SIAM J. IMAGING SCIENCES
Vol. 11, No. 1, pp. 293–324

Local Image Structure and Procrustes Metrics*

Jan Koenderink† and Andrea van Doorn‡

Abstract. The differential geometry of images (one of numerous applications but certainly an important one) involves singly isotropic space, rather than Euclidean space, the reason being that image intensity is not commensurate with the dimensions of the image plane. The Procrustes root mean square *nonuniformity measure* in such spaces immediately leads to a principled definition of curvedness, shape index, and orientation of local second-order structure. However, it is categorically different from the conventional *curvature-based* measures. One obtains a natural Euclidean metric of shape space that is readily extended to the cubic and quartic orders of approximation. In this setting, it is simple to derive the marginal probability densities for curvedness, shape index, and orientation for isotropic and anisotropic Gaussian random fields. For slight anisotropies, the marginals are much closer to those empirically found in natural images than the conventional formalisms suggest. The main articulation fits the edge and ridge structure imposed by the linear and quadric orders. The cubic and quartic structures contribute in a natural manner to these edge and ridge structures, whereas their higher-order saddle structures contribute little to the variance and can generally be ignored in applications.

Key words. curvature, Procrustes metric, invariants, local image structure

AMS subject classifications. 68U10, 68U05, 53B99

DOI. 10.1137/17M1136079

1. Introduction.

In image processing, local image structure is mainly of interest in the design and analysis of various feature detectors, the generic instance perhaps being "edge finders" [51]. In vision science, local image structure is of major interest because most of the processing in the front-end visual system is of a local nature. The generic elements are so-called receptive fields of various complexity. Among the simplest receptive field structures, one recognizes generic instances of edge finders. Indeed, from their early discovery, receptive fields have been identified as local feature detectors, the features being elementary geometrical configurations like edges or bars.

A principled formalism that accounts for a large fraction of the receptive fields in the visual front end is the theory of scale-space jets [33, 29, 30, 44, 14, 58, 45, 46]. Virtually all local processing identified in the primary visual cortex can be described that way. It is an attractive

*Received by the editors June 26, 2017; accepted for publication August 25, 2017; published electronically February 6, 2018.
http://www.siam.org/journals/siims/11-1/M113607.html

†Department of Psychology, Justus Liebig University Giessen, Giessen 35390, Germany; Department of Experimental Psychology, University of Leuven (KU Leuven), Leuven 3000, Belgium; and Department of Experimental Psychology, Helmholtz Institute, Utrecht University, 3584 CS Utrecht, The Netherlands (koenderinkjan@gmail.com, http://www.gestaltrevision.be/en/about-us/contact/all-contacts/45-jan-j-koenderink).
‡Department of Psychology, Justus Liebig University Giessen, Giessen 35390, Germany, and Department of Experimental Psychology, Helmholtz Institute, Utrecht University, 3584 CS Utrecht, The Netherlands (andrea.vandoorn@gmail.com).

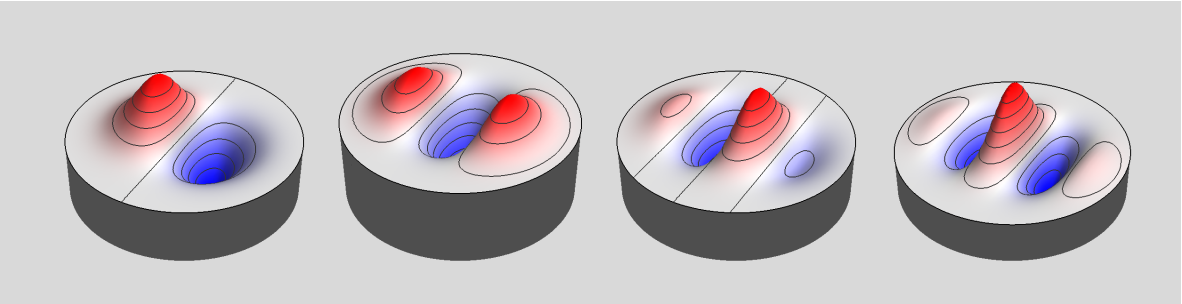

Figure 1. *The derivatives of images at a given scale are exactly calculated using derivatives of Gaussian filters. This allows stable computation of essentially exact partial derivatives of any order at any scale. The implementation in the human brain involves overcomplete bases of directional derivatives. One needs only one type of filter per order; a complete basis for order n involves $n+1$ independent directions* [29, 30]. *A "cortical pinwheel" implements essentially all orientations for the order 2 ("line finder"* [21]*) derivatives. (For the sake of clarity, the positive and negative lobes of the receptive fields have been color coded red and blue here, the zero level being represented as gray.)*

description because its image processing power is well proven by the numerous applications in, for instance, medical image processing. From a conceptual perspective, the description is attractive because it maps local image structure seamlessly on differential geometry. Thus, standard mathematical accounts immediately map on technical implementations and double as likely frameworks for parts of brain theory.

The fundamental structures of differential geometry are the local jets, which are complete sets of partial derivatives up to some maximum order. At any given order, there is a freedom of representation; for instance, one might implement the second order in terms of the partial derivatives $\{\partial_{xx}, \partial_{xy}, \partial_{yy}\}$ in Cartesian coordinates $\{x, y\}$. This is indeed the standard representation adopted in image processing. However, the brain uses a different implementation in that it maintains an overcomplete basis of directional derivatives (like ∂_{xx} for the x-direction) in *any* direction.

In the case of image processing, all expressions used in this paper can be implemented *exactly* by substituting convolution with Gaussian derivative filters for the Cartesian partial derivatives [33, 44, 14, 58] (see section 10). In the case of brain theory, one would rewrite all expressions in terms of unidirectional derivatives (see Figure 1).

The present discussion skips such issues because many apparently different implementations are equivalent under linear combination. It is assumed here that the visual cortex maintains (over-)complete representations of jets up and including order 4 at any point. This involves 14 DOF (two for the linear, three for the quadric, four for the cubic, and five for the quartic order), which is slightly less than a 4×4-pixel icon. Moreover, it is assumed that this involves a continuous range of spatial scales.[1] Thus, "local" can imply any *size*. The local processing is the same at all scales.

[1] In practice, discretely sampled, of course.

An image is a cross section of a fiber bundle $\mathbb{E}^2 \times \mathbb{I}$ (here $\mathbb{I}$ denotes the isotropic line), not $\mathbb{E}^3$, because pixel intensities[2] and image plane distances are mutually incommensurable, as is generically the case with *graphs*. Evidently, $\mathbb{E}^2 \times \mathbb{I}$ is quite distinct from $\mathbb{E}^3$; thus, the notion of "surface normal" loses its meaning, and the classical (Gaussian) differential geometry [15, 57] fails to apply.

A convenient formalism that allows one to retain much of the flavor of the classical differential geometry of surfaces is to treat images as *surfaces* (more properly "cross sections") in singly isotropic space $\mathbb{E}^2 \times \mathbb{I}$, where $\mathbb{I}$ is the isotropic line [39, 63]. Assuming Cartesian coordinates $\{x, y, z\}$, with $\{x, y\} \in \mathbb{E}^2$ and $z \in \mathbb{I}$, the degenerate metric is

$$\mathrm{d}s^2 = \mathrm{d}x^2 + \mathrm{d}y^2. \tag{1}$$

In the generic case that $x_1 \neq x_2$ or $y_1 \neq y_2$, the mutual distance d_{12} of two points is $\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$. In case $d_{12} = 0$ (and only in *that* case!), one says that their *special distance* δ_{12} is $z_1 - z_2$, where the latter is a signed quantity [39, 63]. The special distance remains undefined for points that are separated in the image plane. Points at zero distance that have nonzero special distance are said to be mutually parallel points in analogy to the case of planes [39, 63].

There exists a complete treatment of the differential geometry of curves and surfaces in graph spaces that is largely analogous to the classical differential geometry of surfaces [55]. Whereas it is not exceptional to see expressions taken from the latter applied to images regarded as surfaces, such a practice is actually meaningless because it implicitly mixes mutually incommensurable dimensions. Moreover, it complicates matters unnecessarily because the analogous expressions in singly isotropic space are invariably much simpler than those for Euclidean space [39, 63].

This paper focusses mainly on surface shape in the infinitesimal neighborhood of a point. Thus, it is natural to concentrate on truncated Taylor expansions

$$z(x, y) = \sum_{k=0}^{n} \left(\frac{1}{k!} \sum_{i=0}^{k} a_{k-i\,i} x^{k-i} y^i \right), \tag{2}$$

where it will rarely be advantageous to go beyond order $n = 4$ (see section 8). Orders 0 and 1 will generally be disregarded, the reason being that order 0 specifies location, whereas order 1 specifies spatial attitude, both in xyz-space. Either is immaterial to *shape*.

In fact, the group of similarities of singly isotropic space is[3] [37, 55]

$$\begin{aligned} x' &= \tau_x + \rho(x \cos\varphi - y \sin\varphi) \\ y' &= \tau_y + \rho(x \sin\varphi + y \cos\varphi) \\ z' &= \mu_z + \chi_x x + \chi_y y + \sigma z. \end{aligned} \tag{3}$$

[2] Notice the use of $\mathbb{I}$ instead of $\mathbb{R}^+$. It is assumed that the jets are applied to the *logarithm* of the intensity, which is the natural scale-invariant way to treat nonnegative quantities. This is implemented in the visual system. In image processing, one has a choice. However, one often is forced to use JPEG (etc.) instead of RAW images, in which cases in-camera compression usually implements a roughly logarithmic mapping.

[3] This is readily obtained from the usual Euclidean expressions, remembering that $\sin \varepsilon p = \varepsilon p$ and $\cos \varepsilon p = 1$ (generally, $f(p + \varepsilon q) = f(p) + \varepsilon f'(p)\, q$ [62]).

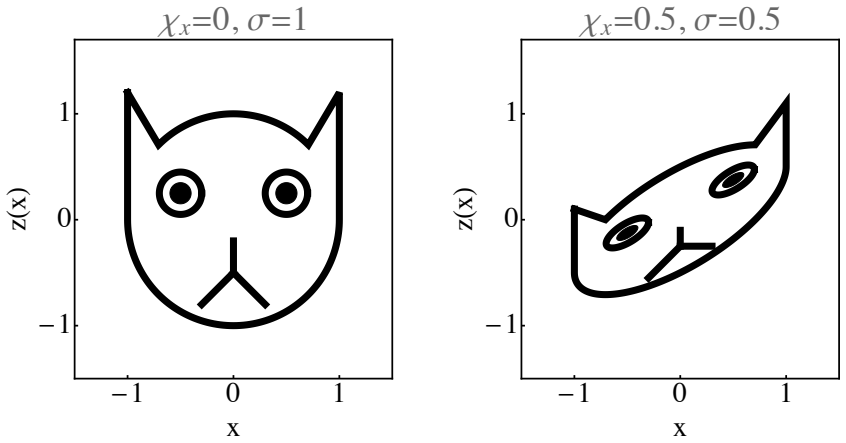

Figure 2. *The effect of an isotropic similarity in the $y = 0$ plane. Left: original configuration; right: after an isotropic rotation over an angle* 0.5 *(remember that the isotropic angle is not periodic; the range is $\mathbb{R}$) and a similarity "of the second kind"* [63] *("first kind" a Euclidean similarity in the image plane) of magnification* 0.5*. Notice that all movements are in the image intensity (z) direction. Thus, image intensity and distance in the image plane never "mix," which is as it should be, but has the consequence that the proper geometry is non-Euclidean.*

The vector $\boldsymbol{\tau}$ denotes a translation in the image plane, φ a rotation in the image plane, and ρ a similarity in the image plane. In this paper, these are ignored; thus, $x' = x$, $y' = y$. One is left with (3) (see Figure 2). The parameter μ parameterizes an isotropic shift; it will generally be ignored too.

The vector $\boldsymbol{\chi}$ parameterizes an *isotropic rotation*, whereas σ parameterizes an *isotropic scaling*. For $\sigma = 1$, one has the important "proper movements" or congruences that allow the formal definition of local *shape*.

It is entirely possible to develop a theory of shape in terms of *curvature*; indeed, Strubecker has done so [38, 37, 55]. However, here we explore an alternative (and unconventional) route to shape, based on the root mean square (RMS) deviation from planarity. This has the advantage of not requiring any smoothness assumptions; mere integrability suffices. Thus, the resulting theory also applies to fractal surfaces and so forth.

2. The RMS deviation from planarity. The primary objective is to set up a "shape space" for the second order. This could be done following the usual path as first traced by Gauss for the Euclidean case and specialized by Strubecker for the singly isotropic space. However, this approach fails to yield a *shape metric*; it only yields a *shape description*.

An alternative approach seeks to define a general measure of any type of deviation from planarity, the plane being considered "shapeless." One uses the variance of the deviations from, say, planarity. Such a measure could then easily be generalized to a full *shape metric*. Of course, this will not yield much of a shape *description*. However, by focusing on a single order (say, order 2), this is easily amended by drawing on the conventional approach. In essence, the variance of deviations from planarity is a generalization of the Procrustes distance as pursued by Kendall for the case of point configurations [26, 27, 5, 24, 16]. However, this generalization involves some novel twists because it depends upon the size of the region of interest (ROI).

Consider a (monochrome) image $z(x, y)$, assuming (for the moment and without loss of generality) $z(0,0) = z_x(0,0) = z_y(0,0) = 0$, where z is taken as log intensity. Notice that gamma transformations [23] become scalings, leaving the shape invariant. Evidently, the variance of a general quadric

$$z(x,y) = \frac{1}{2}\left(a_{20}x^2 + 2a_{11}xy + a_{02}y^2\right) \tag{4}$$

must *diverge* when taken over the full plane. However, the interest is in local regions of interest. Accordingly, the spatial weighting function (of unit integral weight)

$$w(x,y;s) = \frac{\mathrm{e}^{-\frac{x^2+y^2}{2s^2}}}{2\pi s^2} \tag{5}$$

is introduced, where s is a measure of the linear size of the ROI. The variance of the deviations from planarity is then defined as $\overline{z^2} - \overline{z}^2$, where

$$\overline{f}(s) = \langle f(x,y)\rangle = \int_{-\infty}^{+\infty} f(x,y)w(x,y;s)\,\mathrm{d}x\,\mathrm{d}y. \tag{6}$$

This variance is evidently a function of s, the size of the ROI (see below).

