Visual saliency and texture segregation without feature gradient

Ohad Ben-Shahar*

Department of Computer Science and the Zlotowski Center for Neuroscience, Ben-Gurion University, Beer Sheva 84105, Israel

Edited by Dale Purves, Duke University Medical Center, Durham, NC, and approved August 17, 2006 (received for review May 31, 2006)

A central notion in the study of texture segregation is that of feature gradient (or feature contrast). In orientation-based texture segregation, orientation gradients have indeed played a fundamental role in explaining behavioral results. Here, however, we show that general, smoothly varying, orientation-defined textures (ODTs) exhibit striking perceptual singularities that are completely unpredictable from orientation gradients. These singularities defy not only popular texture segregation theories but also virtually all computational segmentation methods, and they confound previous behavioral studies with smoothly varying ODTs. We provide psychophysical evidence that perceptual singularities in smooth ODTs are salient visual features consistent across observers and with significant effect on the perception and segregation of oriented textures. We further show that, although orientation gradients cannot predict them, perceptual singularities in smooth ODTs emerge directly from, and can be spatially localized by, two ODT curvatures. Given the traditional role of feature gradients in early vision, the significance of these findings extends well beyond orientation-based texture segregation to issues ranging from curve integration and fragment grouping, through the perception of 3D shape, to the functional organization of the primary visual cortex.

curvature | perceptual organization | perceptual singularity measure

The visual perception of texture plays a fundamental role in many aspects of vision from figure–ground segregation to 3D shape and depth perception. The ability to effortlessly segregate texture stimuli into coherent parts has long been attributed to the changes in the spatial distribution of elementary features, sometimes called textons (1). Of these features, one that has been studied extensively is orientation, because of its neurophysiological basis (2), its central role in perceptual organization (1, 3), and its close relationship to shape perception (4, 5).

Orientation-defined textures (ODTs) in visual stimuli result from pattern formation processes that cover surfaces in the real world. These pattern formation processes often enjoy smoothness properties, and segregation is usually reported when the observed orientation changes abruptly, an event occurring typically because of occlusion or surface discontinuities. Following such early insights, and motivated by Julesz's influential theories that emphasize feature gradients as a source for segregation (1), orientation-based texture segregation (OBTS) research for the last two decades (e.g., refs. 6–8) has concentrated on quantifying segregation based on the orientation contrast between regions of constant orientation (up to noise). Although this exploration is valid in a sense of employing simplified stimuli for use in laboratory experiments, it is such a gross oversimplification of general ODTs as to obscure key aspects of the visual process itself. Interestingly enough, regardless of the formation process of ODTs, in natural stimuli, they are seldom constant or piecewise constant because this outcome requires an accidental match between the surface geometry, the texture formation process, and the observer's viewpoint. Part of our goal in this article is to bring this larger context of smoothly varying ODTs to the fore, both psychophysically and in terms of computational modeling.

Among the bulk of work on OBTS, few studies have considered orientation variations in an attempt to examine how orientation gradients within perceptually coherent regions ($\Delta\theta_{\text{within}}$) affect the minimal orientation gradient between them ($\Delta\theta_{\text{between}}$) that is needed to facilitate segregation (e.g., see refs. 9 and 10). The more general model that emerged is that $\Delta\theta_{\text{between}}$ must be significantly larger than $\Delta\theta_{\text{within}}$ for segregation to occur. These results were further extended to show dependency on the relative orientation between the ODT and the texture boundary (8, 11) and, most recently, to exhibit sensitivity to curvature discontinuities as well (12).

One explicit prediction from existing segmentation models is that, without significant feature gradient, textures are perceptually coherent. As implied by perceptual and computational modeling studies (e.g., refs. 6, 7, 9, 10, 13, and 14), an ODT described by the orientation function $\theta(x, y)$ is therefore likely to exhibit perceptual singularities or perceptual boundaries along curves and points where the magnitude of its orientation gradient, i.e., $|\nabla\theta|$, crosses some threshold. Indeed, this gradient-based approach accurately predicts perceptual singularities for the well studied piecewise constant ODTs.

Unfortunately, however, the predictions of this approach become utterly wrong in the general case of ODTs with smoothly varying orientation. As we show in this paper, 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, and they segregate the ODT into perceptually distinct regions. Importantly, these perceptual singularities have no apparent relationship to the orientation gradient of the ODT. For example, the pattern in Fig. 1*a* has constant orientation gradient across the entire image. In other words, nowhere in this pattern does the orientation change 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 (1), and not merely their local contrast, determines the saliency of different regions and the segregation of the pattern. Similarly, the ODT in Fig. 1*b* exhibits a salient double spiral structure that has no apparent relationship to the orientation gradient of this pattern. It is therefore clear that, to model these saliency differences and segregation correctly, something beyond orientation gradients must be considered.

That salient perceptual singularities in smooth ODTs are independent of orientation gradients is easily demonstrated by creating perceptually different ODTs with identical orientation gradients. In particular, if $\theta_0(x, y)$ is the orientation function of an ODT, it is guaranteed that its orientation gradient is identical to that of the ODT defined by the orientation function $\theta_1(x, y) = \theta_0(x, y) + \Delta\theta$, where $\Delta\theta$ is an arbitrary phase shift. One such phased pair is demonstrated in Fig. 1*c* and shows how drastically different the

Author contributions: O.B.-S. designed research, performed research, analyzed data, and wrote the paper.

