

Saliency and Segregation *Without* Feature Gradient: New Insights for Segmentation from Orientation-Defined Textures

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

Abstract

The analysis of texture patterns, and texture segregation in particular, are at the heart of perceptual organization. In this paper we question the widely accepted view that the detection (both perceptual and computational) of salient perceptual singularities between perceptually coherent texture regions is tightly dependent upon feature gradients. Specifically, we study smooth orientation-defined textures (ODTs) and show that they exhibit striking perceptual singularities even without any outstanding gradients in their defining feature, namely orientation. We further show how these generic singularities are not only unpredictable from the orientation gradient, but that they also defy popular segmentation algorithms and neural models. We then examine smooth ODTs from a (differential) geometric point of view and develop a theory that fully predicts their perceptual singularities from two ODT curvatures. The computational results exhibit striking correspondence to segregation performed by human subjects and provide a conclusive evidence for the role of curvature in texture segregation. Extensions and implications of our results are developed for various aspects of visual processing.

1. Introduction

The ability to effortlessly segregate texture stimuli into coherent parts has long been attributed to rapid changes (or high contrast) in the spatial distribution of elementary visual features - a notion that had become known in the study of texture segregation as *feature gradients* [1, 32, 16, 31, 36], and which motivated some of the most popular texture segregation algorithms to-date (e.g. [25]). In this work we revisit the role of feature gradients in the context of *orientation-define textures* (ODTs) and *orientation-based texture segmentation* (OBTS). While ODTs are frequent in natural and artificial visual stimuli, textures are rarely characterized solely by orientation. Nevertheless, understanding the effect of orientation on texture segregation is essential due to its

neurophysiological basis [14], its central role in perceptual organization [16, 17], and its close relationship to shape perception [40, 41].

Just like feature gradients in general, *orientation gradients* were shown to play a key role in OBTS and in explaining behavioral results in human observers (e.g., [29, 30, 22]). However, much of this insight is due to the fact that the overwhelming majority of these studies have concentrated on ODTs of *piecewise-constant* orientation, where presumed perceptually coherent regions were defined by constant orientation, while salient *perceptual singularities*¹ emerged from high orientation contrasts between these regions (e.g. [28, 22, 45, 21, 23]). While such exploration is valid in a sense of employing simplified stimuli for use in lab experiments, it is such a gross oversimplification of general ODTs as to obscure key aspects of the visual process itself. Interestingly enough, regardless of their formation process, ODTs in natural stimuli are seldom constant or piecewise constant since this requires an accidental match between the surface geometry, the texture formation process, and the observer's view-point. Furthermore, perspective projection dictates that even completely parallel lines in the world are likely to give rise to a non-constant ODT in the image. Indeed, as we discuss in this paper, the larger context of piecewise *smooth* ODTs reveals entirely new aspects and much wider scope for OBTS, both perceptually and computationally. Broader implications for perceptual organization and other aspects of vision are discussed in Sec. 6.

2. New phenomenological motivation

One prediction from existing segmentation models is that without significant feature gradient/contrast textures are perceptually coherent (i.e., neither segregated into distinct regions nor exhibit any other perceptual singularities). As implied by perceptual studies, as well as by computational algorithms (e.g. [25]), an ODT described by the orienta-

¹We use the notion of *perceptual* singularities to emphasize the difference from genuine singularities in the visual signal.




Figure 1. Perceptual singularities in piecewise constant ODTs are well predicted by orientation gradients. (a) A typical piecewise constant ODT. This ODT, as well as most other textures illustrated in this paper, are visualized and generated using the Line Integral Convolution (LIC) method [6]. (b) The map of orientation gradient magnitude (depicted as intensity) corresponds well to the segregation of the texture by human observers.

tion function $\theta(x, y)$ is therefore likely to exhibit perceptual inhomogeneities, perceptual singularities, or perceptual boundaries, along curves and points where $|\nabla(\theta * G_\sigma)|$ crosses some threshold, where the wrapped Gaussian G_σ is a blurring kernel and the shift in scale space is required to allow a well behaved gradient everywhere. Indeed, as shown in Fig. 1, this gradient-based approach accurately predicts perceptual singularities for the well studied ODTs of piecewise constant orientation.

Unfortunately, however, the predictions of this approach become utterly wrong in the general case of ODTs with (piecewise-) *smoothly varying* orientation. Such smoothly varying ODTs almost always exhibit striking perceptual singularities – curve-like structures that are significantly more salient than any other region of the ODT. These singularities are always tangent to the ODT (i.e., they coincide with certain streamlines, or segments thereof) and they are responsible for the segregation of the ODT into perceptually distinct regions. Importantly, this segregation, and these perceptual singularities, have *no apparent relationship to the orientation gradient* of the ODT. For example, the pattern in Fig. 2a has *constant* orientation gradient across the entire image (Fig. 2b). In other words, nowhere in this pattern the orientation changes more rapidly or differently from any other part. Still, virtually all observers segregate this pattern into diagonal bands separated by perceptually salient border lines. Evidently, something in the spatial relationship between the orientation textons [16], and not merely their local contrast, determines the saliency of different regions and the segregation of the pattern². Similarly, the ODT in Fig. 2c exhibits a salient double spiral structure that has no apparent relationship to the orientation gradient of this pattern (shown in Fig. 2d). It is therefore clear that in order to model these saliency differences and segregation correctly, something beyond orientation gradients (or feature gradients

²Note also that the perceived perceptual singularities have nothing to do with the wrapping around of the orientation itself since they are invariant to Euclidean transformations of the image.

in general) must be considered.