It is convenient to notice that (the integrations being elementary) $\langle x^{n-m}y^m\rangle = 0$ when n or m odd; otherwise,

$$\langle x^{n-m}y^m\rangle = \frac{2^{\frac{n}{2}}\Gamma(\frac{m+1}{2})^{\frac{n-m+1}{2}}}{\pi}s^n, \tag{7}$$

as this allows one to find the variance of arbitrary truncated Taylor expansions by simple algebraic combination. Up to order $n = 8$, the results gathered in Table 1 cover all cases discussed in this paper.[4]

Table 1

The values of $\langle x^{n-m}y^m\rangle$ for the orders up to 8 and $n \geq m$ and $s = 1$.

	$m=0$	2	4	6	8
$n=0$	1				
2	1	1			
4	3	1	3		
6	15	3	3	15	
8	105	15	9	15	105

Throughout this paper, the size s of the ROI is set to unity. However, it is important to take due notice of the s-dependence of the variances of the various orders. As the size of the ROI is increased, the importance of the higher orders increases accordingly. This is a well-known effect; e.g., if one is restrained to a limited region, the earth will appear flat, and if one increases the scope, the importance of the second-order terms may dominate the lower orders, and the curvature of the globe becomes noticeable.

[4]We consider Taylor expansions, including quartic terms; thus, because we need to consider variances, we have to reckon with orders as high as 8.

The Procrustes metric can also be used *incrementally*, as will often be advantageous in applications. For instance, the cubic order then yields the distance to the quadric order approximation and so forth. The RMS deviation is really a *nonuniformity* measure.

The variance for the *zeroth order* vanishes, as expected, because the zeroth-order approximation is the featureless plane.

Although the *first order* is not of importance because it can always be removed, it is still of some interest to see how the variance works out for this simple case. The variance for the first order is[5]

$$\Gamma_{(1)} = a_{10}^2 + a_{01}^2, \tag{8}$$

that is, the squared modulus of the gradient, a differential invariant. This makes obvious intuitive sense. The gradient is a vector $\{p, q\}$ (the reason for this notation will become clear in section 3) and thus transforms as

$$\begin{pmatrix} p' \\ q' \end{pmatrix} = \begin{pmatrix} \cos\varphi & \sin\varphi \\ -\sin\varphi & \cos\varphi \end{pmatrix} \begin{pmatrix} p \\ q \end{pmatrix}, \tag{9}$$

here written in a form that readily generalizes to the higher orders (see section 3).

In most cases involving shape, one simply annuls the contribution due to the gradient through an isotropic rotation. The case is only mentioned here because it illustrates that the Procrustes measure leads to the natural result that "gradient space" has the metric of $\mathbb{E}^2$. More generally, we use a congruence to ensure that $\langle f(x,y)\rangle$, $\langle f_x(x,y)\rangle$, and $\langle f_y(x,y)\rangle$ are all zero at $\{0,0\}$ for the given (assumed fixed) scale.

This defines a reference plane that comes in place of the conventional *tangent plane*. The reference plane is defined whenever the integral equation (6) exists. It naturally depends upon the scale. It is a "flat earth" reference.

The variance of the second order is

$$\Gamma_{(2)} = \frac{1}{2}(a_{20}^2 + 2a_{11}^2 + a_{02}^2), \tag{10}$$

often called "bending energy" [59], which can be shown to equal $(\kappa_1^2+\kappa_2^2)/2$, where $\kappa_{1,2}$ denote the principal curvatures [57]. Thus, the variance equals the so-called Casorati curvature [10] squared, a well-known differential invariant (for instance, see [20, 49]). However, the nature of the derivation is quite different from that conventionally associated with either "bending energy" or "Casorati curvature."

Felice Casorati (1835–1890) introduced this "curvature"[6] as a more commonsense curvature measure as compared to the Gaussian curvature or the mean curvature, his arguments being that it is nonnegative and only vanishes for planes [10]. He was severely chastised for this rather arbitrary choice, and his measure was soon forgotten. Koenderink and van Doorn [32]

[5]Notice the notation $\Gamma_{(1)}$. For the RMS deviation at *any* order Γ will be used, when necessary in order to avoid ambiguity a subscript (like in $\Gamma_{(1)}$) will be used.

[6]The quotes indicate that Casorati's measure is not a curvature in the conventional sense, that is, a rate of change of spatial attitude, like the rotation of the normal when moving in a certain direction in Gauss's classical treatment [15].

reintroduced it a century later (being unaware of Casorati's paper at that time) and suggested the term "curvedness," which is now in regular use. Casorati's intuition was evidently very close to the nature of the present derivation; he was primarily interested in the deviation from planarity.

Notice that (10) is formulated in a somewhat opaque (although conventional) form. Reason is that various important facts are not explicitly displayed. For instance, $a_{20} + a_{02}$ is a differential invariant (the Laplacian), and one would like to see that appear explicitly. A coefficient like a_{20} has no geometrical meaning, whereas the Laplacean $a_{20} + a_{02}$ has. We consider more appropriate formulations below.

The present explanation as the RMS deviation from planarity yields an independent and intuitive alternative to the differential geometry. It is essentially a Procrustes metric in Kendall's [26, 27] sense. Notice that the notion of reference plane (as different from "tangent plane") does not imply smoothness and likewise that the notion of nonplanarity (as different from "curvature") does not involve a change of spatial attitude with displacement. This is not at all "differential geometry" in the proper sense, although it—of course, intentionally—appears that way.

For the pure *third order*, one considers the general cubic

$$z(x,y) = \frac{1}{3!}\left(a_{30}x^3 + 3a_{21}x^2y + 3a_{12}xy^2 + a_{03}y^3\right). \tag{11}$$

This is a pure cubic, "pure" meaning that one uses a reference quadric instead of a reference plane. One finds that the cubic variance is

$$\Gamma_{(3)} = \frac{1}{12}(5(a_{30}^2 + a_{03}^2) + 6(a_{21}a_{03} + a_{12}a_{30}) + 9(a_{21}^2 + a_{12}^2)). \tag{12}$$

Likewise, for the pure (analogous definition) *fourth order*, that is,

$$z(x,y) = \frac{1}{4!}\left(a_{40}x^4 + 4a_{31}x^3y + 6a_{22}x^2y^2 + 4a_{13}xy^3 + a_{04}y^4\right), \tag{13}$$

one obtains the variance

$$\Gamma_{(4)} = \frac{1}{12}(2(a_{40}^2 + a_{04}^2) + 3a_{22}(a_{40} + a_{04}) + 5(a_{31}^2 + a_{13}^2) + 6(a_{22}^2 + a_{31}a_{13})). \tag{14}$$

These are evidently differential invariants with potential applicability. However, they do not seem to have found much use. In order to be able to make intuitive—that is, geometrical—sense of such expressions, one needs to put the various orders of the Taylor series in a form that exhibits their relation to changes of orientation in the picture plane in an explicit form. This is discussed in the next section.

When computing the variance of the full Taylor expansion ((2), that is, referenced to the plane), one finds that the terms in s^4 involve a certain mixture of the linear and cubic orders, namely,

$$\Gamma_{(1,3)} = a_{01}(a_{03} + a_{21}) + a_{10}(a_{12} + a_{30}), \tag{15}$$

where the notation should be self-explanatory. Notice that the mixed term can assume negative values.

Likewise, one finds that the terms in s^6 involve a certain mixture of the quadric and quartic orders, namely,

$$\Gamma_{(2,4)} = \frac{1}{2}\left(a_{20}(a_{40}+a_{22}) + 2a_{11}(a_{31}+a_{13}) + a_{02}(a_{04}+a_{22})\right). \tag{16}$$

The significance of this will become clear later in the discussion (section 8). It derives from correlations between the first and the third and between the second and the fourth orders. Notice that there is no such interaction between even and odd orders. Thus, $\Gamma_{(1,3)}$ can be understood as a third-order "correction" on the first order that is irrelevant for small regions because of going with s^4 instead of s^2. An analogous reasoning applies to $\Gamma_{(2,4)}$. This illustrates the importance of the size of the ROI.

3. Canonical forms for orders up to 4.

3.1. The zeroth order. The zeroth order plays no role because it can always be transformed away. However, it may be instructive to see what one would obtain in case one kept it.

The zeroth order is, of course, trivial. The variance vanishes as expected and does not depend upon the width s. After all, $z(x,y) = a_{00}$ is just the reference plane up to an isotropic translation.

3.2. The linear order. Since an isotropic rotation will annul the slope at a fiducial point, the variance of the linear order is usually of no interest. Nevertheless, it is still of some interest to mention this case in order to see how it reveals properties of the Procrustes metric that will appear in much greater complexity at the higher orders.

The space of linear bivariate polynomials is usually known as "gradient space." The variance is simply the squared distance to the origin $\Gamma^2 = p^2 + q^2$ (notice that we succinctly write Γ for $\Gamma_{(1)}$ and $\{p,q\}$ for $\{a_{10}, a_{01}\}$, a practice that we will continue for the higher orders), and the mutual distance of two instances is $\sqrt{(p_1-p_2)^2+(q_1-q_2)^2}$; thus, gradient space is a Euclidean plane in the Procrustes metric.

The unit variance locus is the unit circle $\mathbb{S}^1$. Each point on that locus can be moved to any other by a Euclidean rotation about the origin, so all points on the circle have the same shape: The shape space is degenerated into a point. The habitus is

$$z(x,y) = \Gamma(x\cos\varphi + y\sin\varphi), \tag{17}$$

where φ is a pure directional parameter, there Being no need for a shape descriptor. The factor Γ is only a size parameter, an amplitude—it does not affect shape. *The* canonical linear shape is simply $z(x,y) = x$.

A smooth relief maps into gradient space as a (locally) two-dimensional patch. The generic singularities are folds and cusps ([61, 1], Figure 3), the former images of the zero locus of the quadric discriminant (see below) and the latter involving the cubic terms.

3.3. The quadrics. A simple reordering of terms[7] in (4),

$$a_{20}x^2 + 2a_{11}xy + a_{02}y^2 = p\frac{x^2+y^2}{2} + \left(q\frac{x^2-y^2}{2} + r\,xy\right), \tag{18}$$

[7]Notice that the notation $p, q, r, s, \ldots$ is used for the canonical coefficients of all orders. This will not lead to confusion in this paper since the analysis is done per order. However, the reader should beware.

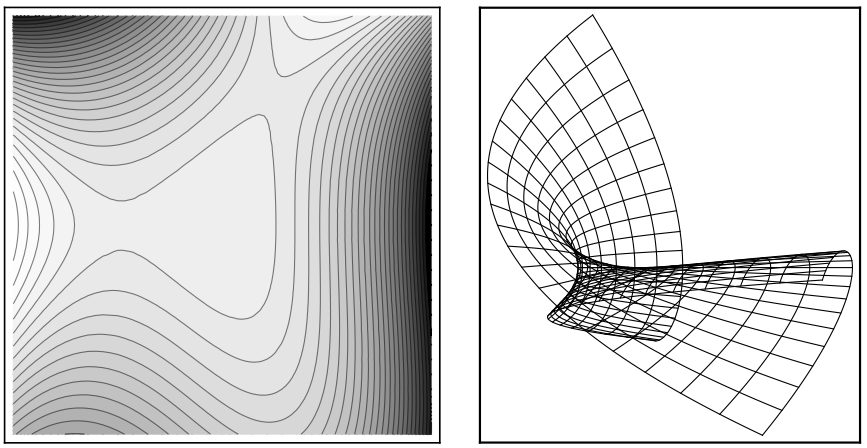

Figure 3. *Left: an image (or relief); right: its map in gradient space. Notice the folds and cusps. These are the generic singularities. This is a simple example: Typical images will give rise to hundreds or thousands of cusps in some particular network. Some statistical measure then often makes sense.*

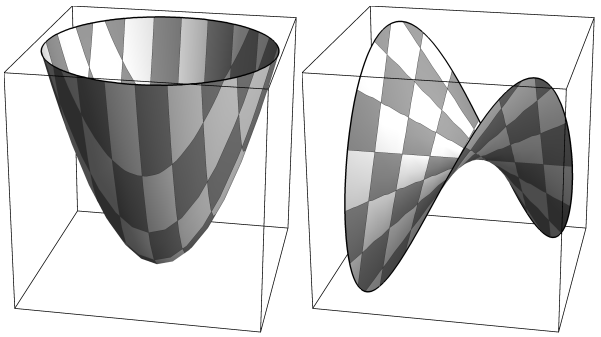

Figure 4. *The quadrics are linear combinations of $(x^2 + y^2)/2$ (left) and $(x^2 - y^2)/2$ or xy (right). The former is rotationally invariant, and the latter repeats after a half-turn.*

setting $p = a_{20} + a_{02}$, $q = a_{20} - a_{02}$, and $r = 2a_{11}$, yields immediate insight into the structure of the second order as it presents the form as the sum of an isotropic part and an anisotropic part that repeats after a half-turn (see Figure 4). Apparently, the parameters $\{q, r\}$ have no individual relevance but transform as a pair; thus, it is meaningful to write[8]

$$\{q, r\} = \Lambda\{\cos 2\Phi, \sin 2\Phi\} \tag{19}$$

with $\Lambda^2 = q^2 + r^2$. The parameter Φ is just the orientation of the largest principal curvature, whereas the shape and Casorati curvature (an amplitude, or size, parameter) will depend upon the pair $\{p, \Lambda\}$. Indeed, a short calculation reveals that the squared Casorati curvature is $\Gamma^2 = p^2 + \Lambda^2 = p^2 + q^2 + r^2$, whereas the shape index (Figure 5) is $\Psi = \arctan(p/\Lambda)$. One has

[8]Notice the use of parameters Γ, Λ, Ψ, ..., a notation that will be used for *all* orders. This should not lead to ambiguity since the discussion is per order. Where necessary, one may use a subscript indicating the order, like in $\Lambda_{(2)}$ and so forth.