The author declares no conflict of interest.

This article is a PNAS direct submission.

Abbreviations: ODT, orientation-defined texture; OBTS, orientation-based texture segregation; PSM, perceptual singularity measure; LIC, line integral convolution; Exp., experiment.

*E-mail: ben-shahar@cs.bgu.ac.il.

www.pnas.org/cgi/doi/10.1073/pnas.0604410103

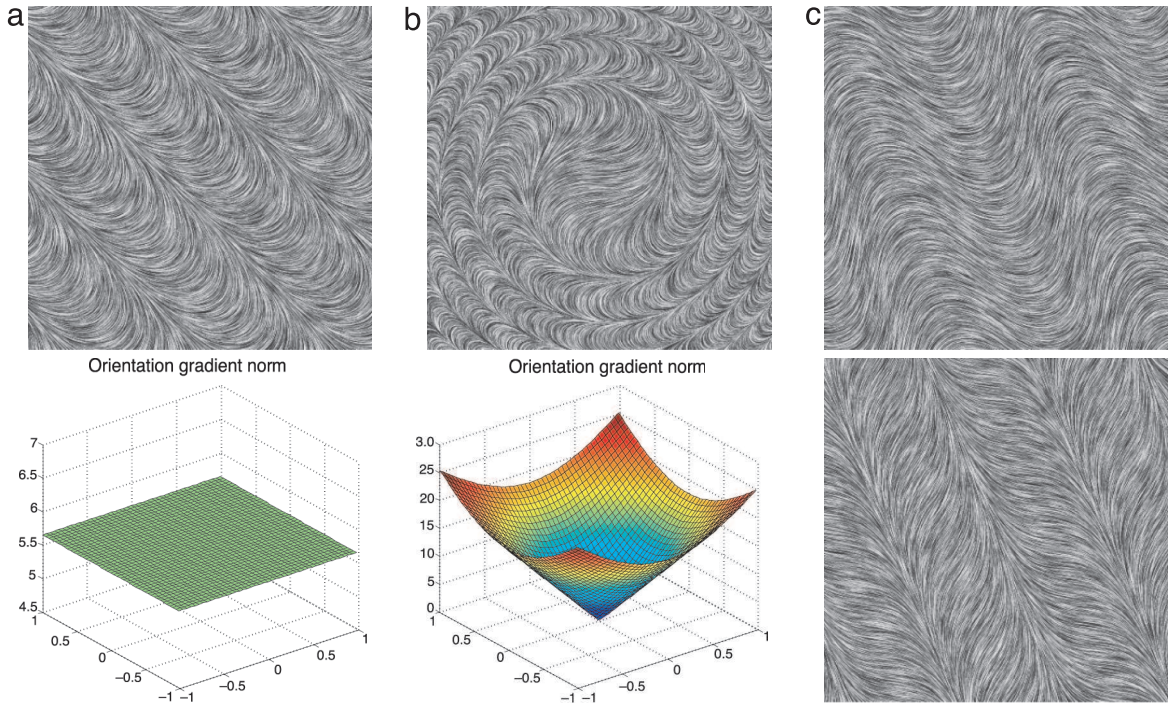

Fig. 1. Perception singularities in smoothly varying ODTs. (*a*) A smoothly varying ODT defined by the function $\theta(x, y) = c{\cdot}x + c{\cdot}y$ and the graph of its orientation gradient magnitude. Clearly, there is a fundamental gap between the (inhomogeneous) perceptual outcome and the (constant) orientation gradient. (*b*) A smooth ODT that exhibits salient double spiral structure and its orientation gradient magnitude. (*c*) A phased pair of smooth ODTs that are different only by a constant phase shift: $\theta_{top}(x, y) = \pi/8 + \sin(5x + 2y)$, where $\theta_{bottom}(x, y) = \theta_{top}(x, y) + \pi/2$. Virtually all observers report no segregation of θ_{top} into distinct regions whereas θ_{bottom} is consistently segregated, into diagonal bands along salient straight perceptual singularities.

structure of perceptual singularities in textures of identical feature gradient can be.

It should be mentioned that perceptual singularities in smoothly varying ODTs were already observed (but ignored) in at least two previous OBTS studies. Nothdurft (9), who was first to experiment with ODTs of piecewise-varying orientation, reported "apparent flow pattern with a pronounced micro structure" (see p. 1074, in ref. 9). To minimize the effect of these structures on his experiment, Nothdurft was forced to randomly change the sign of the background orientation shift and therefore to compromise the link between the dependent and independent variables. More recently, Ben-Shahar and Zucker (12) reported occasional "kinks" in their piecewise-smooth ODT stimuli (12) but observed no confounding effect on the results in their particular experiment. Similarly, unexplained salient configurations were reported but were left as an open question in the contour integration literature (see figure 7 in ref. 15). More than a decade later, in this article, we offer a solution.