That salient perceptual singularities in smooth ODTs are independent on orientation gradients is easily demonstrated by creating different ODTs with identical orientation gradient map and comparing the perceptual outcome. In particular, if $\theta_0(x, y)$ is the orientation function of a given ODT, it is guaranteed that its orientation gradient is identical to that of the ODT defined by the orientation function $\theta_1(x, y) = \Delta\theta + \theta_0(x, y)$, where $\Delta\theta$ is an arbitrary phase shift. One such a pair is demonstrated in Fig. 2e,f to show how drastically different the structure of perceptual singularities in textures of identical feature gradient can be.

The perceptual singularities exemplified in Fig. 2 are extremely consistent across observers³. They are also very robust and insensitive to the specific visualization method of the ODT, and they persist even in sparse representations (Fig. 3). It may be surprising, then, that these singularities, and the resultant segregations, are poorly predicted by most, if not all, segmentation algorithms available to-date. We have thoroughly tested a variety of algorithms on images of smooth ODTs and show selected results in Fig. 4. Clearly, none of these algorithms is able to replicate human perception in these cases. In part, however, this result was expected. The notion of feature gradient, and quite often that of orientation gradient, are grounded in many segmentation methods, either explicitly, via a gradient-like measure, or implicitly, via a pixel or region similarity factor. Since, as we have shown, orientation gradients are not a reliable determinant for perceptual singularities in ODTs, it is unlikely that segmentation methods that rely on these gradients could provide the correct segmentation results.

Given all these observations and the introduction of the

³In an extensive behavioral study we have examined a variety of psychophysical aspects of the perception of singularities in smooth ODTs. The results are omitted from this paper but reported elsewhere [2].




Figure 3. Perceptual singularities in piecewise constant ODTs are robust to the specific visualization method. Here we show the same ODT depicted as a LIC dense texture [6] and as a sparse array of jittered oriented segments (texels). The same global structure mediated by the perceptual singularities is equally salient in both. The sparse, texel-based representation is particularly intriguing in this sense since its salient structures are neither better aligned nor less curving than certain non salient configurations (e.g., between the black dots). Evidently, saliency is not determined by element to element matching based on orientation difference or good continuation [10]. But if not, then how?