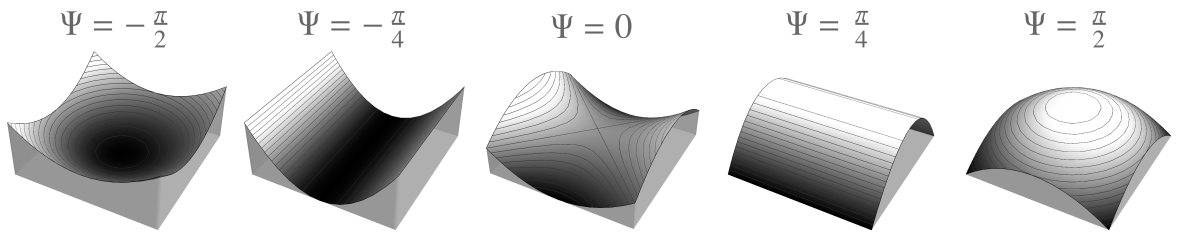

Figure 5. *The shape index scale of the quadrics. It is a continuum; here, only the major landmarks are shown. The limits of the scale are convex and concave umbilical, and the center is the symmetric saddle. In between are the parabolic ridge and rut.*

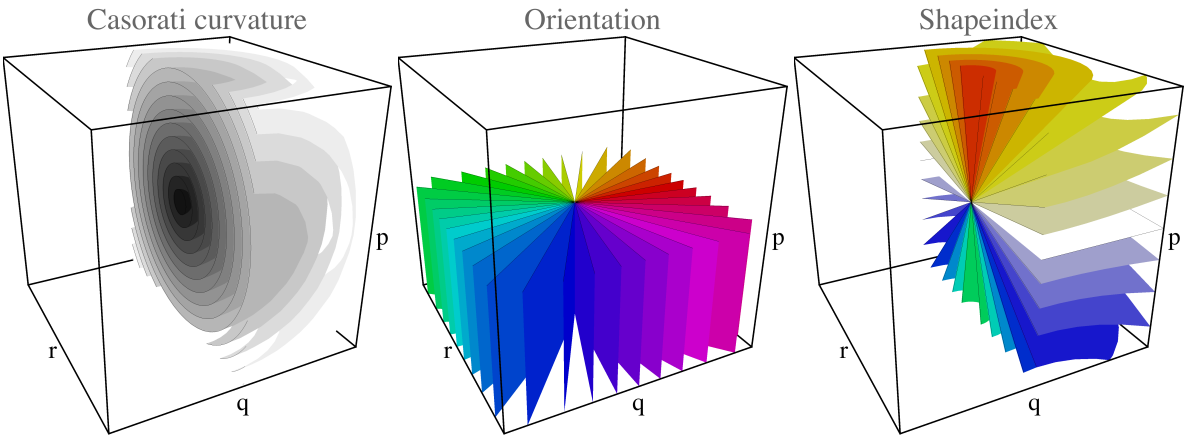

Figure 6. *The foliation of the quadrics shape space by the concentric spheres of constant Casorati curvature* Γ *(left), the half-planes of constant orientation* Φ *(center), and the semicones of constant shape index* Ψ *(right). In each case, half of the space was cut out so as to render the structure more intuitively legible.*

$$\begin{pmatrix} p' \\ q' \\ r' \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos 2\varphi & \sin 2\varphi \\ 0 & -\sin 2\varphi & \cos 2\varphi \end{pmatrix} \begin{pmatrix} p \\ q \\ r \end{pmatrix}, \tag{20}$$

an equation similar to (9).

Notice that the variance of the difference of two quadrics is simply their Euclidean distance in the $\{p, q, r\}$-space. This quadric shape-space has a natural Euclidean metric.

It is immediate to show that

$$\begin{aligned} \Gamma \sin \Psi &= p, \\ \Gamma \cos \Psi \cos 2\Phi &= q, \\ \Gamma \cos \Psi \sin 2\Phi &= r. \end{aligned}$$

Thus, the Casorati curvature Γ, the shape index Ψ, and the orientation Φ form a polar coordinate system for the shape space (see Figures 6 and 7). The locus of unit variance is $\mathbb{S}^2$. Antipodal points are related by a z-inversion, and points on small circles concentric with the umbilical axis can be superimposed through a Euclidean rotation about the origin.

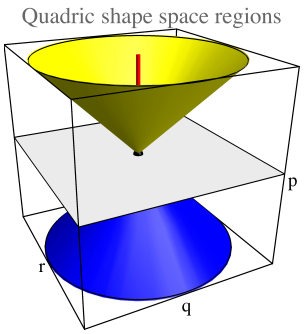

Figure 7. *This is a very coarse-grained partition of the quadrics shape space. Notice that it is axially symmetric about the umbilical axis (red line partly visible at top). The interior of the yellow semicone contains convexities, and that of the blue semicone contains concavities. The space outside the cones is the region of saddle shapes. The central plane is the locus of symmetrical saddles; it divides the saddles into ridge-like and rut-like.*

Thus, all distinct shapes are exhausted by a single meridian of the unit sphere, being mutually distinguished by the latitude, which is the conventional shape index. The habitus (see Figure 4) is conveniently written as

$$z(x,y) = \Gamma\left(\cos\Psi(x^2+y^2) + \sin\Psi\left((x^2-y^2)\cos(2\Phi) + 2xy\sin(2\Phi)\right)\right), \tag{21}$$

which can be simplified to $z(x,y) = \cos\Psi(x^2+y^2) + \sin\Psi(x^2-y^2)$ by factoring out the orientation and size so as to obtain the pure shape dependence.

This discussion reveals that the usual definition of the shape index is not just a convenience, as it is often made out to be [32], but that Casorati curvature, shape index, and orientation have an intimate relation and need to be defined as a triple instead of individually.

For constant Casorati curvature, one obtains a "shape sphere." Notice that the attitude runs at double speed, though, making the Shape sphere formally analogous to the Poincaré sphere in optics [53, 6] or the Bloch sphere [4] in quantum mechanics. Disregarding the orientation, the shape space modulo size and attitude reduces to a meridian of the sphere. This explains why the shape index space is a finite segment $\mathbb{I}^1$, really one half of $\mathbb{S}^1$. The endpoints—the poles of the sphere—represent the umbilical points and the center the symmetric saddle (Figure 5).

More generally, the p-axis is the locus of umbilicals, the plane $p = 0$ the locus of symmetrical saddles, and the right circular cone about the p-axis with semi-top-angle $\pi/4$ the locus of parabolic, or cylindrical, points (the zero locus of the quadric discriminant; see Figure 7). Moreover, the spheres concentric with the origin are loci of constant Casorati curvature. The origin itself represents the planar point.

A general surface will map on a smooth surface in shape space. From this, one immediately obtains a number of general observations. For instance,

- flat points generally do not occur on smooth reliefs;
- umbilical points are generically isolated;
- parabolic point generically lie on smooth curves;
- the surface parcellates in elliptical and hyperbolical regions, with parabolic curves as their common boundaries;

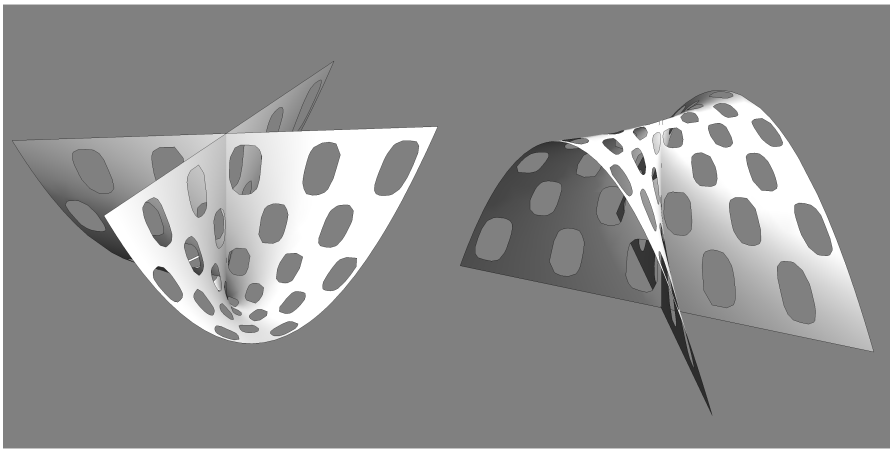

Figure 8. *The Whitney umbrella* [60, 17]*, also known as a "pinch point." It is a locus where a surface starts to intersect with itself.*

- the elliptical regions are either convex or concave; they cannot be mixed and do not share a common boundary.

A pure quadric maps on a point in shape space, cubic terms extend that to a planar patch, and quartic terms yield a curved surface patch. Thus, one needs at least the four lowest orders to even start a description of the nature of the relief in a local neighborhood, which is why the present discussion includes the cubic and quartic terms.

The generic singularities of the image in shape space are Whitney umbrellas [60, 17].[9] These are related to the self-intersections of the surface, and the Whitney umbrella (Figure 8) represents a pinch point (or cuspidal point).

3.4. The cubics. For the general cubic, one rewrites the relief of (11) as (setting $p = 3(a_{12} + a_{30})/4$, $q = 3(a_{03} + a_{21})/4$, $r = (a_{30} - 3a_{12})/4$, and $s = (3a_{21} - a_{03})/4$)

$$\frac{1}{3!}\left[(p+r)x^3 + (q+3s)x^2y + (p-3r)xy^2 + (q-s)y^3\right]. \tag{22}$$

This is indeed natural because the parameters $\{p, q, r, s\}$ transform as

$$\begin{pmatrix} p' \\ q' \\ r' \\ s' \end{pmatrix} = \begin{pmatrix} \cos\varphi & \sin\varphi & 0 & 0 \\ -\sin\varphi & \cos\varphi & 0 & 0 \\ 0 & 0 & \cos 3\varphi & \sin 3\varphi \\ 0 & 0 & -\sin 3\varphi & \cos 3\varphi \end{pmatrix} \begin{pmatrix} p \\ q \\ r \\ s \end{pmatrix}. \tag{23}$$

Thus, any cubic is the sum of two canonical cubics (a "shoe surface" (Figure 9, left) and a "monkey saddle" (Figure 9, right)) that mutually transform differently, one turning three times as fast as the other.

[9] Try $z(x, y) = (2x^2 + y^2)/2 + (x^3 + x^2y + xy^2 + y^3)/6 + (x^4 + y^4)/24$ to see a Whitney umbrella in shape space. A parametric model of the Whitney umbrella is $x(u, v) = uv$; $y(u, v) = u$; $z(u, v) = v^2$, or, as an implicit equation, $x^2 = y^2z$, where the negative z-axis is the *handle* of the umbrella.

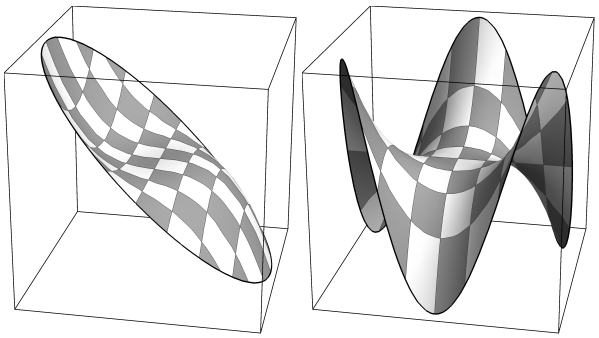

Figure 9. *The cubics are linear combinations of the shoe surface (left) and the monkey saddle (right). The shoe surface repeats after a full turn and the monkey saddle after a third turn. Notice that the shoe surface looks—in a global view over the ROI—much like a planar slope. Indeed, this part of the cubics tends to correlate strongly with the first order. The monkey saddles rarely play a role in applications.*

Notice that the coefficients p and q are the cubic terms occurring in the mixed variance (15). This apparently happens because they transform just like the gradient. (See (9) and also section 8.)

Define the parameters

$$\Gamma = \sqrt{p^2 + q^2 + r^2 + s^2} \tag{24}$$

$$\Lambda = \arctan(\sqrt{p^2 + q^2}, \sqrt{r^2 + s^2}) \tag{25}$$

$$\Psi = \arctan(p, q) \tag{26}$$

$$\Phi = \arctan(r, s) \tag{27}$$

where $\Gamma \in (0, \infty)$, $\Lambda \in (0, \pi/2)$, $\Psi \in (0, 2\pi/3)$, and $\Phi \in (0, 2\pi)$. This allows one to parameterize the 3-sphere in terms of hyperspherical coordinates

$$\begin{pmatrix} p \\ q \\ r \\ s \end{pmatrix} = \Gamma \begin{pmatrix} \cos\Lambda\cos\Phi \\ \cos\Lambda\sin\Phi \\ \sin\Lambda\cos 3\Psi \\ \sin\Lambda\sin 3\Psi \end{pmatrix}. \tag{28}$$

One has $p^2 + q^2 = \Gamma^2\cos^2\Lambda$, $r^2 + s^2 = \Gamma^2\sin^2\Lambda$. For any Λ, one obtains a Clifford torus (or flat[10] torus, a topological product of two circles $\mathbb{S}^1 \times \mathbb{S}^1$). The 3-sphere $\Gamma = 1$ foliates into such flat tori [40, 28, 7], each torus parameterized by $\{\Phi, \Psi\}$. One parameter is just the orientation, so fixing either Φ or Ψ, one obtains a circle. The analog of the shape index is thus a pair like $\{\Lambda, \Phi\}$ (say, fixing Ψ), where Λ selects a certain torus and Φ a point on the circle. This geometry can be explicitly visualized through stereographic projection from $\{1, 0, 0, 0\}$ (Figure 10).

The full cubic shape space is $\mathbb{E}^4$. Ignoring Casorati curvature (the original definition suitably extended), one obtains a shape sphere $\mathbb{S}^3$. This sphere naturally foliates into flat tori,

[10] "Flat" implies that the surface can be rolled out over the plane. Indeed, the $\{\Phi, \Psi\}$-coordinate patch is just a rectangle with its opposite sides identified.