From Geometry to Saliency in ODTs

By their definition, an abstract representation for all ODTs would make explicit the orientation at each point and therefore would consist of either an orientation function $\theta(x, y)$ in the image plane or a unit length vector field $\vec{E}_T(x, y)$ tangent to the ODT at each point (x, y). Because an oriented texture is a texture flow (16), here we follow the theoretical foundations developed in our earlier work to derive a computational measure that makes accurate predictions of the perceptual outcome.

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 (17). More specifically, by attaching a second unit vector $\vec{E}_N(x, y)$ to each point (x, y) such that it is perpendicular to $\vec{E}_T(x, y)$, one obtains a frame at each point of 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-centered view (Fig. 2). Perhaps the most important vectors (other than the frame vectors themselves) that beg such an object-centered representation are the covariant derivatives of $\vec{E}_T$ and $\vec{E}_N$. These covariant derivatives, $\nabla_V\vec{E}_T$ and $\nabla_V\vec{E}_N$, represent the initial rate of change of the frame in any given direction V, a quantity that, in the $\{\vec{E}_T, \vec{E}_N\}$ coordinates, is captured by Cartan's connection equation (17):

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

The coefficient $w_{12}(V)$ is a function of the direction vector V, which reflects the fact that the local behavior of the flow depends on the direction along which it is measured (see Fig. 2). Fortunately, $w_{12}(V)$ is linear (17), which allows us to fully represent it by its value along two linearly independent directions. The freedom in selecting these directions is naturally resolved by adopting the natural directions implied by the frame itself, which lead to the following two scalars:

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

These two scalars, defined at each point of the ODT, are called its tangential curvature (κ_T) and normal curvature (κ_N), respectively (16), and they represent the initial rate of change of the ODT orientation in its tangential and normal directions, respectively (Fig. 2). In practical terms, the two ODT curvatures can be evaluated at each point as the coefficients of the orthogonal expansion (17) of the orientation gradient vector of the ODT based on the frame itself:

$$\begin{aligned}\kappa_T &= \nabla\theta{\cdot}\vec{E}_T = \nabla\theta{\cdot}(\cos\theta, \sin\theta)\\ \kappa_N &= \nabla\theta{\cdot}\vec{E}_N = \nabla\theta{\cdot}(-\sin\theta, \cos\theta).\end{aligned} \qquad [3]$$

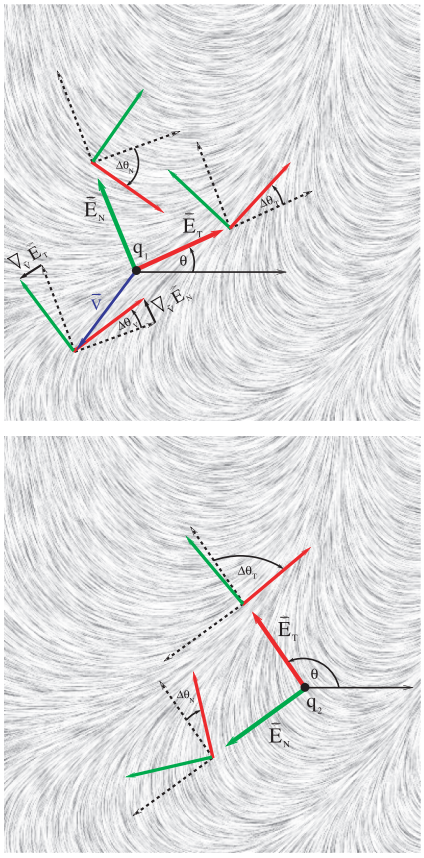

Fig. 2. The frame field theory in action on a blown-up (and brightened) section of the spiral-like stimuli from Fig. 1*b*. Tangent and normal vectors are shown in red and green, respectively. (*Upper*) An infinitesimal translation of the frame from point q_1 along an arbitrary direction $\vec{V}$ (blue vector) rotates it (counterclockwise in this case) by an amount related to the covariant derivatives $\nabla_{\vec{V}}\vec{E}_T$ and $\nabla_{\vec{V}}\vec{E}_N$. The initial rate of rotation along the directions $\vec{E}_T$ and $\vec{E}_N$ themselves defines the two curvatures κ_T and κ_N. Here, a ''tangential translation'' induces a small counterclockwise rotation $\Delta\theta_T$ relative to the original pose (dashed) whereas the ''normal translation'' induces large clockwise rotation $\Delta\theta_N$. The corresponding (small) κ_T and (large) κ_N translate to a high PSM (Eq. **4**), which peaks at the perceptual singularity that passes around q_1. (*Lower*) Similar analysis around q_2 shows a large ''tangential rotation'' $\Delta\theta_T$ and small ''normal rotation'' $\Delta\theta_N$. The corresponding curvature values translate to a small PSM (Eq. **4**) that predicts correctly the lack of saliency around q_2.

We note that, although Eq. **3** provides signed curvatures, as indeed is possible in the plane, in the following, we will be interested in their absolute values (i.e., in $|\kappa_T|$ and $|\kappa_N|$) because the orientation of the ODT is determined only up to 180° and because the sign of curvature seems to play no role in the segregation process.