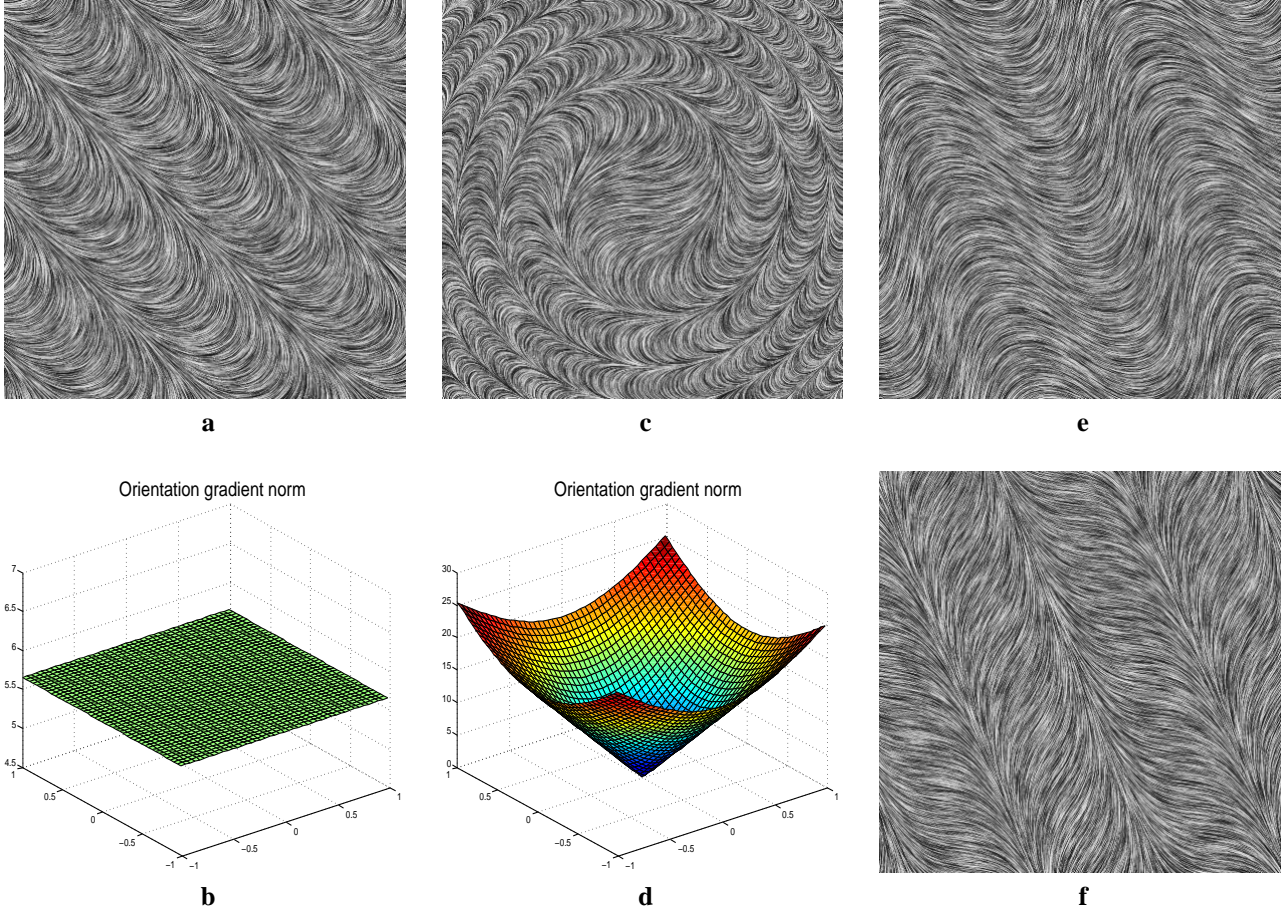


Figure 2. perceptual singularities in *smoothly varying* ODTs are poorly predicted by orientation gradients. **(a)** A smoothly varying ODT defined by the function $\theta(x, y) = c \cdot x + c \cdot y$. Virtually all observers report segregation of this pattern into multiple diagonal bands separated by salient border lines. **(b)** The graph of the orientation gradient magnitude of the pattern in **a**. Clearly, there is fundamental gap between the (inhomogeneous) perceptual outcome and the prediction based on the (constant) orientation gradient. **(c)** This smooth ODT exhibits salient double spiral structure that is easily segregated by human observers. **(d)** The graph of the orientation gradient magnitude of the pattern in **c**. Again, the perceptual outcome clearly is dissociated from the feature gradient. **(e+f)** This pair of smooth ODTs are different only by a constant phase shift: $\theta_e(x, y) = \pi/8 + \sin(5x + 2y)$ where $\theta_f(x, y) = \theta_e(x, y) + \pi/2$. Despite having an identical feature (i.e., orientation) gradient across the pattern, the perceptual outcome is drastically different. In fact, virtually all observers report no segregation for θ_e while θ_f is consistently segregated into diagonal bands along salient straight perceptual singularities. (Note: to avoid aliasing and other artifacts, please print the ODT images on a high resolution color printer or use the electronic version to view on a monitor after extending each to decent size).

hitherto overlooked singularities in smoothly changing textures, we next examine the (geometrical) nature of smooth ODTs, we explore possible factors that effect their local saliency perceptual singularities, and we devise a computational method that accurately predicts and detects them. The contribution of this paper is therefore twofold, both for better understanding perceptual organization at the perceptual level, and for improving performance of computational methods to handle a novel aspect of segmentation and grouping.

3. Geometrical foundations

Consider an ODT in the image plane. An abstract representation for this ODT would make explicit the orientation at each point, hence an ODT can be described as an orien-

tation function $\theta(x, y)$ in the image plane or as a unit length vector field $\vec{E}_T(x, y)$ tangent to the ODT at each point. In this sense an ODT is an oriented pattern [19], an oriented texture field [35], or a texture flow [4]. Here we follow the theoretical foundations developed in the latter to derive a computational measure that makes accurate predictions of the perceptual outcomes.

An extension of the vector field representation $\vec{E}_T(x, y)$ that makes tools from differential geometry readily available is the *frame field* representation [33]. More specifically, by attaching a second unit vector $\vec{E}_N(x, y)$ to each point (x, y) such that $\vec{E}_T \cdot \vec{E}_N = 0$ one obtains a frame at each point on the ODT, and therefore a *frame field* $\{\vec{E}_T, \vec{E}_N\}$ in the image domain. This frame field is not simply a redundant representation; it also provides a *local coordinate system* in which all other vectors can be represented in a natural, object




Figure 4. Perceptual singularities in *smoothly varying* ODTs are poorly detected by existing segmentation algorithms. This table illustrates selected segmentation results compared to manual segregations by human observers. In all cases, the algorithmic results bear no resemblance to the perceptual one. (a) Original image (b) Typical manual segregation by a human observer. In these behavioral studies (see [2] for details) stimuli were presented with fading margins to minimize potential boundary effects. User segmentations therefore correspond to the viewable regions only. (c) Segmentation result using normalized cuts [39]. (d) Segmentation result using mean shift segmentation [7] with parameters tuned to handle best the texture regions of our ODTs ($h_s = 8, h_r = 6.5, M = 100$). (e) Boundary probability by multi cue segmentation due to [26]. All segmentations were produced from publicly available code released by the authors of the different methods. All shown images were scaled down significantly for space reasons. To appreciate these results it is recommended to zoom in on each image using the electronic version of the paper.