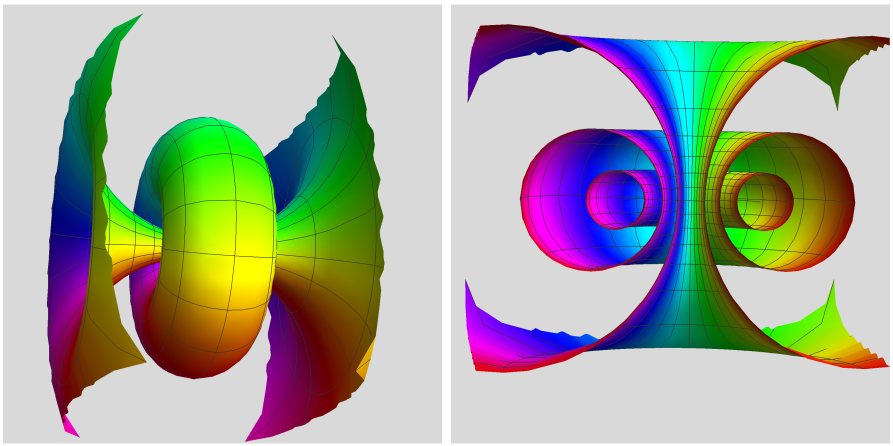

Figure 10. *The flat tori in cubic shape space shown in stereographical projection from* $\{1,0,0,0\}$. *The parameter* Λ *is set to a number of constant values, and the hue reveals the value of parameter* Ψ. *Right: half of the space has been cut out, thus showing the foliation by flat tori to good advantage.*

thus explaining 1 DOF of the cubic shape index. Any torus foliates into circles parameterized by orientation. Ignoring orientation, one obtains a circle representing the remaining DOF of the canonical cubic shape. This explains the geometrical structure of cubic shape space to the extent required for the present discussion.

The cubic habitus as a function of the shape index pair is

$$z(x,y) = \frac{1}{3!}\left(\cos\Lambda\left(x^2+y^2\right)\left(x\cos\Phi + y\sin\Phi\right) + x\sin\Lambda\left(x^2-3y^2\right)\right), \tag{29}$$

where the parameter range should be restricted so as to avoid congruent repeats. This construction yields a cubic scale space in Cartesian and polar coordinates, with a simple Euclidean Procrustes metric.

The monkey saddle part (Figure 9, right) is only of minor interest in applications involving natural images (see below), whereas the shoe-surface part occurs as a refinement on the gradient, that is, the edge character of the relief[11] (see section 8).

3.5. The quartics. Introducing the parameters $p{=}3(a_{04}{+}2a_{22}{+}a_{40})/8$, $q{=}(a_{40}{-}a_{04})/2$, $r = a_{13} + a_{31}$, $s = (a_{04} - 6a_{22} + a_{40})/8$, and $t = (a_{31} - a_{13})/2$ in (13), the quartic part can be rewritten as

$$z(x,y) = \frac{p}{4!}\left(x^2+y^2\right)^2 + \frac{q}{4!}\left(x^2-y^2\right)\left(x^2+y^2\right) + \frac{r}{4!}2xy\left(x^2+y^2\right) + \frac{s}{4!}\left(\left(x^2-y^2\right)-2xy\right)\left(\left(x^2-y^2\right)+2xy\right) + \frac{t}{4!}4xy\left(x^2-y^2\right).$$

[11] Remember that a generic "edge detector" [51, 9, 30, 58] is nothing but a first-order directional derivative at some finite scale [30].

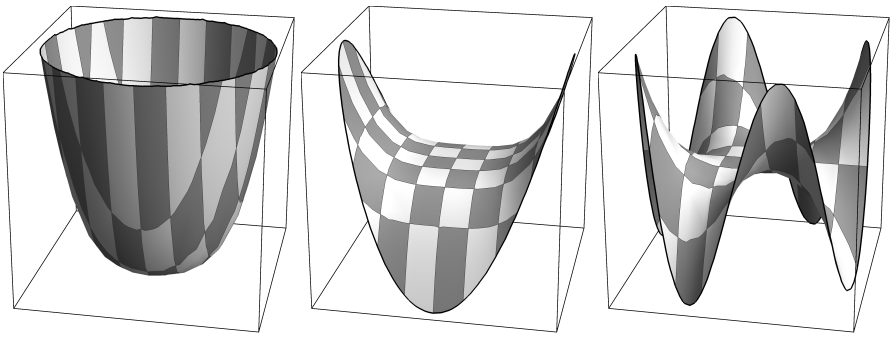

Figure 11. *The quartics are a linear combination of these three surfaces. The leftmost is rotationally symmetric, the center one repeats after a half turn, and the rightmost one repeats after a quarter turn. Notice that the former two are very similar—in a global view of the* ROR*—to the quadrics. Indeed, this part of the quartics tends to correlate strongly with the quadrics (Figure 4). The higher-order saddles rarely play a role in most applications.*

This is again natural because the parameters $\{p, q, r, s, t\}$ transform as

$$\begin{pmatrix} p' \\ q' \\ r' \\ s' \\ t' \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \cos 2\varphi & \sin 2\varphi & 0 & 0 \\ 0 & -\sin 2\varphi & \cos 2\varphi & 0 & 0 \\ 0 & 0 & 0 & \cos 4\varphi & \sin 4\varphi \\ 0 & 0 & 0 & -\sin 4\varphi & \cos 4\varphi \end{pmatrix} \begin{pmatrix} p \\ q \\ r \\ s \\ t \end{pmatrix}. \tag{30}$$

Thus, any quartic is the sum of three canonical quartics that mutually transform differently, and one turns two times and another four times as fast as the first one (Figure 11).

Notice that the mixed variance (16) can be expressed as a linear combination of terms p, q, r, which indeed transform just like the quadric case.[12] See also section 8.

The quartic variance satisfies $\Gamma \propto \frac{5}{3}p^2 + q^2 + r^2 + s^2 + t^2$. The full quartic shape space is $\mathbb{E}^5$, the Procrustes metric—deviation from planarity—again induces a natural Euclidean representation. Disregarding Casorati curvature (the original definition suitably extended), one obtains a shape sphere $\mathbb{S}^4$. Again, there are various loci on the shape sphere that represent congruent shapes.

The quartic shape index has 3 DOF. It can be studied in a similar way as done for the cubics. However, this will rarely be of interest because the high-order saddle part has no relevance to most applications involving natural images (see below).

4. The edge and ridge character of the relief. As an explicit polar representation, it is convenient to write

$$\begin{aligned} z(\varrho, \varphi) = p_{00} + p_{20}\frac{\varrho^2}{2} + p_{40}\frac{\varrho^4}{24} + \varrho\left((p_{11} + p_{31}\frac{\varrho^2}{6})\cos\varphi + (q_{11} + q_{31}\frac{\varrho^2}{6})\sin\varphi\right) + \\ \frac{\varrho^2}{2}\left((p_{22} + p_{42}\frac{\varrho^2}{12})\cos 2\varphi + (q_{22} + q_{42}\frac{\varrho^2}{12})\sin 2\varphi\right) + R(\varrho, \varphi), \end{aligned} \tag{31}$$

[12] A direct calculation shows that the mixed variances are $\Gamma_{(13)} = \frac{4}{3}\left(p_{(1)}p_{(3)} + q_{(1)}q_{(3)}\right)$ (15) and $\Gamma_{(2,4)} = \frac{1}{6}\left(4p_{(2)}p_{(4)} + 3q_{(2)}q_{(4)} + 3r_{(2)}r_{(4)}\right)$ (16), where we use $\{p_{(i)}, q_{(i)}, r_{(i)}, \ldots\}$ for the parameters of the orders i.

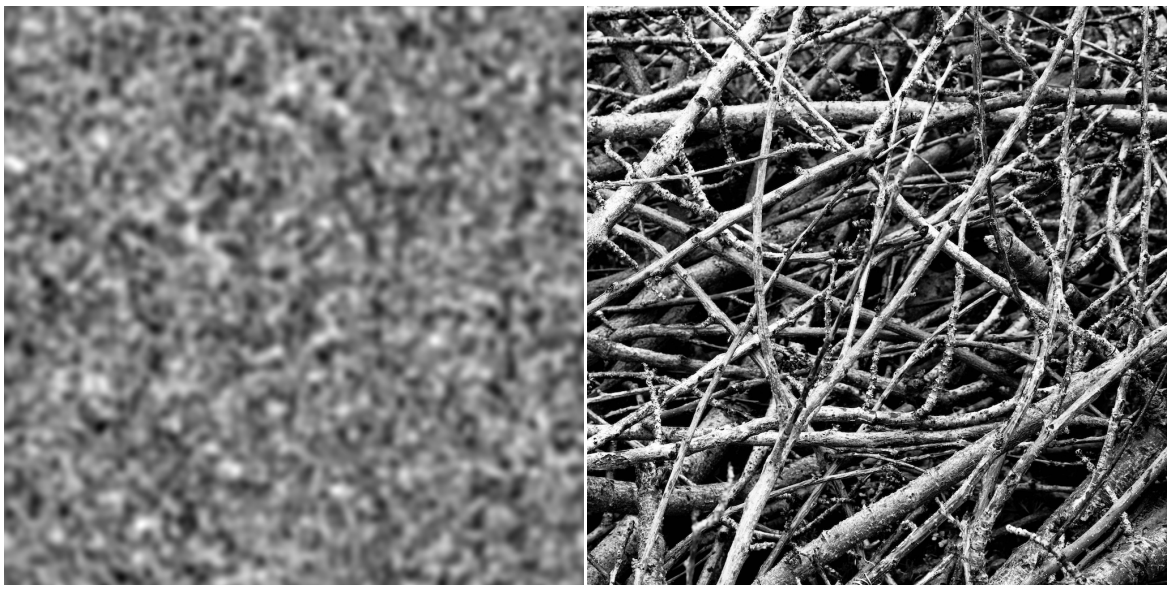

Figure 12. *Two images used to illustrate various facts. Left: an isotropic random Gaussian image with a Gaussian autocorrelation function (henceforth called "noise image"); right: the "twigs image," which is a texture-like natural image.*

where

$$R = (p_{33}\sin 3\varphi + q_{33}\cos 3\varphi)\frac{\varrho^3}{6} + (p_{44}\sin 4\varphi + q_{44}\cos 4\varphi)\frac{\varrho^4}{24} \tag{32}$$

represents the contributions of the higher-order saddles. It is immediate to see how this form can be extended to progressively higher orders. However, this is unlikely to be of interest in applications.

Here the $\{p_{nm}, q_{nm}\}$ are the coefficients of

$$\{\cos m\varphi, \sin m\varphi\}\frac{\varrho^n}{n!}. \tag{33}$$

The part without R explicitly describes the edge ($\{\cos\varphi, \sin\varphi\}$) and ridge[13] character ($\{\cos 2\varphi, \sin 2\varphi\}$) of the relief [19, 18, 41, 42, 43] up to order 4.

This representation immediately reveals the angular properties of directional derivatives of any order (here limited to 4, but the series can be continued indefinitely, of course). It might be said to represent the angular spectrum of the relief explicitly. Although one would prefer Cartesian coordinates in technical image processing, it is well known that the visual cortex uses (overcomplete) bases of directional derivatives [21]. Here (31) is the natural representation of the "orientation pinwheel" [22], the neural implementation of the polar bases [29].

In "natural images" (Figure 12), one finds empirically that the edge components of the linear and cubic orders tend to be aligned [19, 18, 41, 42, 43]. Likewise, one finds that the

[13]The famous instance is the Hubel and Wiesel [21] "line detector." Formally, this is just a second-order directional derivative at some finite scale [29].

ridge components of the quadric and quartic orders also tend to be aligned [19, 18, 41, 42, 43]. That is,

$$\left|\frac{p_{11}q_{31} - p_{31}q_{11}}{p_{11}p_{31} + q_{11}q_{31}}\right| \ll 1 \tag{34}$$

and

$$\left|\frac{p_{22}q_{42} - p_{42}q_{22}}{p_{22}p_{42} + q_{22}q_{42}}\right| \ll 1. \tag{35}$$

For Gaussian random images, this can be shown to apply on the average (see section 8).

5. Distributions in relief-shape space. Consider the second order. A natural way to sample from the shape space is to draw from an isotropic normal distribution centered at the origin of *pqr*-space [3] with, say, variance s^2. In polar coordinates, the probability density function is

$$P(\Gamma, \Lambda, \Phi, s) = \frac{\Gamma^2 e^{-\frac{\Gamma^2}{2s^2}} \cos(\Psi)}{2\sqrt{2}\pi^{3/2}s^3}, \tag{36}$$

from which one easily derives the marginals

$$P_\Gamma(\Gamma) = \frac{\sqrt{\frac{2}{\pi}}\Gamma^2 e^{-\frac{\Gamma^2}{2s^2}}}{s^3}, \qquad P_\Psi(\Psi) = \frac{\cos\Psi}{2}, \qquad P_\Phi(\Phi) = \frac{1}{2\pi}. \tag{37}$$

Thus, the distribution of the Casorati curvature follows a Maxwell–Boltzmann distribution, whereas the distribution of orientations is uniform, both much as expected. The distribution for the shape index is zero at the umbilical endpoints of the scale and peaks at the symmetric saddles (see Figure 13, left).

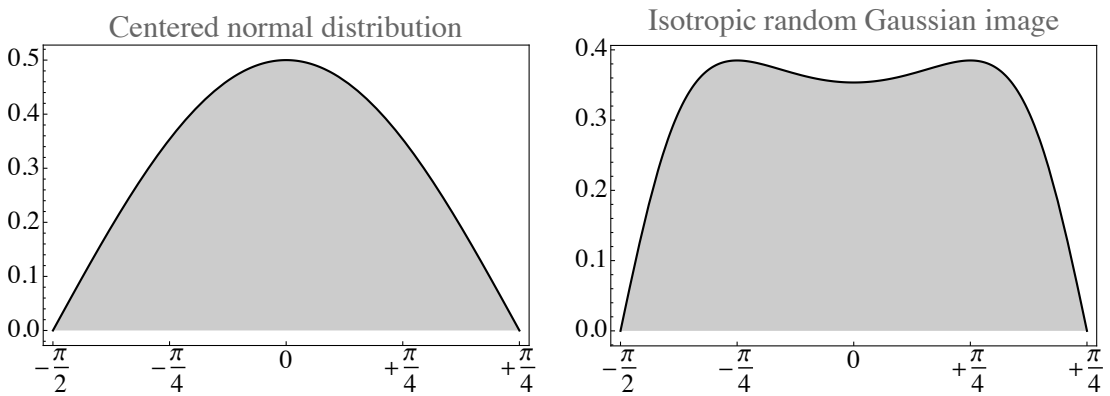

Figure 13. *Left: the marginal distribution of the quadric shape index for a normal, centered distribution in shape space. It peaks at the symmetric saddle and vanishes at the limits of the scale, reflecting the fact that umbilics occur only as isolated points. Right: the distribution for points on an isotropic random Gaussian image (see (44)). This distribution peaks at the parabolic points.*

However, this ignores the fact that the samples should *derive from an image*. A generic example is the isotropic random Gaussian image (Figure 12, left). In this case, the spatial derivatives and thus the coefficients of the quadric are mutually correlated [48, 3, 50]. One has

$$\overline{\left(\frac{\partial^p\partial^q z}{\partial x^p y^q}\right)\left(\frac{\partial^{p'}\partial^{q'} z}{\partial x^{p'} y^{q'}}\right)} = (-1)^{\frac{1}{2}(p+p'+q+q')} m_{(p+p')(q+q')}, \tag{38}$$

where $m_{rs} = M_{r+s}\mu_{rs}$ with $r = p + p'$, $s = q + q'$. The "circular moments" M_n depend only on the sum of the orders, whereas the angular distributions μ_{rs} are captured by averages over directions.

