From boundaries to bumps: When closed (extremal) contours are critical

Benjamin Kunsberg

Applied Mathematics, Brown University,
Providence, RI, USA ✉

Steven W. Zucker

Computer Science, Biomedical Engineering,
Yale University, New Haven, CT, USA ✉

Invariants underlying shape inference are elusive: A variety of shapes can give rise to the same image, and a variety of images can be rendered from the same shape. The occluding contour is a rare exception: It has both image salience, in terms of isophotes, and surface meaning, in terms of surface normal. We relax the notion of occluding contour and, more accurately, the rim on the object that projects to it, to define closed extremal curves. This new shape descriptor is invariant over different renderings. It exists at the topological level, which guarantees an image-based counterpart. It surrounds bumps and dents, as well as common interior shape components, and formalizes the qualitative nature of bump perception. The invariants are biologically computable, unify shape inferences from shading and specular materials, and predict new phenomena in bump and dent perception. Most important, working at the topological level allows us to capture the elusive aspect of bump boundaries.

Introduction

Intuitively, shape inferences seem uniform, accurate, and holistic: Glancing at an image of a bump, one would "see" the bump (almost) regardless of how it was rendered or lit, and would probably have the impression that the bump percept was physically veridical. Such intuition has driven computational approaches to shape-from-shading, which seek maps from *the* given image directly to *the* surface that gave rise to it. Thus far, this goal has been elusive. Instead, to find a well-defined image-to-surface map, computational researchers constrain the problem by, for example, limiting reflectance (say, to Lambertian materials), limiting surfaces (say, to elliptical patches), or training deep neural networks with limited data sets (references below). We propose an alternative approach. Instead of limiting the problem in this fashion, we seek a more general solution. To achieve this, we sacrifice the unique surface as our goal and instead postulate an intermediate stage between the image and the surface. This intermediate stage is topological in nature and different from previous approaches. With it we are able to identify those areas of the image that are most informative about shape and on which shape inferences could be anchored. This leads to a novel invariant between collections of images and collections of surfaces. The images could derive from different material, lighting, and rendering physics or from different observers. The invariant provides an explanation for certain aspects of shape psychophysics. Perhaps the most radical aspect of our invariant, for the perception community, is that it is defined over a neighborhood—a portion of the image or surface—that surrounds certain types of curves. These are the extremal contours mentioned in the title and which are developed in this article.

Since our characterization of the invariant is abstract and involves mathematics not normally used in perceptual modeling, we include background material and try to make the development conceptual rather than rigorous; see Appendix B. The article is a corollary to—and extends to bumps—an earlier model of *critical contours* (Kunsberg & Zucker, 2018). Since the earlier work was more formal, the first part of this article is an informal review of the relevant material. The hope is this will make the material more accessible to the perception community. The new technical contribution in this article is the specialization of our invariant to defining bumps and dents (i.e., protrusions up to the convex/concave ambiguity).[1] This specialization is nontrivial, since it involves showing that the definition of a bump is generic. While bumps are classical components of shapes, and we all seem to "know one when we see one," finding a definition for them has been problematic. In fact, there is even disagreement around where the boundary of a bump should be located (Watt & Morgan, 1983; Georgeson et al., 2007; Morgan, 2011; Subedar & Karam, 2016). We embrace this ambiguity; in fact, it is precisely what our topological invariant captures.
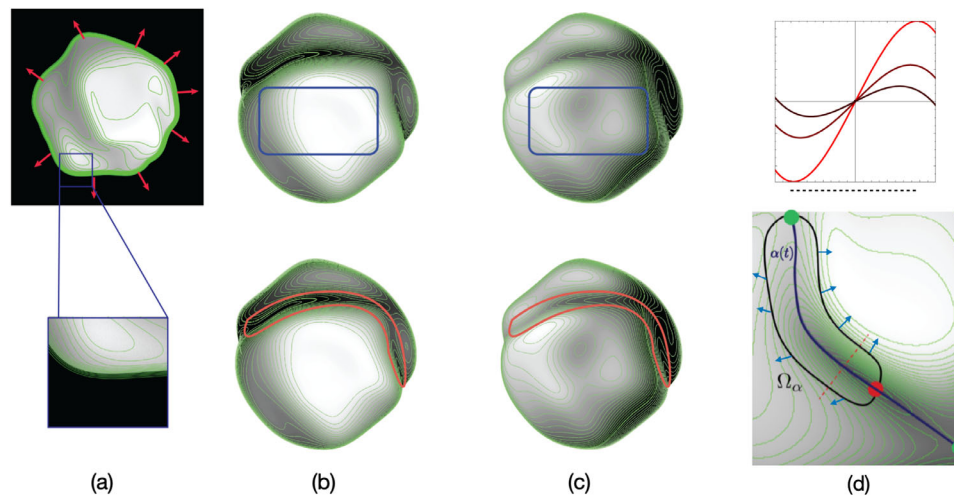
Figure 1. Isophotes concentrate near features of interest. (a) The occluding contour. Notice how the isophotes (green), an image property, are parallel to the occluding contour on the object side, while the (projected) normal (red), a surface property, is orthogonal to it. (b, c) Isophote behavior in interior regions of a Lambertian shape. When the illumination shifts from "in front" (b) to "above" (c), notice how the isophotes change drastically in the blue region but remain stable in the red "sausage." (d) This sausage is at the heart of our theory. Mathematically, the (red) sausage Ω surrounds a curve α (blue), a critical contour. Ω encloses a region within which the isophotes follow a pattern; notice how the image gradient (blue arrows) point away and how the critical contour α cuts directly across the isophotes. This pattern is consistent with certain surface configurations. Possible cross sections of the surface along the red dashed contour (d, bottom) could be among the profiles shown (d, top). Because of this variability, we employ topological methods that constrain the ridge-like profile to be within the sausage. A curve on the surface—the extremal contour—is the analog of the critical contour (not shown in this figure). It lives within a similar sausage. Figures (b) and (c) adapted from Kunsberg et al. (2018); (d) from Kunsberg and Zucker (2018).

Occluding contours

The occluding contour (Appendix A) provides motivation for what an invariant should comprise. It has an image signature (parallel isophotes), and it has a corresponding contour on the surface, the rim contour. Thus, the occluding contour links an image salient property to a surface property (the normal; see Figure 1a). This link between the two-dimensional (2D) image domain and the three-dimensional (3D) shape domain is formal and generic, in the sense that it holds across different rendering and viewing situations (Huggins et al., 2001; Huggins & Zucker, 2001). It is perceptually useful (Palmer & Ghose, 2008; Mooney et al., 2019; Ghose & Peterson, 2021), and computationally, it provides a kind of "anchor" structure for shape reconstruction: Regardless of the shape, it must lie within the occluding contour or its deformation (Wang et al., 2018). But all of this obtains because—and only because—occluding contours are mathematically special. *Occluding contours are the projections of rim curves of maximal slant*. But therein also lies their limitation; occluding contours exist only at exterior boundaries and self-occlusions, that is, at very special locations within the image. Our goal in this article is to generalize such constructs to interior shape features, specifically bumps (and dents). In effect, we seek to "extend" the occluding contour into the interior

of a shape. Just as the occluding contour bounds shapes, we study how to bound bumps (and dents and ridges, etc.). Importantly, the well-defined crispness of the occluding contour is relaxed into a neighborhood on the shape and thus into a neighborhood in the image. This is specifically where our sacrifice (from the opening paragraph) manifests: It is within these neighborhoods that the image and surface salient properties are defined.

Critical contours

To emphasize: Bump boundaries are, mostly, fuzzy; occluding contours are crisp. Another subtlety: While traditional measures of shape, such as normals and curvatures, are local and defined pointwise, bumps extend over a region. (This is the holistic comment in the opening sentence.) In generalizing the occluding contour to bumps, then, we shall need to capture this transition from sharp to fuzzy and from a point to a neighborhood. Our proposal to do this extends the geometry of surfaces with constructions from topology. Instead of seeking *the* shape, the uniqueness requirement is relaxed (partly). Bumps, in our theory, have a boundary but it is only implicit; it may differ for individuals, for renderings, and for viewings, but it always "lives" within a band of uncertainty—a

neighborhood. This neighborhood has a small diameter (an ϵ-neighborhood of the bump boundary); for illustration, see Figure 1b–d. The red "sausage" is the heart of our approach. Regardless of how the ridge bends, there is a sausage that surrounds it. But there is more: The sausage has a special geometric structure, through which contours can be defined. When the sausage is over the image, we call the contour a *critical contour*. Its counterpart on the surface is an extremal curve of slant. The sausage contains an image pattern relating to a surface property that is invariant of the rendering function. In the spirit of how an artist's drawing can convey the impression of shape rather than its exact coordinates, the critical contour constrains the possible shapes. Taken in cross section, the surface is ridge-like but only partly constrained (see Figure 1d, top). As the rendering changes a little, or the viewpoint changes a little, the critical contour will always be in a similar red sausage. As the ridge becomes steeper, the sausage gets thinner; the limiting process is defined in Kunsberg and Zucker (2018). The image in the interior of the sausage may change a little with changes in material or lighting, but the existence of the boundary and the presence of the sausage are invariant. In effect, the meaning derives from the sausage and what it implies. The technical contribution in this article studies what happens when the critical contours are closed (i.e., when they surround bumps).

Because of its importance, we underline the difference between the blue and the red regions in Figure 1. Others have observed that *the isophotes change drastically with changes in lighting or reflectance (e.g., Todd et al., 2014). While this observation is true in many—perhaps most—places in an image, it is not true in the red sausage neighborhood around critical contours.* This is the image side of the invariant we shall be developing.

Put in psychophysical terms, our approach is motivated by the many observations that shape inferences are qualitatively similar, but not quantitatively identical, across subjects and conditions. That is, different individuals infer similar—but not identical—3D surfaces from the same image. There is a huge literature confirming this, sometimes emphasizing material or lighting influences and sometimes emphasizing over/underestimation of surface normals (Sun & Schofield, 2012; Mamassian & Kersten, 1996; Mingolla & Todd, 1986; Todd et al., 2014; Egan & Todd, 2015; Christou & Koenderink, 1997; Egan & Todd, 2015; Seyama & Sato, 1998; Curran & Johnston, 1996; Khang et al., 2007; Koenderink et al., 1996; Marlow et al., 2015; Nishida & Shinya, 1998; Nefs et al., 2006; Vangorp et al., 2007; Ohara et al., 2020; Bernhard et al., 2016; Faisman & Langer, 2013). Since we are studying shape descriptions, we are working at a scale different from textures and surface coverings, and we are not seeking to estimate lighting directions

(cf. Koenderink et al., 2004) or material versus blur properties (Mooney et al., 2019). Nevertheless, the psychophysical results in Mooney et al. (2019), which explore how boundary orientations compare to shading gradients, are close in spirit to what we are attempting, although they do not yet capture the global, qualitative content in our topological approach. Further along these lines, we hope that our research will help to convey the qualitative aspect of shape perception discovered by the psychophysics community to the computational community.

We were particularly motivated by an example from the laboratory of J. Todd, which we review here (Figure 2a, b). Subjects viewed a pattern of bumps and their 3D percepts were estimated. Notice the loose but not exact perceptual agreement (Figure 2b) between them. This variation is not arbitrary; subjects agreed about what we might call the "bumpiness", but not the magnitudes and precise locations. In our terms, we will say that the subjects agreed about the topology but not about the geometry. Subjectively, it is as if the variation across subjects were constrained by a "band of uncertainty" (Figure 2c) around each bump. This band of uncertainty prefigures the difficulty in specifying precisely where a bump begins—somewhere within the band—and is why we were attracted to a topological approach. The difference between viewers is effectively a "stretch" between their inferences, but not a completely arbitrary stretch. Rather, this "stretch" band is localized to a neighborhood (the sausage) and thus is central to our definition of a bump. For an invitation to the relevant topological ideas, see Appendix B.

We now consider the sausage from a photometric perspective. It is well known that the intensity distribution will change as the rendering, material, lighting, and viewing angle change, sometimes by a large amount. The computational challenge, as it is classically applied, is to "explain" (i.e., model) all of these changes. Perhaps it is because these changes can be so extreme that it explains why one is tempted to simplify the problem to, say, Lambertian reflectance. But it is still difficult to define priors and regularizers so that they ameliorate all of these changes *for the entire object.* Here topology again provides the payoff: The isophote[2] structure within the sausage follows a particular template for a wide range of renderings. While the intensity distribution changes *almost* everywhere, it does not change everywhere. It follows a plan within the sausage, as is illustrated by the blue and red regions in Figure 1 and first characterized in Kunsberg and Zucker (2018), Kunsberg et al. (2018). Notice in particular how the (image) gradient points away from the sausage on both sides (Figure 1d). Consistency in this gradient orientation will be important in linking different sausages together and in finding the generic surface configuration corresponding to it.
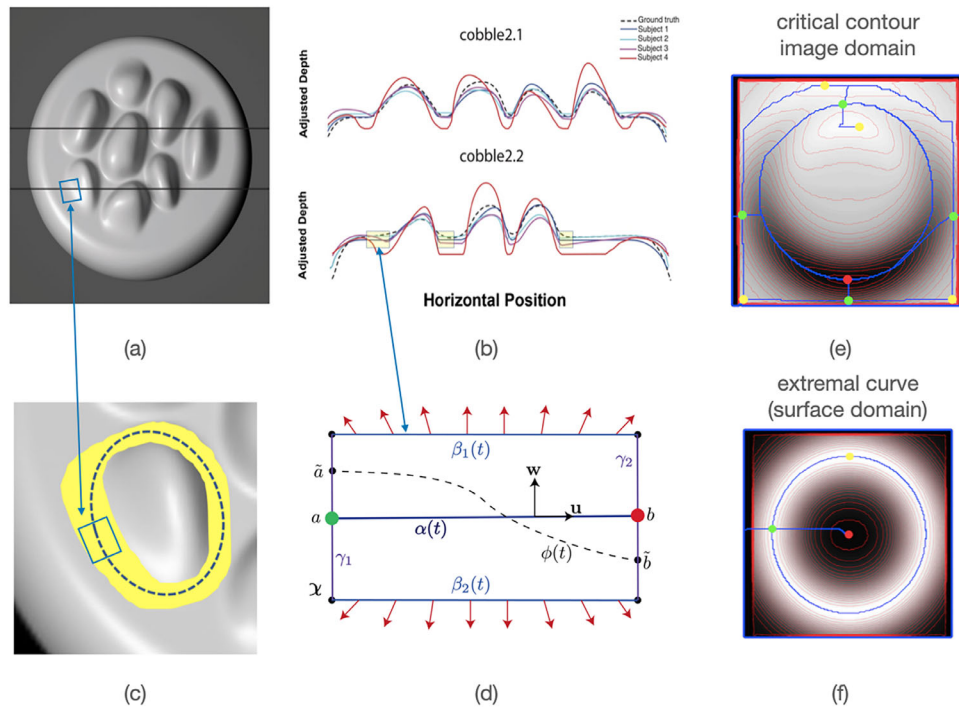
Figure 2. Overview of results. (a) Example image from an experiment in the Todd Lab. (b) Reconstructions of surface profiles along the scan lines in (a). Note the differences between subjects. (c) Zoom on a single bump from (a) showing (an estimate of) the sausage (yellow); the blue arrow indicates corresponding locations. The sausage surrounds a possible critical contour (dashed). (d) Mathematical abstraction of (a piece of) a sausage. Note the critical contour $\alpha(t)$ starts at a and ends at b. Nearby is another contour $\phi(t)$ that derives from the surface; its starting and ending points are close to a and b. In general, as the sides of the bump become steeper, the sausage box in (d) gets narrower and the curves approach each other. This diagram is explained more fully in Extremal contours: Image salience. (e) Example of a critical contour $\alpha(t)$ surrounding a bump; note how it cuts directly across the isophotes. (f) Example of an extremal contour $\phi(t)$ on the slant function of the bump in (e). Notice how the circular critical contour and the circular extremal contour are close. As we will explain, the red dot is minima and the green dots are maxima of either the intensity or the slant functions. Such maxima and minima are central to the topological approach. Figures (a) and (b) from Nartker et al. (2017); (d) from Kunsberg and Zucker (2018).