centered view (Fig. 5). Perhaps the most important vectors (other than the frame vectors themselves) that asks for such an object centered representation are the *covariant derivatives* of \vec{E}_T and \vec{E}_N . These covariant derivatives represent the initial rate of change of the frame in any given direction \vec{V} , a quantity which in the $\{\vec{E}_T, \vec{E}_N\}$ coordinates is captured by Cartan's *connection equation* [33]:

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

The coefficient $w_{12}(V)$ is a function of the tangent vector V , which reflects the fact that the local behavior of the flow depends on the direction along which it is measured. Fortunately, $w_{12}(V)$ is a *1-form* [33] and thus linear. This allows us to fully represent it with two scalars at each point since

$$w_{12}(V) = w_{12}(a \vec{E}_1 + b \vec{E}_2) = a w_{12}(\vec{E}_1) + b w_{12}(\vec{E}_2).$$

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

$$\begin{aligned} \kappa_T &\triangleq w_{12}(\vec{E}_T) \\ \kappa_N &\triangleq w_{12}(\vec{E}_N) \end{aligned} \quad (2)$$

These two scalars, defined at each point of the ODT, are called its *tangential curvature* (κ_T) and *normal curvature* (κ_N), respectively [4], and they represent the initial rate of change of the ODT orientation in its tangential and normal

directions, respectively. Practically, these two curvatures can be evaluated at each point as the coefficients of the orthogonal expansion [33] of the orientation gradient of the ODT (relative to a fixed coordinate system) based on the frame itself

$$\begin{aligned} \kappa_T &= \nabla \theta \cdot (\cos \theta, \sin \theta) \\ \kappa_N &= \nabla \theta \cdot (-\sin \theta, \cos \theta). \end{aligned} \quad (3)$$

We note that although Eq. 3 provides signed curvatures, as indeed is possible in the plane (e.g. [8, p. 21]), in the following we will be interested in their *absolute values* (i.e., in $|\kappa_T|$ and $|\kappa_N|$) since the orientation of the ODT is determined only up to π and since the sign of curvature appears to play no role in the segregation process.




Figure 5. The intrinsic local geometry of smooth ODTs is best captured by its representation as a differentiable frame field which is everywhere tangent and normal to the direction of the flow (see pairs of red vectors). An infinitesimal translation of the frame in a direction V rotates it by some angle determined by the connection form of the frame field. Since the connection form is a linear operator, it is fully characterized by two numbers obtained by orthogonal expansion. The natural expansion based on the frame itself yields the two curvatures κ_T and κ_N .

4. A theory of saliency in ODTs

One interesting property of ODT curvatures is that unlike orientation gradients, or their components, in general these curvatures cannot be simultaneously constant in a neighborhood of the ODT, however small. This constraint is a consequence of a more general differential constraint expressed as follows [4]:

$$\nabla \kappa_T \cdot \vec{E}_N - \nabla \kappa_N \cdot \vec{E}_T = \kappa_T^2 + \kappa_N^2. \quad (4)$$

Eq. 3 further suggests that even if the orientation gradient $\nabla \theta$ remains constant, the two ODT curvatures will change in a *periodic* manner. This formal observation becomes particularly significant once Fig. 2a is scrutinized. In this figure $\nabla \theta$ is constant across the entire pattern but perceptually this ODT exhibits *periodic* singularities. Is this link accidental or does it indicate a deeper relationship between abstract geometry, and in particular curvatures, and perceptual organization in the human visual system?

Further insight into this possibility is revealed by examining the actual behavior of the two curvatures (in absolute value) of this ODT, as shown in Fig. 6a-c. Evidently, in this case, the straight line-like perceptual singularities are parallel to the levelsets of the two curvatures. Strikingly, aligning the texture pattern with the two curvature maps shows that its perceptual singularities in fact *coincide* with $|\kappa_T(x, y)|$'s zeros on one hand, and with $|\kappa_N(x, y)|$'s maxima on the other. Does this relationship represent a universal computational rule for detecting perceptual singularities in smooth ODTs?

Unfortunately, $|\kappa_T(x, y)|$'s zeros cannot serve as such a universal measure. Although this criterion is appealing because of its explicit link to much research into the role of *collinearity* in perceptual organization and biological vision (e.g. [11, 18, 37]), there are at least two reasons that renders it less useful than expected. First, $\kappa_T(x, y)$ may vanish along a perceptual singularity only if the latter is straight (as in Fig. 6a). In general, however, since perceptual singularities in smooth ODTs bend and curve (as in Figs. 2c and 3), the tangential curvature along them will not vanish. Second, streamlines of zero tangential curvature often are the *least* salient structure in ODTs, as is demonstrated in Fig. 6d. Blindly detecting structures of zero tangential curvature is therefore prone to provide false positives in these cases.