For the Gaussian autocorrelation function $\exp(-(x^2 + y^2)/(2\sigma^2))$, one has

$$M_n = \pi^2 2^{\frac{n+5}{2}} \Gamma\left(\frac{n+3}{2}\right) \sigma^{-n-1} \tag{39}$$

and

$$\mu_{rs} = \frac{((-1)^r + 1)\,((-1)^{r+s} + 1)\,\Gamma\left(\frac{r+1}{2}\right)\Gamma\left(\frac{s+1}{2}\right)}{4\pi\Gamma\left(\frac{1}{2}(r+s+2)\right)}. \tag{40}$$

For other autocorrelation functions, one simply obtains an overall constant factor; thus, the actual form of autocorrelation function is irrelevant here.

The covariance matrix Σ_2 for the derivatives $\{\partial_{xx}, \partial_{xy}, \partial_{yy}\}$ reveals correlations between the partial derivatives, whereas the covariance matrix Ξ_2 for the canonical coefficients $\{p, q, r\}$ is diagonal:

$$\Sigma_2 = \frac{M_4}{8}\begin{pmatrix} 3 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 3 \end{pmatrix}, \quad \Xi_2 \propto \begin{pmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{41}$$

Thus, the canonical coefficients are mutually uncorrelated, and the power is equally distributed between p^2 and $q^2 + r^2$.

Likewise, in the cubic case for the derivatives $\{\partial_{xxx}, \partial_{xxy}, \partial_{xyy}, \partial_{yyy}\}$ and the canonical coefficients $\{p, q, r, s\}$, one has

$$\Sigma_3 = -\frac{M_6}{16}\begin{pmatrix} 5 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 5 \end{pmatrix}, \quad \Xi_3 \propto \begin{pmatrix} 9 & 0 & 0 & 0 \\ 0 & 9 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{42}$$

from which it follows that the distribution of power between $p^2 + q^2$ and $r^2 + s^2$ is 9:1.

For the quartic case for the derivatives $\{\partial_{xxxx}, \partial_{xxxy}, \partial_{xxyy}, \partial_{xyyy}, \partial_{yyyy}\}$ and the canonical coefficients $\{p, q, r, s, t\}$, the covariance matrix is

$$\Sigma_4 = \frac{M_8}{128}\begin{pmatrix} 35 & 0 & 5 & 0 & 3 \\ 0 & 5 & 0 & 3 & 0 \\ 5 & 0 & 3 & 0 & 5 \\ 0 & 3 & 0 & 5 & 0 \\ 3 & 0 & 5 & 0 & 35 \end{pmatrix}, \quad \Xi_4 \propto \begin{pmatrix} 18 & 0 & 0 & 0 & 0 \\ 0 & 16 & 0 & 0 & 0 \\ 0 & 0 & 16 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \tag{43}$$

so the distribution of power is very uneven, $p^2 : q^2 + r^2 : s^2 + t^2 = 9 : 16 : 1$.

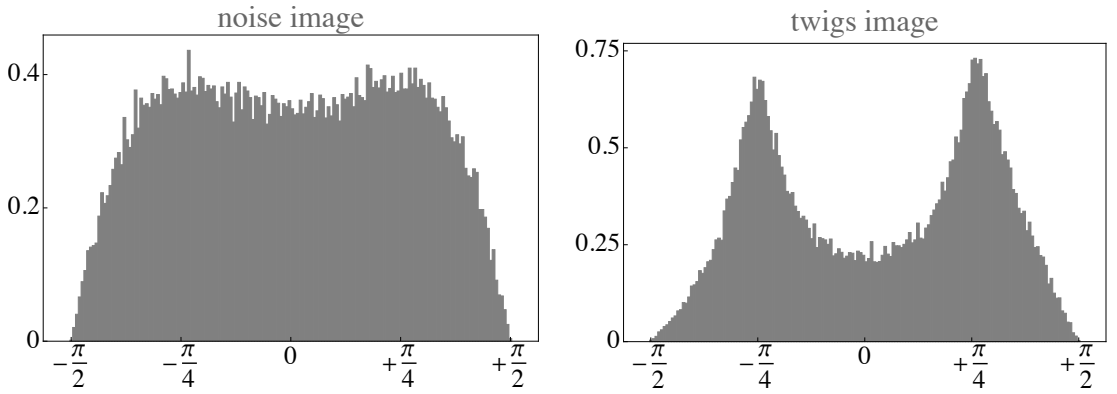

Figure 14. *Empirical shape index marginal distributions for the two images (Figure 12). As is the case for almost any instance, the distributions are strongly peaked at the ruts and ridges. Moreover, there is a marked asymmetry, no doubt due to factors of ecological optics. Compare this to Figure 13, right.*

These correlations determine the various marginal distributions for Gaussian random surfaces. Most importantly, the power of the higher-order saddles is very low, only 10% for the cubic and 4% for the quartic parts. This is indeed trivially confirmed by a numerical investigation involving Gaussian random surfaces, but, more importantly, essentially the same statistics prove to apply[14] to the case of natural images.

In the case of the quadrics, the marginal that yields the PDF of the shape index is (the integrations over the Casorati curvature and the orientation being elementary)

$$Q_\Psi(\Psi) = \frac{2^{3/2}\cos\Psi}{(3+\cos 2\Psi)^{3/2}}, \tag{44}$$

an expression that has been derived earlier [36, 43]. It has the result that parabolic points become relatively more numerous (Figure 13, right).

In order to quantify this, the shape index range may be divided into the ranges $E^+ = (-\frac{\pi}{2}, -\frac{3\pi}{8})$ (elliptic convex), $P^+ = (-\frac{3\pi}{8}, -\frac{\pi}{8})$ (ridges), $H = (-\frac{\pi}{8}, +\frac{\pi}{8})$ (saddles), $P^- = (\frac{\pi}{8}, \frac{3\pi}{8})$ (ruts), and $E^- = (\frac{3\pi}{8}, \frac{\pi}{2})$ (elliptic concave). These are the Voronoi cells for the special points $\pm\frac{\pi}{2}$ (umbilics), $\pm\frac{\pi}{4}$ (parabolic points), and 0 (the symmetric saddle).

Granted the Voronoi cells, one counts the probability of encountering elliptic p_E, parabolic p_P and hyperbolic p_H cases. This allows the definition of the *ridginess* $R = p_P$ and *ellipticity* $E = p_E/(p_E+p_H)$. For the distribution P_Ψ, one finds $R = 0.541\ldots$, $E = 0.166\ldots$, and for the distribution Q_Ψ, one finds $R = 0.582\ldots$, $E = 0.328\ldots$. Apparently, the ridginess increases when one considers points sampled from an image instead of shape space proper.

It has been noticed [43] that the empirical distribution for natural images is markedly more peaked at the ruts and ridges than what is seen for the Gaussian random surfaces (example in Figure 14). One likely reason appears to be the violation of isotropy. In order to check this, one has to derive the probability density function for the shape index of a nonisotropic Gaussian random surface (no result appears to be recorded in the literature). The covariance matrix for the derivatives $\{\partial_{xx}, \partial_{xy}, \partial_{yy}\}$ is readily available [3]:

[14] At least in the (many) instances tried more or less at random.

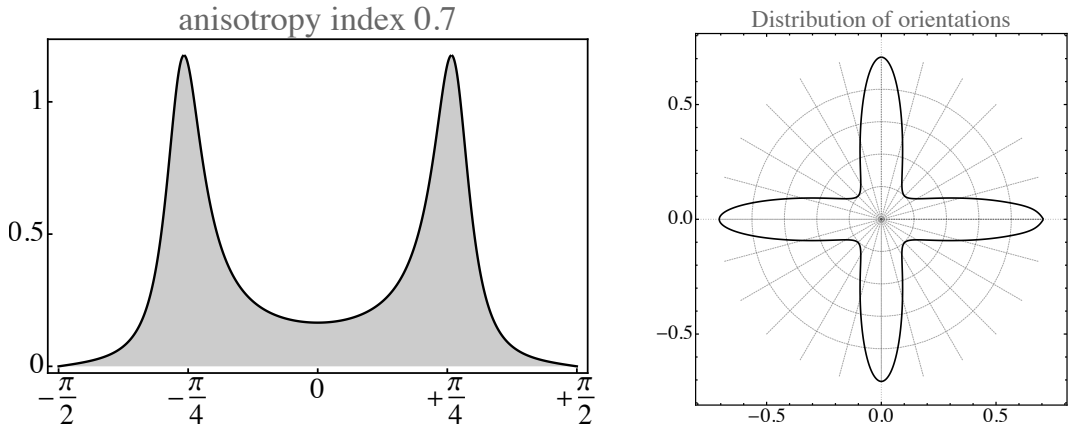

Figure 15. *The marginal distribution of the quadric shape index (left) for a Gaussian random surface with anisotropy index* 0.7*. It peaks strongly at the parabolic points; thus, the surface consists mainly of ruts and ridges. For the ideal Gaussian random surface, ruts and ridges are equally likely. In natural images, one often finds asymmetries, no doubt due to various physical factors: The radiometry works out quite differently for concavities (effects of vignetting and interreflections) and convexities (mainly shading proper). At the right is the marginal distribution of the orientations for the same anisotropy. For a natural image, one would not expect to see this since the orientation of the anisotropy would vary from region to region; thus, the articulation seen here would average out.*

$$\frac{M_4}{8}\begin{pmatrix} 3\xi^{-4} & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 3\xi^{4} \end{pmatrix}, \tag{45}$$

where ξ denotes the stretch of the x-axis, whereas ξ^{-1} is the corresponding shrink of the y-axis. A convenient index of anisotropy is $s = (\xi^4 - 1)/(\xi^4 + 1)$.

However, although marginalization with respect to the Casorati curvature is immediate, the marginalization with respect to the orientations leads to integrals that cannot be done in closed form (at least, Mathematica gives up). However, numerical integration is easily and stably done to any desired accuracy. As expected, the PDF peaks rather strongly at the parabolic shapes, even for moderate degrees of anisotropy (see Figure 15).

The anisotropy for the natural images is not necessarily reflected in the overall ridge orientation distribution (see Figure 16). In such images, the local anisotropies are not lined up. In the case of the twigs image, the distribution of orientations is almost flat, although there are evidently very strong local anisotropies (Figure 17).

Right from the beginning [36, 43], it was noticed that real images yield results that differed from the predictions. The reason is clear by now: In real images, one deals with a *spectrum of anisotropy* depending upon the size of the ROIs. This effectively decouples the histogram of shape indices from that of the orientations.

The ridginess increases monotonically with the degree of anisotropy, reaching a fraction of 0.9 for an anisotropy 0.760 (see Figure 18). The ellipticity falls monotonically with the anisotropy: It is down to 10% of its initial value for anisotropy of 0.712. Thus, (near) umbilicals become really scarce, and most of the surface is ridge or rut, the remainder being mainly

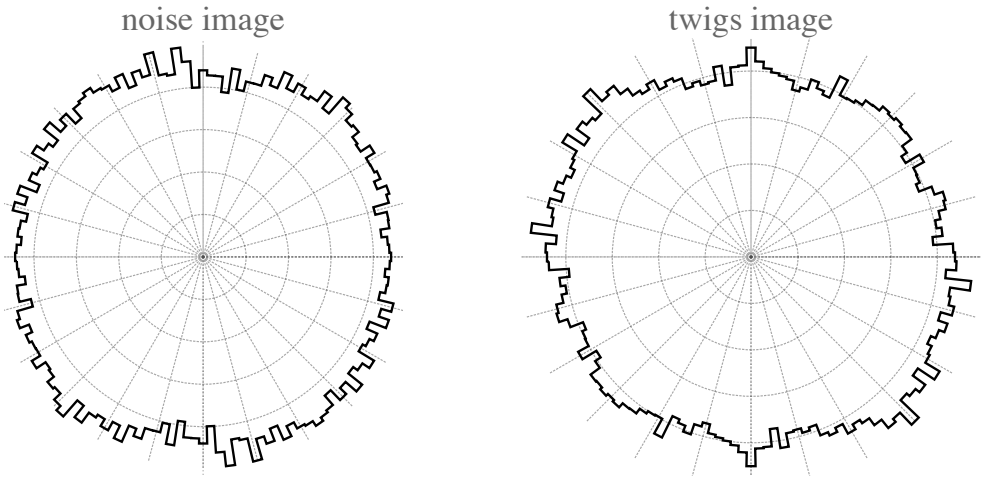

Figure 16. *Polar histograms for the quadric ridge orientation of the two images (Figure 12). Notice that local anisotropy has a strong influence on the shape index distribution but is not necessarily evident in the overall orientations (see Figure 17).*

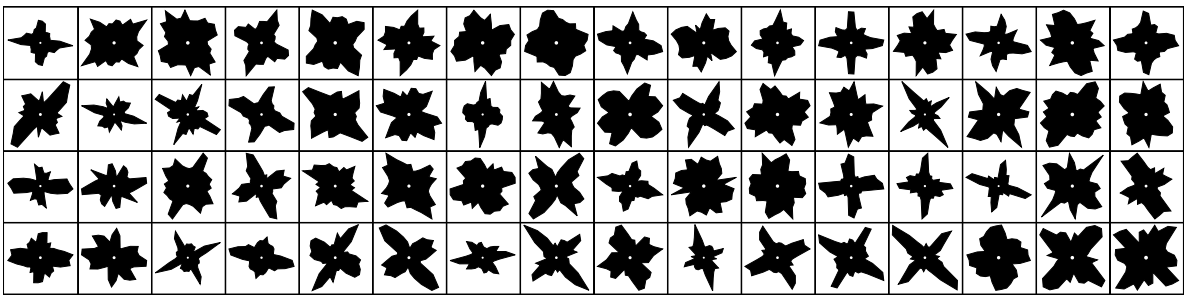

Figure 17. *Here the twigs image was cut in eight-by-eight pieces and an orientation histogram made for each piece. Notice that there is very strong local anisotropy, which averages out over the whole picture. This explains why the shape index histogram (Figure 15) corresponds to significant anisotropy, whereas the orientation histogram (Figure 16 right) appears to suggest otherwise (it "should" look much like Figure 15, right).*

hyperbolic. (Figure 18.) All these predictions are perfectly reflected in simulations using random Gaussian image samples.