The sausage surrounds critical contours through a consistent image orientation flow. Above we talked about variation across subjects, implying that different subjects, if asked, might draw different contours around the bumps. Earlier theoretical activities sought to explain where (minima vs. inflections, etc.) the bump boundary might be placed (Watt & Morgan, 1983; Georgeson et al., 2007; Morgan, 2011; Subedar & Karam, 2016), not to capture its variability. In addition, different contours could also arise from lighting or rendering changes. For any of this to be useful, there must be a connection back to the scene domain. Here is the key point: The critical contour can also arise directly from properties of the surface and hence indicates surface properties. In particular, if surface slant is used as a rendering function, an image is created. Since there are then curves from the image (say, from a Lambertian rendering) and curves from the surface (say, from the slant renderer), we have our invariant: Both curves are within the sausage. To be concrete,

consider the bump in Figure 2. In the image domain, there is a special contour—a *critical contour*—that cuts across the isophotes in a particular fashion (Figure 2e). This could be a boundary of the bump as seen in the image. In the scene domain, there is a corresponding contour—what we shall later describe as an *extremal curve of slant*—that anchors the bump in surface terms (Figure 2f). Although in this artificial example, the critical contour in the image and the extremal curve of slant on the surface correspond closely, in general, they are only formally guaranteed to be contained in the sausage (Figure 2d).[3] Importantly, the presence of either one implies the existence of the sausage, which implies the existence of the other. Changing the rendering solely moves the critical contour around inside the sausage. All of this is proved in Kunsberg and Zucker (2018). We summarize this informally with the slogan:

> **critical contour $\Longleftrightarrow$ sausage $\Longleftrightarrow$ extremal curve of slant**

The sausage is thus the heart of our theory. It signifies the existence of regions in images that are especially informative about shape and implies the existence of structures that can anchor our shape percepts. But the sausage does not need to be explicitly computed. Computationally, we work instead with the critical contours and the extremal curves of slant. These can be estimated for individuals and can be calculated. While this intermediate sausage is implicit, the contours establish the intermediate stage from which the rest of the shape has to be interpolated or filled in. They serve as a kind of anchor, just as shapes are perceptually "filled in" from artists' drawings.

Before moving on, we note another delicacy in Figure 2e, f. The slant function is a kind of opposite to the image function; namely, bright image patches can correspond to low-slant portions of the surface. This duality is inherent in the topology and will be discussed further in Critical contours and qualitative shape representations.

Comparisons with other approaches

In this section, we provide a brief review of earlier approaches to shape-from-shading, to compare and contrast with our current proposal. We set the stage with classical computer vision approaches and then proceed to recent deep network approaches.

Traditional computer vision approaches were either differential or integral (Zhang et al., 1999). Mach (1965) and Horn and Brooks (1989) set the stage with the differential approach; for expository reasons, we illustrate the second one, formulated as a minimization. The problem is to find the normal $N(x, y)$ to the surface at every point that best matches the image $I(x, y)$ for a known light source at L subject to regularization constraints involving the partial derivatives of the normal:

$$\min_{N(x,y)} \int \underbrace{||I(x, y) - N(x, y) \cdot L||^2}_{1} \, dx \, dy$$

$$+ \lambda \int \underbrace{||N_x + N_y||}_{2} dx \, dy \qquad (1)$$

The normals enter through a Lambertian model, and the argument in Term 1 pushes them to agree with the image. The argument in Term 2 is a smoothness term. λ is a parameter for trading off data fidelity (Term 1) with regularization (Term 2). Without regularization, the problem is ill-posed; regularization is needed to reduce the possible solutions to a unique one. Loosely, minimizing the partial derivatives reduces surface undulations. Since the regularization and norms are evaluated over the entire image (or surface), each point matters and smoothness is reduced uniformly. This approach assumes the model is known, and estimating its parameters (e.g., L) can be problematic. Clearly, the model is far from general and the regularization can be extreme, especially in areas of high curvature. Some generalization was attempted by placing probability measures over models and parameters (e.g., Barron & Malik, 2012), but this led to a much higher dimensional minimization. Such computational problems are delicate and the results are extremely fragile.

Learning rather than explicit modeling is the current preferred approach for many vision problems (Saxena et al., 2005), often using deep neural networks (DNNs) (Lindsay, 2021). Instead of postulating a specific image formation model, DNNs function in (semi)supervised fashion, using training data given as pairs, (x_n, y_n), $n = 1, 2, ..., N$. The deep network can be viewed as a function f that "learns" the relationship $y_n \approx f(x_n)$; that is, it learns to interpolate between images and surfaces by, for example, solving a variational problem of the form:

$$\min_{f} \sum_{n=1}^{n=N} loss(f(x_n), y_n) + \lambda ||f||^2 \quad (2)$$

Note that, again, there are two terms, one for data fidelity and another for regularization (Parhi & Nowak, 2021). For shape-from-shading networks, x_n might be an (RGB) input image and y_n is the associated depth map (Eigen et al., 2014) or surface normal map (Tang et al., 2012; Wang et al., 2020). Although this follows the traditional model of seeking surface normals at every point, now it is the function f that is being solved for. Some loss functions add in Lambertian terms (Tang et al., 2012) (or other physical rendering models; Wu et al., 2015). Others argue that these models amount to latent variables and should be learned (Storrs & Fleming, 2021), introducing constraint through selected training data. Examples include, for example, "faces" (Sengupta et al., 2018) or "chairs" or "dormitory rooms" (Kulkarni et al., 2015); review in Breuß et al., 2021 In any case, the DNN architecture is tuned to interpolate the given data, so that the resulting algorithms can be brittle outside of it. Progress in generalizing from, say, *chairs* to *cars* is limited (Sitzmann et al., 2019). While particular solutions engineered for particular applications can be useful (Parhi & Nowak, 2021), especially when the data set is acquired from natural scenes, it can be difficult to understand why the interpolation works (Hutson, 2018). One might speculate that "faces" work because there is an underlying invariance—a template (Blanz & Vetter, 1999)—that guides the interpolation. In any case,the robustness of our visual systems remains elusive.

While we applaud DNN researchers' advances in building practically useful vision systems, we remain convinced of the importance of a deeper

level of understanding for achieving the generality and robustness of biological vision systems. The first artificial neural network for the shape-from-shading problem (Lehky & Sejnowski, 1988) provides a cautionary tale. This three-layer network was trained to indirectly classify shaded images of elliptical surface patches. That is, an image generation model was used to build the data set – Lambertian reflectance of elliptical paraboloids with different curvature – and the network was trained to classify curvature. The connections learned by the network yielded receptive fields resembling those for endstopped neurons. At the time, this result was surprising, because endstopping was a property thought to be associated with edge detection, not with shape-from-shading. In retrospect, however, endstopped neurons signal curvature generally, so the network could be understood as signaling the isophote curvature (Dobbins et al., 1987). In effect, the training data had the "result" built in directly: The elliptical patches were readily separable by the isophote curvatures, and the curvatures also specified the surface patches. The danger, to put it more generally, is that such networks learn to approximate the inverse to the data generator. To paraphrase Storrs (Storrs & Fleming, 2021), to "learn about the world by learning about images" *depends on which images are chosen*. Elliptical patches were a special case where a few curvature parameters specified the solution.

To avoid confounds such as this, a deeper understanding of the relationship between images and surfaces is required. We hope that our invariant can contribute to this.

Overview of the article

The article is organized as follows. We begin with background material, first introducing critical contours and the topological framework from which they derive (Critical contours and qualitative shape representations). Here we provide a pictorial introduction to the Morse-Smale complex (M-S complex). As background, we include two appendices to this article, an invitation to the general fields of differential topology (Appendix B) and a brief discussion of the M-S complex (Appendix C). Following this, we begin the new contributions in this article. In Extremal curves of slant, we extend the introductory discussion of the occluding contour to the boundary of bumps. Since bump boundaries are not occluding, we relax the global maximum slant requirement. The key derivation is within Figure 4. The result is a schematic definition of a bump as consisting of an extremal curve of slant that must be closed. It passes from a maximum of slant through a saddle. Examples illustrate how critical contours in the image

domain relate to extremal curves of slant in the surface domain.

Since we are only working with slant and not with a full description of the surface, the tilt coordinate remains unspecified. To constrain it, we use a notion of a generic surface, Extremal contours: Surface meaning, and show how this constrains the tilt so that the normals point consistently toward (or away from) the bump. Otherwise, a small change in viewpoint could lead to a large change in the image, a nongeneric situation.

To argue for the biological plausibility of our approach, we briefly discuss two approaches to computing critical contours (Computing extremal contours). The first is abstract; the second leads to a round, closed image flow; critical contours are integral curves through it, and its extent suggests the sausage from the Introduction. The importance of such circular flows has been known for some time, but the connection to bump shapes is, we believe, new. Examples illustrate its application to differently rendered bumps, from Lambertian (with different light sources) to specular (with different scene backgrounds).

Finally, we end with several phenomenological observations. First, since the extremal curves can provide a form of segmentation, we show how different bump constructs can spontaneously flip from convex to concave, without affecting their surrounds. Second, we extend the generic argument to show how normals must follow a consistent orientation, like the way figure and ground remain organized along contours. Although both the computational experiments and the psychophysical demonstrations are only proofs-of-concept, we hope that they will stimulate others to explore the rich possibilities that arise in applying (differential) topological ideas to perception.

An earlier version of this material was presented at (Zucker & Kunsberg, 2019).

Critical contours and qualitative shape representations

"Quantitative," as the term is typically used, connotes accurate, numerical, and formal, while "qualitative" can suggest vague and "informal." We use the term "qualitative" in a formal, mathematical sense, to denote topological mixed with geometrical ideas. Topology studies bumps and valleys; differential geometry studies curvatures. Topology is global (invariance over rubber sheet deformations), and differential geometry is concerned more with local properties, such as curvature at a point. We shall exploit this idea of smooth deformations when we define extremal contours.

Just as a sketch is a union of thin black lines on white paper, *critical contours* (Kunsberg & Zucker, 2018) are
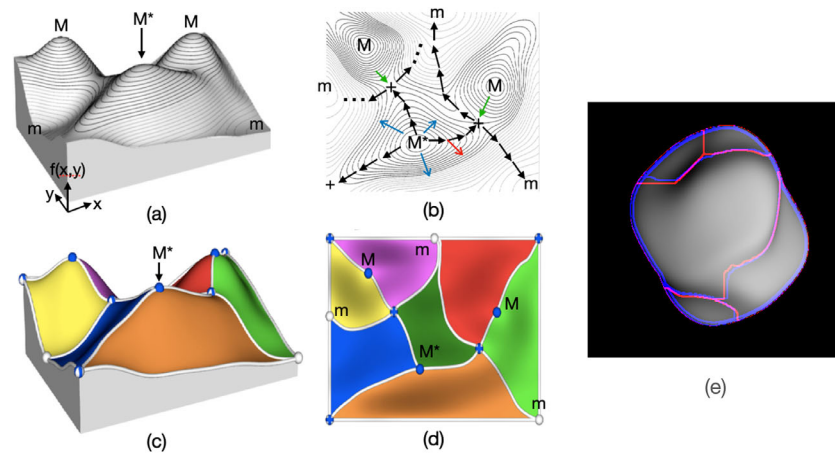
Figure 3. Brief introduction to the Morse-Smale complex. (a) Consider a scalar function $f(x, y)$ as a smooth mountain range (top view in (b)). Note how the level sets are nested and circle the maxima. (b) Imagine pouring water on the central peak (M*); most will flow downhill to nearby minima m along the gradient (blue arrows), cutting across the level sets. A special flow would follow distinguished paths along ridge lines (black arrows) from the maxima to saddle points (+) and then to minima. Should the flow deviate from the path, it will fall directly toward a minimum (red arrow). Other maxima may flow into these saddles (green arrows) so that in general, for each saddle, there are two entering and two leaving paths. (c, d) The Morse-Smale complex is a graph, in which the nodes (0-cells) are extrema (maxima, minima, and saddles) and the edges (1-cells) connect maxima to saddles and saddles to minima, thereby linking the black arrows. The graph forms quadrilaterals, called 2-cells, that provide a tesselation of the domain into components. (e) Illustration of a simplified M-S complex for a random figure: Red is computed from the intensities and blue from the slant function. Notice how they are close but do not coincide exactly. Figures (a) to (d) modified from Gyulassy (2008); see further discussion in Kunsberg and Zucker (2018), Kunsberg et al. (2018), and the Appendix.

a concentration of shading. Imagine a drawing of a ridge: A thin line, perhaps drawn by an artist, would be the limiting case in which the ridge has infinitely steep intensity "walls" surrounding it (see Figure 7 in Kunsberg et al., 2018). Informally, critical contours can be viewed as a sketch of the shading inside a shape, just as an occluding contour is a sketch of the boundary of a shape. Qualitative judgments, such as relative depths, can still be made in some circumstances (Koenderink et al., 2015). Formally, the critical contours are those edges (technically, 1-cells) of the M-S complex that have large transverse second derivatives.

Before explaining these words, a little of the historical background provides intuition. Maxwell (1870), building on Cayley, started the topological description of landscapes in the 19th century by examining how water flows downhill from mountaintops to valleys, along ridge lines and courses, and settles in minima. In brief, water flows down the gradient to minima, except when it follows a ridge line exactly and flows from a maximum to a saddle (Figure 3a, b). These ridge lines are special because they separate the flows toward one minimum from those toward another; any deviation off the ridge line and the water is drawn toward a minimum rather than the saddle. Of course, once the saddle is encountered, the flow is then toward a minimum. Note, in particular, that the flow is along the gradient of the surface; this is perpendicular to the level lines everywhere and that it

would remain relatively unchanged (i.e., qualitatively similar) if the (x, y) domain were smoothly stretched or compressed.

In the 20th century, Poincaré, and later Morse, developed the idea that the maxima, minima, and saddles—the critical points of a 2D function in 3D—could inform enormously the topology of the surface on which they were located. This idea is one of the foundations of modern topology and is introduced in Appendix B. Considering not the function, but the gradient of the function, was Smale's contribution. This gradient of the function could be thought of as a vector field on the surface, parts of which are illustrated in Figure 3b.[4]

The modern form of these observations is elegantly captured in the Morse-Smale complex, a rigorous way to qualitatively describe a scalar field via a set of curves (Biasotti et al., 2008; Smale, 1961); see Figure 3c, d and Appendix C.[5] It results, in effect, in an abstract, graphical version of what was just described. The mountain range becomes the value of a scalar function $f(x, y)$. The nodes of the graph (0-cells) are the extrema of this scalar function, or places where its derivative is 0; edges in the graph (1-cells) connect maxima to saddles and saddles to minima. Notice, in particular, how cycles of four edges (2-cells) are formed, connecting a maximum, a minimum, and two saddles in alternating order. These 2-cell quadrilaterals segment the mountain range into characteristic domains. We

shall shortly be modifying these components to develop the abstract definition of a bump as a special type of domain.

