

The Second Order Local-Image-Structure Solid

Lewis D. Griffin

Abstract—Characterization of second order local image structure by a 6D vector (or jet) of Gaussian derivative measurements is considered. We consider the affect on jets of a group of transformations—affine intensity-scaling, image rotation and reflection, and their compositions—that preserve intrinsic image structure. We show how this group stratifies the jet space into a system of orbits. Considering individual orbits as points, a 3D orbifold is defined. We propose a norm on jet space which we use to induce a metric on the orbifold. The metric tensor shows that the orbifold is intrinsically curved. To allow visualization of the orbifold and numerical computation with it, we present a mildly-distorting but volume-preserving embedding of it into euclidean 3-space. We call the resulting shape, which is like a flattened lemon, the second order local-image-structure solid. As an example use of the solid, we compute the distribution of local structures in noise and natural images. For noise images, analytical results are possible and they agree with the empirical results. For natural images, an excess of locally 1D structure is found.

Index Terms—Scale space, image derivatives, feature analysis, noise, natural images.

1 INTRODUCTION

THIS paper is concerned with the analysis of local structure of 2D scalar functions, especially those, such as images, that are the result of physical measurement. The general approach to local structure used herein is characterization by derivatives; as compared to, for example, local Fourier analysis [1]. The focus of the paper is indicated in Fig. 1, which shows example image patches whose local structure is dominated by first order structure, by second order structure and, equally, by first and second order structure. While the varieties of pure first and pure second order local structure are already well-understood [2], [3], the case of mixed first and second is not. The aim of this paper is to derive a principled framework within which mixed first and second order structure can be expressed and investigated.

In the remainder of this section, we establish a mathematical framework for local image analysis (Section 1.1), discuss the idea of transformations that leave local structure unchanged (Section 1.2), and review previous work in this area (Section 1.3). In Section 2, we derive the structure of an orbifold that represents local structure, as measured by derivatives, but with the effect of a group of structure-preserving transformations factored out. The orbifold is locally 3D but intrinsically curved. In Section 3 we present an embedding of the orbifold as a bounded solid in 3D space that allows us to visualize its shape (look ahead to Fig. 4 for a preview). In Section 4, we present, as an example use of the orbifold and its embedding, an analysis of the distribution of local structures present in different types of noise and in natural images. An electronic supplement, in the form of a Mathematica [4] notebook contains all derivations, computations, and diagrams for the reader who wants additional detail.

- The author is with the Computer Science Department, University College London, Malet Place Engineering Building, London, WC1E 6BT UK. E-mail: L.Griffin@cs.ucl.ac.uk.

Manuscript received 25 Jan. 2006; revised 9 June 2006; accepted 20 Sept. 2006; published online 18 Jan. 2007.

Recommended for acceptance by J. Weickert.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0045-0106. Digital Object Identifier no. 10.1109/TPAMI.2007.1066.

1.1 Measurement of Local Image Structure

Mathematical ideas of “local structure” are intimately tied to the concept of derivative; formally defined by the limit $I'(0) = \lim_{\varepsilon \rightarrow 0} \varepsilon^{-1}(I(\varepsilon) - I(0))$ and, thus, inapplicable to functions, such as images, that are the result of measurement. Scale Space analysis [5], [6], [7], [8], [9], [10], [11], [12] proposes a two-step solution to the problem of operationalizing image derivative measurement. First, it defines a way of changing the inner scale (roughly the size of the smallest resolvable detail) of an image by convolution (denoted \otimes) of the image with Gaussian kernels. Gaussian kernels of scale $\sigma \in \mathbb{R}^+$ are defined as

$$G_{\sigma}(x) := \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}}, \quad G_{\sigma}(x, y) := G_{\sigma}(x)G_{\sigma}(y).$$

Because of the excellent localization of the Gaussian both in space and in frequency [13], the rescaling operation $I_{\sigma} = G_{\sigma} \otimes I$ can be performed easily and stably even if the raw image (I) is the result of physical measurement, and so discretely sampled [5]. The second step of the Scale Space approach to calculating image derivatives is via the property $(I_{\sigma})' = G'_{\sigma} \otimes I$, which means that the derivative of a rescaled image can be obtained by convolving the original image with a derivative of Gaussian (DtG). DtGs are defined as follows:

$$G_{\sigma}^{(u)}(x) := \frac{d^u}{dx^u} G_{\sigma}(x), \\ G_{\sigma}^{(u,v)}(x, y) := G_{\sigma}^{(u)}(x)G_{\sigma}^{(v)}(y), \quad u, v \in \mathbb{Z}^+.$$