The average density of umbilical points has been derived by Berry [3]. It is $M_6/(4\pi M_4)$. It varies only very slightly with the anisotropy (less than 8% over the full anisotropy range). Thus, although almost all of the area becomes eventually occupied by parabolic points, the number of umbilics is not really affected by that.

6. Distribution of power over the orders. A major distinction between single scale, isotropic Gaussian random noise and "natural images" appears when comparing the distribution of power over the orders. For the noise image, the distribution of powers ($\Gamma_{(2)}$, $\Gamma_{(3)}$, $\Gamma_{(4)}$; (10), (12), (14)) is

$$\text{quadric} : \text{cubic} : \text{quartic} = 50\,(33, 63) : 31\,(16, 51) : 16\,(10, 23), \tag{46}$$

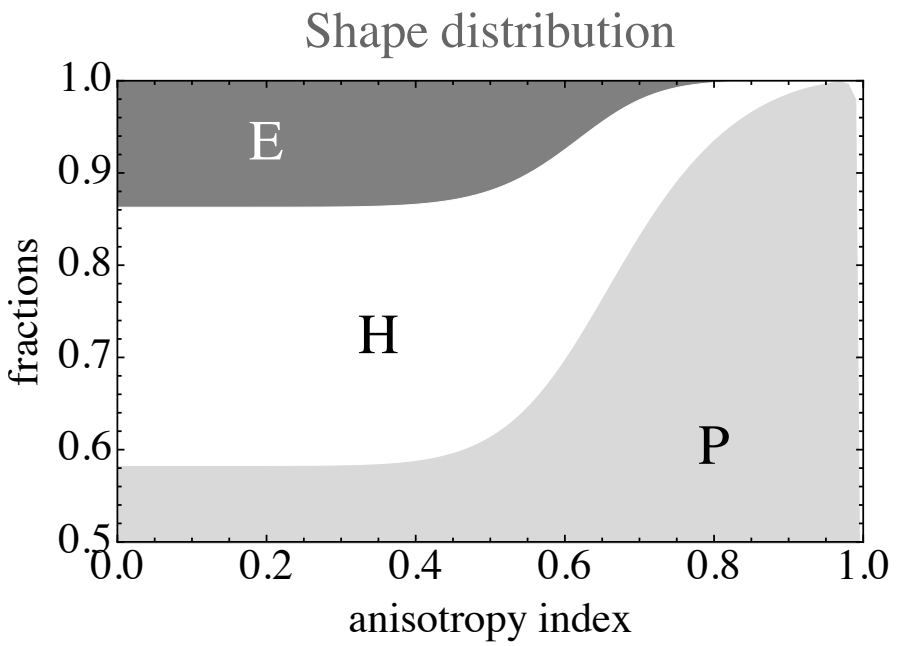

Figure 18. *The distribution of surface type (elliptic (E), parabolic (P), or hyperbolic (H)) as a function of the anisotropy index. As the anisotropy grows, the proportion of parabolic points grows, and the ratio of elliptic to hyperbolic points diminishes.*

where the interquartile ranges are given in brackets.

Another way to look at this is to count pixels where either the quadric, cubic, or quartic power dominates. One finds

$$\text{quadric} : \text{cubic} : \text{quartic} = 63 : 35 : 2; \tag{47}$$

thus, almost two-thirds of the area is dominated by the quadric terms and one-third by the cubic terms, whereas only about 2% is dominated by the quartic order.

This is *very* different for the twigs image, where one finds

$$\text{quadric} : \text{cubic} : \text{quartic} = 27\,(15, 43) : 30\,(13, 53) : 34\,(19, 49) \tag{48}$$

and

$$\text{quadric} : \text{cubic} : \text{quartic} = 25 : 38 : 37, \tag{49}$$

the quartic contribution being increased by more than an order of magnitude (see Figure 19).

7. Remarkable landmarks. "Remarkable landmarks" are various singularities, points, or curves as the case may be [56, 11, 52, 8, 25, 12, 54, 3, 35]. We consider the lowest-order cases, but, of course, one may extend the analysis arbitrarily.

Singularities of the first order are the points where the gradient magnitude vanishes. These are nodes or saddles, their nature being decided by the sign of the quadric discriminant, that is, $K = a_{20}a_{02} - a_{11}^2$, which is the intrinsic (often called "Gaussian") curvature in the singly isotropic (graph) space. Notice that K is the magnification of the map to gradient space.

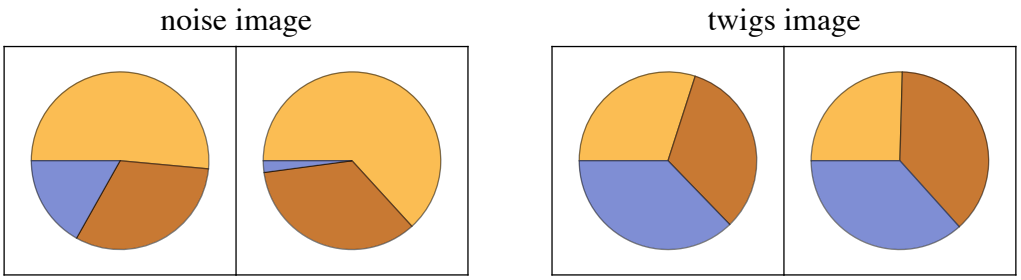

Figure 19. *In each frame, one has at left the distribution of the powers $\Gamma_{(2)}$ (yellow), $\Gamma_{(3)}$ (brown), and $\Gamma_{(4)}$ (blue) and at right the distribution of the most significant order (quadric yellow, cubic brown, quartic blue) over all pixels. The distributions are very different for an isotropic Gaussian random field (frame at left, the "noise image," Figure 12, left) as compared to a natural image (frame at right, the "twigs image," Figure 12, right). The result for the latter image is typical for that of many "natural images."*

The familiar "gradient space" is the plane $\{z_x, z_y\}$. The surface $z(x, y)$ is mapped in gradient space with as generic singularities folds and cusps. The folds are the loci $K = 0$; on the surface, these are the parabolic curves.

The cusps are singular points [2] on the folds where the gradient of K is tangent to the parabolic curve. Their nature is determined by the cubic and quartic terms.

The singular points of the quadrics are umbilical points [3]. These are generically isolated points on the surface; they occur in three topologically distinct kinds, their nature being determined by the cubic terms.

The surface may also be mapped into the space of quadrics. The generic singularities are Whitney umbrellas, which are points where the surface starts to intersect itself, often called "cross-caps." At such points, one needs to draw on the quartic terms.

In all cases generic local models are readily available.

8. Geometrical meaning of mixed order differential invariants. In most cases, the lowest orders (first and second) carry most of the interesting structure. The generic case can always be brought into the canonical form

$$z(x, y) = x + \frac{1}{2}(Px^2 + 2Qxy + Ry^2) + \dots \tag{50}$$

with a movement in the singly isotropic (graph) space. As expected, the mixed differential invariants $\{P, Q, R\}$ have immediate geometrical significance [54, 35], namely (Figures 20, 21, and 22), the following:

- P is proportional to the isotropic curvature of a section $y = 0$ at the origin;
- Q is proportional to the Euclidean curvature of the slope line (in the x-direction) at the origin;
- R is proportional to the Euclidean curvature of the isohypse $z = 0$ at the origin,

as will be illustrated below.

Indeed, the mixed invariant R (Figure 22) is commonly used as a "curvature detector" [34] or even a "corner detector," whereas the mixed invariant P (Figure 20) has been used as a "cliff detector" or a "subcontour detector." So far, the mixed invariant Q (Figure 21) seems to have found little application.

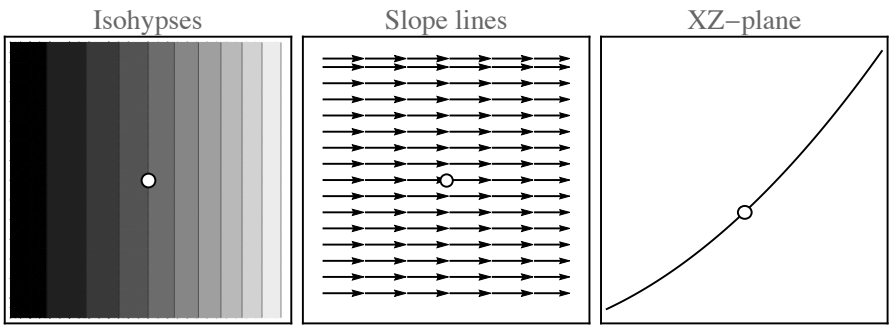

Figure 20. *The effect of the mixed first- and second-order differential invariant P is well illustrated by the relief $z(x, y) = x + x^2/2$. Shown are the isohypses of the relief (left) and the slope lines (center) in the XY-plane as well as the surface section in the XZ-plane (right). Notice that the isohypses are straight but unevenly spaced and that the flow of slope lines is rectilinear and uniform but that the section in the $y = 0$ plane is curved.*

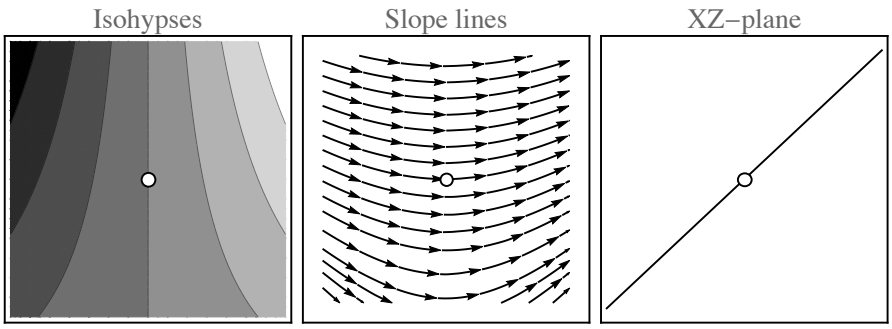

Figure 21. *The effect of the mixed first- and second-order differential invariant Q is well illustrated by the relief $z(x, y) = x + xy$. Shown are the isohypses of the relief (left) and the slope lines (center) in the XY-plane as well as the surface section in the XZ-plane (right). Notice that both the isohypse $z = 0$ and the section in the plane $y = 0$ are straight but that the slope line through the origin is curved.*

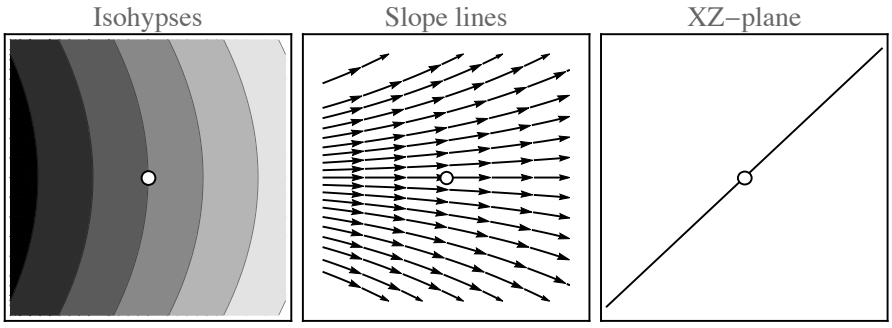

Figure 22. *The effect of the mixed first- and second-order differential invariant R is well illustrated by the relief $z(x, y) = x + y^2/2$. Shown are the isohypses of the relief (left) and the slope lines (center) in the XY–plane as well as the surface section in the XZ–plane (right). Notice that both the slope line through the origin and the section in the plane $y = 0$ are straight but that the isohypse $z = 0$ is curved.*

The isohypses are integral curves of the equation $z_x \mathrm{d}x + z_y \mathrm{d}y$, which is exact since $(z_x)_y = (z_y)_x$. Thus, the isohypses can be naturally labeled by some monotonic function of z.

The slope lines are integral curves of the equation $z_y \mathrm{d}x - z_x \mathrm{d}y$, which is *not* exact since generically $z_{xx} \neq -z_{yy}$. Using an integrating divisor, say, $\vartheta(x, y)$, one may seek for a function $w(x, y)$ that labels the slope lines up to a monotonic transformation. However, the labeling cannot bridge the zero loci of the integrating factor(s).

An example is the case shown in Figure 22, illustrating $z(x, y) = x + y^2/2$. Here the slope equation is $y\mathrm{d}x - \mathrm{d}y$, which allows an integrating divisor $\vartheta(x, y) = y$. The line $y = 0$ is a common asymptote of the areas $\mathrm{sign}(y) = \pm 1$; the consistent labeling of slope lines $w(x, y) = x - \log|y|$ cannot cross the asymptote $y = 0$. This is important in the proper definition of the "water divides" and "river courses" of a relief. (Indeed, the line $y = 0$ in Figure 22 is evidently a "divide" since all slope lines originate from it. Inverting the flow it becomes a "course" since any element of the flow eventually joins it.)

