\mathbf{E}}_T(q)=(1,0)$ or $\theta(q)=0$, a model that reflects good continuation for the locally parallel structure of the levelsets has already been proposed in the literature [1] and takes the form

$$\theta(x,y)=\tan^{-1}\left(\frac{K_T x+K_N y}{1+K_N x-K_T y}\right) \tag{4}$$

where $K_T=\kappa_T(q)$ and $K_N=\kappa_N(q)$. When viewed as a surface in the space $XY\theta$ (i.e., $\mathbf{R}^3$ whose Z axis represents orientation), this function is known as a *right helicoid*. Several instances of this model, depicted as texture patterns stripped of any value structure, are shown in Fig. 4.

Such a helicoidal model possesses several geometrical properties that associate it with good continuation. Perhaps most importantly, 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

$$\frac{\kappa_T(x,y)}{\kappa_N(x,y)}=\text{const}=\frac{K_T}{K_N}\qquad \forall(x,y)\in N(q). \tag{5}$$

Therefore, unlike common ways to achieve "good behavior" of 2D (or higher dimensional) structures, typically via anisotropic diffusion and deformable models (e.g., [31, 30]), the helicoidal model emerges from considerations of the behavior of curvatures in the image plane and thus it is a closer in spirit to the methodology that was employed in the study of good continuation of curves (e.g., [32, 22, 21, 17]). 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.,[26, 12]) and therefore are clearly inappropriate for a model for coherence.

Although it emerges from explicit good continuation considerations in the image plane, the helicoidal model does enjoy properties that link it to scale space and diffusion techniques. In particular, this model (as an orientation function) was also proved to be both p-harmonic for all values of p and a minimal surface in $XY\theta$. 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 model for 2D functions

The helicoidal model defines only the levelset 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 express the helicoidal model directly in terms of a function $I(x,y)$ whose levelsets obey the helicoidal orientation function and then explore the constraints that this derivation entails. From the identity $\theta=\tan^{-1}\frac{-I_x}{I_y}$ we can derive $\nabla\theta$ in terms of I's derivatives

$$\theta_x=\frac{I_x I_{xy}-I_y I_{xx}}{I_x^2+I_y^2}\qquad \theta_y=\frac{I_x I_{yy}-I_y I_{xy}}{I_x^2+I_y^2} \tag{6}$$

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

$$\begin{aligned}\kappa_T &= \frac{I_x^2 I_{yy}-2I_x I_y I_{xy}+I_y^2 I_{xx}}{(I_x^2+I_y^2)^{3/2}}\\ \kappa_N &= \frac{I_{xy}((I_x^2+I_y^2)+I_x I_y(I_{xx}-I_{yy})}{(I_x^2+I_y^2)^{3/2}}\end{aligned} \tag{7}$$

By applying the constraint in Eq. 5 we can therefore obtain a constraint that forces $\hat{I}$ to have helicoidal levelsets

$$\begin{aligned}0 = {} & I_{xx}(K_N I_y^2-K_T I_x I_y)+\\ & I_{xy}(K_T(I_x^2-I_y^2)-2I_x I_y)+\\ & I_{yy}(K_N I_x^2+K_T I_x I_y)\,.\end{aligned} \tag{8}$$

While this second order PDE defines a necessary condition which $\hat{I}$ must satisfy, its nonlinear nature suggests that a closed form solution may be difficult to obtain. A more productive way to generalize the levelsets' helicoidal model to a general 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 derive a closed form expression for $\hat{I}$'s levelsets 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 [16] which is used to parameterize the entire solution as a dense collection of levelsets. This parameterization 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. Levelsets closed form