Unfortunately still, using $|\kappa_N(x, y)|$'s maxima to detect perceptual singularities also presents some difficulties. First, these maxima come at different values that vary both within and between ODTs. Detection via global thresholding is therefore prone to fail. More critical is that $|\kappa_N(x, y)|$'s maxima actually fail to localize correctly certain salient singularities. Both phenomena are exemplified in Fig. 6e-h using the ODT defined by

$$\theta(x, y) = \frac{\pi}{4} + \left(a + \frac{b}{1 + e^{\frac{c-(x-y)}{d}}} \right) (x - y)$$

where a, b, c, d are constants. The magnitude of orientation gradient of this ODT has a sharp sigmoid-like profile in the direction

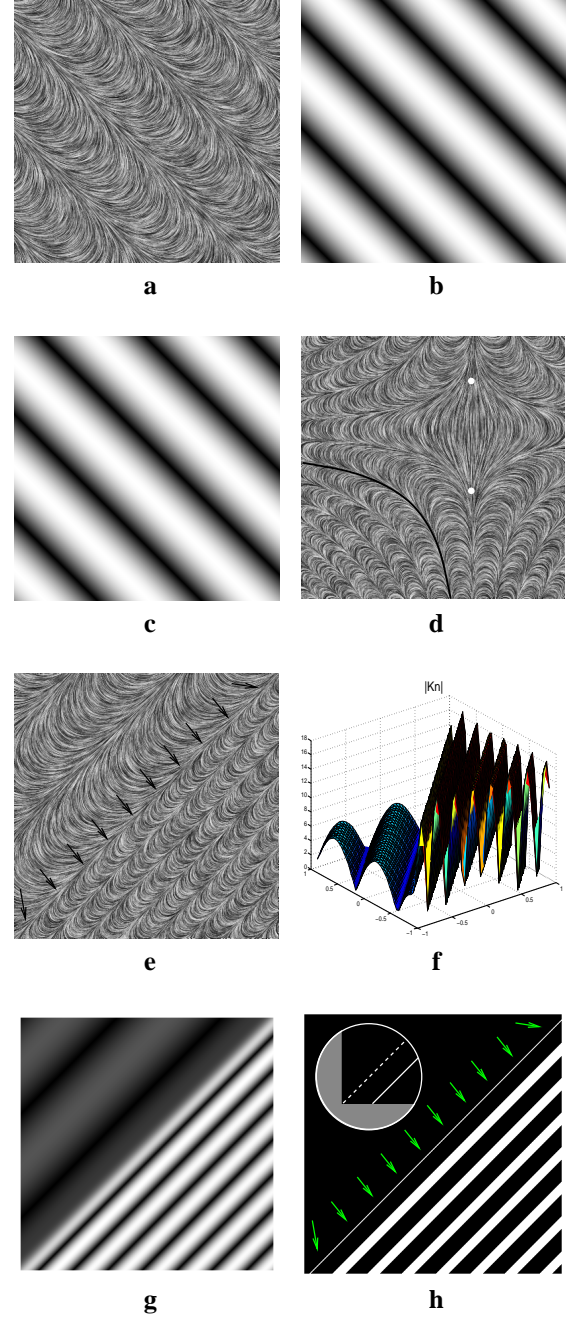


Figure 6. Preliminary observations about the relationship between ODT curvatures and perceptual singularities. (a) ODT of constant orientation gradient (same as Fig. 2a). (b) $|\kappa_T(x, y)|$ of the ODT from panel a, depicted as intensity. Note how the perceptual singularities coincide with $|\kappa_T(x, y)|$'s zeros. (c) $|\kappa_N(x, y)|$ of the same ODT. Note the correspondence of the singularities coincide with $|\kappa_N(x, y)|$'s maxima. Compare to manual segregation in Fig. 4. (d) $|\kappa_T(x, y)|$'s zeros cannot serve as a universal saliency measure. In general, singularities will emerge along curving lines (black solid curve) while curves of zero curvatures often will not be salient at all (streamline between two white dots). (e) ODT whose orientation gradient has a sigmoid-like profile parallel to the main diagonal (top left to bottom right). Note the singularity that stretches corner to corner (marked with arrows) between regions in which additional singularities occur in different frequencies. (f) $|\kappa_N(x, y)|$ of the ODT in panel e. Note the different maxima values. (g) $|\kappa_N(x, y)|$ plotted as intensity. (h) $|\kappa_N(x, y)|$ thresholded just below the maximum value of the central singularity. Note how this maximum curve (solid line in the inset) is offset relative to the perceptual singularity (dashed line in the inset).

of the main diagonal, and constant zero profile in the direction of the other diagonal. Hence, this smooth ODT provides an opportunity to examine perceptual singularities that separate regions of significantly different orientation gradients, as is the case with the central singularity in this figure. Indeed, one side effect of this configuration is the false localization of this central singularity by $|\kappa_N(x, y)|$'s maxima.