Remark 1. The Morse-Smale complex is a topological description of a function; it makes certain of its shape features explicit but does not specify the precise function values everywhere. Values on the complex can be used to get a "weak" representation of the original function.

We use the Morse-Smale complex as the foundation for our shape description. Just as a sketch is a union of thin black lines on white paper, *critical contours* (Kunsberg & Zucker, 2018) are a concentration of shading. Imagine a drawing of a ridge: A thin line, perhaps drawn by an artist, would be the limiting case in which the critical contour has infinitely steep intensity "walls" surrounding it (see Figure 7 in Kunsberg et al., 2018). Informally, critical contours can be viewed as a sketch of the shading inside a shape, just as an occluding contour is a sketch of the boundary of a shape. Formally, the critical contours are those edges (technically, 1-cells) of the Morse-Smale complex that have large gradients surrounding them. Critical contours are computable from the image, and a main result of that theory is critical contours are abstractly invariant to changes in the rendering function. They are contained within the sausage and, in effect, they define a type of scaffold on which a shape can be built.

The 1-cells of the M-S complex lie along the gradient flow, which is orthogonal to the level sets everywhere. A delicacy arises because, although many have suggested that the shading flow (along level sets) is the foundation for shading analysis (Koenderink & van Doorn, 1980; Breton & Zucker, 1996), others have observed that the level sets change drastically with changes in lighting or reflectance (e.g., Todd et al., 2014). While this observation is true in many places, it is not true in a neighborhood around critical contours (Figure 2); see also the detailed examples in Kunsberg et al. (2018). Second, across neighborhoods like this, the intensities change rapidly (the image gradient is large), from dark along the critical contour to bright in either direction normal to it. These two observations illustrate the basis for critical contours, and they hold generically for a wide class of intensity and surface variations (Kunsberg & Zucker, 2018).

While the Morse-Smale complex was built for a general scalar function, we here emphasize the two scalar functions utilized in our work. First, there is the image scalar function, the height of which is the brightness. Level sets, curves of constant brightness, are the isophotes. The gradient flow cuts across the isophotes. The second scalar function is the slant function on the surface. Assuming orthogonal projection, it can be aligned with the image intensities. The level sets are curves of constant slant. The gradient cuts across these, so it specifies the direction in which slant is changing most rapidly. An example is shown in Figure 3e, where the red complex is computed from the slant and the blue from the intensities. Notice how they are close but not in exact agreement (i.e., are in the same sausage neighborhood).

Remark 2. Critical contours provide a topological signature of key interior shape components, stable under generic lighting and rendering variations.

We illustrate these constructs for shape in detail after we specialize them to bumps.

Extremal curves of slant

We now begin the technical contributions in this article. At the outset, we note a difference in direction. Our previous development (Kunsberg & Zucker, 2018) began with the image, developed the limiting process for shading, and then connected it to surface properties. We now work in the other direction. Extremal curves of slant are motived from the rim, a surface property, and analysis of the slant function. We then develop image signatures for them.

The main challenge is how to define an analogous rim in the interior of the shape, so that it could surround a bump. This requires a relaxation of the slant along the occluding contour to something less than the global maximum value at every point, since the tangent plane never rotates enough to include the line-of-sight for a smooth bump. (This would cause an occlusion.) One might naively try to relax from the global maxima to some version of local maxima of slant along a contour, but this is impossible for technical reasons. (The surface has to be generic.) For Morse functions, this requires critical points to be nondegenerate (their Hessian, or second, derivatives must be full rank). Mathematically, for a generic Morse function, there are no appropriate contours consisting entirely of local maxima of slant.

There is a modification for the general surface shown in Figure 3 that does work. The closest analogue to the maximal slant contour is a Morse Smale 1-cell that includes maximal critical points while also existing for generic surfaces. This is why we call it an extremal curve of slant. In effect, we are seeking a curve around a bump that is steep everywhere and such that every step across it is a large step toward the top.

Before proceeding, it is helpful to recall from the Introduction that the Morse-Smale complex can be defined over any scalar field. We are working with two scalar fields here, one for the slant (extremal curves) and one for the image (critical contours). The main result is that they coincide, at least to the level of the sausages overlapping, so what is said about structure in one domain has an analogue in the other.
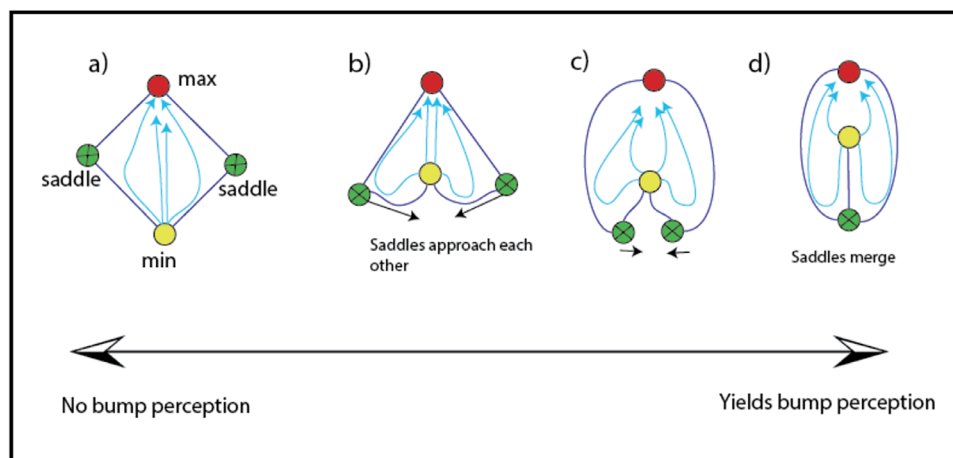
Figure 4. Development of the M-S complex for bumps on the slant function. (a) Starting with the standard MS cycle, (b) smoothly deform the domain so the saddles begin to (c) approach one another until (d) they merge. Notice how bringing the saddles together forces the minimum (in slant) toward the inside; this slant min is the top of the bump. Furthermore, the integral curves (cyan) deform along with the saddles. These will become important when studying this configuration in the image domain, because they then denote isophotes.

To start, we focus on extremal curves of slant. Recall (Figure 3) that the 2-cell, the basic building block of the M-S complex, consists of a region surrounded by a special flow from a maximum to a pair of saddles and then to a minimum. While all flows are along the gradient, the 1-cells (the edges of the 2-cells) are special—they delimit the building blocks with the flows into saddles and then into minima.

We now seek the appropriate generalization from the occluding contour to an interior curve, that is, the closed extremal curves that could surround a bump (more technically, the closed curves that anchor the sausages). These are the contours following the gradient flow that go through local maxima and saddles, but with a modification of the MS-complex, in which the saddles have been joined. To derive these (Figure 4), begin with the standard 2-cell from the M-S complex for a scalar function, say the slant, and deform it by allowing the saddles to approach one another until they merge. This deformation is smooth and generic, and amounts to a distortion of the surface. It is helpful to imagine this deformation starting from the configuration in Figure 3. In the end, when the saddles merge, the associated minimum is forced into the center. Importantly, this configuration remains generic. Furthermore, corresponding to this evolution of the slant function, there is a corresponding evolution of the image function. This latter has been considered in the computer vision literature (Nackman, 1984; Griffin & Colchester, 1995) before, but not the slant version. Nor, to our knowledge, has anyone previously exploited its topological implications for perception. It gives rise to a distinct type of 2-cell typical for a bump.

It is instructive to analyze an image of a bump (Figure 5) in detail. First, we illustrate part of the definition of a Morse function. Recall that these were functions whose critical points were isolated. This would discount flat planes or ridges, for example. So, to make the surface Morse, we added slight undulations; these are easily visible in the surface on which the bump sits.[6] Now, viewed from above, notice how the slant is minimal (Figure 5f), then increases to its maximum on the steep sides, and then decreases again. While this maximum is not $\pi/2$, it is large. The M-S complex for this is shown in Figure 5j. A blue extremal curve passes through a maximum (of the slant function) (yellow dot) and a (green) saddle point. The circular curve is the extremal line of slant, because this is the only part of the M-S complex that has "steep sides" of the slant function. The minimum is somewhere inside it. The template for this pattern (Figure 5b) schematizes this; in effect, this template is the definition of a bump in the slant domain. By passing from the max to saddles, it illustrates how the extremal curve is a relaxation of the occluding contour. The M-S complex on the slant function shows how the extremal contour encircles the bump, with a slant minimum (and no other maximum) inside it.

Also shown in Figure 5c, d, e are three different renderings of the bump. The first two are Lambertian, and the third is a nonlinear and somewhat artificial mix of lights. Each image of the bump shows a different intensity distribution and, necessarily, different isophote arrangements. Notice how the (image) critical contour is the same circular curve, however; moving the light source or changing the reflectance moves the max position and the isophotes, but not how the critical contour cuts through them. This cut, the gradient flow through the level sets along a distinguished path, remains essentially invariant over lighting and rendering
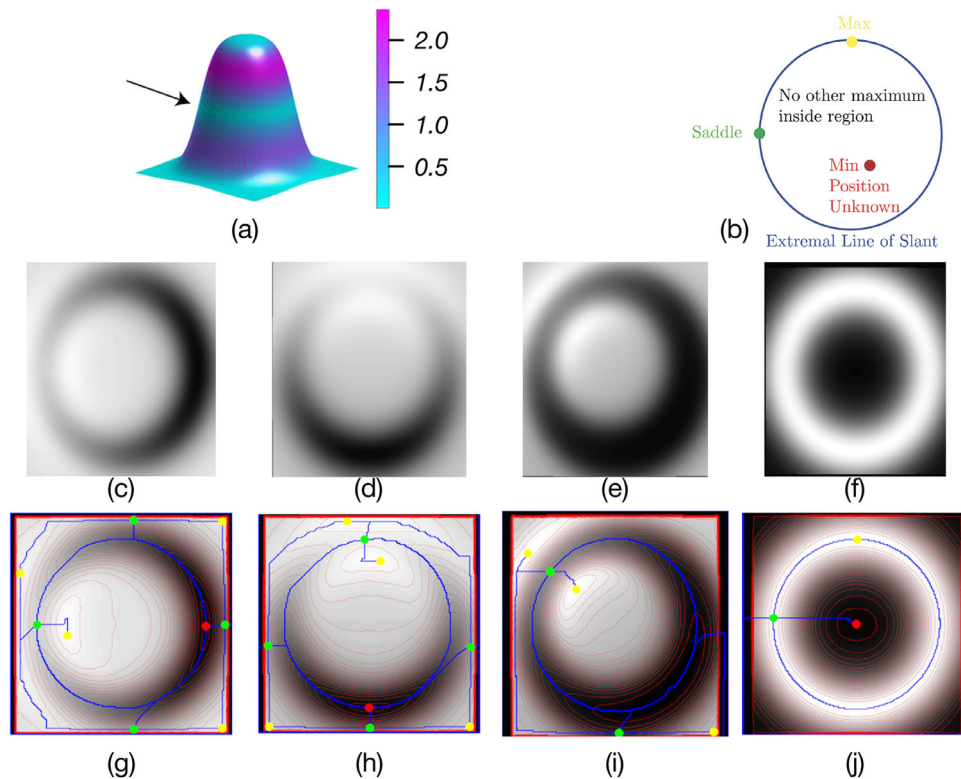
Figure 5. An artificial bump on a slightly bent (i.e., generic) surface. (a) The bump is colored by the Gaussian curvature; the arrow points to the parabolic curve along which the Gaussian curvature is zero. (b) The defining template for a bump, shown in the slant domain so that an extremal curve surrounds a minimum. (c, d) A Lambertian reflectance function on the bump illuminated from two different positions. (e) An artificial rendering function based on a nonlinear combination of Lambertian reflectances. (f) The slant function for the bump viewed from above. (g–i) The isophotes and M-S complex for the three images and (j) for the slant function.

changes. The cut defines an integral curve through the gradient flow, and this curve anchors the sausage neighborhood around it.

Comparing the intensity images to the slant figure reveals a reverse in sign: The brighter parts of the image correspond to darker parts of the slant. This makes total Lambertian sense: The bright part of the bump points toward the light source, while the slant is minimal. In topological terms, this is handled by the "stable" and "unstable" components of the M-S complex. They are related, informally, by taking the flow along the scalar function, where the water runs downhill, or the flow along "minus" the scalar function, where the water runs uphill. That is, maxima and minima are interchanged. These are the "ascending" and "descending manifolds" technically, and the M-S complex is their "transverse intersection." For our purposes, however, it is useful to consider them separately.

Our next example is an image from the Todd database (Figure 6). Beginning with the slant function, notice how each bump is surrounded by an extremal curve that passes through a maximum and a saddle; this maximum is the "steepest" part of the bump. As we saw above for the Lambertian example, interior to the bump is a minimum in slant that often corresponds to

a maximum in intensity. Note that the extremal curve clings to the high-slant portion of the function and that this bright region (high slant) encloses—and is surrounded by—darker, low-slant neighborhoods.

In multibump images such as these, the M-S complex provides some constraint on the "space" between bumps (Figure 6c). This space can undulate slowly; it supports a number of saddle points that can connect the bumps together. Unlike the extremal curves, these 1-cell/saddle connectors are not stable; they can move a good distance with changes in lighting or material. The intensity change across them is not as extreme as with critical contours. As we show later, however, they do provide some constraint on the bump/dent or protrusion/intrusion ambiguity.

Stating the previous discussion concisely, we have:

Definition 3.1. Slant extremal contours are the saddle-maxima 1-cells of the M-S complex of the slant function.

Slant extremal contours are interior contours, defined on the surface, that pass through the maxima of the slant function, follow the gradient flow, and, at the same time, segment the surface into regions. For brevity in the remainder of this article, we will
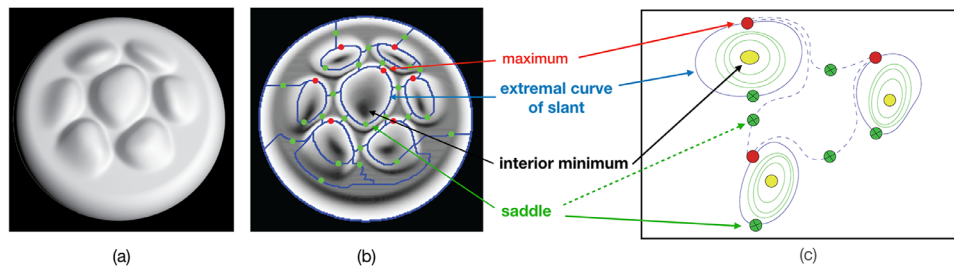
Figure 6. Illustration of a bump max, saddle, and unknown interior minimum on a Todd example. (a) Intensity image. (b) Slant function for the surface. Typically, the slant function has a minimum somewhere in the interior, while the intensity function often has a corresponding maximum (e.g., for Lambertian renderings). Notice how the extremal curve passes through the bright region denoting high values of slant; the portions of the M-S complex other than the closed extremal curves are not critical contours. (c) A sketch of the topological network (part of an M-S complex) corresponding to several bumps. Although this sketch was not built directly from the more complex example in (b), notice how there are saddles that link maxima of slant on nearby bumps together. These 1-cells, shown as dashed contours, do not have significant slant changes across them, so are not as stable as the extremal curves surrounding each bump. The iso-slant contours surrounding each bump minimum will be discussed shortly.

drop the word "slant" and call them simply "extremal contours."