Let $I(x,y)$ an arbitrary smooth scalar function over the image plane. The orientation of I's levelsets satisfy (see Eq. 3)

$$\theta=\tan^{-1}\left(\frac{-I_x}{I_y}\right) \tag{9}$$

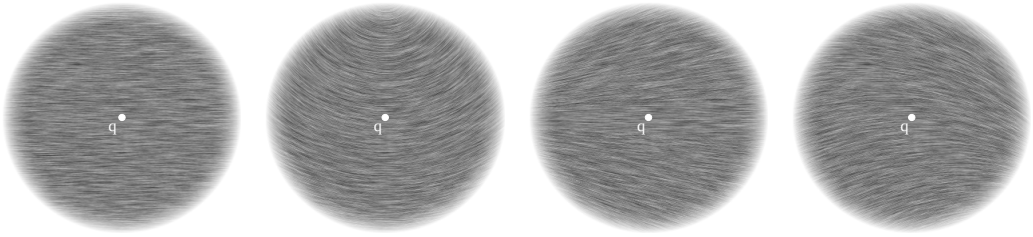

Figure 4. Stripped of its value (say, shading) structure, underlying $\hat{I}(x,y)$ - the model for good continuation - is the geometry of its levelsets, which as we argue should have good continuation qualities of its own. For this purpose we utilize an already proposed model for the good continuation of locally parallel structure - a right helicoid in $XY\theta$ [1]. Shown here are just several examples of helicoidal levelset 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 [5].

and therefore, once we are given an orientation function $\theta(x,y)$ for I's levelsets, $I(x,y)$ must satisfy the PDE

$$I_x + I_y \tan\theta(x,y) = 0\ . \tag{10}$$

Now, since $\theta(x,y)$ will be based on the helicoidal model for good continuation of locally parallel structure [1], from Eq. 4, we obtain the following PDE that our model should satisfy

$$(1 + K_N x - K_T y)I_x + (K_T x + K_N y)I_y = 0\ . \tag{11}$$

The characteristic curves [16] 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

$$\begin{aligned} x(t) &= e^{tK_N}(c_1\cos(tK_T) - c_2\sin(tK_T)) - \tfrac{K_N}{\xi^2} \\ y(t) &= e^{tK_N}(c_1\sin(tK_T) + c_2\cos(tK_T)) + \tfrac{K_T}{\xi^2} \\ I(t) &= c_3 \end{aligned} \tag{12}$$

where $\xi^2 = K_T^2 + K_N^2$ and the coefficients c_1, c_2, and c_3, are determined from the initial data

$$\left.\begin{array}{l} x(t=0) = x_0 \\ y(t=0) = y_0 \\ I(t=0) = I_0 \end{array}\right\} \Rightarrow \left\{\begin{array}{l} c_1 = x_0 + \frac{K_N}{\xi^2} \\ c_2 = y_0 - \frac{K_T}{\xi^2} \\ c_3 = I_0 \end{array}\right. . \tag{13}$$

Given any point in the image plane, we are now able to trace its levelsets via Eqs. 12 and 13. Deriving such a closed form for the levelsets should not be taken for granted. In fact, many levelset 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 model for $\hat{I}(x,y)$ as a whole.

4.2. Solution parameterization

Since Eqs. 12 and 13 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. As long as Γ is not a characteristic curve (and as long as other minor anomalies are avoided), the substitution on Γ into Eqs. 12 and 13 yields the following parametric surface

$$\begin{aligned} x(t) &= \tfrac{-K_N + e^{tK_N}[(x_0(s)\xi^2 + K_N)\cos(tK_T) + (K_T - y_0(s)\xi^2)\sin(tK_T)]}{\xi^2} \\ y(t) &= \tfrac{K_T + e^{tK_N}[(x_0(s)\xi^2 + K_N)\sin(tK_T) + (y_0(s)\xi^2 - K_T)\cos(tK_T)]}{\xi^2} \\ I(t) &= I_0(s) \end{aligned} \tag{14}$$

Given this closed form, 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 "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 venue is to select a $\Gamma_{xy}(s)$ that is everywhere orthogonal to the characteristic curves. If the latter are computed by tracing the direction tangential to the levelsets, i.e., the direction of $\hat{\mathbf{E}}_T$, constructing a curve that is everywhere orthogonal will be based on integration along $\hat{\mathbf{E}}_N$. Repeating the process from Sec. 4.1 for the system

$$\tfrac{dx}{ds} = -K_T x - K_N y \qquad \tfrac{dy}{ds} = 1 + K_N x - K_T y$$

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

$$\begin{aligned} x_0(s) &= \tfrac{-K_N + e^{-sK_T}(K_N\cos(sK_N) + K_T\sin(sK_N))}{\xi^2} \\ y_0(s) &= \tfrac{K_T + e^{-sK_T}(-K_T\cos(sK_N) + K_N\sin(sK_N))}{\xi^2}\ . \end{aligned} \tag{15}$$