Explicit equations for DtGs can be found using the following formula that expresses their form as the product of a 0th order Gaussian multiplied by a Hermite polynomial (H_u) [14]:

$$G_{\sigma}^{(u)}(x) := \left(\frac{-1}{\sigma\sqrt{2}}\right)^u H_u\left(\frac{x}{\sigma\sqrt{2}}\right) G_{\sigma}(x).$$

The Scale Space approach thus allows the computation of image derivatives of any order at any scale. If derivatives across the entire image are required the convolution formalism should be used; if derivatives are required at a single location only, then an inner product formalism ($\langle \cdot | \cdot \rangle$)




Fig. 1. Examples of different types of local structure.

is more convenient. For instance, we define image derivatives of scale σ , at the origin to be

$$c_{uv} := (-1)^{u+v} \langle G_{\sigma}^{(u,v)} | I \rangle := (-1)^{u+v} \int_{x,y \in \mathbb{R}} G_{\sigma}^{(u,v)}(x,y) I(x,y).$$

Note that the derivative measurements c_{uv} are dependent on the measurement scale σ , though to prevent cluttered equations we do not indicate this with a superscript. The inner-product formulation shows the similarity between the DtG approach and the operation of the simple cell neurons of area V1 of the mammalian cerebral cortex [15]. To avoid any incorrect presumptions, we note that the DtGs are not an orthonormal set; for example, $\langle G_{\sigma}^{(2,0)} | G_{\sigma}^{(0,2)} \rangle = (16\pi\sigma^6)^{-1}$.

In this paper, we are concerned with structure up to second order, the measurement of which requires a total of six DtG filters (Fig. 2). The vector of derivative measurements is referred to as a jet [16], [17], and we will later make use the perspective that regards such jets as being points in a jet space [18], [19], [20]. Formally, the second order jet at the origin is given by $\vec{j} := (c_{00} \ c_{10} \ c_{01} \ c_{20} \ c_{11} \ c_{02})^T$.

1.2 Image Similarity Transforms

In geometrical problems, one typically chooses some group of transformations and then studies structure that is invariant to them. The group (\mathbf{T}) that we will study here consists of the following and their compositions: (\mathbf{T}_1) translation of the image, (\mathbf{T}_2) rotation of the image, (\mathbf{T}_3) reflection of the image, (\mathbf{T}_4) addition of a constant intensity, and (\mathbf{T}_5) multiplication of intensities by a positive factor. We motivate these in turn.

1.2.1 Translation, Rotation, and Reflection (\mathbf{T}_1 , \mathbf{T}_2 , and \mathbf{T}_3)

In particular imaging systems, for example, those with space-variant resolution, prior knowledge may guide us to treat different image locations or orientations differently. However, for an uncommitted imaging system, we require that our analysis of local image structure be invariant to translation, rotation, reflection, or their combination. Translation is dealt with by choosing the point to be analyzed as the origin of our coordinate system.

Rotation is more difficult. Consider the effect on the jet of a rotation about the origin of angle θ . The 0th order term of the jet (c_{00}) is unaffected. The first order terms transform according to




Fig. 2. The derivatives of the 2-D Gaussian kernel (DtGs).

$$\begin{pmatrix} c_{10} \\ c_{01} \end{pmatrix} \rightarrow \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} c_{10} \\ c_{01} \end{pmatrix}$$

and the second order terms according to

$$\begin{pmatrix} c_{20} \\ c_{11} \\ c_{02} \end{pmatrix} \rightarrow \frac{1}{2} \begin{pmatrix} 1+c & 2s & 1-c \\ -s & 2c & s \\ 1-c & -2s & 1+c \end{pmatrix} \begin{pmatrix} c_{20} \\ c_{11} \\ c_{02} \end{pmatrix},$$

where $s \equiv \sin 2\theta$, $c \equiv \cos 2\theta$. So, first order structure requires a full 2π rotation to return to the starting values, while second order structure returns after a rotation by π .

For reflection, consider the effect on the jet of a reflection in the line $y = x$. Again, the 0th order term of the jet is unaffected. The first order terms transform according to: $c_{10} \leftrightarrow c_{01}$; the second order terms according to: $c_{20} \leftrightarrow c_{02}$.