In general, extremal contours will not be closed. However, motivated by the previous examples, we are especially interested in the special case when an extremal contour connects back to itself.

Definition 3.2. Extremal rings are closed slant extremal contours.

Putting these pieces together, we have (up to a concave/convex reversal):

Definition 3.3. A bump/valley is an interior region within an extremal ring.

Extremal rings and occluding contours

We now investigate extremal contours and extremal rings and show that these features define the surface topologically while also being visually salient for many rendering functions. It will formalize the previous observations. In particular, we show that, in all likelihood, these extremal contours have the same two properties that the occluding contour had: *surface meaning* and *image salience*, provided the surface and rendering are generic.

Extremal contours: Surface meaning

We first show that extremal contours represent boundaries of bumps and valleys of the surface. The argument will be based on the conclusion that the surface normal field near the extremal contour (i.e., within the sausage) should generically point uniformly

to the interior (or exterior) of the region bounded by the extremal contours. Our argument is based on a generic prior for vision, in which small changes (e.g., in lighting) should imply bounded changes in the image (cf. Freeman, 1994). This section is slightly technical and may be perused quickly to get the gist.

Consider Figure 7, and focus on the enlarged portion of the extremal ring. This was defined to cut across the level sets of the slant function but, by definition, it does not inform the tilt function. Here we show that, for generic surfaces, the tilt must be constrained along the extremal ring to prevent self-occlusion and image instability. There are two basic possibilities: In the first case, the surface behaves like a bump boundary, where the surface normal on the extremal ring points consistently; the second case leads to rather wild surfaces with "crazy" curvatures. While this second case is a mathematical possibility (developed next), it violates our requirement that the surface be generic in the perceptual sense. That is, a small change in the lighting would lead to a drastically different image (Figure 8). This is typically not the case for natural surfaces, since their image does not change drastically for small changes in viewpoint. For the wild surfaces, the extensive changes happen because the normal points in many different directions in very small neighborhoods, covering a very large portion of the Gauss map (introduced in Appendix D).

More technically, let $\alpha(t)$ denote an extremal ring that bounds a region R, and let $\{x, y\}$ be image coordinates under orthographic projection. Let $\sigma(x, y)$ represent the slant in a neighborhood of a point (x_0, y_0) on $\alpha(t)$. Rotate the frame so that the slant gradient (tangent direction to $\alpha(t)$) points locally in the y direction. We compare two possible solutions for the local surface depth $S(x, y)$ by defining two pairs of Taylor expansions for the slant and tilt functions. Let $\sigma_1(x, y) = \sigma_2(x, y) = c_1 + c_2 y + \sigma_{xx} x^2 + \sigma_{xy} xy + \sigma_{yy} y^2$.
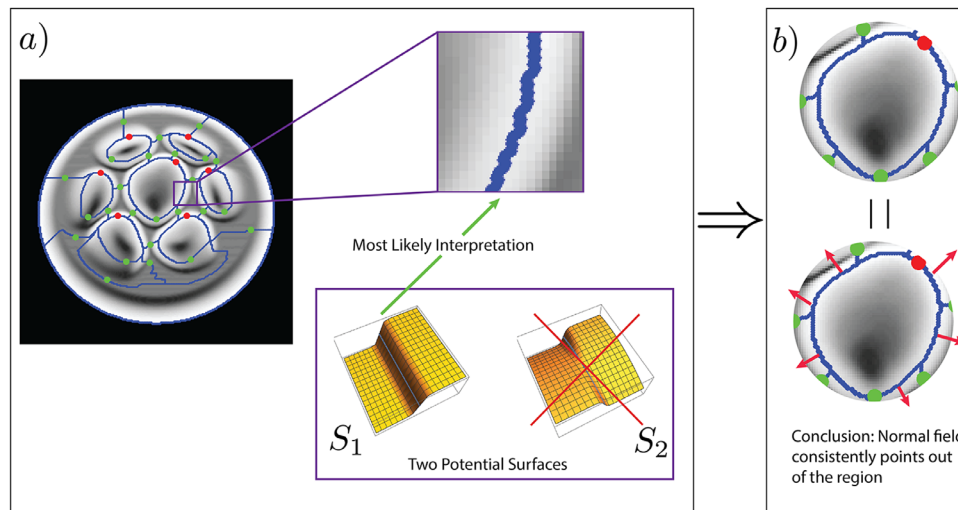
Figure 7. Extremal curves and surface normals. (a) Consider a portion of a maximal curve of slant. Two distinct types of local surfaces could have caused this local slant function, one with no twist and another with substantial twist. The first surface, without twist, is much more likely than the second surface. Note the normal field for the first surface consistently points to the same side (left) of the extremal curve. (b) Taking the most likely interpretation for each portion of the extremal curve, the normal field must point uniformly outside (or inside) the entire contour.
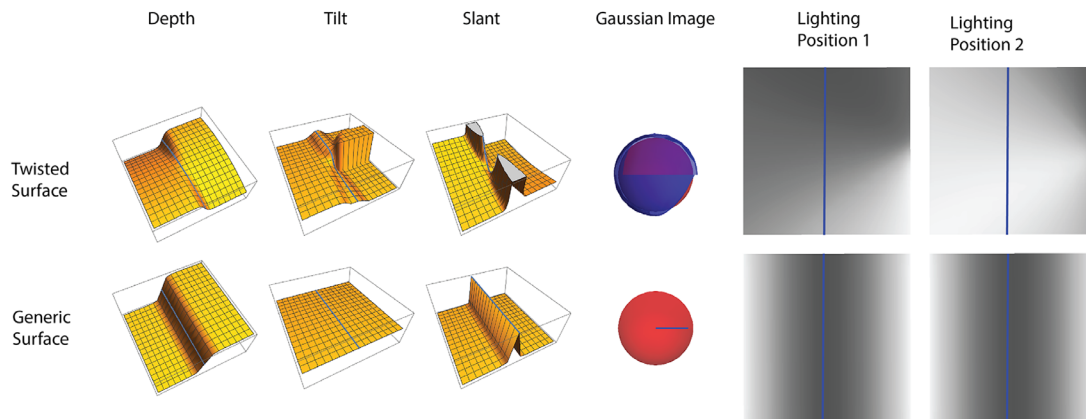

Figure 8. Comparing possible surface explanations for a maximal slant curve. Each row depicts different properties of the surfaces shown in the first column. Although both are technically solutions, the twisted surface covers most of the Gauss map (Gaussian image shown) and, when rendered under slightly different lightings, gives rise to drastically different images. The generic surface solution occupies a small portion of the Gauss map and renders almost the same image under different light sources. Both arguments show that the twisted surface (top row) is much less likely.

(There is no linear x term here due to the slant gradient pointing along the y axis.) Let $\tau_1 = x$ and $\tau_2 = y$. Now, $\{\sigma_1, \tau_1\}$ defines a surface S_1 and $\{\sigma_2, \tau_2\}$ defines an alternative surface S_2. Which of these is more likely?

Both S_1 and S_2 have the same magnitude tilt gradient. However, S_1 (generic surface) has the tilt change along the contour while S_2 (twisted surface) has a tilt change perpendicular to the extremal contour.

Let $N_1(x, y)$, $N_2(x, y)$ represent the normal fields for each of these two solutions. We compare the relative probability of each of these surfaces by considering the term $T_i = \int_\Omega \det(DN_i DN_i^T)$ for each surface.

This is essentially the Gaussian curvature of each solution integrated over the patch.[7] The solution with higher Gaussian curvature will be the solution that is less smooth, more dependent on lighting direction (Freeman, 1994), and with a higher chance of occlusion. We compare the relative likelihood of the two solutions $\frac{L_1}{L_2}$ by considering the inverse of the ratio $\frac{T_2}{T_1}$. A simple calculation shows

$$\frac{L_1}{L_2} \propto \frac{T_2}{T_1} \propto \frac{\sigma_{xx}^2}{\sigma_y^2} \quad (3)$$

where σ_{xx} is the transversal second derivative of the slant across $\alpha(t)$ while σ_y is the gradient along $\alpha(t)$. Since the slant on α is extremal, its gradient will necessarily be small. In addition, since the slant is changing rapidly across α, σ_{xx} will be large. Thus, for a slant patch as shown in Figure 7a, the ratio $\frac{\sigma_{xx}^2}{\sigma_y^2}$ will be large. This statement is illustrated by the comparison in Figure 8. We conclude that the image patch derives from a surface patch with a surface normal field that is not twisted. In other words, it is most likely that the normal points to a single side of the curve in the entire Taylor expansion.

Applying the above argument completely around the extremal ring $\alpha(t)$ shows that the most probable interpretation of the surface normal field along $\alpha(t)$ is that it does not have a twist and therefore must point uniformly outside (or inside) the region bounded by $\alpha(t)$. It then follows that the region R is "higher" (or "lower") than the surrounding area. More precisely, one can see that the R must be an ascending or descending manifold of depth; in other words, the region is a bump or a valley.

Remark 3. The surface normal along an extremal contour points consistently to the interior or exterior, almost always. When closed, the normal points consistently to a bump or a dent.

Extremal contours: Image salience

We now show that the second important property of the occluding contour, image salience, also exists for extremal contours. Since extremal contours are related to critical contours (Kunsberg & Zucker, 2018), image salience will follow from the image invariance of critical contours, which we now review.

A given surface, when rendered differently (e.g., with a different light source), can yield drastically different images overall. This was illustrated in Figures 1 and 5. However, many rendering functions such as Lambertian shading, specular shading, texture, and so on, as well as line-drawing algorithms (DeCarlo et al., 2003; Judd et al., 2007), all involve the surface normal field in order to create an image.

The surface normal field $N(x, y)$ can be defined as a map from the image domain $\mathbb{R}^2$ to the unit sphere S^2. It is then natural to define a rendering function F as a smooth map from the unit sphere to the real line, that is, $F : S^2 \rightarrow \mathbb{R}$. The image is then expressed as the combined map $I(x, y) = F(N(x, y))$. The orientation field (for a smooth rendering function) is then computed perpendicular to the gradients ∇I.

In Kunsberg and Zucker (2018), we defined critical contours that were computable from the image. Generally,

Definition 4.1. Critical contours are gradient flows in the image with large transversal second derivatives.

Remark 4. Critical contours are computable from the image gradients, whereas extremal contours are computable from the slant gradients.

In Kunsberg and Zucker (2018), we showed critical contours had an invariance to the choice of F. More precisely, to show image invariance of a critical contour for a wide class of rendering functions $\mathcal{C}$, we proved:

Theorem 1. *Given a surface normal field $N(x, y)$ and any two choices of generic rendering functions $F_1, F_2 \in \mathcal{C}$, construct $I_1 = F_1(N(x, y))$, $I_2 = F_2(N(x, y))$. If a critical contour is present in I_1, then there is a arbitrarily close critical contour in I_2.*

A corollary of this statement (Corollary 10 in Kunsberg & Zucker, 2018) can be restated:

Corollary 4.1. An extremal contour must lie in the tubular neighborhood of a critical contour and have the same endpoints.

This tubular neighborhood is the formal version of the red sausage that has appeared throughout the article, and the above statement is a formal version of Figure 2d. Informally, the tubular neighborhood is a region of points on the manifold displaced from the critical contour in the orthogonal direction.

As extremal contours lie near the critical contours, the invariance statements from Kunsberg and Zucker (2018) show that extremal contours will nearly always be salient from the image, regardless of rendering function. Extremal contours delineate the surface features while the critical contours delineate image features. We can observe a critical contour, infer an extremal contour next to it, and then use the previous section to attribute surface protrusions to image regions. The sausage links the two; in the limiting case, when the sausage collapses, the extremal contour coincides with the occluding contour. This is what lies behind the slogan in the Introduction.

Thus far, we have only considered single bumps. But since the M-S complex is global, multiple bumps can be identified from portions of the M-S complex. We provide an example of this in Figure 9, a collection of figures from the Todd experiment. Although closed critical contours are evident, variations are clearly present. Part of this variation is due to rendering differences, such as around the large bump on the far right. As just explained, these variations are to be expected and live in the sausage. Another part is due to the manner in which critical contours are computed; examples of this are evident in the breaks connecting to the exterior bump. We elaborate on computational issues in the next section.
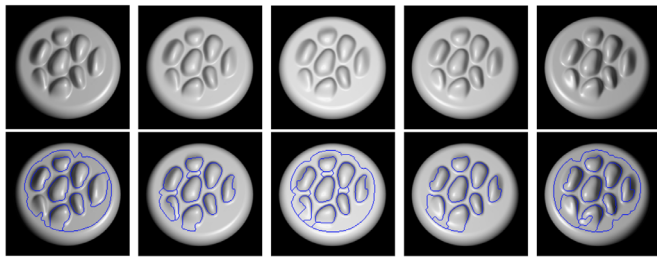
Figure 9. Critical contours on differently rendered images by persistence simplification. Top row: A sequence of images of a surface rendered with specular and Lambertian shading; the light source moves from left to right. Bottom row: Closed critical contours on the image; our theory implies that these will contain "bumps." Note that some bumps are "connected" by 1-cells. Although these are part of the simplified M-S complex, they are not critical contours. See text for further discussion. Top figures courtesy J. Todd.
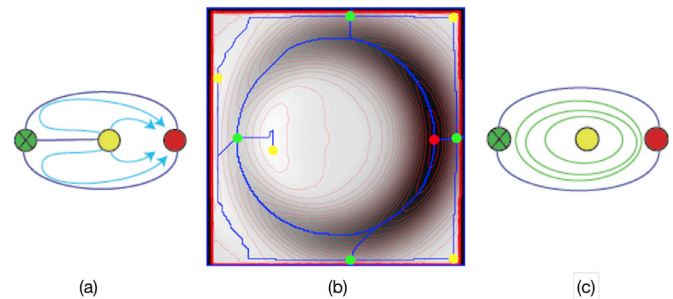

Figure 10. Concentration of isophote tangent orientation around bumps. (a) The slant extremal contour from Figure 4. Note the gradient integral curves in cyan. (b) A corresponding image from Figure 5. Notice how the closed critical contour (corresponding to the slant extremal contour) cuts across the isophotes. Moving slightly interior, the isophotes form concentric rings; this is cartooned in (c) by the concentric iso-slant contours in green. Such families of nested level sets can serve as an image signature for the identification of bumps.

Computing extremal contours

Topological constructs, such as the Morse-Smale complex, are defined using continuous mathematics. They involve real numbers. But when noise and sampling are introduced, technical problems can arise in assessing topological signatures. For example, since pixels are discrete, what is the precise location of a singularity? To compute the gradient at a fine scale, in a practical sense, we are restricted to computing differences in the image coordinate directions: Will algorithms converge to the actual gradient? Finally, is there a real tiny hole in the shape or are values missing because of noise or quantization? The field of computational topology is being developed to answer questions such as these and to develop algorithms that run correctly on discrete domains such as images.

There are two different ways in which this can be approached (Günther et al., 2014), one topological and the other based on the gradient flow. We used the first approach in Kunsberg and Zucker (2018), and we used it to compute all of the examples in this article. The second approach is more relevant to biological realizations. We discuss both in turn.

Persistence simplification (Edelsbrunner & Harer, 2010; Carlsson, 2009) is one of the important developments in computational topology. Basically, take a discrete structure, cover it with a smooth object, and then use it to calculate topological features such as critical points. By the above arguments, sampling and noise can introduce singularities that are irrelevant to the actually continuous structure. Just as blurring can smooth over tiny holes due to noise in an image, persistence simplification is a globally consistent way to reveal overall structure while removing those tiny holes and "irrelevant" noisy details that derive from quantization and discretization.