Substituting Eq. 15 into Eq. 14 results with the following expression for our sought after model

$$\begin{aligned} x(s,t) &= \tfrac{-K_N + e^{(tK_N - sK_T)}[K_N\cos\tilde{K}(s,t) + K_T\sin\tilde{K}(s,t)]}{\xi^2} \\ y(s,t) &= \tfrac{K_T + e^{(tK_N - sK_T)}[K_N\sin\tilde{K}(s,t) - K_T\cos\tilde{K}(s,t)]}{\xi^2} \\ I(s,t) &= I_0(s) \end{aligned} \tag{16}$$

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

Eq. 16 is a key result of special importance to this paper. It provides an explicit parameterization of the desired good continuation model based on its desired levelset structure (i.e., the helicoidal good continuation model). It therefore parameterizes 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. 5). 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 levelset geometry can constrain this distribution, we will need to incorporate additional good continuation constraints, as discussed below. However, before doing so, one needs to realize that the parameterization provided by Eq. 16 is not Cartesian. This is a major limitation because eventually the good continuation model should be used with images provided in Cartesian coordinates. Of particular concern is the fact that once a model is available, it should be applied to image patches whose spatial extend is defined using the L_2 norm to have circular shape. Can we bridge this parameterization gap easily? Is it possible to re-parameterize Eq. 16 based on Cartesian coordinates?

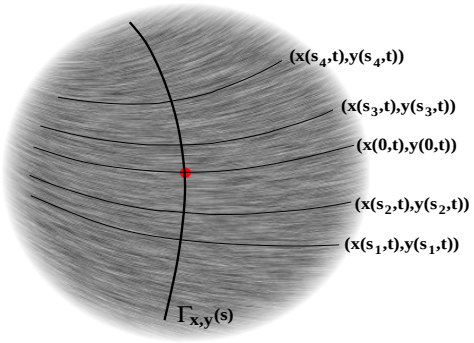

Figure 5. Eq. 16 parameterizes the plane as a collection of curves, each of which is a projection of a levelset curve on the image plane. Shown here is the generator curve $\Gamma(s)$ and several of the generated levelset curves, superimposed on a helicoidal pattern of $K_T = 0.1$ and $K_N = 0.1$.

Given the gap between the available and desired parameterizations, and the nontrivial form of Eq. 16, it is remarkable that the last question can be answered in the affirmative. 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 Eq. 16 makes both $x(\tilde{s}(x, y), \tilde{t}(x, y))$ and $y(\tilde{s}(x, y), \tilde{t}(x, y))$ the identity transformation. 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 Eq. 16 is reversible and that every coordinate (x, y) can be traced back to the pair (s, t) with the following transformation

$$\begin{aligned}\tilde{s}(x,y) &= \frac{\text{sign}(K_T x + K_N y) K_N \cos^{-1}\omega(x,y) - K_T \log \mu(x,y)}{\xi^2} \\ \tilde{t}(x,y) &= \frac{\text{sign}(K_T x + K_N y) K_T \cos^{-1}\omega(x,y) + K_N \log \mu(x,y)}{\xi^2}\end{aligned} \tag{17}$$

where

$$\mu(x,y) = \sqrt{1 + 2K_N x - 2K_T y + \xi^2(x^2+y^2)}$$

$$\omega(x,y) = \frac{1 + K_N x - K_T y}{\mu(x,y)} \quad .$$

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

Having Eq. 17, and an explicit way to switch back and forth between the levelset-centered and the Cartesian parameterization, we are now able to plug $\tilde{s}(x, y)$ to the $I(x, t)$ component of Eq. 16 to obtain an explicit description of our desired model in terms of Cartesian coordinates

$$I(x,y) = I_0\left(\frac{\text{sign}(K_T x + K_N y) K_N \cos^{-1}\omega(x,y) - K_T \log \mu(x,y)}{\xi^2}\right) \tag{18}$$

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 levelsets).