The trigonometric functions in Eq. **3** suggest 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. 1*a* is scrutinized. In this figure, $\nabla\theta$ is constant across the entire pattern, but perceptually this ODT exhibits periodic singularities, just like the behavior of the curvatures in such patterns. Based on these observations and additional pilot experiments, we assert that the relationship between saliency in smooth ODTs and abstract geometry, and in particular the ODT curvatures, is not accidental. In fact, we state the following (see ref. 18 for an extended computational discussion):

Proposition 1. *The locus of salient perceptual singularities in a smoothly varying ODT of orientation function* $\theta(x, y)$ *corresponds to the ridges of the curvature-based local differential measure*

$$PSM(x,y) = \bigwedge_{\kappa_T^2+\kappa_N^2>\tau^2}\left[\frac{\kappa_N(x,y)^2}{\kappa_T(x,y)^2+\kappa_N(x,y)^2}\right], \qquad [4]$$

where $\kappa_T(x, y)$ *and* $\kappa_N(x, y)$ *are defined by Eq.* **2** *and computed by Eq.* **3** *and the operator* $\bigwedge$ *rectifies its argument by the condition specified.*

It should be noted that, because the proposed perceptual singularity measure (PSM) is defined as the quotient of curvatures, without the rectification it is "blind" to the amount of local orientation variation that is present in the ODT. In other words, without proper thresholding, the PSM might predict spurious perceptual singularities in stimulus regions where the underlying orientation function is changing very slowly or includes noisy fluctuations. One example for this behavior can be seen in Fig. 1*b*. Although the PSM without rectification would predict a double spiral that continues almost to the center point (see ref. 18), applying the rectification guarantees that the predicted structure stretches no further than the perceptual one. One of our goals in this article is to perform a psychophysical evaluation of the proposed PSM for saliency in smooth ODTs, including a measurement of the threshold τ that is the only perceptual parameter involved in it.

Results

To explore perceptual singularities in smoothly varying ODTs, we examined (*i*) their perception in a free viewing experiment, (*ii*) their discrimination in preattentive visual sequences, and (*iii*) their segregation in rapid displays. Our goal was to examine consistency across observers, the dissociation of perceptual singularities from orientation gradients, and the validity of our saliency measure and theory.

Experiment (Exp.) 1. We first examined qualitatively the degree to which perceptual singularities in smooth ODTs are consistent across observers. Subjects (N = 7) were asked to trace perceptual singularities in line integral convolution (LIC) ODTs (19) presented on a computer screen (Fig. 3). Experimental sessions incorporated 24 different stimuli, which included both smoothly varying ODTs and two types of control stimuli: (*i*) constant ODTs with no perceptual singularities and (*ii*) piecewise constant ODTs with two main perceptual segments defined by 90° orientation contrast and a constant texture-boundary orientation difference of 45°. Observers were shown one such piecewise constant pattern during their briefing and were asked to trace salient curves and boundaries. No smoothly varying ODT was presented during the instructions. Nevertheless, virtually all subjects traced multiple curves on the smoothly varying ODTs, with results being extremely consistent across observers. A sample of results for one such ODT is demonstrated in Fig. 3.

Exp. 2. Having a qualitative confirmation that perceptual singularities in smooth ODTs are a significant and consistent perceptual phenomenon, we used a second experiment to examine our hypothesis about the dissociation of these singularities from orientation gradients. In this experiment, observers performed a forced-choice discrimination task in which pairs of LIC ODTs were flashed rapidly one after another (separated by an interstimulus mask) and subjects were asked to decided whether they were similar or not. Within each trial, the two ODTs presented could be either (*i*) the exact same ODT, (*ii*) completely different ODTs (i.e., ODTs with unrelated orientation function), or (*iii*) a phased pair, i.e., two ODTs that only differed by a 90° local orientation shift and therefore had an identical orientation gradient function. Each trial type was equally frequent, and their order was preselected randomly and fixed for all subjects. The experiment consisted of 20 different ODT stimuli that formed 10 pairs of phased stimuli. Seven of these 10 pairs were smoothly varying (see *Methods*) whereas the other three were control stimuli of piecewise constant orientation

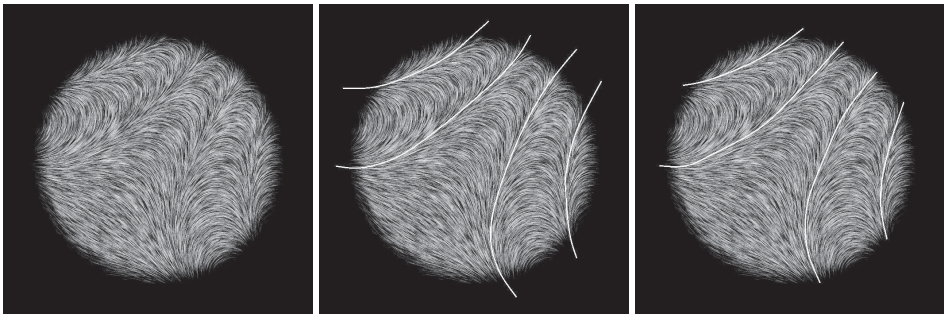

Fig. 3. Manual segregation of smoothly varying ODTs. (*Left*) One smooth ODT used in Exp. 1. (*Center* and *Right*) Manual segregation by two different subjects. Results of the other five subjects that participated in this experiment were virtually identical.

with strong orientation gradients of 90° along designated curves and a fixed texture-boundary orientation difference of 45°.