1.2.2 Affine Scaling of Intensity (\mathbf{T}_4 and \mathbf{T}_5)

In this paper, we assume the most general model of intensity values, that they are unconstrained real numbers. Different analyses may be appropriate if it can be assumed that intensities are nonnegative and/or not exceeding a maximum value [21], [22]. We require that our analysis of image structure is unaffected by adding a constant value (α) to all intensity values. Such a change can come about, for example, by there being a veiling glare over a scene or by a change in the dark current of a CCD. Only the 0th order term of the jet is effected and it simply transforms as $c_{00} \rightarrow \alpha + c_{00}$.

We also require that image structure be invariant to the multiplication of all image intensities by a nonzero, positive factor (ρ). Such a change can come about by adjusting the overall level of illumination of a scene, exposure time, or the gain of a CCD readout mechanism. The effect on the jet is simply that all terms are multiplied by the factor, i.e., $\vec{j} \rightarrow \rho \vec{j}$. Note that we do not require invariance to multiplication of intensities by a negative factor. This is because we cannot conceive of a physical change (that we would wish to ignore!) that could cause this. This illustrates that the transformation group that we choose to be invariant to does not arise inevitably from some mathematical argument, but rather is an attempt to model the changes that arise by altering aspects of the imaging setup rather than the scene itself.

1.3 Previous Work

The literature on image feature analysis is now substantial, so inevitably we will only mention work that shares a similarity in approach with that presented here. Related work is further considered in the discussion section.

First, we mention approaches that assume that images are like Morse functions—that is to say, infinitely-differentiable functions containing only generic structure that is

resilient to infinitesimal perturbations. The possible local structures of such functions are documented in [3], [23] and, further, in [24]. In brief, Morse functions may be locally slope-like or one of three critical point types—maximum, minimum, and saddle point. Slope points may be further analyzed by considering the curvature of the isophote and flowline (curve of steepest descent) through the point [25], [26]. These curvatures are intrinsic quantities as they are invariant to rotations of the image and monotonic transformations of the intensity.

A second type of approach is to focus, not on itemizing the possibilities for individual points but, rather, on the distribution of types of point that occur in a given class of images [27]. The seminal reference for this is [28], in which expressions for the joint distribution of derivative measurements for noise images with a fixed Fourier amplitude spectrum but random phase spectrum are derived. More recently, the joint distribution of derivatives (up to third order) of noise and natural images has been determined [29], [30], [31], [19], as has the second order (only) derivatives of noise and natural images [2]. The frequency of occurrence of image patches has also been investigated using a 3×3 pixel representation, modulo intensity scaling and rotation, rather than a derivative one [32].

The final family of approaches relevant to this work are those that consider the distribution of local image patches conditioned on the local derivative structure [33], [34], [35]. These approaches show that the underlying structure of natural images is typically piece-wise constant even though the derivatives of the blurred intensity are of course continuous.

2 INTRINSIC GEOMETRY OF THE LOCAL STRUCTURE ORBIFOLD

In this section, we derive the core results of the paper. The first step (Section 2.1) is to choose a norm, and thus a metric, for jet space. Next (Section 2.2), we present a nonlinear reparameterization of jet space that separates out coordinates that change and not do change under the action of the transformation group. In Section 2.3, we show that the unchanging parameters are a coordinate system for the orbifold which results from factoring the jet space by the transformation group. Finally, in Section 2.4, we bring the preceding sections together to derive a metric on the orbifold.

2.1 A Norm on Jet Space

A norm ($\|\cdot\|$) on jet space is a function that associates with each element a positive number that expresses, in some sense, the magnitude of the structure measured by the jet. This norm will play a fundamental role in our analysis, but we are unaware of any preexisting work that makes a convincing case for what the norm should be; so, in the following, we argue for a particular choice. We start by stating some requirements that characterize the jet space norm that we wish to specify:

- R₁.** The norm should satisfy the standard axioms for a seminorm.
- R₂.** The norm should be invariant to translation, rotation, or reflection of the image domain.
- R₃.** The norm should be unaffected by adding a constant to image intensities.