Notice that partial derivatives of different orders on a Gaussian random surface are often correlated [48]. Whereas even and odd orders are mutually uncorrelated, the first order is correlated with the third,

$$\frac{M_4}{8} \begin{pmatrix} 3 & 0 & 1 & 0 \\ 0 & 1 & 0 & 3 \end{pmatrix}, \tag{51}$$

and the second with the fourth,

$$-\frac{M_6}{16} \begin{pmatrix} 5 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 5 \end{pmatrix}, \tag{52}$$

where the partial derivatives occur in their natural order, as in (41), (42), (43). This is a necessary constraint to take into account when computing marginal statistics for landmarks of mixed orders.

The correlations for the orders 1 through 4 for the images (Figure 12) can be determined numerically. (We state them as percentages.) They are (the "noise image" at the left, the "twigs image" at the right)

	2	3	4
1	-1	66	-2
2		-1	75
3			-1

,

	2	3	4
1	-4	57	2
2		0	71
3			6

.

Notice that the twigs image yields correlations that are very similar to those of the noise image.

The above constraints have remarkable consequences. For instance, using these correlations, one easily verifies that the edge contributions of the linear and cubic orders line up in the average, and so do the ridge contributions of the quadric and quartic orders (see (34), (35)). On the other hand, a similar computation involving averages of products of partial derivatives reveals that there is no average relation between the orientations of the higher-order saddles and those of the linear and quadric orders. This regularity has been empirically found in natural images. It is hardly a specific property of natural images, though, since it can

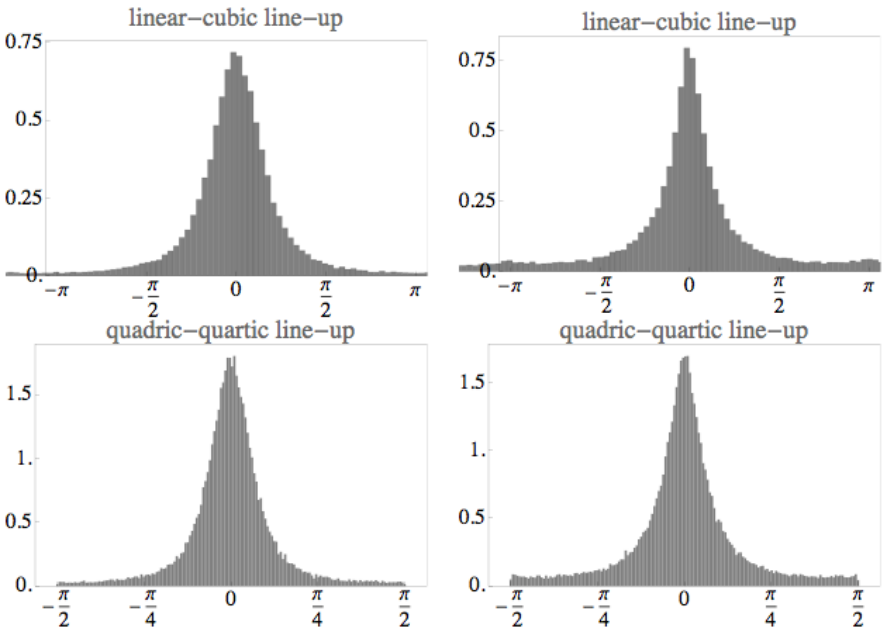

Figure 23. *Lineups computed from the two images (left: "noise image"; right: "twigs image"). In the top row are the deviations between the gradient and cubic edge directions, and in the bottom row are the deviations between the quadric and quartic ridge orientations. Results obtained for "natural images" are similar to that for the noise. A useful measure of angular spread is the semi-interquartile width. For the noise image, we find* 24° *for the edge and* 10° *for the ridge spreads; for the twigs image, the values are* 27° *and* 11°.

be proven to apply to random Gaussian images. It is due to the correlation between partial derivatives for smooth random surfaces (see Figure 23).

This is also of importance in the correlation of derivatives of a single scale at (slightly) different locations. For instance, an order 1 derivative at a slightly different location can be expressed as a local Taylor expansion involving the higher orders. Thus, the correlation over space translates into a correlation over orders. The methods discussed here readily allow one to investigate such important relations.

In neurophysiological investigations, the system is typically stimulated through retinal activation. Many of the correlations detected in the visual cortex will then be due to the structures discussed here; in many cases, there will be no immediate need to invoke hypothetical neural connections.

9. Very local structure. So far, much of the emphasis has been on the overall statistics of local structures. This is of interest for brain science since much of the available data are exactly of that nature. Because of that, there has been much interest in the statistical structure of natural images as opposed to such simple cases as isotropic, single-scale, Gaussian images. However, in any case, involving image interpretation, statistics is only of limited interest. The statistics is useful to characterize the "texture" of (usually large) image regions that surround a main center of interest. The center of interest tends to be a singular location, and the ROI is such that the number of relevant DOF is very limited. Such centers of interest are the local images that *focal vision* deals with at any moment.

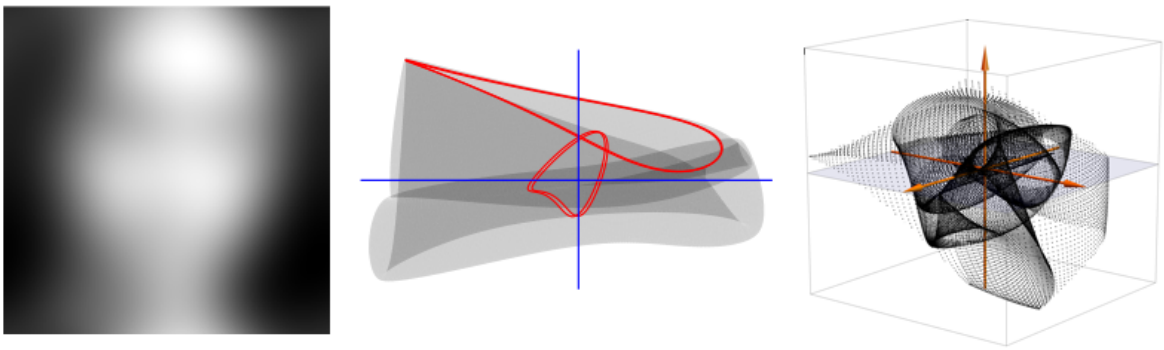

Figure 24. *This is an example of a "local structure," in this case a distant face. At this scale (left figure), one can just make out the characteristic "bar code" structure of a face. At the center is the image in gradient space. Even for a simple image like this one has various folds and cusps in a complicated pattern. The red curve is the image of the (square) boundary of the image at left. The axes are centered at the origin. At the right is the image in quadric shape space. The axes are again centered at the origin. The horizontal plane is the locus of symmetrical saddles. Notice that the map roams through all octants. There are several Whitney umbrellas, although they cannot be distinguished from cusps due to the projection here.*

The term "local" should be understood relative to the current scale. What is meant is an image extent (or ROI) that offers just a very limited number of DOF for the current scale. Such images look like very low resolution icons. Figure 24 shows an example.

For such cases, the overall statistics are not relevant at all. One looks at singular structures instead of statistical measures.

For the case of the (very typical) example of Figure 24, one notices that the image in gradient space is dominated by singularities, both folds and cusps, and that the same observation holds for the image in quadric shape space. It is very evident that even such extremely blurred pictures offer geometrical complexities that are beyond the reach of simple analysis, whereas they reveal insufficient articulation for a statistic description to make much sense.

This is the type of structure the visual cortex is working with at a white spectrum of nested scales.

10. Local differential geometry and image processing. The simple, local differential geometry discussed in this paper can immediately be applied to local operations in image processing. The reason is that spatial derivatives can be implemented exactly, except for the discretization implied by pixelation. The one factor to consider is that any derivation implies a *scale*. No *actual* object can be considered "differentiable," but the mode of observing them (or addressing them) can. This is the use of the familiar "scale-space" [33, 44, 47, 14, 58] representation.

Confronted with a twofold extended distribution of intensity $\mathcal{I}$, say, the radiance at the eye due to scattering of sunlight by the scene in front of you, one considers its representation at finite scale σ. The representation $\mathcal{R}$ is obtained through a physical sampling that is formally modeled by a convolution with a Gaussian kernel g_σ as $\mathcal{R} = g_\sigma \circ \mathcal{I}$. The point here is that the representation $\mathcal{R}$ is an *observation* of an ideal object $\mathcal{I}$ that can only be known through its representations. The object $\mathcal{I}$ cannot even be known; thus, it can hardly be considered

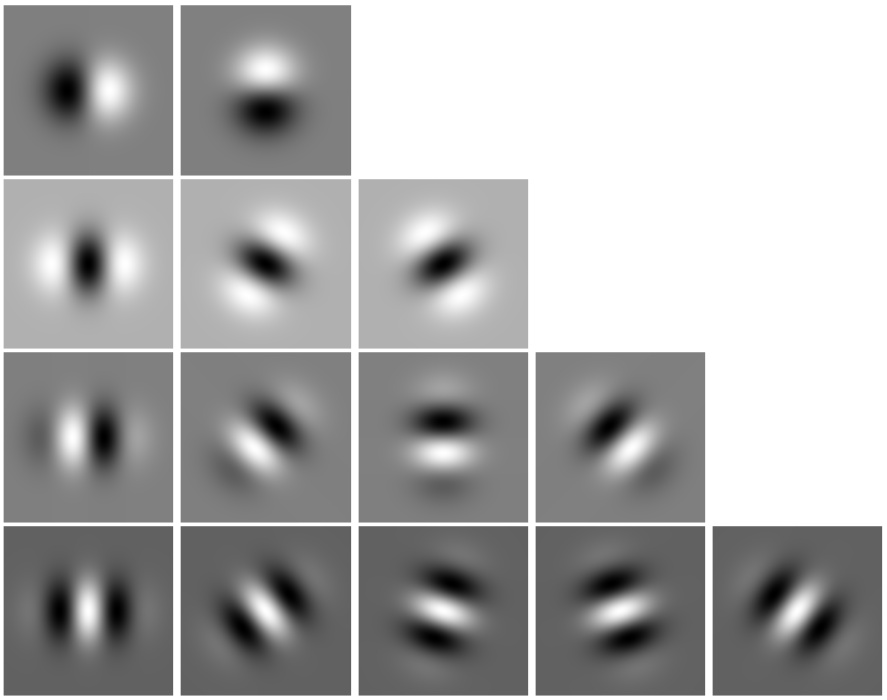

Figure 25. *This is the jet of partial derivatives up to (and including) order* 4*, shown as a complete basis. (In the leftmost column, one has* ∂_x*,* ∂_{xx}*,* ∂_{xxx} *and* ∂_{xxxx}*.) It involves four RF-profiles that occur in a number of directions or orientations (the cortex uses overcomplete, continuous bases). All relations discussed in this paper can be explored using this basis of image operators. Notice that these operators involve a scale. In applications, one would use a range of scales.*

differentiable! But the *observation is known* (indeed, all that is initially known about $\mathcal{I}$), but its derivatives can also be known—in the sense of being observed—because (at least formally)

$$\partial\mathcal{R} = \partial(g_\sigma \circ \mathcal{I}) = (\partial g_\sigma) \circ \mathcal{I} = g_\sigma \circ (\partial\mathcal{I}), \tag{53}$$

where g_σ is differentiable because it is an analytic function.[15] Neither $\partial\mathcal{I}$ nor $\partial\mathcal{R}$ exists: the former because $\mathcal{I}$ is not known and the latter because it is an observation and thus necessarily only approximately known. Only $(\partial g_\sigma) \circ \mathcal{I}$ makes sense as yet another *observation*.

In the same vein, all derivatives of $\mathcal{I}$ are observable *at some finite scale*. Collecting all partial derivatives up to some maximum order lets one observe a jet bundle as the representation of $\mathcal{I}$ (see Figure 25).

[15] The partial derivative of order n in the coordinate x for the scale σ has the one-dimensional receptive field profile

$$g_n(x,\sigma) = \frac{\partial^n}{\partial x^n}\left(\frac{\mathrm{e}^{-\frac{x^2}{2\sigma^2}}}{\sigma\sqrt{2\pi}}\right) = \left(\frac{-1}{\sigma\sqrt{2}}\right)^n H_n\left(\frac{x}{\sigma\sqrt{2}}\right)\frac{\mathrm{e}^{-\frac{x^2}{2\sigma^2}}}{\sigma\sqrt{2\pi}},$$

where H_n denote the Hermite polynomial of order n. The two-dimensional receptive fields $g_{nm}(x,y,\sigma)$ corresponding to mixed derivatives of order nm are simply products $g_n(x.\sigma)g_m(y,\sigma)$ of these because the Gaussian is (uniquely!) separable in Cartesian coordinates. In practice, such operators are implemented in the Fourier domain because one invariably needs the convolution with an image. In the brain, the receptive fields are implemented spatially though, there is no global Fourier transform in the cortex.

Images are given as brute facts. All their derivatives can be observed and thus represented at scales coarser than their pixelation. The kernels $\partial^n g_\sigma$ are best considered as implementations of derivative operators. They satisfy all properties of the derivatives in the infinitesimal domain [33, 44, 14, 58], but their use does not require assumptions about the smoothness of an image. Thus, scale-space theory allows all relations discussed in the text to be implemented *exactly* at some finite scale σ.

10.1. Local differential geometry and the brain. The model of local differential geometry discussed above may immediately double for a model of signal processing in the primary visual cortex V1. The well-known principal receptive field types [22] may readily be interpreted in terms of derivative operators [29, 30]. For instance, the well-known "line finder" discovered by Hubel and Wiesel [21] may naturally be interpreted as a second-order directional derivative operator.

The various differential invariants discussed in this paper are algebraic combinations of derivatives of various orders (for instance, see section 8). Any expression can immediately be compiled into a small neural network [34, 30, 31], probably largely on the dendritic level. This yields a principled model of local cortical processing. It is a formal, conceptual model, similar to objects like the "ideal gas" in physics [13]. There is no such a thing as an ideal gas. Yet any specific gas is first and for all understood as an instance of the ideal gas.

We propose that a productive way to deal with the complexity of the brain is in terms of conceptual models of the ideal gas type. The local differential geometry discussed here is meant as a likely candidate for such a model.