We analyzed discrimination performance both by the different trial types (same, different, phased) and by the different stimulus type (smoothly varying vs. piecewise constant). The critical trials were those with phased pairs because, by being different ODTs of identical orientation gradient, these pairs could be used to examine the role of orientation gradients in the segregation of smooth ODTs.

The results in terms of discrimination accuracy are shown in Fig. 4. Discrimination of different ODTs, as well as detection of identical ODT pairs, was equally accurate at ≈95% and were not influenced by stimulus type (Fig. 4*a*). However, a qualitatively different performance was revealed for phased pairs (Fig. 4*b*): discrimination of discontinuous phased pairs was below chance level with 28.00% accuracy (SE = 6.02%) whereas smoothly varying phased pairs were discriminated correctly above chance level with 60.60% accuracy (SE = 6.32%). These differences were highly significant ($P < 0.00005$), indicating qualitative differences in the perceptual experience obtained from discontinuous, piecewise-constant ODTs vs. smoothly varying ones.

Why would subjects find it significantly more difficult to discriminate between phased pairs of discontinuous, piecewise-constant ODTs than phased pairs with smoothly varying orientation function? The fact that the discrimination accuracy of the former is well below chance level suggests that most subjects do not use, or are unable to observe, the microstructure of briefly presented ODTs and therefore rely on some global perceived structure conveyed by salient perceptual singularities that are effected by the high orientation contrasts along the discontinuities. Because discontinuous, piecewise-constant phased pairs have identical perceptual singularities, they become almost indistinguishable in brief presentations. As our experiment reveals, however, under identical conditions, the discrimination between smoothly varying phased ODTs is significantly above chance level, suggesting that their perceptual singularities are indeed quite different despite having identical orientation gradients. We therefore conclude that, in general, orientation gradients are not the main cause for saliency and perceptual singularities in ODTs, which confirms our hypothesis above.

Exp. 3. With the dissociation of perceptual singularities from orientation gradients established in Exp. 2, we turned to a psychophysical evaluation of our saliency theory and our PSM (Eq. **4**). Qualitatively, our computational theory can be shown to produce extremely accurate predictions in comparison with the free viewing task of Exp. 1 (see ref. 18), but here we sought psychophysical confirmation in a preattentive setup, which could also allow the calibration of the rectification operator $\wedge$ in Eq. **4**.

To achieve these goals, we had to design a set of stimuli and an experimental procedure that would test segregation according to predicted perceptual singularities in a manner independent of orientation gradients. Natural candidate stimuli for such an exploration are ODTs of constant orientation gradient (of the sort shown in Fig. 1*a*) whose predicted (and observed) perceptual singularities are straight lines and are therefore easy to manipulate and control. We generated a large set of such stimuli and incorporated them in a preattentive segregation task in which subjects were asked to decide whether or not two salient probes that were embedded in each stimulus were within the same perceptual segment or not. The degree to which a subject's response would correspond to predicted singularities could then serve as a validation or refutation of the proposed theory.

All stimuli in this experiment were ODTs of the form $\theta(x, y) = \theta_0 + d \cdot g \cdot (x \pm y)$, plotted as LIC textures in the spatial range

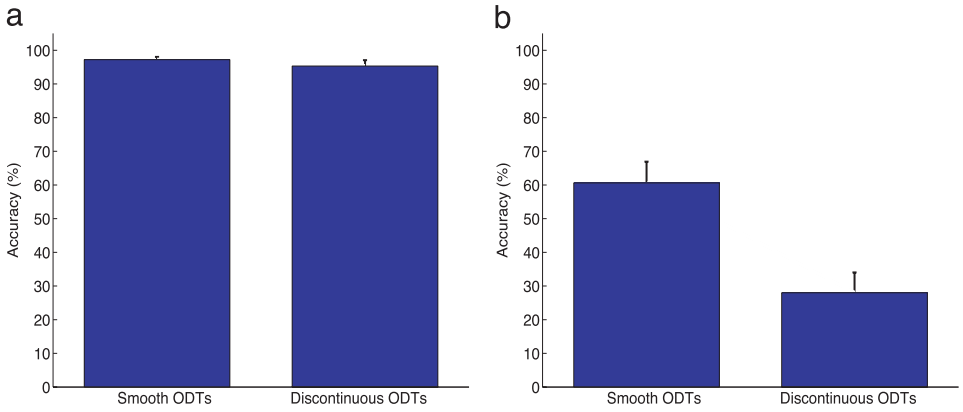

Fig. 4. ODT discrimination accuracy in Exp. 2. (*a*) Detection accuracy of identical pairs was not influenced by the ODT type: 97.22% for smoothly varying ODTs and 95.31% for discontinuous ones. (*b*) Detection accuracy of phased pairs was significantly affected by ODT type: 60.60% for smoothly varying ODTs and 28.00% for discontinuous ones ($P < 0.00005$). Error bars represent ±1 SE.