If either κ_T or κ_N accurately predicted the salient regions and perceptual singularities in ODTs, it would be a strong indication that the other curvature is less important for perceptual organization in both man and machine. In this sense, one perhaps could expect that the normal curvature κ_N would prove insignificant, since despite its theoretical justification [4], no psychophysical or physiological evidence currently supports its involvement in early vision. It is therefore a striking result that a computational measure (henceforth called the Perceptual Singularity Measure, or PSM) that accurately pinpoints perceptual singularities and salient curves in ODTs involves *both* curvatures:

Proposition 1 *The locus of perceptual singularities in smoothly varying ODTs of orientation function $\theta(x, y)$ is defined by the ridges of the local differential measure*

$$PSM(x, y) = \prod_{\kappa_T^2 + \kappa_N^2 > \tau} \left[\frac{\kappa_N(x, y)^2}{\kappa_T(x, y)^2 + \kappa_N(x, y)^2} \right] \quad (5)$$

where $\kappa_T(x, y)$ and $\kappa_N(x, y)$ are defined by Eq. 2 and the operator \prod rectifies its argument by the condition specified.

It is important to realize that the point-wise normalization by the total variation $\kappa_T^2 + \kappa_N^2$ not only re-localizes detected singularities and pinpoints them to the perceptual outcome (see Fig. 7 and results below), but it also provides a normalized measure that is restricted to the interval $[0, 1]$ and therefore easier to use. The rectifier requires in order to avoid detection of spurious singularities in regions where the ODT changes too slowly to trigger any segregation. The threshold for this operation is determined psychophysically [2]. It should be noted that the normalization also suggests that an equivalent way to perceptually organize ODTs is to extract the *valleys* of the coupled measure

$$\overline{PSM}(x, y) = \prod_{\kappa_T^2 + \kappa_N^2 > \tau} \left[\frac{\kappa_T(x, y)^2}{\kappa_T(x, y)^2 + \kappa_N(x, y)^2} \right] \quad (6)$$

thereby placing both curvatures on completely equal footing in terms of their contribution to the perceptual organization process.

5. Experimental results

We have exhaustively tested the proposed computational measure for saliency and perceptual singularities in smooth ODTs on a variety of ODT patterns. Textures were generated from randomly selected quadratic orientation functions for which curvatures, $PSM(x, y)$, and other differential properties, were computed numerically using Matlab. The results on a variety of other types of ODTs are omitted for space considerations. Despite much debate in the literature on the proper way to compute ridges (e.g., [12, 43, 27, 20, 9], here we use a local method by López et al. [24] that provides excellent numerical stability.




Figure 7. Our Perceptual Singularity Measure applied on the asymmetric ODT from Fig. 6e. (a) $PSM(x, y)$. Compare to Fig. 6f. (b) $PSM(x, y)$ depicted as intensity. Compare to Fig. 6g. (c) $PSM(x, y)$ ridges. Note how this time the central singularity (marked with arrows) is localized accurately. Ridges in this case are particularly easy to detect since they coincide with the points where $PSM(x, y) = 1$.

Results of our theory on selected ODTs are demonstrated in Fig. 8 (also in Fig. 7) and illustrate the outstanding match between perception and the proposed computational theory. Performance of our saliency measure was compared to manual segregations by performed human observers who viewed LIC visualizations [6] of the same ODTs. While our psychophysical work examine multiple aspects of segmentation without feature contrast, and incorporate numerous experiments using several experimental paradigms, in the current computational paper we only bring samples of manually drawn segregations for comparison with the computational method. Readers interested in the behavioral aspects of this work are kindly referred to [2].

Strikingly, the same saliency measure works also for *piecewise constant* textures (after infinitesimal shift in scale space) and therefore fully generalizes existing theories based on feature gradients. Although a theoretical account on this generalization is part of our forthcoming work, here we briefly illustrate it (Fig. 9) for one piecewise constant texture and one ODTs that combines both types of singularities (with and without feature gradient)

6. Summary, implications, and future work

In this paper we seek to make several contributions to computational perceptual organization based on a novel examination of texture segregation by the human visual system. First and foremost, we argue that perceptual organization and the detection of perceptual singularities in textures cannot be determined reliably by feature gradients. Although high feature gradients often do signal perceptual singularities, *the lack of the former does not imply perceptual coherence*, as is vividly demonstrated by the hitherto neglected smooth ODTs presented here. Second, we have shown that the accurate detection of perceptual singularities in ODTs, both *with* and *without* feature gradients, emerges directly *and solely* from their intrinsic geometry and curvature properties. This not only provides better computational results than existing segmentation methods, but it also provides (1) a strong evidence for the advantage of a theoretical framework suggested recently in the literature [4] and (2) novel insights for better understanding of segregation processes in the human visual system.