4.3. Derivation of the final model

Eq. 18, and the function $I_0(s)$ on which it is built, represent the backbone on which any scalar function with helicoidal levelsets 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 [15] or by minimization of some variation measure through anisotropic diffusion [30, 31]. In this sense, the helicoidal model for the levelsets 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 in, e.g., [31]) for all values of p and to be a minimal surface (used in, e.g., [7]) in $XY\theta$ [1]. 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 levelset and value structures. Fortunately, Eq. 18 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,

$$\mathcal{F}(I, I_x, I_y, I_{xx}, I_{xy}, I_{yy}; x, y) = 0 \;\Rightarrow\; \mathcal{G}(I_0, \frac{\partial I_0}{\partial s}, \frac{\partial^2 I_0}{\partial s^2}; s) = 0 \,,$$

then solve $\mathcal{G}$ for $I_0(s)$ (if solvable) and plug back into Eq. 18 to obtain the final explicit solution.

Eq. 18 can be differentiated for an arbitrary differentiable $I_0(s)$ using the chain rule

$$\begin{aligned} I_x &= \frac{\partial I_0}{\partial \tilde{s}}\frac{\partial \tilde{s}}{\partial x} \\ I_y &= \frac{\partial I_0}{\partial \tilde{s}}\frac{\partial \tilde{s}}{\partial y} \end{aligned} \qquad \begin{aligned} I_{xx} &= \frac{\partial^2 I_0}{\partial \tilde{s}^2}\left(\frac{\partial \tilde{s}}{\partial x}\right)^2 + \frac{\partial I_0}{\partial \tilde{s}}\frac{\partial^2 \tilde{s}}{\partial x^2} \\ I_{xy} &= \frac{\partial^2 I_0}{\partial \tilde{s}^2}\frac{\partial \tilde{s}}{\partial x}\frac{\partial \tilde{s}}{\partial y} + \frac{\partial I_0}{\partial \tilde{s}}\frac{\partial^2 \tilde{s}}{\partial x \partial y} \\ I_{yy} &= \frac{\partial^2 I_0}{\partial \tilde{s}^2}\left(\frac{\partial \tilde{s}}{\partial y}\right)^2 + \frac{\partial I_0}{\partial \tilde{s}}\frac{\partial^2 \tilde{s}}{\partial y^2} \end{aligned} \tag{19}$$

and the derivatives of $\tilde{s}(x, y)$ are readily computed from Eq. 17 (with some special care involving the sign function):

$$\begin{aligned} \frac{\partial \tilde{s}}{\partial x} &= -\frac{K_T x + K_N y}{\mu^2(x,y)} \\ \frac{\partial \tilde{s}}{\partial y} &= \frac{1 + K_N x - K_T y}{\mu^2(x,y)} \\ \frac{\partial^2 \tilde{s}}{\partial x^2} &= \frac{\mu^2(x,y) K_T - 2(\xi^2 x + K_N)(K_T x + K_N y)}{-\mu^4(x,y)} \\ \frac{\partial^2 \tilde{s}}{\partial x \partial y} &= \frac{\mu^2(x,y) K_N - 2(\xi^2 y - Kt)(K_T x + K_N y)}{-\mu^4(x,y)} \\ \frac{\partial^2 \tilde{s}}{\partial y^2} &= \frac{-\mu^2(x,y) K_T - 2(\xi^2 y - Kt)(1 + K_N x - K_T y)}{-\mu^4(x,y)} \,. \end{aligned} \tag{20}$$

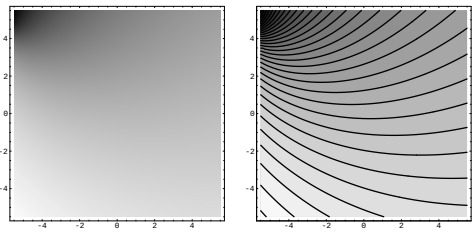

Figure 6. One instance of the dual harmonic good continuation model of Eq. 21, both as a shading patch and as a set of levelsets (highly quantized for clarity of display). This instance corresponds to $\theta(0,0) = 0$, $K_T = 0.1$, and $K_N = 0.1$, where the origin is at the center of these plots.