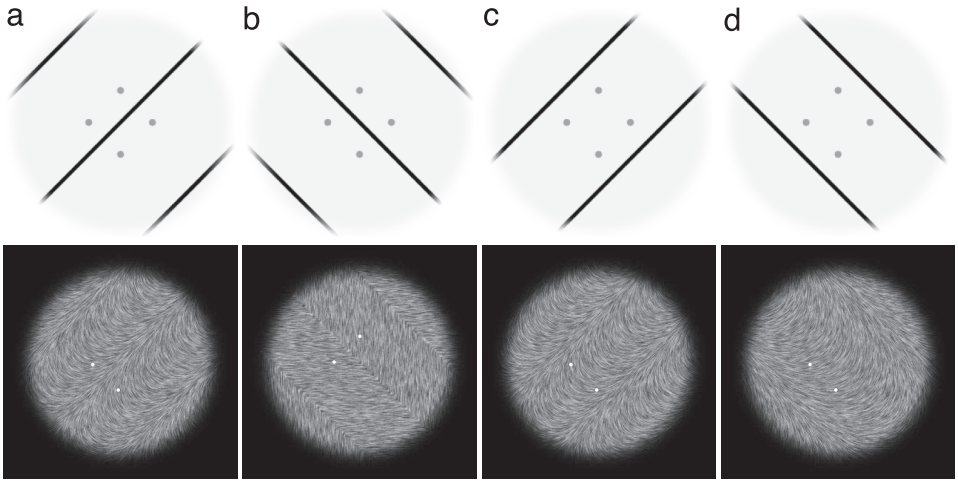

Fig. 5. Sketches and examples of the stimuli used in Exp. 3. All these stimuli had (predicted) perceptual singularities (shown as dark lines in the sketches, *a–d Upper*) oriented either 45° (*a* and *c*) or −45° (*b* and *d*), which occurred either along an image diagonal (*a* and *b*) or offset from it (*c* and *d*). Each stimulus was presented with two circular probes plotted on two neighboring corners of a virtual diamond (sketches show all four corners). Note that the stimuli shown in *a*, *c*, and *d* are all smoothly varying ODTs whereas *b* shows a discontinuous, piecewise-constant ODT.

$[-1, 1] \times [-1, 1]$. The parameter g was assigned values such that the magnitude of orientation gradient in each stimulus was taken from the set {5, 10, 15, 20, 25, 30, 35, 40} in units of degrees per visual angle (from viewing distance of 80 cm). By definition, g reflects the total local variation of the ODT at each point, which relates directly to the two ODT curvatures by $g = \sqrt{\kappa_T^2 + \kappa_N^2}$. The lower (5°/visual angle), and higher (40°/visual angle), bounds were set below and above the perceptual threshold as determined in a pilot experiment.

For each value of g, two possible values for θ_0 were used such that the resultant ODT would have a (predicted) perceptual singularity either along the image diagonal or offset from it (see Fig. 5). For each value of g, we therefore had 4 stimuli that counterbalanced the direction and position of the predicted singularities. This set of 32 stimuli was further duplicated to a total of 64 to counterbalance the direction of variation of the microtexture in the pattern, a parameter that is expressed by the value of the parameter d: either 1 or −1. An additional set of 8 stimuli, with similar perceptual singularities but using a piecewise constant orientation function, was added for control (see Fig. 5 for a few selected stimuli). Importantly, during the instructions, subjects were shown only piecewise constant ODTs (e.g., as in Fig. 5*b*) with which the task was exemplified. No smoothly varying ODT was presented during the instructions.

By design (and assuming no rectification; see Eq. **4**), the two probes (see *Methods*) occurred on two different sides of a predicted perceptual singularity in 25% of all trials having any particular g value, and in 50% of all trials with a predicted singularity that coincides with the stimulus diagonal (e.g., as in Fig. 5 *a* and *b*). The analysis of this latter set of trials, in terms of how well the model prediction is reflected in subjects' performance, is plotted against g in Fig. 6. Note that, for ODTs that vary very slowly ($g \leq 15$), performance is at the 50% chance level. Indeed, subjects did not segregate these patterns into distinct regions and perceived the two probes as always belonging to the same perceptual unit. This behavior changes drastically once the total variation $g = \sqrt{\kappa_T^2 + \kappa_N^2}$ exceeds some threshold. From that point on, subjects' response depends on the geometric relationship between the perceptual singularities and the probes position, and their performance clearly crosses the (conservative) 75% prediction threshold that validates the theory, and up to 93.5% on average. This accuracy comes very close to the 97% ceiling performance (dashed red line in the graph), which was defined by segregation accuracy on the control stimuli with discontinuous, piecewise-constant ODTs.