Critical points that are extremely fine scale and local are eliminated in a kind of structure-preserving smoothing. Algorithms to compute the Morse-Smale complex from discretized images (i.e., a mesh) have been developed by Reininghaus and Hotz (2011), Sahner et al. (2008), and Weinkauf et al. (2010), among others. We use the algorithm of Reininghaus and Hotz (2011) in all of our experiments. However, because persistence only involves how much smoothing a critical point can survive, the explicit "steep sides" in the critical contours definition are not always incorporated. This explains the variation in the cobblestone image in Figure 9. Some artifacts (noise/discretization loops) were eliminated, and by and large, almost all bumps are localized in every example. But the result is not perfect and more work would be required to fully implement the critical contours "steep sides" criterion.

Instead, a more biological approach to computing extremal contours would work directly on the orientation flow rather than the discrete image, and builds on the fact that cells in visual cortex are orientation selective. We illustrated earlier how isophotes concentrate near the occluding contour (Figure 1); similar properties hold around bumps. In particular, focusing on the level sets during the evolution shown in Figure 4 makes clear that they concentrate in an analogous fashion near the critical contour; see Figure 10. In effect, as the saddles converge, they "pull" the isophotes with them. This leaves a special signature where the isophotes form a concentric flow within and around the critical contour. Here, discretization works in our favor, because exceptions could be regularized away with co-circularity (Ben-Shahar & Zucker, 2004).

This observation connects to another body of literature. Perceptual sensitivity to closed contours (Kovacs & Julesz, 1993; Elder & Zucker, 1993) and circular textures (e.g., Wilson & Wilkinson, 1998;
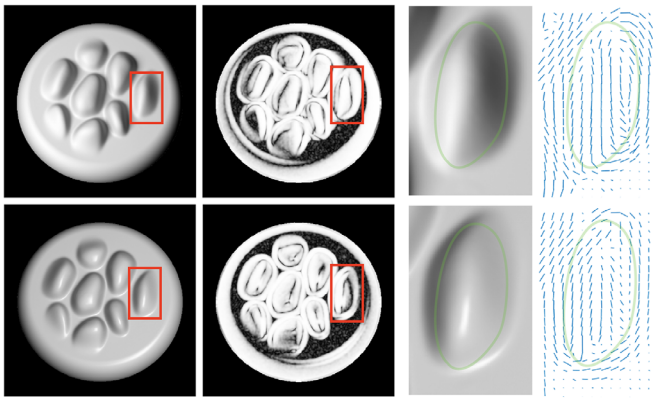
Figure 11. Max slant cycles—closed critical contours—can be found via maximum compression in shading images. Across each of the rows, we show original image, "flow compression," an enlarged region, and the "flow direction." Notice the similarities in the flow compression and flow direction across the different renderings. The width of the flow provides an indication of the red sausage.

Dumoulin & Hess, 2007; Dakin & Bex, 2002) is well known. What remains to be explained is why this sensitivity is so common. Our answer is that these concentric flows are the image-domain signature of bumps and dents. Importantly, just as concentric flows are not particularly sensitive to positioning and other perturbations (e.g., Achtman et al., 2003), the details required for accurate computational algorithms in topology are regularized away. Furthermore, there is an accumulating body of evidence that orientation structure is at the basis of related shape perception (Li & Zaidi, 2000; Fleming et al., 2011; Marlow et al., 2019) and has been interpreted as a component of shape representation at an intermediate stage (Gallant et al., 2000). In total, all of this evidence has a particular advantage from our perspective: Our result about concentric flows formally connects an image salient property with 3D shape interpretation. To speculate even further, this connection suggests how intermediate visual areas, such as V4, could be sensitive to 3D structure (of bumps). Some preliminary evidence in support of this is presented in Kunsberg and Zucker (2018a).

In Figure 11, we show the oriented shading flow around two pebble examples, plus an integral curve[8] through it. Notice how this curve surrounds the bumps just as the extremal ring did. The extent of the parallel flow correlates with the sausage around the contour, which could provide a means for estimating it. Since these flows are a type of oriented texture, we speculate that working with nested flows also supports a generalization from shape-from-shading to shape-from-texture. Formally, this follows because texture compression induces similar flows (Cholewiak et al., 2014). Finally, we show a number of specular examples, in which the oriented flows arise from
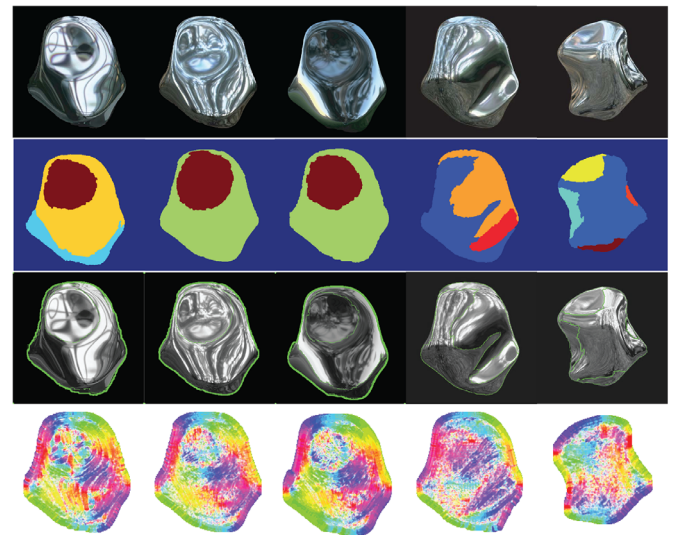

Figure 12. Critical contours on specular images (from Fleming et al., 2004). (Row 1) Examples of five specular images. Note the organized flows apparent around the bumps and dents. (Row 2) A segmentation of the images according to integral curves through the structure tensor direction. (Row 3) The integral curves. (Row 4) The colorized orientation map. In Extremal contour relational constraints, we suggest why these examples appear as dents rather than bumps.

compression of the visual scene around the object (Fleming et al., 2004; Mooney & Anderson, 2014); see Figure 12.

In all of these examples, the flow compression was computed by a classical technique in image processing, the structure tensor (Rao & Schunck, 1991; Bigün et al., 1991), from which the primary direction was derived. Compression was assessed by the ratio of eigenvalues. In effect, the structural tensor is an ellipse, and the compression is a measure of elongation. There are more biological ways to compute these flows (Ben-Shahar & Zucker, 2004). We use the structure tensor here only as an easy proof-of-concept. Although more experimental work is needed, we hope that these examples illustrate how topological ideas could inform different ways of thinking about relatively well-understood biological constructs (namely, orientation selectivity and its uses) and how orientation flows could make abstract topological ideas more concrete, more computable, and more robust. Several other directions for future study are described next.

Directions for future study

We have just illustrated how the theory can be used to find "bumps" in images with unknown rendering and illumination, and how there is at least the potential for biological realization. We now move up to the phenomenological level. First, we show how the theory
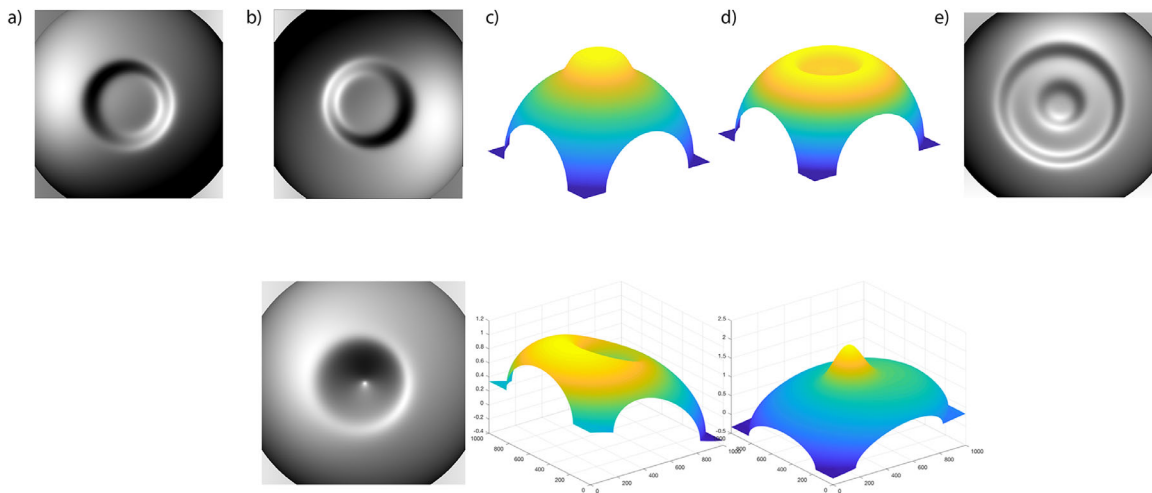
Figure 13. Multistable bumps and dents. Top: (a, b) Images with bistable perceptions governed by an extremal curve segmentation. The inner disk is carved out by an extremal contour and can be seen as coming out or going into the page. At first glance, (a) often looks like an indentation and (b) like a protrusion; as in the crater illustration, (b) is the inversion of (a). Staring at either one shows that it can flip to the other. (c, d) Possible surface interpretations of images in (a) and (b). (e) Nested bumps provide a higher-order multistable illusion. Again, components can be flipped independently. Bottom: An extension in which the bump is not directed at the viewer and the object is asymmetric. Bistability still maintains.

predicts a novel 3D bistable illusion. Following, we show an illustrative example applying the theory to nonclosed extremal curves; a constraint labeling problem arises. Finally, we show a demonstration comparing the importance of these extremal contours over other image regions. While none of these excursions is complete, our goal, in this section, is to illustrate directions for, and perhaps stimulate, future study.

Bistable dimples and bumps

The definitions of extremal contours and closed rings were in terms of the M-S complex, and involved only singularities and their (global) relationships. However, since the singularities are organized through their indices (Appendix C), maxima and minima can enjoy a complementary relationship. For example, we showed in Figure 5 how maxima in the image can correspond to minima in slant, which follows from Lambertian (and other) rendering functions. Formally, we discussed how this arose in the context of ascending and descending manifolds. We here show how this duality is perceptually present.

Closed extremal contours, and their image-equivalent closed critical contours, can each signal the existence of bumps and valleys, which immediately brings the convex/concave illusion to mind (Ramachandran, 1988b). Set up properly, this is a classical instability in shape perception, often disambiguated by the common light-source-from-above heuristic. Since an extremal

ring perceptually creates either a bump or a valley, that is, does not resolve the convex/concave ambiguity, it follows that similar instabilities should be demonstrable using extremal curves as well.

Here is a key difference from changing the overall light direction: The extremal curve provides a different kind of organization. Since the extremal contour actually segments the bump or valley, it follows that these segmented parts should be independent of other portions of the image. This raises a prediction: that the individual parts of an image should also be subject to the multistability *individually*.

To confirm this prediction, that is, to show that the individual parts—and not the full image—can be flipped separately, we created the bistable images in Figure 13 with an extremal ring in the center whose interior region is a disk. The extremal ring creates a perceived segmentation but does not specify whether the interior disk is a bump or a valley. The normal vectors along the ring, which we showed had to point consistently once projected into the image, can, in this case, point into or outside the ring. Thus, the interior part can be seen as a bump (pointing out of the page) or a dimple (into the page); because it is ambiguous, they can be perceptually flipped. Normally, one might expect this to be solely an effect of lighting (Ramachandran, 1988a; Morgenstern et al., 2011, 2014), but this example emphasizes different structural aspects. An asymmetric version of this illusion is shown in Figure 13 (bottom). The concept of segmentation facilitates describing this perceptual phenomenon; it is cleaner than representing the two solutions as independent
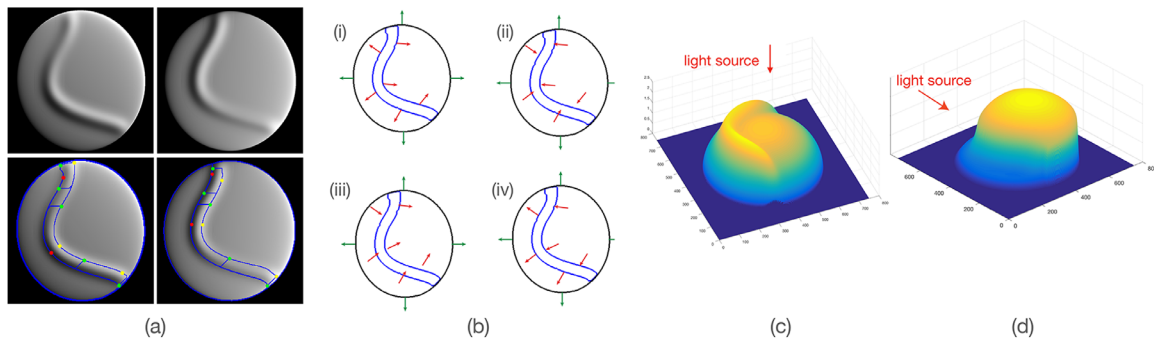
Figure 14. From bumps to extended ridges. (a) Top row: Images of the same surface illuminated from different positions. Bottom row: Persistence-simplified Morse-Smale complexes showing the extremal curves in blue. (The short horizontal segments should be ignored; they are not extremal contours.) (b) Four possible labelings of the surface normal field on these curves. (b,i) A ridge above the surface; (b,ii) a valley; (b,iii; b,iv) two extended slopes. (c) The surface corresponding to (b,i). (d) The surface corresponding to (b,iv). Note, in this case, the gradient points inconsistently, so that the surface continues to decline. (See profile at edge of ridge.)

depth fields. The bistable bumps can be nested as well.

Remark 5. This is not the concave/convex ("hollow face") illusion in disguise because the surface portion outside the extremal ring is stably perceived as convex in depth. This is due to the portions of the occluding contour that are visible at the edges of the image. This example could be generalized to include N extremal "rings" leading to possible 2^N ambiguous perceptions. Rather than a single concave/convex ambiguity on the global object, we could have concave/convex ambiguity on individual parts governed by the extremal curves.

Extremal contour relational constraints

The bistable illusion above arises from the binary choice of the normal field on the extremal ring—it can either point consistently toward the "inside" (as in a dent) or the "outside" (as in a bump). We previously argued that, to be generic, such consistency is required almost everywhere along an extremal contour. We now elongate the bumps to show how a combinatorial logic of normal constraints develops. This logic may indicate why concave/convex ambiguities are less frequent in natural vision than they might be.

Like closed extremal contours, individual nonclosed slant extremal curves can also provide information about protrusions and other features within a shape. The result in Extremal contours: Surface meaning can be summarized as the normal field along an extremal curve must point to one side of the curve (almost always). This can be thought of as a labeling of the curve, just as the "border" side of a curve is indicated by the Gestalt notion of "border ownership" (Zhou et al., 2000). If extremal curves are extended toward an occluding contour without introducing other structure, then the logic can be applied. For each of
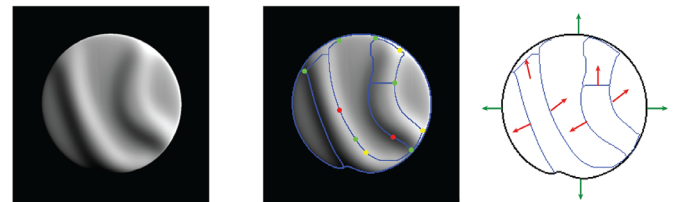

Figure 15. An extension of Figure 14 in which the critical contours partition the surface into several parts. These can then be interpreted as ridges and valleys through a consistent labeling scheme.

the two extremal curves in the example surface shown in Figure 14a, we get the four qualitatively different labelings depicted in Figure 14b. The red arrows depict the binary choice for the surface normal on the curve and the green arrows depict the necessary normal constraints due to the occluding contour. These labelings are simply a choice of orientation for the surface normal on each of the two blue curves.