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 $\mathcal{G}$. Remarkably, asking $I(x,y)$ to have a vanishing Laplacian (the special case of a p-Laplacian with $p = 2$, or a critical point of the harmonic energy $\iint ||\nabla I||^2 dxdy$) 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 $I_0(s)$ takes the form

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

Combining this form with Eq. 18 gives rise to the following

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

$$\hat{I}(x,y) = I_0(0) + \frac{I_0'(0)}{\xi^2}\left[\mathit{sign}(K_T x + K_N y) K_N \cos^{-1}\omega(x,y) - K_T \log \mu(x,y)\right] \tag{21}$$

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

It should be noted that although expression 21 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 Eq. 21 represents a most unique object. It is a scalar function whose levelset geometry *and* it value structure are both harmonic functions at the same time! 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, by construction, is based on those levelset-centered local features which we discussed in Sec. 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). Fig. 6 shows one instance of this model, with and without its levelsets emphasized.

To conclude this section, and the theoretical part of this paper, it should be mentioned that an analysis similar to the one presented above proves that "dual p-harmonic" good continuation models exist for *all values of* p. No solution exists, however, for a model that its value structure is a minimal surface. For space considerations, a formal discussion and proofs of these last statements are described elsewhere.

5. Contextual inference and results

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* [14, 24], *recurrent neural networks* [13], and *belief propagation networks* [23]. Here we present results using a relaxation labeling network whose nodes $i = (x,y)$ are the image pixels and the labels λ at each node are drawn from the set of 5-tuples $\Lambda = \{(I, ||\nabla I||, \theta, K_T, K_N)\}$, where $I \in [0,1]$, $||\nabla I|| \in [0, \nabla_{max}]$, $\theta \in [-\pi, \pi)$, and $\kappa_T, \kappa_N \in [-K, K]$ (after appropriate quantization of these domains).

Initially, the measurement process assigns a confidence value, or probability $p_i^0(\lambda)$, to each possible label such that at each node $\sum_{\lambda \in \Lambda} p_i(\lambda) = 1$. The relaxation process itself drives this initial confidence distribution $p_i^0(\lambda)$ to a final (possibly ambiguous) distribution $p_i^\infty(\lambda)$. What governs the dynamics of this process, and ultimately its convergence state, is the compatibility relationships $r_{ij}(\lambda, \lambda')$ between different labels at different nodes. In our case, these compatibilities represent the degree to which two nearby pixels have consistent values based on the model for consistency developed above (Proposition 2). For lack of space we omit the rest of the technical details and refer the reader to the many papers on the subject that describes its theoretical foundations (e.g., [14, 24]).

We have examined the performance of this approach on a variety of shading images, both synthetic and natural. Three of the benchmark images we used are presented in Fig. 7. Since both values and levelset geometry are important, we show both aspects in all figures.

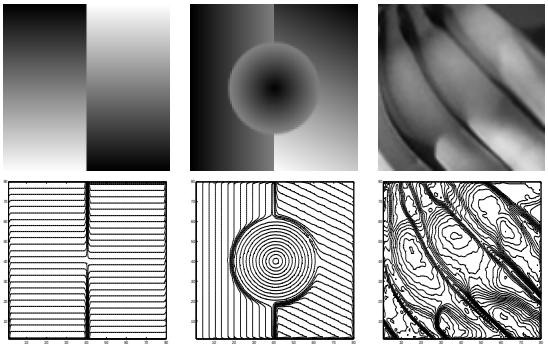

Figure 7. Test images (gradient edge, gradient patches, and bananas) and their levelsets geometry.