The performance shown in Fig. 6 not only supports our theory but also provides the orientation variation threshold τ that controls the rectification for the detection of perceptual singularities (Eq. **4**). Clearly, these singularities emerge and trigger robust segregation as soon as $g = \sqrt{\kappa_T^2 + \kappa_N^2} \geq 18°$ per visual angle. From that point on, saliency and perceptual singularities in smoothly varying ODT are mainly effected by ODT curvatures and can be spatially localized by our PSM measure. For a qualitative comparison of the computational measure and human performance in free viewing conditions, the reader is referred to ref. 18.

Discussion

This work attempts to make several fundamental contributions to perceptual organization and texture segregation by the human visual system. First and foremost, we argue that perceptual organization, saliency, and the detection of perceptual singularities in textures cannot be determined reliably by feature gradients. Although outstanding 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 smoothly varying ODTs presented here. To that end, we have shown that the accurate detection of perceptual singularities in

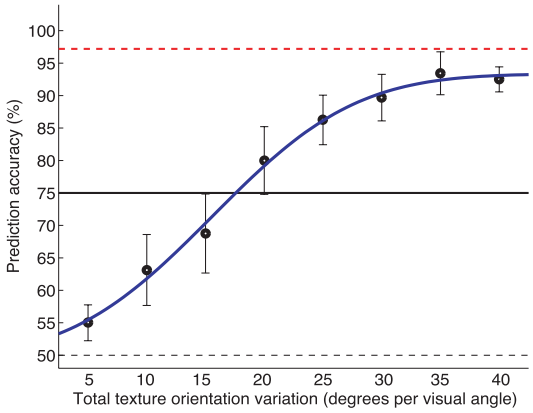

Fig. 6. Results of Exp. 3. Model prediction accuracy reveals agreement with subjects' response as soon as the total orientation variation exceeds a perceptual threshold of $g \geq 18°$ per visual angle. The bottom dashed line shows chance level. The top dashed line shows ceiling performance observed from the discontinuous, piecewise constant stimuli. The center solid line shows threshold level. Error bars are ±1 SE, and mean accuracy is fitted by a cumulative normal distribution.

smoothly varying ODTs emerges directly and solely from their intrinsic geometry and curvature properties. We proposed a computational measure that is able to pinpoint the spatial location of these perceptual singularities, and validated this theory psychophysically in a preattentive segregation experiment. Strikingly, the same saliency measure works also for the classical piecewise constant textures, and therefore it fully generalizes existing theories based on feature gradients. The same measure may also explain recent findings about an unexpected salience of curvature in grouping by proximity in dot-sampled structured grids (20). Furthermore, smoothly varying ODTs also have been used in studies of sensitivity to orientation modulations (e.g., refs. 21 and 22). The results presented here therefore suggest that such studies could benefit from considering ODT curvatures or the proposed PSM as additional independent variables.

Although this article focuses on ODTs and OBTS, we believe that its implications are much broader in scope. Texture representation by the human visual system is believed to rely on orientation maps (14, 7, 13) whose response pattern can be thought of as multiple ODT layers. We therefore predict that saliency fluctuations and perceptual singularities should emerge in multiply oriented and general textures as well, despite uniform or constant feature gradient. Furthermore, orientation is explicit not only in ODTs but also in motion and optical flows, and therefore the theory developed in this article is directly applicable to motion-based segregation. More importantly, orientation is implicit in other visual cues that take part in the segregation process, in particular, shading (23) and color (24). Therefore, the results presented here imply that the role of these other cues in early vision and segregation should be revisited.

The results presented in this article also are linked to aspects of vision not traditionally related to texture segregation and in particular to the perception of 3D shape. Often, when viewed from a slanted viewpoint 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 valleys that correspond to the salient perceptual singularities signaled by our PSM. It is therefore possible that our study may be placed in the context of the shape-from-surface contours and shape-from-texture flow (e.g., refs. 4, 5, and 25–28). Unfortunately, most such predictive studies focus on specific cases such as lines of curvatures (e.g., refs. 4 and 26), geodesics (e.g., refs. 25 and 28), or developable surfaces (e.g., refs. 26 and 28), none of which seem to apply directly to our case. Closer examination of this important link is therefore part of our future work.

Finally, the current study also raises intriguing questions concerning the functional organization of primary visual cortex. For example, much of the texture segregation and perceptual organization literature has linked psychophysical results to anatomical and physiological findings in V1. Although long-range horizontal connections are among the main biological structures that are thought to facilitate these processes by means of contextual modulations (e.g., refs. 29–31), hardly any findings connect them directly to curvature tuning (for a recent exception, see ref. 32), as is explicitly suggested by our study. Although some research has indicated curvature coding at the neural level (32–35), no work has been done on curvature in the context of texture in general or the representation of normal curvature (κ_N) in particular. Our work suggests that these and similar related issues deserve systematic investigation.

Methods

Exp. 1. All stimuli were intensity modulated to create a fuzzy circular aperture to minimize possible interaction with the image boundaries (Fig. 3). Subjects could trace either polygonal or spline curves using a mouse and were not constrained by time or viewing distance.

Exp. 2. All smoothly varying ODTs were generated from quadratic orientation functions of the form $\theta(x, y) = a_{00} + a_{10}x + a_{01}y + a_{20}x^2 + a_{11}xy + a_{02}y^2$, where the coefficients a_{ij} were preselected randomly from the range $[-5, 5]$. All stimuli were LIC textures with a fuzzy circular aperture (of the type shown in Fig. 3) and effective size of 7° in radius from viewing distance of ≈80 cm.