While this paper focuses on ODTs and OBTS, we believe that its implications are much broader in scope. Orientation is explicit not only in ODTs but also in motion and optic flows and therefore the results developed in this paper are directly applicable to




Figure 8. Experimental results of our saliency measure and detection of perceptual singularities in smooth ODTs. (a) ODT. (b) Manual segregation by human observer (stimuli were presented with fading margins to minimize potential boundary effects. User segregations therefore correspond to the viewable regions only). (c) Magnitude of the orientation gradient of the pattern. Note the lack of correspondence with the perceptual outcome and manual segregations. (d) PSM of the pattern (without rectification). (e) PSM ridges (with rectification). Compare to manual segregations in column b.

motion-based segmentation. Orientation is also implicit in other visual cues that take part in the segmentation process, in particular shading [5] and color [3]. The results presented in this paper therefore imply that the role of these other cues in early vision and segmentation should be revisited. Our forthcoming work concentrates on handling these visual cues based on the insights discussed in this paper and combining them all into a novel approach to the segmentation of general textures and natural images.

The results presented in this paper also carry implications for another central issue of perceptual organization, namely that of contour integration. Indeed, grouping edge elements into salient curve structures has been motivated by Gestalt principles like proximity and good continuation, both in behavioral studies (e.g. [10, 42]) and computational ones (e.g. [38, 34, 23]). But just like the study of texture segregation has been dominated by piecewise constant feature distributions, curve integration has mostly gained insights from curve-like structures embedded in *noisy and cluttered* background (e.g. [38, 13, 44]). Once the noisy background is replaced with a structured one, classical saliency results become invalid and 1D good continuation fails to explain or replicate correctly perceptual grouping [10, pp. 191-192]. One such example was illustrated in the texel array in Fig. 3 for which existing curve integration theories would falsely prioritize the diagonal collinear configuration between the black dots in the bottom left quadrant.

The solution to this difficulty lies in the theory presented in our current work, which both solves an open question from the perception literature ([10, pp. 191-192]) and suggests ways to handle it computationally.

Interestingly enough, the results presented in this paper are also linked to aspects of vision not traditionally related to perceptual organization, for example to the perception of 3D shape. Often, when viewed from a slanted view point or when forced to make a shape interpretation, observers of smooth ODTs of the kind shown here report the perception of 3D terrains with narrow ridges and valleys⁴ that correspond to the salient perceptual singularities signaled by our computational measure. Why this is the case may correspond to how shape is perceived from surface contours [40, 41], or due to generic edge-flow configurations that define occluding contours [15]. Closer examination of these links are part of our future work.

Acknowledgements: This work was supported by the Toman fund, the Frankel fund, and the Paul Ivanier center for Robotics Research at Ben-Gurion University. I thank Brian Cabral and Leith Leedom for making their LIC visualization code available, and Antonio López for sharing his ridge computation code.

⁴Try viewing the ODT in Fig. 2c after slanting the page.




Figure 9. Segregation results on *piecewise* continuous textures based on our proposed saliency measure. (a) Piecewise constant texture of the sort presented in Fig. 1. (b) Piecewise smooth ODT that combines both types of singularities (with and without feature gradient). In both cases detected perceptual singularities are marked as white curves.