Fig. 8 illustrates the results of the inference of coherent structure on the input images, and compares these results to state-of-the-art anisotropic diffusion processes [25, 30] to show how this computation is able to provide superior organization that effectively eliminates spurious measurements and noise while preserving both singularities and much of the underlying levelset geometry. All relaxation results are shown after 5 iterations. Diffusion results are shown after 200 iterations on clean images (to illustrate the qualitative effect of the process on typical visual structure), after 100 iterations on images with additive noise (which was near the minimal number of iteration required to get rid of most, but not necessarily all of the noise), and after 500 iteration on images with Salt and Pepper noise (which was never enough to eliminate the noise). Observe the near flawless behavior of the good continuation process in comparison to these contemporary diffusion schemes, not only in terms of noise removal and preservation of perceptual edges, but more so in terms of preserving and reconstructing the general levelset geometry. In comparison, diffusion of the noisy images distorts these structures long before noise is removed (though this is true in general, it is particularly evident with Salt and Pepper noise). Results on clean images are shown in order to illustrate qualitatively the effect of all these processes on typical visual structure which without noise should be preserved as long as possible.

6. Summary

This paper presents a theory, a formal derivation, and a method for the organization of coherent structure of general 2D scalar visual features based on the principle of good continuation. We have shown how geometrical good continuation may be applied to general functions by considering both their levelsets and the distribution of values and have developed a formal model in which both aspects are harmonic functions *simultaneously* (extendable to p-harmonic for any value of p). We have used this model in a relaxation labeling computation and demonstrated how this perceptual organization-based approach provides superior performance in eliminating spurious measurements while preserving singularities and much of the underlying levelset geometry.

Acknowledgements: This work was supported by the Toman and Frankel funds at Ben-Gurion University, and the Offi ce of Naval Research (ONR). We thank Chen Keasar for valuable CPU time on his BGU cluster and Brian Cabral and Leith Leedom for making their LIC visualization code available.

References

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

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

[3] W. Bosking, Y. Zhang, S. B., and D. Fitzpatrick. Orientation selectivity and the arrangement of horizontal connections in the tree shrew striate cortex. *J. Neurosci.*, 17(6):2112–2127, 1997.

[4] P. Breton, L. Iverson, M. Langer, and S. Zucker. A new approach to shape from shading. In G. Carpenter and S. Grossberg, editors, *Neural Networks for Vision and Image Processing*, chapter 5, pages 111–131. The MIT Press, 1992.

[5] B. Carbal and L. Leedom. Imaging vector fields using line integral convolution. In *Proc. of SIGGRAPH*, pages 263–270, 1993.

[6] V. Caselles, B. Coll, and J.-M. Morel. Geometry and color in natural images. *J. of Math. Imaging and Vision*, 16:89–105, 2002.

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

[8] M. Fischler. The perception of linear structure: A generic linker. In *IUW*, volume 2, pages 1565–1579, 1994.

[9] G. Guy and G. Medioni. Inferring global perceptual contours from local features. *Int. J. Comput. Vision*, 20(1/2):113–133, 1996.

[10] G. Guy and G. Medioni. Interence of surfaces, 3d curves, and junctions from sparse, noisy 3d data. *IEEE Trans. Pattern Anal. Machine Intell.*, 19(11):1265–1277, 1997.

[11] R. Hess and D. Field. Integration of contours: New insights. *Trends Cogn. Sci.*, 3:480–486, 1999.

[12] D. Hoffman. *Visual Intelligence*. W.W. Norton & Company, 1998.

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

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

[15] A. Jain. *Fundamentals of Digital Image Processing*. Prentice-Hall International, Inc., 1989.

[16] F. John. *Partial Differential Equations*. Springer-Verlag New York Inc., 1982.

[17] B. Kimia, L. Frankel, and A. Popescu. Euler spiral for shape completion. In *IEEE Computer Society Workshop on Perceptual Organization in Computer Vision*, 1999.

[18] J. Koenderink and A. van Doorn. Image processing done right. In *ECCV*, pages 158–172, 2002.

[19] K. Koffka. *Principles of Gestalt Psychology*. Routledge & Kegan Paul Ltd., 1935.

[20] Z. Li. A neural model of contour integration in the primary visual cortex. *Neural Comput.*, 10(4):903–940, 1998.

[21] D. Mumford. Elastica in computer vision. In B. Chandrajit, editor, *Algebric Geometry and it's applications*. Springer-Verlag, 1994.

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

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

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

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

[26] W. Richards, J. Keonderink, and D. Hoffman. Inferring 3d shapes from 2d silhouettes. *J. Opt. Soc. Am.*, 4:1168–1175, 1987.