Although we have not done formal psychophysics on this, the authors perceive labeling (a) although one can construct surfaces with precisely the same images in Figure 14a from the other labelings also; see Figure 14d. Similar labeling constraints can be applied to more complex imagery (Figure 15), and we would speculate that the reason the dents appear fixed in Figure 12 is because the "arrow" along the occluding boundary must point outward. Further examination of these issues is warranted, at least to extend our understanding of generic lighting and surface interactions.

Critical contours as 3D shape anchors

Our basic hypothesis is that critical contours are key to 3D shape perception. As part of the Morse-Smale complex, they are like the pencil lines that might be drawn by an artist. In both cases, they provide a scaffold
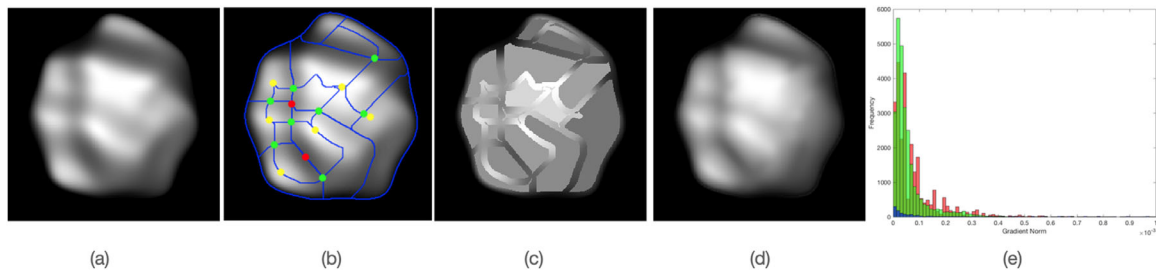
Figure 16. (a) A random blob with a vivid shape percept. (b) Its (simplified) Morse-Smale complex. (c) A flattened image. The intensities on the M-S complex match the image intensities while, for each region within them, the intensity is replaced by the average. (d) A pseudo-heat equation blur of image (c). Notice how this shape appears similar to the original (a), even though the image intensities differ. (e) Histogram of intensities in the different images. Red corresponds to (a), blue corresponds to (c), and green corresponds to (d).

on which qualitative shape inferences could be based. This radical hypothesis—that only selected areas in an image anchor shape perception—needs far more exploration. In a related article (Kunsberg et al., 2018), we exploited a color-shape interaction to demonstrate that the neighborhood around critical contours, and not the space between them, sufficed to ground 3D inferences. We now provide additional support for this observation with two new displays. Both exploit the blur/shape ambiguity developed at length in Mooney et al. (2019), which should be consulted for discussion and many references.

The placement of critical contours is key to 3D shape perception, as we now demonstrate. We start by illustrating the relative unimportance of shading values away from the critical contours by modifying images. We first isolate the critical contours, extend a small sausage around them, and then set the intensity within this to the underlying value of the critical contour. The areas between critical contours (the 2-cells) are set to a constant, which is the mean value of intensity within it. Finally, the resultant discontinuous image is smoothed by a pseudo-heat equation, thereby implementing a blur operation. While the intensities in the resultant image differ from those in the original by a nonlinear function, the shapes appear about the same; see Figure 16. Importantly, many of the pixel intensities differ—see histogram—while, at least informally, the shapes do not.

The second demonstration in Figure 17 complements the above and builds explicitly upon the blur investigated in Mooney et al. (2019). We begin with a set of contours of constant intensity against a neutral background. Blurring these provides a rich 3D percept, not the impression of blurred lines, with the same multistability properties discussed previously. Importantly, it is the arrangement of the starting contours that matters—notice in particular that they are placed in threesomes with alternating contrast. After blurring, these replicate the isophote structure that must surround critical contours.
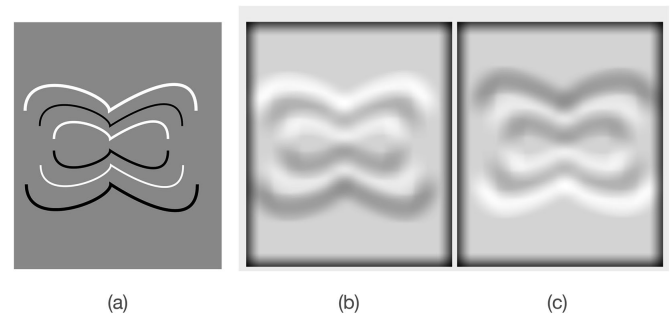

Figure 17. (a) A set of image contours in black and white on a gray background arranged as nearly parallel triples. (b) A simple blurring operation on these contours yields a rich 3D percept, not the flat appearance of blurred lines. Note in particular the appearance of ridges and bumps, how the curvature of the contours induces curvature in the ridges, and how the discontinuities are smoothed over. (c) When the original contours are inverted (black to white and vice versa), the percept is also inverted.

While these two demonstrations are only suggestive, we include them to (hopefully) stimulate others to examine these questions more rigorously. Modifications of the example in Figure 17 could include different orderings, arrangements, thicknesses, and contrasts of the original curves. These could provide further insight into the necessary structure of the sausage and of the perceptual basis for critical contours.

Conclusions

Our visual systems readily infer aspects of 3D shape across renderings, lightings, and contexts. Although the solutions seem veridical, the ill-posedness of the problem suggests this is unlikely. The standard approach, both in classical computer vision and in deep networks, is to introduce (directly or indirectly) some

form of regularization and to solve for parameters (e.g., surface normal) over the entire image. We have argued that, instead, a qualitative 3D surface representation is helpful. We formalized this in topological terms and were able to use it to link measurable 2D image features invariantly with surface properties. We summarized this with a slogan emphasizing how the image-based critical contour is constrained to lie within a sausage; this sausage, in turn, also contains extremal curves of slant. Fundamentally, the slogan works in the other direction as well.

In this article, we focused on bumps and dents, and advocated for the use of slant extremal curves as the candidate surface representation. To summarize the key points in our development, we argued:

1. *Bumps exist in the interior* of a shape. The occluding boundary delimits the full extent of an object. We showed that bumps have a description that is a relaxation of the occluding contour in a precise mathematical sense.
2. *Bumps are qualitative.* While bumps have a boundary, precisely where it lies is less clear. Evidence suggests that shape perception is qualitative: That is, while different subjects agree on certain basic properties of shape, they disagree on quantitative details. We encompass this qualitative aspect of shape inferences in a topological representation. The uncertainty in precision of the bump boundary was represented as a topological sausage around the critical contour.
3. *Bumps are global objects*; they are defined not at a point, such as the curvature at the peak, but over a neighborhood. Like a mountain, a bump is a collection of material that builds to a peak; they can be climbed from many sides. We characterized the global nature of bumps with the Morse-Smale complex, a global topological descriptor for surfaces. The M-S complex was attractive because the level sets in climbing a bump are nested and increasing.
4. *Bumps have both image and scene signatures.* To be perceivable, there must be some image signature to the bump. To define this, we built on previous theoretical ideas of critical contours. Bumps were defined using extremal curves of slant. We prove invariance to many aspects of lighting and material changes.
5. *Bumps are distinct parts* of a shape. As such, they are bounded from one another and should be manipulable separately. We introduced several visual illusions that illustrate this multistability.
6. *Bumps have consistent normals* The M-S complex also indicates how surface normals fit together, providing a constraint system that operates on the surface normal and can be extended to the occluding contour.

Image structure is highly dependent on the object material and rendering function. For example, a simple change in illumination will cause the vast majority of image pixels to change. In addition, rendering function changes (e.g., Lambertian vs. specular) can lead to unpredictable pixel changes in most of the image. Instead, by focusing on extremal slant curves, we ignore many of the "nuisance" positions in the image and many of the "nuisance" parameters in different rendering functions. By doing so, the computation becomes more straightforward and biological: We showed that there is a direct route from local image orientations to global surface parts without having to solve a particular differential equation, for example.

Much work remains to be done. First, the theory of critical contours and extremal curves leads to psychophysical tests; we mention the predictions inherent in the multistable displays and the extended ridges. Initial demonstrations show that those regions away from extremal curves do not carry much 3D information (Figures 16, 17 and Kunsberg et al., 2018). Second, computational issues remain. The image extremal contours may be incomplete due to noise and occlusion; some type of contour completion may be needed. Also, as shown in Figure 14, a labeling algorithm may be developed to compute the "side" to which the normal field points. Finally, the transition from the critical contour "scaffold" to a full surface realization needs to be studied. Initial experiments on reconstruction of the entire slant field from its extremal curves are promising (Kunsberg & Zucker, 2018), and related work exists in the computational literature (e.g., Allemand-Giorgis et al., 2014; Weinkauf et al., 2010). Whether this is analogous to the filling-in processes around Kanizsa figures, neon color spreading, and so on (van Lier et al., 2009; Pessoa & De Weerd, 2003; Bressan et al., 1997) remains an enticing possibility.

In summary, we have shown a path toward 3D shape inference based on stable and simply expressed intermediate image features. It captures the qualitative nature of shape perception while guaranteeing a degree of constancy even when the rendering function is unknown.

Acknowledgments

Corresponding author: Steven W. Zucker.
Email: steven.zucker@yale.edu.
Address: Computer Science, Biomedical Engineering, Yale University, New Haven, CT, USA.

Footnotes

[1] For brevity, we will ignore the convex/concave ambiguity until later in the article.

[2] Isophotes yield a smooth flow—their tangent map—when sampled by visual cortex (Ben-Shahar & Zucker, 2004; Kunsberg et al., 2018), so they are appropriate as a model of how image structure is represented in visual cortex.

[3] Note: The sausage in Figure 2d is the mathematical abstraction of the sausage in Figure 1d.

[4] As a tangential aside, but to build intuition, recall the "hairy ball theorem." Suppose that there is a closed surface covered in hair, just as the gradient covers our function. Now, having a smooth gradient field on a ball amounts to being able to "comb the hair" on the ball all over. If the ball is one-dimensional (i.e., a circle), then this is easy. However if the ball is 2D, this is impossible. Consider how the North and South poles are singularities of the lines of longitude and latitude. That is, one-dimensional (1D) spheres can support smooth vector fields, and 2D spheres must have critical points. In general, these results extend to higher (even and odd) dimensions, showing how critical points inform what we can deduce about surface topology.

[5] Shortly we shall instantiate this scalar field in two ways, either as an image, or brightness at a point, or as the slant function, or the value of slant at a point.

[6] It is perhaps worth emphasizing how the slightly undulating surface is generic. If one were to add a bit more undulation, or to move the undulations around, the topological description would not change. Adding a tiny undulation to a plane, however, changes it significantly.

[7] For an explanation of these terms, see Holtmann-Rice et al. (2018).

[8] An integral curve through a vector field is a curve whose tangent agrees with the vector at each point.

[9] http://www.math.stonybrook.edu/Videos/IMS/Differential_Topology/.

References

Achtman, R. L., Hess, R. F., & Wang, Y.-Z. (2003). Sensitivity for global shape detection. *Journal of Vision, 3*(10), 4, https://doi.org/10.1167/3.10.4.

Allemand-Giorgis, L., Bonneau, G.-P., & Hahmann, S. (2014). Piecewise polynomial re-construction of functions from simplified Morse-Smale complex. *IEEE Visualization Conference 2014, SciVis Posters, Nov 2014*, Paris, France. IEEE.

Banchoff, T. F. (1990). *Beyond the third dimension*. New York, NY: Scientific American Library.

Barron, J. T., & Malik, J. (2012). Shape, albedo, and illumination from a single image of an unknown object. In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, pp. 334–341.

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

Bernhard, M., Waldner, M., Plank, P., Soltészová, V., & Viola, I. (2016). The accuracy of gauge-figure tasks in monoscopic and stereo displays. *IEEE Computer Graphics and Applications, 36*(4), 56–66.

Biasotti, S., De Floriani, L., Falcidieno, B., Frosini, P., Giorgi, D., Landi, C., Papaleo, L., . . . Spagnuolo, M. (2008). Describing shapes by geometrical-topological properties of real functions. *ACM Computing Surveys, 40*(4), 12:1–12:87.

Bigün, J., Granlund, G. H., & Wiklund, J. (1991). Multidimensional orientation estimation with applications to texture analysis and optical flow. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (8), 775–790.

Blanz, V., & Vetter, T. (1999). A morphable model for the synthesis of 3D faces. In *SIGGRAPH '99: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 187–194, https://doi.org/10.1145/311535.311556.

Bott, R. (1980). Marston morse and his mathematical works. *Bulletin of the American Mathematical Society, 3*(3), 907–950.

Bott, R. (1988). Morse theory indomitable. *Publications Mathématiques de l'IHÉS, 68*, 99–114.

Bressan, P., Mingolla, E., Spillmann, L., & Watanabe, T. (1997). Neon color spreading: A review. *Perception, 26*(11), 1353–1366.

Breton, P., & Zucker, S. (1996). Shadows and shading flow fields. *IEEE Conference on Computer Vision and Pattern Recognition*, 782–789.

Breuß, M., Bruckstein, A., Maragos, P., & Wuhrer, S. (2021), *Perspectives in shape analysis*. New York: Springer.

Carlsson, G. (2009). Topology and data. *Bulletin of the American Mathematical Society, 46*(2), 255–308.

Cholewiak, S. A., Kunsberg, B., Zucker, S., & Fleming, R. W. (2014). Predicting 3D shape perception from shading and texture flows. *Journal of Vision, 14*(10), 1113.

Christou, C. G., & Koenderink, J. J. (1997). Light source dependence in shape from shading. *Vision Research, 37*(11), 1441–1449.

Curran, W., & Johnston, A. (1996). The effect of illuminant position on perceived curvature. *Vision Research, 36*(10), 1399–1410.

Dakin, S., & Bex, P. (2002). Summation of concentric orientation structure: Seeing the glass or the window? *Vision Research, 42*(16), 2013–2020.

DeCarlo, D., Finkelstein, A., Rusinkiewicz, S., & Santella, A. (2003). Suggestive contours for conveying shape. *ACM Transactions on Graphics (Proc. SIGGRAPH), 22*(3), 848–855.

Dobbins, A., Zucker, S. W., & Cynader, M. S. (1987). Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature, 329*(6138), 438–441.

Dumoulin, S. O., & Hess, R. F. (2007). Cortical specialization for concentric shape processing. *Vision Research, 47*(12), 1608–1613.

Edelsbrunner, H., & Harer, J. (2010). *Computational topology: An introduction*. American Mathematical Soc: New York, Springer, 2016.

Egan, E. J. L., & Todd, J. T. (2015). The effects of smooth occlusions and directions of illumination on the visual perception of 3-D shape from shading. *Journal of Vision, 15*(2), 24.

Eigen, D., Puhrsch, C., & Fergus, R. (2014). Depth map prediction from a single image using a multi-scale deep network. arXiv preprint arXiv:1406.2283.

Elder, J., & Zucker, S. (1993). The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research, 33*(7), 981–991.

Faisman, A., & Langer, M. S. (2013). Qualitative shape from shading, highlights, and mirror reflections. *Journal of Vision, 13*(5), 10, https://doi.org/10.1167/13.5.10.