References

- [1] J. Beck. Effect of orientation and the shape similarity on perceptual grouping. *Percept. Psychophys.*, 1:300–302, 1966. 1
- [2] O. Ben-Shahar. Perceptual singularities in smooth orientation-defined textures: Segregation without feature gradient. In *VSS*, 2006. 2, 4, 6
- [3] O. Ben-Shahar and S. Zucker. Hue fields and color curvatures: A perceptual organization approach to color image denoising. In *Proc. Computer Vision and Pattern Recognition*, pages 713–720, 2003. 7
- [4] 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. 3, 4, 5, 6
- [5] P. Breton and S. Zucker. Shadows and shading flow fields. In *Proc. Computer Vision and Pattern Recognition*, pages 782–789, 1996. 7
- [6] B. Carbal and L. Leedom. Imaging vector fields using line integral convolution. In *Proc. of SIGGRAPH*, pages 263–270, 1993. 2, 6
- [7] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Machine Intell.*, 24(5):603–619, 2002. 4
- [8] M. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Inc., 1976. 4
- [9] D. Eberly. *Ridges in Image and Data Analysis*. Kluwer Academic Press, 1996. 6
- [10] D. Field, A. Hayes, and R. Hess. Contour integration by the human visual system: Evidence for a local “association field”. *Vision Res.*, 33(2):173–193, 1993. 2, 7
- [11] W. Geisler, J. Perry, B. Super, and D. Gallogly. Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res.*, 41(6):711–724, 2001. 5
- [12] R. Haralik. Ridges and valleys on digital images. *CVGIP*, 22:28–38, 1983. 6
- [13] L. Herault and R. Horaud. Figure-ground discrimination: A combinatorial optimization approach. *IEEE Trans. Pattern Anal. Machine Intell.*, 15(9):899–914, 1993. 7
- [14] D. Hubel and T. Wiesel. Functional architecture of macaque monkey visual cortex. In *Proc. R. Soc. London Ser. B*, volume 198, pages 1–59, 1977. 1
- [15] P. Huggins and S. Zucker. Folds and cuts: how shading flows into edges. In *ICCV*, 2001. 7
- [16] B. Julesz. Textons, the elements of texture perception, and their interactions. *Nature*, 290(12):91–97, 1981. 1, 2
- [17] G. Kanizsa. *Organization in Vision: Essays on Gestalt Perception*. Praeger Publishers, 1979. 1
- [18] M. Kaschube, F. Wolf, T. Geisel, and S. Löwel. The prevalence of colinear contours in the real world. *Neurocomputing*, 38-40:1335–1339, 2001. 5
- [19] M. Kass and A. Witkin. Analyzing oriented patterns. *CVGIP*, 37:362–385, 1987. 3
- [20] J. Koenderink and A. van Doorn. Local features of smooth shapes: Ridges and courses. In *Proc. of the SPIE*, volume 2031, pages 2–13, 1993. 6
- [21] L. Kwan and D. Regan. Orientation-tuned spatial filters for texture-defined form. *Vision Res.*, 38:3849–3855, 1998. 1
- [22] M. Landy and J. Bergen. Texture segregation and orientation gradient. *Vision Res.*, 31(4):679–691, 1991. 1
- [23] Z. Li. A neural model of contour integration in the primary visual cortex. *Neural Comput.*, 10:903–940, 1998. 1, 7
- [24] A. López, F. Lumberras, J. Serrat, and J. Villanueva. Evaluation of methods for ridge and valley detection. *IEEE Trans. Pattern Anal. Machine Intell.*, 21(4):327–335, 1999. 6
- [25] J. Malik and P. Perona. Preattentive texture discrimination with early vision mechanisms. *J. Opt. Soc. Am.*, 7(5):923–932, 1990. 1
- [26] D. Martin, C. Fowlkes, and J. Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Trans. Pattern Anal. Machine Intell.*, 26(1):530–549, 2004. 4
- [27] O. Musin. Topographic structure of image. In *Proceedings of the International Conference on Computer Analysis of Images and Patterns*, volume 719 of *Lecture Notes in Computer Science*, pages 24–30. Springer-Verlag, 1993. 6
- [28] H. Nothdurft. Orientation sensitivity and texture segmentation in patterns with different line orientation. *Vision Res.*, 25(4):551–560, 1985. 1
- [29] H. Nothdurft. Sensitivity for structure gradient in texture discrimination tasks. *Vision Res.*, 25(12):1957–1968, 1985. 1
- [30] H. Nothdurft. Texture segmentation and pop-out from orientation contrast. *Vision Res.*, 31(6):1073–1078, 1991. 1
- [31] H. Nothdurft. The role of features in preattentive vision: Comparison of orientation, motion, and color cues. *Vision Res.*, 33(14):1937–1958, 1993. 1
- [32] R. Olson and F. Attneave. What variables produce similarity grouping? *American Journal of Psychology*, 83:1–21, 1970. 1
- [33] B. O’Neill. *Elementary Differential Geometry*. Academic Press, 1966. 3, 4
- [34] P. Parent and S. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Trans. Pattern Anal. Machine Intell.*, 11(8):823–839, 1989. 7
- [35] A. Rao and R. Jain. Computerized flow field analysis: Oriented texture fields. *IEEE Trans. Pattern Anal. Machine Intell.*, 17(7):693–709, 1992. 3
- [36] D. Sagi. The psychophysics of texture segmentation. In T. Pappas, C. Chubb, A. Gorea, and K. E., editors, *Early Vision and Beyond*, pages 69–78. The MIT Press, 1995. 1
- [37] 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. 5
- [38] 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. 7
- [39] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Machine Intell.*, 22(8):888–905, 2000. 4
- [40] K. Stevens. The visual interpretation of surface contours. *Artif. Intell.*, 17:47–73, 1981. 1, 7
- [41] J. Todd and F. Reichel. Visual perception of smoothly curved surfaces from double-projected contour patterns. *J. Exp. Psych.: Human Perception and Performance*, 16(3):665–674, 1990. 1, 7
- [42] T. Tversky, W. Geisler, and J. Perry. Contour grouping: Closure effects are explained by good continuation and proximity. *Vision Res.*, 44:2769–2777, 2004. 7
- [43] L. Vincent and P. Soille. Watersheds in digital spaces: An efficient algorithm based on immersion simulations. *IEEE Trans. Pattern Anal. Machine Intell.*, 13(6):583–598, 1991. 6
- [44] L. Williams and K. Thornber. A comparison of measures for detecting natural shapes in cluttered backgrounds. 34(2/3):81–96, 2000. 7
- [45] S. Wolfson and M. Landy. Discrimination of orientation-defined texture edges. *Vision Res.*, 35(20):2863–2877, 1995. 1