[27] R. Šára. Isophotes: the key to tractable local shading analysis. In *CAIP*, pages 416–423, 1995.

[28] K. Schmidt, R. Goebel, S. Löwel, and W. Singer. The perceptual grouping criterion of colinearity is reflected by anisotropies in the primary visual cortex. *Eur. J. Neurosci.*, 9:1083–1089, 1997.

[29] A. Sha'ashua and S. Ullman. Structural saliency: The detection of globally salient structures using a locally connected network. In *ICCV*, pages 321–327, 1988.

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

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

[32] S. Ullman. Filling in the gaps: The shape of subjective contours and a model for their creation. *Biol. Cybern.*, 25:1–6, 1976.

[33] M. Wertheimer. Laws of organization in perceptual forms. In W. Ellis, editor, *A source book of Gestalt Psych.*, pages 71–88. Routledge & Kegan Paul Ltd., 1955.

[34] L. Williams and D. Jacobs. Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Comput.*, 9(4):837–858, 1997.

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

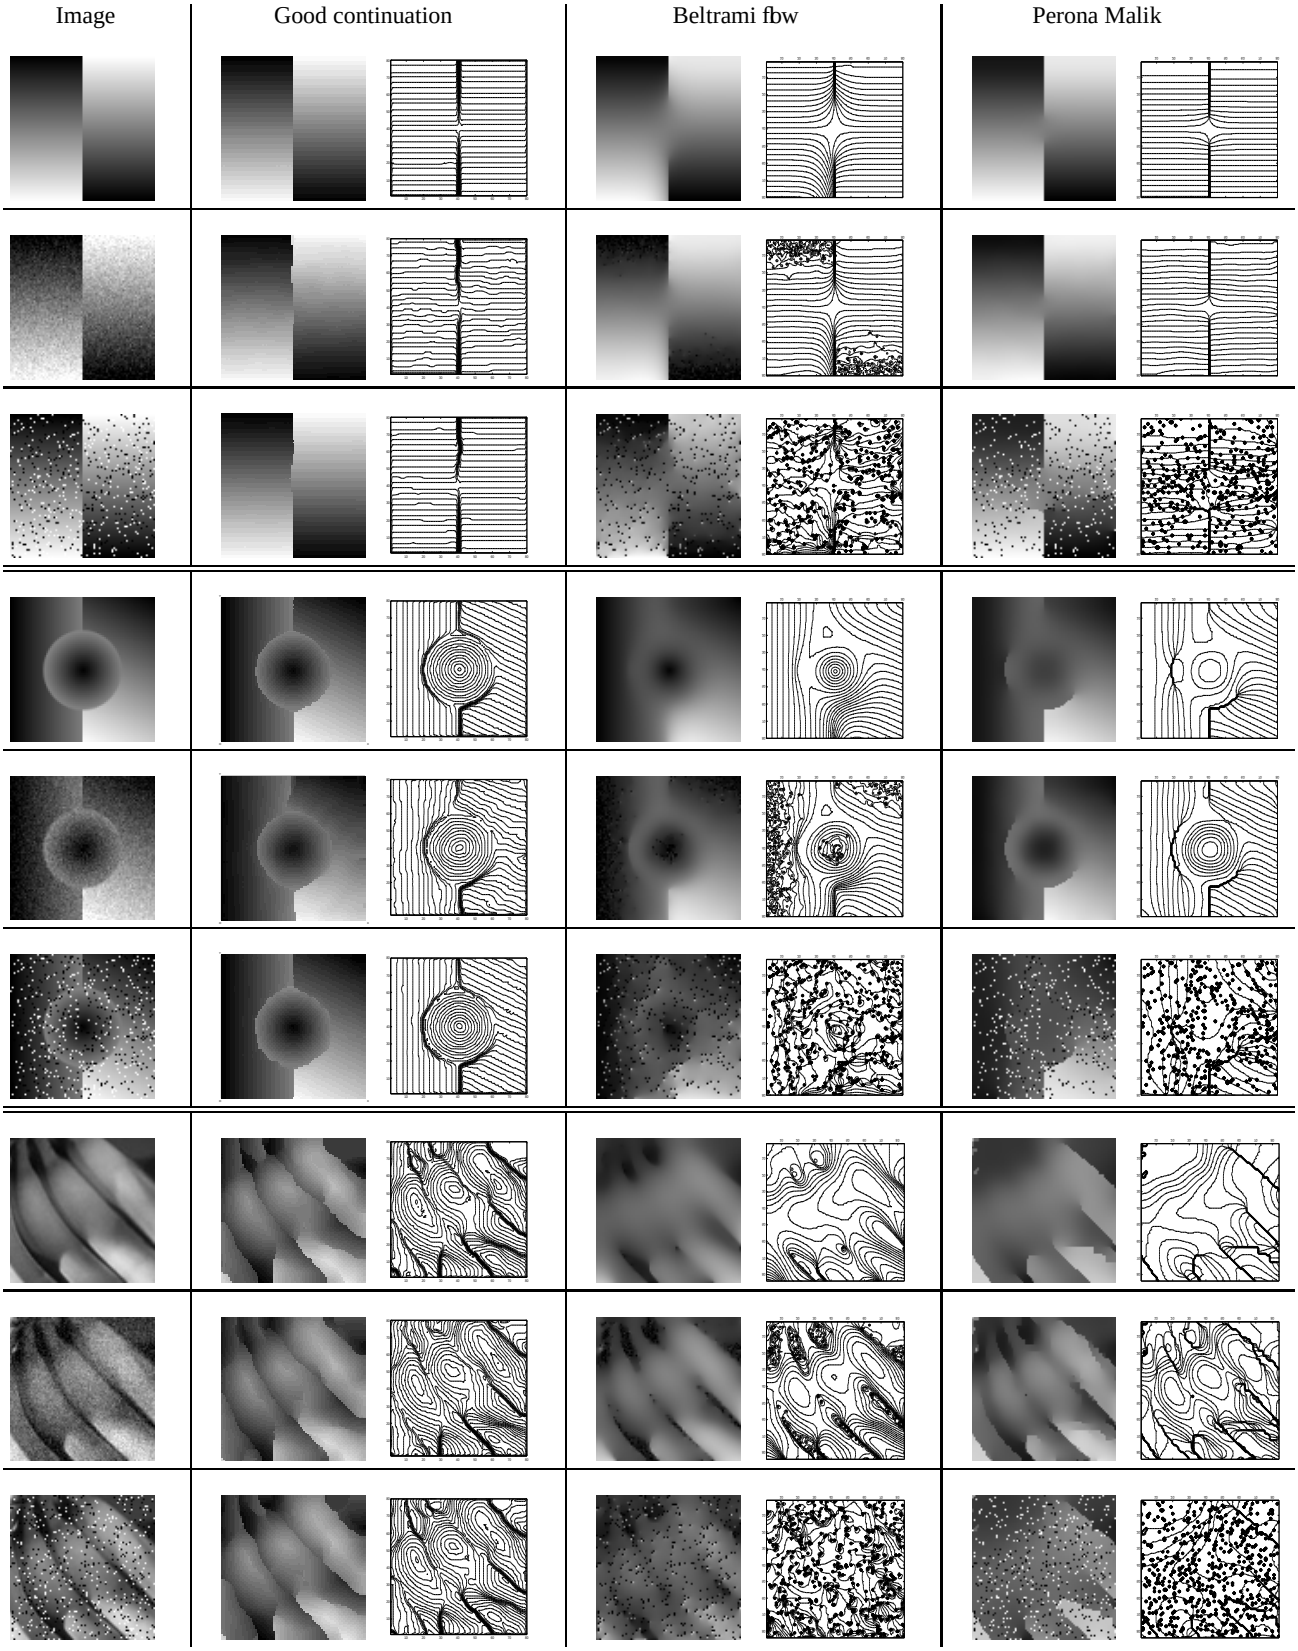

Figure 8. Results of applying the good continuation inference process on a synthetic and natural shading patterns (unless printed on a quality printer, details in several images may be difficult to observe. In this case, please refer to the electronic version and zoom in on individual panels). Different 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. Please refer to the main text for additional explanations.