11. Conclusions. Numerous applications—image processing being a key example—involve differential geometry in graph spaces, that are singly isotropic spaces of the kind $\mathbb{E}^2 \times \mathbb{I}$, where the line $\mathbb{I}$ is isotropic. This involves formalisms that are quite distinct from the familiar differential geometry of curves and surfaces immersed in Euclidean three-space $\mathbb{E}^3$.

Almost without exception, the resulting formalism will be much less involved in the graph spaces than in Euclidean spaces. The reason is that the distance metric in Euclidean space is parabolic, whereas the angle metric is elliptic. In contradistinction, in graph spaces, both the distance and angle metrics are parabolic, leading to a full symmetry between planes and points in three-dimensional space. Many of the annoying *exceptions* that plague Euclidean geometry [63] simply fail to occur.[16]

The notion of graph spaces is not generally taught, with the unfortunate result that it is not at all uncommon to see Euclidean differential geometry being applied to images considered as surfaces. Although this often "works" in applications—no doubt the reason why the problem remains unidentified—this involves comparison of apples and oranges that is the mixing of incommensurate dimensions.

The notions of "curvedness" and "shape index" were originally introduced independently and essentially in an ad hoc manner [32]. Curvedness can be identified with the Casorati curvature in Euclidean differential geometry, but it was never accepted as a valid curvature measure because Casorati [10] used mainly informal arguments that left considerable room

[16] For this very reason, Strubecker [39] proposed that the teaching of geometry in primary schools should be based on singly isotropic rather than Euclidean geometry. Although he had a valid point, this proposition—as was to be expected—was never taken up, no doubt to the detriment of the kids.

for ambiguity. In the differential geometry of graph spaces, it can be shown that curvedness, shape index, and orientation occur as different partial aspects of a single entity. It is the notion of the Procrustes measure [27] *RMS deviation from planarity* that leads to this insight. This is interesting because of a statistical rather than differential geometric nature. This perspective effectively lifts the arbitrariness of the original definitions.

Curvedness and shape index, are widely used in applications in various fields of endeavor. Research as to the statistical structure of natural images has led to an interest in the marginal probability densities of curvedness, shape index, and orientation. The distributions are quite different for a natural distribution in shape space—say, an isotropic normal distribution centered at its origin (the planar point)—and distributions of shape over a Gaussian random surface. Moreover, it has been noticed that the empirical distributions often deviate from the theoretical ones [36, 43]. As it turns out, this is at least to some extent due to the constraint of isotropy. If the Gaussian random image is only slightly anisotropic, the shapes become concentrated on the ridges and ruts, most points become hyperbolic, and approximately umbilical points become very rare.

In natural images, the statistics are often anisotropic on local but approximately isotropic on global scales. Lack of appreciation of this fact has led to doubts as to the relevance of geometrical predictions, as empirical estimates of, for instance, the shape index marginal differed greatly from the formal expectations. The scale dependence of the degree of anisotropy is apparently an important property of "natural images."

Another interesting point that arises from the present discussion is the importance of the global structure of the relief. This implies that the statistics of features at a certain order may depend upon the fact that they occur in the context of an *image* (a planar distribution) [48, 3]. The shape index is an obvious example. Its statistics cannot be derived from the structure of the quadric shape space alone. It depends crucially on the fact that partial derivatives in an image are mutually correlated. This immediately leads to the peaking of the shape index probability density at the ridges and ruts. The present discussion generalizes that to the case of anisotropic Gaussian images.

Another implication of this are the mutual alignments of edge directions in the linear and cubic orders, the mutual alignments of ridge orientations in the quadric and quartic orders, and the very small contribution to the Procrustes metric of the cubic and quartic saddles. Such correlations are not specific to natural images as is sometimes suggested; they also occur for isotropic random Gaussian surfaces and are due to the correlation of partial derivatives in such fields, having nothing to do with image content as such.

Acknowledgments. The work was supported by the DFG Collaborative Research Center SFB TRR 135 headed by Karl Gegenfurtner (Justus-Liebig Universität, Giessen, Germany) and by the program by the Flemish Government (METH/14/02) awarded to Johan Wagemans. Jan Koenderink was supported by the Alexander von Humboldt Foundation.

REFERENCES

[1] V. Arnold, *Catastrophe Theory*, Springer, New York, 1992.
[2] T. Banchoff, T. Gaffney, and C. McCrory, *Cusps of Gauss mappings*, in Research Notes in Mathematics, no. 55, Pitman, London, 1982.

[3] M. V. Berry and J. H. Hannay, *Umbilic points on gaussian random surfaces*, J. Phys. A Math. Gen., 10 (1977).
[4] F. Bloch, *Nuclear induction*, Phys. Rev., 70 (1946), pp. 460–474.
[5] F. Bookstein, *Morphometric Tools for Landmark Data*, Cambridge University Press, Cambridge, 1991.
[6] M. Born and E. Wolf, *Principles of Optics*, 7th ed., Cambridge University Press, Cambridge, 1999.
[7] V. Borrelli, S. Jabrane, F. Lazarus, and B. Thibert, *Flat tori in three-dimensional space and convex integration*, Proc. Natl. Acad. Sci., 109 (2012), pp. 7218–7223.
[8] M. J. Boussinesq, *Sur une propiété remarquable des points où les lignes de plus grande pente d'une surface ont leurs plans osculateurs verticaux, est sur la différence qui existe généralement, à la surface de la terre, entre les lignes de faîte ou de thalweg et celles le long desquelles la pente du sol est un minimum*, C. R. Sèances Acad. Sci., 73 (1871), pp. 1368–1371.
[9] J. Canny, *A computational approach to edge detection*, IEEE Trans. Pattern Anal. Mach. Intell., 8 (1986), pp. 697–698.
[10] F. Casorati, *Mesure de la courbure des surfaces suivant l'idée commune. ses rapports avec les mesures de courbure gaussienne et moyenne*, Acta Math., 14 (1890), pp. 95–110.
[11] A. Cayley, *On contour and slope lines*, London Edinb. Dublin Philo. Mag. J. Sci., 18 (1859), pp. 264–268.
[12] M. B. de Champ, *Mémoire sur les lignes de faîte et thalweg que l'on est conduit à considérer en topographie*, J. Math., 3 (1877), pp. 99–114.
[13] R. Feynman, R. Leighton, and M. Sands, *The Feynman Lectures on Physics*, Vol. 1, Library of Congress Catalog Card No. 63-20717, 1964.
[14] L. M. J. Florack, *Image Structure*, Kluwer, Dordrecht, 1997.
[15] C. F. Gauss, *Disquisitiones generales circa superficies curvas*, Commentattones Soci. Regiae Sci. Gottingensis, 6 (1827), pp. 4–50.
[16] J. Gower and G. Dijksterhuis, *Procrustes Problems*, Oxford University Press, Oxford, 2004.
[17] A. Gray, *Modern Differential Geometry of Curves and Surfaces with Mathematica*, 2nd ed., CRC Press, Boca Raton, FL, 1997.
[18] L. Griffin and M. Lillholm, *Hypotheses for image features, icons and textons*, Int. J. Comput. Vis., 70 (2006), pp. 213–230.
[19] L. Griffin, M. Lillholm, and M. Nielsen, *Natural image profiles are most likely to be step edges*, Vis. Res., 44 (2004), pp. 407–421.
[20] L. D. Griffin, *The second order local-image-structure solid*, IEEE Trans. Pattern Anal. Mach. Intell., 29 (2007), pp. 1355–1366.
[21] D. Hubel and T. Wiesel, *Receptive fields and functional architecture of monkey striate cortex*, J. Physiol., 195 (1968), pp. 215–243.
[22] D. H. Hubel and T. N. Wiesel, *Brain and Visual Perception*, Oxford University Press, New York, 2005.
[23] R. W. G. Hunt, *The Reproduction of Colour*, 6th ed., Wiley, Chichester, 2004.
[24] I. L. Dryden and K. Mardia, *Statistical Shape Analysis*, Wiley, Chichester, 1998.
[25] M. C. Jordan, *Nouvelles observations sur les lignes de faîte et de thalweg*, C. R. Sèances Acad. Sci., 75 (1872), pp. 1023–1025.
[26] D. G. Kendall, *Shape manifolds, procrustean metrics, and complex projective spaces*, Bull. London Math. Soc., 16 (1984), pp. 81–121.
[27] D. G. Kendall, *A survey of the statistical theory of shape*, Stat. Sci., 4 (1989), pp. 87–99.
[28] Y. Kitagawa, *Embedded flat tori in the unit* 3*-sphere*, J. Math. Soc. Jpn., 47 (1995).
[29] J. Koenderink and A. van Doorn, *Receptive-field families*, Biol. Cybernet., 63 (1990), pp. 291–297.
[30] J. Koenderink and A. van Doorn, *Generic neighborhood operators*, IEEE Trans. Pattern Anal. Mach. Intell., 14 (1992), pp. 597–605.
[31] J. Koenderink and A. van Doorn, *Receptive field assembly specificity*, J. Vis. Comm. Image Represent., 3 (1992), pp. 1–12.
[32] J. Koenderink and A. van Doorn, *Surface shape and curvature scales*, Image Vis. Comput., 10 (1992), pp. 557–564.
[33] J. J. Koenderink, *The structure of images*, Biol. Cybernet., 50 (1984), pp. 363–370.
[34] J. J. Koenderink and W. Richards, *Two-dimensional curvature operators*, J. Opt. Soc. Amer. Ser. A, 5 (1988), pp. 1136–1141.

[35] J. J. Koenderink and A. J. van Doorn, *The structure of relief*, Adv. Imaging Electron Phys., 103 (1998), pp. 65–150.
[36] J. J. Koenderink and A. J. van Doorn, *Local structure of gaussian texture*, IEICE Trans. Inform. Syst., 86 (2003), pp. 1165–1171.
[37] K. Strubecker, *Differentialgeometrie des isotropen Raumes iii: Flächentheorie*, Math. Z., 48 (1942), pp. 369–427.
[38] K. Strubecker, *Überblick über die Differentialgeometrie des isotropen Raumes*, Jahresbr. Dtsch. Math.-Verein., 54 (1951), p. 34.
[39] K. Strubecker, *Geometrie in einer isotropen Ebene*, Math.-Naturwiss. Unterr., 15 (1962), pp. 297–306, 343–351, 385–394.
[40] G. Lemaître, *Quaternions et espace elliptique*, Acta Pontif. Acad. Sci., 12 (1948), pp. 57–80.
[41] M. Lillholm and L. Griffin, *Maximum likelihood metameres for local 2nd order structure of natural images*, in Scale Space and Variational Methods in Computer Vision, Vol. LNSC, Vol. 4485, Springer, New York, 2008, pp. 394–405.
[42] M. Lillholm and L. D. Griffin, *A category system on the shape index descriptor of local image structure induced by natural image statistics*, Perception, 35 (2006), pp. 48–49.
[43] M. Lillholm and L. D. Griffin, *Statistics and category systems for the shape index descriptor of local 2nd order natural image structure*, Image Vis. Comput., 27 (2009), pp. 771–781.
[44] T. Lindeberg, *Scale-Space Theory in Computer Vision*, Kluwer, Dordrecht, 1994.
[45] T. Lindeberg, *A computational theory of visual receptive fields*, Biolo. Cybernet., 107 (2013), pp. 589–635.
[46] T. Lindeberg, *Time-causal and time-recursive spatio-temporal receptive fields*, J. Math. Imaging Vis., 55 (2016), pp. 50–88.
[47] L. M. J. Florack, B. H. Romeny, J. J. Koenderink, and M. A. Viergever, *Cartesian differential invariants in scale-space*, J. Math. Imaging Vis., 3 (1993), pp. 327–348.
[48] M. S. Longuet-Higgins, *The statistical analysis of a random, moving surface*, Philos. Trans. Roy. Soci. London Ser. A, 249 (1957), pp. 321–387.
[49] M. Loog, *The jet metric*, in Scale Space and Variational Methods in Computer Vision, Lecture Notes in Comput. Sci. 4485, F. Sgallari, A. Murli, and N. Paragios, eds., Springer, New York, 2007, pp. 25–31.
[50] B. Markussen, K. Steenstrup Pedersen, and M. Loog, *Second order structure of scale-space measurements*, J. Math. Imaging Vis., 31 (2008), pp. 207–220.
[51] D. Marr and E. Hildreth, *Theory of edge detection*, Proce. Roy. Soci. London Ser. B, 207 (1980), pp. 187–217.
[52] J. C. Maxwell, *On hills and dales*, London Edinb. Dublin Philoso. Mag. J. Sci., 40 (1870), pp. 421–425.
[53] H. Poincaré, *Théorie mathematique de la lumière comparaison de cette théorie avec la théorie élastique*, Georges Carré, Paris, 1892.
[54] R. Rothe, *Zum problem des Talwegs*, Sitzungsber. Berliner Math. Gesellschaft, 14 (1915), pp. 51–69.
[55] H. Sachs, *Isotrope Geometrie des Raumes*, Vieweg, Berlin, 1990.
[56] D. Saint-Venant, *Surfaces à plus grande pente constituées sur des lignes courbes*, Bulle. Soc. Philomath. Paris, (1852), pp. 24–30.
[57] M. Spivak, *A Comprehensive Introduction to Differential Geometry*, Publish or Perish, Boston, MA, 1975.
[58] B. M. ter Haar Romeny, *Front-End Vision and Multi-Scale Image Analysis*, Springer, New York, 2004.
[59] L. van Vliet and P. Verbeek, *Curvature and bending energy in digitized 2d and 3d images*, in SCIA'93, Proceedings of the 8th Scandinavian Conference on Image Analysis, Tromso, Norway, 1993, pp. 1403–1410.
[60] H. Whitney, *The singularities of a smooth n-manifold in* (2n - 1)*-space*, Ann. Math., 45 (1944), pp. 247–293.
[61] H. Whitney, *On singularities of mappings of Euclidean spaces.* I. *Mappings of the plane into the plane*, Ann. Math. 2nd Ser., 62 (1955), pp. 374–410.
[62] I. M. Yaglom, *Complex Numbers in Geometry*, Academic Press, New York, 1968.
[63] I. M. Yaglom, *A Simple Non-Euclidean Geometry and Its Physical Basis: An Elementary Account of Galilean Geometry and the Galilean Principle of Relativity*, Abe Shenitzer, trans., Springer-Verlag, New York, 1979.