Fleming, R. W., Holtmann-Rice, D., & Bülthoff, H. H. (2011). Estimation of 3D shape from image orientations. *Proceedings of the National Academy of Sciences, 108*(51), 20438–20443.

Fleming, R. W., Torralba, A., & Adelson, E. H. (2004). Specular reflections and the perception of shape. *Journal of Vision, 4*(9), 10.

Forman, R. (2002). How many equilibria are there? An introduction to Morse theory. In *Calculus of Variations: The Rice Undergraduate Colloquium*. American Mathematical Society.

Freeman, W. T. (1994). The generic viewpoint assumption in a framework for visual perception. *Nature, 368*(6471), 542–545.

Gallant, J. L., Shoup, R. E., & Mazer, J. A. (2000). A human extrastriate area functionally homologous to macaque v4. *Neuron, 27*(2), 227–235.

Georgeson, M. A., May, K. A., Freeman, T. C., & Hesse, G. S. (2007). From filters to features: Scale–space analysis of edge and blur coding in human vision. *Journal of Vision, 7*(13), 7.

Ghose, T., & Peterson, M. A. (2021). Task set and instructions influence the weight of figural priors: A psychophysical study with extremal edges and familiar configuration. *Attention, Perception, & Psychophysics, 83*(6), 2709–2727, doi:10.3758/s13414-021-02282-5.

Griffin, L. D., & Colchester, A. C. (1995). Superficial and deep structure in linear diffusion scale space: Isophotes, critical points and separatrices. *Image and Vision Computing, 13*(7), 543–557.

Günther, D., Reininghaus, J., Seidel, H.-P., & Weinkauf, T. (2014). Notes on the simplification of the morse-smale complex. In *Topological methods in data analysis and visualization III*. Springer, pp. 135–150.

Gyulassy, A. G. (2008). *Combinatorial construction of Morse-Smale complexes for data analysis and visualization*. PhD thesis, University of California, Davis.

Holtmann-Rice, D. N., Kunsberg, B. S., & Zucker, S. W. (2018). Tensors, differential geometry and statistical shading analysis. *Journal of Mathematical Imaging and Vision, 60*(6), 968–992.

Horn, B., & Brooks, M. (1989). *Shape from shading*. Cambridge, MA: The MIT Press.

Huggins, P. S., Chen, H. F., Belhumeur, P. N., & Zucker, S. W. (2001). Finding folds: On the appearance and identification of occlusion. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, pp. II-718–II-725.

Huggins, P. S., & Zucker, S. W. (2001). Folds and cuts: how shading flows into edges. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001* (Vol. 2, pp. 153–158). IEEE.

Hutson, M. (2018). AI researchers allege that machine learning is alchemy. *Science, 360*(6388), 861.

Judd, T., Durand, F., & Adelson, E. H. (2007). Apparent ridges for line drawing. *ACM Transactions on Graphics, 26*(3), 19.

Khang, B.-G., Koenderink, J. J., & Kappers, A. M. L. (2007). Shape from shading from images rendered with various surface types and light fields. *Perception, 36*(8), 1191–1213.

Koenderink, J., van Doorn, A., & Wagemans, J. (2015). Part and whole in pictorial relief. *i-Perception, 6*(6), 2041669515615713.

Koenderink, J., van Doorn, A., & Wagemans, J. (2015). Part and whole in pictorial relief. *i-Perception, 6*(6).

Koenderink, J. J. (1984). What does the occluding contour tell us about solid shape? *Perception, 13*(3), 321–330.

Koenderink, J. J. (1990). *Solid shape*. Cambridge, MA: The MIT Press.

Koenderink, J. J., & van Doorn, A. J. (1980). Photometric invariants related to solid shape. *Journal of Modern Optics, 27*(7), 981–996.

Koenderink, J. J., van Doorn, A. J., Christou, C., & Lappin, J. S. (1996). Perturbation study of shading in pictures. *Perception, 25*(9), 1009–1026.

Koenderink, J. J., Van Doorn, A. J., & Pont, S. C. (2004). Light direction from shad (ow) ed random Gaussian surfaces. *Perception, 33*(12), 1405–1420.

Kovacs, I., & Julesz, B. (1993). A closed curve is much more than an incomplete one: Effect of closure in figure-ground segmentation. *Proceedings of the National Academy of Sciences, 90*(16), 7495–7497.

Kulkarni, T. D., Whitney, W., Kohli, P., & Tenenbaum, J. B. (2015). Deep convolutional inverse graphics network. *arXiv preprint arXiv:1503.03167*.

Kunsberg, B., Holtmann-Rice, D., Alexander, E., Cholewiak, S., Fleming, R., & Zucker, S. W. (2018). Colour, contours, shading and shape: flow interactions reveal anchor neighbourhoods. *Interface Focus, 8*(4), 20180019.

Kunsberg, B., & Zucker, S. W. (2018a). Models of 3D shape perception. D. Jaeger, & R. Jung (Eds.), *Encyclopaedia of computational neuroscience*. Springer, https://doi.org/10.1007/978-1-4614-7320-6_100661-110.1007/978-1-4614-7320-6_100661-1.

Kunsberg, B., & Zucker, S. W. (2018). Critical contours: An invariant linking image flow with salient surface organization. *SIAM Journal on Imaging Sciences, 11*(3), 1849–1877.

Lawlor, M., Holtmann-Rice, D., Huggins, P., Ben-Shahar, O., & Zucker, S. W. (2009). Boundaries, shading, and border ownership: A cusp at their interaction. *Journal of Physiology–Paris, 103*(12), 18–36.

Lehky, S. R., & Sejnowski, T. J. (1988). Network model of shape-from-shading: Neural function arises from both receptive and projective fields. *Nature, 333*(6172), 452–454.

Li, A., & Zaidi, Q. (2000). Perception of three-dimensional shape from texture is based on patterns of oriented energy. *Vision Research, 40*(2), 217–242.

Lindsay, G. W. (2021). Convolutional neural networks as a model of the visual system: Past, present, and future. *Journal of Cognitive Neuroscience, 33*(10), 2017–2031.

Mach, E. (1965). On the physiological effect of spatially distributed light stimuli. In F. Ratliff (Ed.), *Mach bands: Quantitative studies on neural networks. Retina.* San Francisco, CA: Holden-Day.

Mamassian, P., & Kersten, D. (1996). Illumination, shading and the perception of local orientation. *Vision Research, 36*(15), 2351–2367.

Marlow, P. J., Mooney, S. W., & Anderson, B. L. (2019). Photogeometric cues to perceived surface shading. *Current Biology, 29*(2), 306–311.

Marlow, P. J., Todorović, D., & Anderson, B. L. (2015). Coupled computations of three-dimensional shape and material. *Current Biology, 25*(6), R221–R222.

Marr, D. (1982). Vision: A computational investigation into the human representation and processing of visual information. MIT Press.

Matsumoto, Y. (2002). *An introduction to Morse theory* (Vol. 208). American Mathematical Society.

Maxwell, J. C. (1870). On hills and dales. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 40*(269), 421–427.

Milnor, J. (1963). *Morse theory, annals of mathematics studies*. Princeton, NJ: Princeton University Press.

Milnor, J. (1965). *Topology from the differentiable viewpoint*, Univ. Press of Virginia, Charlottesville, 1990.

Mingolla, E., & Todd, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics, 53*(3), 137–151.

Mooney, S. W., & Anderson, B. L. (2014). Specular image structure modulates the perception of three-dimensional shape. *Current Biology, 24*(22), 2737–2742.

Mooney, S. W., Marlow, P. J., & Anderson, B. L. (2019). The perception and misperception of optical defocus, shading, and shape. *Elife, 8*, e48214.

Morgan, J., Tian, G., & Flow, R. (2007). *The Poincaré Conjecture, Clay Mathematics Monographs*, vol. 3. American Math. Society, Providence, 521.

Morgan, M. J. (2011). Features and the 'primal sketch.' *Vision Research, 51*(7), 738–753.

Morgenstern, Y., Geisler, W. S., & Murray, R. F. (2014). Human vision is attuned to the diffuseness of natural light. *Journal of Vision, 14*(9), 15.

Morgenstern, Y., Murray, R. F., & Harris, L. R. (2011). The human visual system's assumption that light comes from above is weak. *Proceedings of the National Academy of Sciences, 108*(30), 12551–12553.

Morse, M. (2007). Topology and equilibria. *The American Mathematical Monthly, 114*(9), 819–834.

Nackman, L. R. (1984). Two-dimensional critical point configuration graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-6*(4), 442–450.

Nartker, M., Todd, J., & Petrov, A. (2017). Distortions of apparent 3D shape from shading caused by changes in the direction of illumination. *Journal of Vision, 17*(10), 324–324.

Nefs, H. T., Koenderink, J. J., & Kappers, A. M. (2006). Shape-from-shading for matte and glossy objects. *Acta Psychologica, 121*(3), 297–316.

Nishida, S., & Shinya, M. (1998). Use of image-based information in judgments of surface-reflectance properties. *JOSA A, 15*(12), 2951–2965.

Ohara, M., Kim, J., & Koida, K. (2020). The effect of material properties on the perceived shape of three-dimensional objects. *i-Perception, 11*(6), 2041669520982317.

O'Neill, B. (2006). *Elementary differential geometry* (Rev. 2nd ed.). Burlington, MA: Elsevier.

Palmer, S. E., & Ghose, T. (2008). Extremal edge: A powerful cue to depth perception and figure-ground organization. *Psychological Science, 19*(1), 77–83.

Parhi, R., & Nowak, R. D. (2021). What kinds of functions do deep neural networks learn? Insights from variational spline theory. *arXiv preprint arXiv:2105.03361*.

Pessoa, L., & De Weerd, P. (Eds.). (2003). *Filling-in: From perceptual completion to cortical reorganization*. Oxford University Press.

Ramachandran, V. S. (1988a). Perceiving shape from shading. *Scientific American, 259*(2), 76–83.

Ramachandran, V. S. (1988b). Perception of shape from shading. *Nature, 331*(6152), 163–166.

Rao, A. R., & Schunck, B. G. (1991). Computing oriented texture fields. *CVGIP: Graphical Models and Image Processing, 53*(2), 157–185.

Reininghaus, J., & Hotz, I. (2011). Combinatorial 2d vector field topology extraction and simplification. In *Topological Methods in Data Analysis and Visualization* (pp. 103–114). Springer, Berlin, Heidelberg.

Sahner, J., Weber, B., Prohaska, S., & Lamecker, H. (2008). Extraction of feature lines on surface meshes based on discrete Morse theory. In *Computer Graphics Forum* (Vol. 27, pp. 735–742). Oxford, UK: Blackwell Publishing Ltd.

Saxena, A., Chung, S. H., & Ng, A. Y. (2005). Learning depth from single monocular images. In *NIPS* (Vol. 18, pp. 1–8).

Sengupta, S., Kanazawa, A., Castillo, C. D., & Jacobs, D. W. (2018). Sfsnet: Learning shape, reflectance and illuminance of facesin the wild'. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 6296–6305).

Seyama, J., & Sato, T. (1998). Shape from shading: Estimation of reflectance map. *Vision Research, 38*(23), 3805–3815.

Sitzmann, V., Zollhöfer, M., & Wetzstein, G. (2019). Scene representation networks: Continuous 3d-structure-aware neural scene representations. *arXiv preprint arXiv:1906.01618*.

Smale, S. (1961). On gradient dynamical systems. *Annals of Mathematics, 74*(1), 199–206.

Smale, S. (2007). Generalized Poincaré's conjecture in dimensions greater than four. In *Topological library: Part 1: Cobordisms and their applications* (pp. 251–268).

Spivak, M. (2018). *Calculus on manifolds: A modern approach to classical theorems of advanced calculus*. Reading MA: Addison Wesley, Publiishing Company.

Storrs, K. R., & Fleming, R. W. (2021). Learning about the world by learning about images. *Current Directions in Psychological Science, 30*(2), 120–128.

Subedar, M. M., & Karam, L. J. (2016). 3D blur discrimination. *ACM Transactions on Applied Perception (TAP), 13*(3), 1–13.

Sun, P., & Schofield, A. J. (2012). Two operational modes in the perception of shape from shading revealed by the effects of edge information in slant settings. *Journal of Vision, 12*(1), 12.

Tang, Y., Salakhutdinov, R., & Hinton, G. (2012). Deep lambertian networks. *arXiv preprint arXiv:1206.6445*.

Todd, J. T., Egan, E. J. L., & Phillips, F. (2014). Is the perception of 3D shape from shading based on assumed reflectance and illumination? *i-Perception, 5*(6), 497–514.

van Lier, R., Vergeer, M., & Anstis, S. (2009). Filling-in afterimage colors between the lines. *Current Biology, 19*(8), R323–R324.

Vangorp, P., Laurijssen, J., & Dutré, P. (2007). The influence of shape on the perception of material reflectance. In *ACM Transactions on Graphics, 26*(3):77, 9, http://doi.acm.org/10.1145/1239451.1239528.

Wallace, A. H. (2006). Differential topology; first steps. *Benjamin/Cummings series in the life sciences mathematics monograph series* (p. 130). W. A. Benjamin.

Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., & Jiang, Y. G. (2018). Pixel2mesh: Generating 3D mesh models from single RGB images. In *Proceedings of the European Conference on Computer Vision (ECCV)* (pp. 52–67).

Wang, R., Geraghty, D., Matzen, K., Szeliski, R., & Frahm, J. M. (2020). Vplnet: Deep single view normal estimation with vanishing points and lines. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 689–698).

Watt, R., & Morgan, M. (1983). The recognition and representation of edge blur: Evidence for spatial primitives in human vision. *Vision Research, 23*(12), 1465–1477.

Weinkauf, T., Gingold, Y., & Sorkine, O. (2010). Topology-based smoothing of 2D scalar fields with C1-continuity. In *Computer Graphics Forum* (Vol. 29, pp. 1221–1230). Oxford, UK: Blackwell Publishing Ltd.

Weinkauf, T., & Günther, D. Separatrix persistence: Extraction of salient edges on surfaces using topological methods. In *Computer Graphics Forum,* (Vol. 28, pp. 1519–1528). Oxford, UK: Blackwell Publishing Ltd.

Wilson, H. R., & Wilkinson, F. (1998). Detection of global structure in glass patterns: Implications for form vision. *Vision Research, 38*(19), 2933–2947.

Wu, J., Yildirim, I., Lim, J. J., Freeman, B., & Tenenbaum, J. (2015). Galileo: Perceiving physical object properties by integrating a physics engine with deep learning. In *Advances in Neural Information Processing Systems, 28*, 127–135.

Zhang, R., Tsai, P.-S., Cryer, J. E., & Shah, M. (1999). Shape-from-shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 21*(8), 690–706.

Zhou, H., Friedman, H. S., & Von Der Heydt, R. (2000). Coding of border ownership in monkey visual cortex. *Journal of Neuroscience, 20*(17), 6594–6611.

Zucker, S., & Kunsberg, B. (2019). From borders to bumps: Circular flows are invariant across materials. In *PERCEPTION* (Vol. 48, pp. 69–69). London, England: Sage Publications Ltd.

Appendix A: Maximal slant along the occluding contour

We show how the occluding contour relates image and scene domains and implicates isophotes. We begin with standard definitions (see, e.g., Marr, 1982).

The *rim* of an object in $\mathbb{R}^3$ is composed of all noninterior points where the view vector "glances" the object, that is, where the view vector lies in the tangent plane to the surface. The *occluding contour* is defined as the projection onto the image of the rim of the object. A powerful (but often elusive) cue, it has been studied in Koenderink (1984, 1990) and Lawlor et al. (2009), among many others.