Fifteen naive subjects participated in this experiment, where each session consisted of 300 trials. Each trial started with a blank (700 ms) followed by a mask (150 ms). The first ODT was then presented briefly (100 ms), followed by an interstimulus mask (150 ms) and the second ODT (100 ms). The sequence ended with a post-stimulus mask (150 ms) after which subjects entered their response by pressing one of two designated buttons. Trials were presented in blocks of 50 each, and subjects were allowed to take short breaks between these blocks.

Exp. 3. In each trial of the experiment, a fixation point was presented at the center of the screen (700 ms) followed by a mask (130 ms). Then a stimulus was presented for 75 ms followed by a post-stimulus mask for another 130 ms. Each stimulus of the possible 72 was presented with two salient circular probes (0.14° visible diameter) superimposed on it. Probes were set at 1.2° from fixation on two neighboring vertices of a virtual diamond (Fig. 5). The four possible configurations of probe pairs were equally likely, randomized, and counterbalanced with the direction and position of the perceptual singularities throughout the experiment. Subjects' task was to determine whether the two probes belonged to the same perceptually coherent region ("SAME" response) or to two distinct ones ("DIFFERENT" response). Ten naive subjects participated in this experiment where each ran a session of 576 trials, (12 blocks of 48 trials each), with eight repetitions of each stimulus (two repetitions for each of the four possible probe configurations).

This research was partially supported by the Toman and Frankel funds of Ben-Gurion University, by the Zlotowski Center for Neuroscience, and by the Paul Ivanier Center for Robotics Research.

1. Julesz B (1981) *Nature* 290:91–97.
2. Hubel D, Wiesel T (1977) *Proc R Soc London Ser B* 198:1–59.
3. Kanizsa G (1979) *Organization in Vision: Essays on Gestalt Perception* (Praeger, New York).
4. Stevens K (1981) *Artif Intell* 17:47–73.
5. Todd J, Reichel F (1990) *J Exp Psychol Hum Percept Perform* 16:665–674.
6. Nothdurft H (1985) *Vision Res* 25:551–560.
7. Landy M, Bergen J (1991) *Vision Res* 31:679–691.
8. Wolfson S, Landy M (1995) *Vision Res* 35:2863–2877.
9. Nothdurft H (1991) *Vision Res* 31:1073–1078.
10. Mussap A, Levi D (1999) *Vision Res* 39:411–418.
11. Nothdurft H (1992) *Percept Psychophys* 52:255–275.
12. Ben-Shahar O, Zucker S (2004) *Vision Res* 44:257–277.
13. Sagi D (1995) in *Early Vision and Beyond*, eds Papathomas T, Chubb C, Gorea AEK (MIT Press, Cambridge, MA), pp 69–78.
14. Malik J, Perona P (1990) *J Opt Soc Am* 7:923–932.
15. Field D, Hayes A, Hess R (1993) *Vision Res* 33:173–193.
16. Ben-Shahar O, Zucker S (2003) *IEEE Trans Pattern Anal Mach Intell* 25:401–417.
17. O'Neill B (1966) *Elementary Differential Geometry* (Academic, New York).
18. Ben-Shahar O (2006) *Proceedings of the 5th IEEE Computer Society Workshop on Perceptual Organization in Computer Vision* (IEEE Computer Society, Los Alamitos, CA).
19. Carbal B, Leedom LC (1993) *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1993* (ACM Press, New York), pp 263–270.
20. Strother L, Kubovy M (2006) *J Exp Psychol Hum Percept Perform* 32:226–234.
21. Kingdom F, Keeble D, Moulden B (1995) *Vision Res* 35:79–91.
22. Kwan L, Regan D (1998) *Vision Res* 38:3849–3855.
23. Breton P, Zucker S (1996) *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (IEEE Computer Society, Los Alamitos, CA).
24. Ben-Shahar O, Zucker S (2004) *Neural Net* 17:753–771.
25. Knill D (1992) *J Opt Soc Am* 9:1449–1464.
26. Li A, Zaidi Q (2000) *Vision Res* 40:217–242.
27. Todd J, Oomes A (2002) *Vision Res* 42:837–850.
28. Knill D (2001) *J Opt Soc Am A* 18:12–35.
29. Malach R, Amir Y, Harel M, Grinvald A (1993) *Proc Natl Acad Sci USA* 90:10469–10473.
30. Sillito A, Grieve K, Jones H, Cuderio J, Davis J (1995) *Nature* 378:492–496.
31. Bosking W, Zhang YBS, Fitzpatrick D (1997) *J Neurosci* 17:2112–2127.
32. Ben-Shahar O, Zucker S (2004) *Neural Comp* 16:445–476.
33. Dobbins A, Zucker S, Cynader M (1987) *Nature* 329:438–441.
34. Versavel M, Orban G, Lagae L (1990) *Vision Res* 30:235–248.
35. Pasupathy A, Connor CE (2002) *Nat Neurosci* 5:1332–1338.