Two properties are key. First, the occluding contour directly informs the viewer of the local surface normal; since the view vector lies in the tangent plane, it has
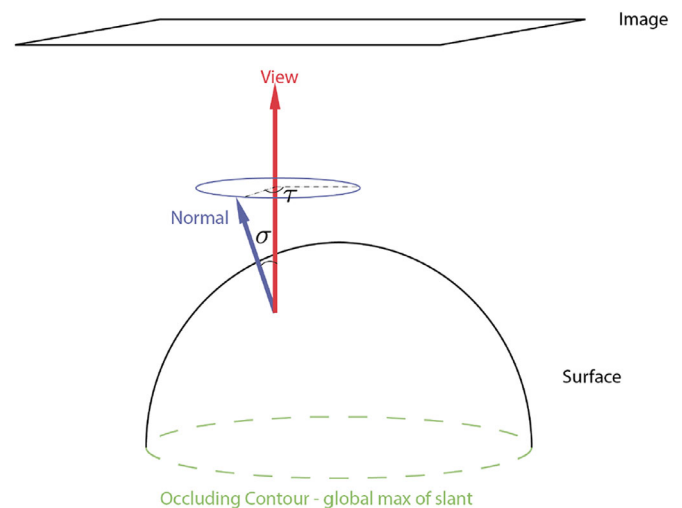

Figure 18. The orientation of a normal vector to the surface at a point can be specified in slant σ and tilt τ coordinates. Since there is a normal at every point, these coordinates define a scalar function over the object. Note how the occluding contour is given by maximal values of the slant.

surface meaning. Second, the occluding contour has a consistent flow "signature," so it also has *image salience*. We develop these in turn, starting with a standard representation for surfaces.

The *slant* $\sigma(x, y)$ is the polar angle between the surface normal and the view direction. The *tilt* $\tau(x, y)$ is the azimuthal angle between the surface normal and the view direction; see Figure 18. Both can be considered scalar functions on the image domain. Of course, these functions are unknown when the surface is unknown.

The surface normal and the view vector are perpendicular at every point on the rim, so the slant on the occluding contour is $\pi/2$. Thus, the occluding contour directly informs the viewer of the slant. This is a rare and unique property for a contour identifiable from the image. From the geometry, we get the important

Remark. The slant achieves a global maximum on the occluding contour, since the slant of the visible surface must always be bounded by $\pi/2$.

Appendix B: An invitation to differential topology

In this article, we exploit techniques and constructions from the field of differential topology, which tends to be unfamiliar to many students of perception. In this appendix, we provide a gentle introduction to some of the ideas that we shall use. The

discussion is informal, intended to give the flavor rather than the formal content, and at least some intuition about ideas and some technical terms. With this basis, we hope the reader will then be able to approach our use of the Morse-Smale complex. This invitation was influenced by the presentations in Milnor (1997) and Forman (2004); the primary reference is, as always, the classic lectures in Milnor (1963); see also Wallace (2006). Two inspiring videos are available from the 1960s, one from Morse (2007) and Milnor's 1965 Earle Raymond Hedrick Lectures.[9] It goes without saying that Morse's actual contributions (Bott, 1980) and their implications (Bott, 1988) are far deeper and wide-reaching than the brief introduction here.

Topology is normally thought of as "rubber sheet" geometry, or the study of which shapes can be made to coincide by stretching and twisting, but not tearing. By this logic, we have the famous example of a coffee cup and a donut being topologically equivalent, because they both have one hole. Stated a little more formally, topology studies which shapes are equivalent under certain types of deformation maps. Extending the coffee cup example, we see that topological equivalence can be characterized by how many handles there are and how they are attached.

Although we focus on smooth shapes here, topology also studies combinatorial ones. They all come together under the idea of constructing complex shapes by joining simpler ones, either for smooth shapes or piecewise linear ones (say, triangles).

We illustrate how this works as follows. A circle drawn on a piece of paper is a collection of tiny drops of lead—think of each as a point. Sets of these points indicate which ones belong together. The collections of these points into sets comprise neighborhoods; some points are neighbors, or neighbors of neighbors, and some are not. Removing a single point from this set of points yields two sets. More familiarly, one set consists of the point and the other an interval, or part of a line ($\mathbb{R}^1$); put back together, they comprise the sphere $S^1 \in \mathbb{R}^2$. We note that the interval is OPEN—loosely, a set that does not include its endpoints—written differently, the set of points indicated by t in $0 < t < 1$, as opposed to CLOSED, which would include its endpoints (the set t such that $0 \leq t \leq 1$). Taking a big leap into the combinatorial side of topology, such components are called CELLS. Their dimension matters: Points are 0-CELLS, lines are 1-CELLS, and so on. Cells comprise basic topological building blocks. The donut, or torus, can be decomposed as follows. First, remove a circle to leave a cylinder. The circle is a 0-cell and a 1-cell. Now, cut the cylinder along a line, which yields a 1-cell and a 2-cell. The result is a description of the torus as a CELL COMPLEX, an object constructed by attaching cells together in a specific manner. Notice that we have not specified the geometric details (e.g., the curvature of the circle or how the 1-cell was "bent").

This is the basic sense in which topology addresses shapes—as neighborhoods of points and how they fit together—and how it differs from (differential) geometry, which is concerned with curvature at a point.

More structure can be added by placing a coordinate system onto it. The sphere $\mathbb{S}^1$ is a MANIFOLD—a generalized curve or surface or higher-dimensional surface—without boundary. You could wander around $\mathbb{S}^1$ forever, feeling like you just walked along an interval of the line $\mathbb{R}$. Going up one dimension, pieces of the sphere $\mathbb{S}^2$ are equivalent to 2D Euclidean space and can be embedded in $\mathbb{R}^3$. Coordinates can be put on the sphere like latitude and longitude. Similarly for the torus. By "equivalent," here we stress one aspect of the coordinate systems: If you select a point p on the manifold, then the points around p on the manifold—a NEIGHBORHOOD of p—are like a neighborhood of points around the origin of a Euclidean space of the same dimension. Locally, then, a manifold is like a surface. The local coordinate patches connect with smooth distortions (diffeomorphisms) and, with a little more structure, one can do calculus on it (Spivak, 2018). This gives rise to the notion of a SMOOTH MANIFOLD.

DIFFERENTIAL TOPOLOGY studies (among other things) which smooth manifolds are equivalent—up to the cell complex—under smooth maps. In addition to supporting a differentiable structure, the manifolds should be COMPACT and BOUNDED, so that intervals contain limit points and cells are complete.

Two example situations are illustrative. First, consider a torus in general position under water, and slowly start to drain the water away. At first, the level just decreases, but then, suddenly, the "top" of the torus appears. Initially, it cuts the water level at a single point, but the cut—the LEVEL SET of points on the torus at the height of the water—grows around it. Soon the level set splits into two components, when the hole in the torus appears. The level sets join at the bottom of the hole and, finally, the surface of water is unbroken again below the torus. The level sets, then, are a series of planar curves cut across the torus; together, they convey information about the shape of the torus. Thinking of "height" as a function on the torus, we note that the "interesting" shape changes occur at the CRITICAL POINTS—places where the derivative of the height function is zero. One immediately identifies the top and the bottom of the torus with the maximum and the minimum of the height function. Between these critical points, the surface could vary somewhat geometrically, but the critical points would not change. For pictures of this, see Chapter 3 in Banchoff (1990).

The second situation takes us a deeper into the critical points and is inspired by Forman (2004). Now imagine the torus is magnetic, and place an iron marble at the top. The tiniest displacement will cause

the marble to roll away under gravity—and it could roll in any direction. That is, it can roll in any of the 2D dimensions on the cap of the torus. Placing the marble at the bottom, it just stays there; there are no dimensions for rolling. It cannot roll up the sides of the bowl. Finally, at the top and bottom of the central hole, there are critical points where there is one free dimension for rolling (downhill) and another dimension along which the marble cannot roll uphill. Thus, the critical points have an INDEX: the degrees of freedom for the marble to roll. While this is informally related to the dynamics of marbles under gravity, technically, the index is the number of negative eigenvalues of the Hessian (matrix of second derivatives).

MORSE THEORY is the study of how the critical points of a function on a smooth manifold inform about the topology of the manifold. As we hope to have shown in the above example using the height, if the function is chosen properly, then critical points inform us about a lot! Critical points are NONDEGENERATE when they are isolated and well behaved (i.e., when the matrix of second derivatives is full rank). Quoting Forman (2004), the main theorem of Morse theory can be stated:

Theorem. *Let M be a closed, compact, smooth submanifold of Euclidean space (of any dimension). Let $E : M \to \mathbb{R}$ be a smooth, real-valued function on M. Suppose that every critical point of E is nondegenerate. Then M can be built from a finite collection of cells, with exactly one cell of dimension i for each critical point of index i.*

With this theorem in mind, let us revisit the sphere. There are two basic singularities of functions on it, such as the height function, one at the top and another at the bottom. This is the content of Reeb's theorem: Each singularity implies a disc—a neighborhood of structure around the singularity—and a topological model of the sphere as two discs "attached" (say) at the equator. That is, a sphere is a cup attached to a bowl; each is signaled by a critical point. These discs, abstract components of the shape, identify parts of the cell complex with parts of the shape, almost regardless of geometric details. Since we shall be looking at aspects of shape around 1-cells of the Morse-Smale complex, one final aspect is needed.

Morse placed smooth functions on manifolds; Smale differentiated them. That is, Smale emphasized dynamical systems within the Morse framework (Smale, 1961). Instead of working with critical points of the function, he worked with critical points of a GRADIENT FLOW—that is, evaluate the gradient of the function at every point on the manifold and think of this as a vector. Placing a vector at every point on the manifold defines a vector field on it such that a marble, if dropped at the any point, will roll in the direction of the gradient to a (very close) nearby point, and it will keep rolling until a critical point is encountered. The vectors define the instantaneous velocity.

Smale's motivation was in solving a fundamental problem in topology (the Poincaré conjecture) (Smale, 2007), so for him, the dynamical systems view was a tool. The story of solving the Poincaré conjecture is itself fascinating (Morgan et al., 2007), but it takes us well beyond what we shall need. The significance (for us) is that we can now talk about trajectories—INTEGRAL LINES—through the flow. These integral lines are generally curved paths through the flow, as would be traced out by the marble. Importantly, some of these trajectories form the 1-cells of the Morse-Smale complex, which is described in the text and the next appendix. As will be seen, it is understanding a neighborhood around 1-cells—what we call a sausage—that drives our main result.

Appendix C: Introduction to the Morse-Smale complex

The Morse-Smale complex is a qualitative representation emphasizing the different stable and unstable regions of a smooth scalar function on a manifold. In this work, we choose the function to be the slant function of the image surface $\sigma(x, y) : \mathbb{R}^2 \to \mathbb{R}$. We will assume σ is a *Morse function*: All its critical points are nondegenerate (meaning the Hessian at those points is nonsingular), and no two critical points have the same function value. Should the surface not be Morse, we can always perturb it slightly to obtain one.

For a smooth surface, the *gradient* $\nabla\sigma = (\partial f/\partial x, \partial f/\partial y)$ exists at every point. A point $p \in \mathbb{R}^2$ is called a *critical point* when $\nabla\sigma(p) = 0$. This gradient field gives a direction at every point in the image, except for the critical points, which are rare (a set of measure zero). Following the vector field will trace out an *integral line*. These integral lines must end at critical points, where the gradient direction is undefined. Thus, one can define an *origin* and *destination* critical point for each integral line.

The type of each critical point is defined by its *index*: The number of negative eigenvalues of the Hessian at that point. For scalar functions on $\mathbb{R}^2$, there are only three types: a maximum (with index 2), a minimum (with index 0), and a saddle point (with index 1).

There are two types of integral lines, depending on the difference in index of the critical points it connects. If the difference is 1, we call the integral line a *1-cell*. It naturally must connect a saddle with either a maximum or a minimum. For example, a *saddle-maximum 1-cell* connects a saddle and a maximum. The set of 1-cells will naturally segment the scalar field into different regions, called *2-cells*. In addition, the scalar values on the 1-cells govern the values on the 2-cells. See Figure 3 for illustration.

Further, for each critical point, its *ascending manifold* is defined as the union of integral lines having that

critical point as a common origin. Similarly, its *descending manifold* is the union of integral lines with that critical point as a common destination.

For two critical points p and q, with the index of p one greater than the index of q, consider the intersection of the descending manifold of p with the ascending manifold of q. This intersection will be either a 1D manifold (a curve called a 1-cell or watershed) or the empty set. For two critical points r and s, with the index of r two greater than the index of s, the intersection of the descending manifold of r with the ascending manifold of s will either be a 2D manifold (a region equivalent to a 2-cell) or the empty set. Thus, the intersection of all ascending manifolds with all descending manifolds partitions the manifold $\mathbb{M}$ into 2D regions surrounded by 1D curves with intersections at the critical points.

The Morse Smale complex is the combinatorial structure (and the corresponding attaching maps) defined by the critical points, 1-cells and 2-cells. It is a structure that relates a set of contours (the 1-cells) to a qualitative function representation. With knowledge only of the slant function at the critical points and 1-cells, one could reconstruct the 2-cells (and thus the entire function) relatively accurately. For some insight, see Allemand-Giorgis et al. (2014) and Weinkauf et al. (2010). In this work, we show how the slant saddle-maxima 1-cell can be used as a model of "bump boundaries" as is relevant to 3D shape perception.

For additional information, see Milnor (1963), Gyulassy (2008), Biasotti et al. (2008), and Matsumoto (2002).

Appendix D: The Gauss map in 1D

We use the Gauss map as an indication of how wildly a surface is varying. We now provide a brief
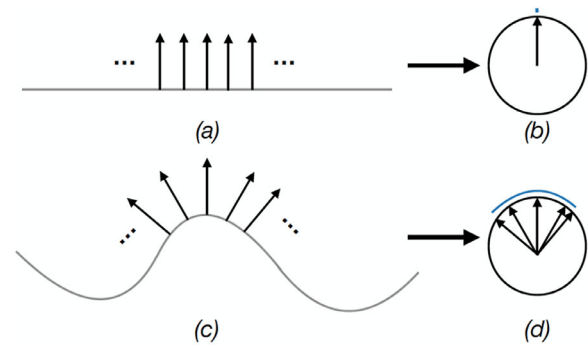

Figure 19. Illustration of the Gauss map, which takes the normals attached to a curve (or surface) (left) and maps them to the unit circle (sphere) (right). (a) A selection of normals attached to a straight line; (b) these map to a single point on the Gauss circle. (c) A selection of normals attached to a curve; (d) these map to a region (in blue) on the Gauss circle.

introduction to it (Figure 19). For a more serious introduction, see O'Neill (2006).

Gauss, working on the foundations of curvature in differential geometry, designed a map that takes the normals to a curve or a surface and maps them, collectively, to a circle or a sphere. Intuitively, the map is accomplished by moving each of the (unit) normals to a single point. Notice how, for the (straight) line, the normals then all overlap, while for the curve, they "spread out" somewhat unevenly. This spreading out can be used as the foundation for a definition of curvature: For a given length of curve, the normals spread out over a portion of the Gauss circle; in the limit as this length of curve approaches zero, the area on the Gauss circle also approaches a limit. The ratio of these two areas is the Gaussian curvature. Since this limit is taken around a point on the original curve, the curvature is a local descriptor. We exploit the measure of the normals over a region (of a surface) to get a global measure of variation.