To generate some ideas for the definition of the jet space norm, we first consider some standard norms on functions rather than jets. The most commonly used norm for functions is the L^2 -norm written, using the inner-product formalism, $\|I\|_2 := \langle I|I \rangle^{\frac{1}{2}}$. A related seminorm, insensitive to the addition of constant values to the function (as required by **R₃**), is: $\|I\|_{\bar{2}} := (\langle 1|I^2 \rangle - \langle 1|I \rangle^2)^{\frac{1}{2}}$. Another common variant for function norms, is to evaluate the function not over the entire domain, but only over some subregion of the domain. For full generality, a weighting function (w) can be used to control how much different parts of the domain are relevant to the norm, i.e., $\|I\|_{\bar{w}} := (\langle w|I^2 \rangle - \langle w|I \rangle^2)^{\frac{1}{2}}$. When the weighting function is a Gaussian, this norm is a natural choice for use in the Scale Space framework. We define the Scale Space norm of an image (I), evaluated at the origin, at scale σ , to be: $\|I\|_{\sigma} := (\langle G_{\sigma}|I^2 \rangle - \langle G_{\sigma}|I \rangle^2)^{\frac{1}{2}}$, i.e., Gaussian-windowed local variance.

The Scale Space norm cannot be applied to jets, as a jet does not specify an image that can be windowed, etc. Jets, in fact, specify metamery classes of all the images that measure to that jet [36]. There is no upper bound on the Scale Space norms of the members of a metamery class, but there is a unique norm minimizer. Our plan then is to define the jet space norm as the minimum of the Scale Space norms of the elements of the metamery class defined by the jet. Below, we prove a theorem on the form of this norm-minimizer.

Theorem. *Of all the functions that measure to a particular jet \vec{j} (of order N), the one that uniquely minimizes the Scale Space norm is: $P_{\vec{j}} = c_{00} + \sum_{1 \leq u+v \leq N} c_{uv} \frac{(2^{-\frac{1}{2}}\sigma)^{u+v}}{u!v!} H_u(\frac{x}{\sigma\sqrt{2}}) H_v(\frac{y}{\sigma\sqrt{2}})$.*

Proof. Let $B : \mathbb{R}^2 \rightarrow \mathbb{R}$ be the unknown member of the metamery class that minimizes the Scale Space norm. Formally, B is defined by the constraints

- (i) $(-1)^{u+v} \langle G_{\sigma}^{(u,v)} | B \rangle = c_{uv}$ for all $0 \leq u + v \leq N$, and
- (ii) it is a minimum of $\|\cdot\|_{\sigma}$.

We combine the constraints using Lagrangian multipliers

$$C(B) = \frac{1}{2} \|B\|_{\sigma}^2 + \sum_{0 \leq u+v \leq N} \lambda_{uv} \left((-1)^{u+v} \langle G_{\sigma}^{(u,v)} | B \rangle - c_{uv} \right).$$

Then, we compute the variation of C , and simplify using (by constraint (i)) that $\langle G_{\sigma} | B \rangle = c_{00}$

$$\begin{aligned} & \frac{\partial}{\partial \varepsilon} C(B + \varepsilon \delta B) \Big|_{\varepsilon=0} \\ &= \langle G_{\sigma} | B \delta B \rangle - \langle G_{\sigma} | B \rangle \langle G_{\sigma} | \delta B \rangle \\ &+ \sum_{0 \leq u+v \leq N} \lambda_{uv} (-1)^{u+v} \langle G_{\sigma}^{(u,v)} | \delta B \rangle \\ &= \left\langle G_{\sigma} B - G_{\sigma} \langle G_{\sigma} | B \rangle + \sum_{0 \leq u+v \leq N} \lambda_{uv} (-1)^{u+v} G_{\sigma}^{(u,v)} \Big| \delta B \right\rangle \\ &= \left\langle G_{\sigma} (B - c_{00}) + \sum_{0 \leq u+v \leq N} \lambda_{uv} (-1)^{u+v} G_{\sigma}^{(u,v)} \Big| \delta B \right\rangle. \end{aligned}$$

Constraint (ii) implies that the left-hand side of the inner product in the expression for the variation is identically zero. So, after rearrangement

$$B = c_{00} - \sum_{0 \leq u+v \leq N} \lambda_{uv} (-1)^{u+v} G_{\sigma}^{-1} G_{\sigma}^{(u,v)},$$

which, using the formula for DtGs in terms of Hermite polynomials (Section 1.1), can be further simplified to

$$B = c_{00} - \sum_{0 \leq u+v \leq N} \lambda_{uv} \left(\frac{1}{\sigma\sqrt{2}} \right)^{u+v} H_u \left(\frac{x}{\sigma\sqrt{2}} \right) H_v \left(\frac{y}{\sigma\sqrt{2}} \right).$$

Finally, using the fact that Hermite polynomials are orthogonal with respect to a Gaussian weighting function, we can solve to find values of the Lagrangian multipliers such that constraint (i) is satisfied, resulting in

$$P_{\vec{j}} = B = c_{00} + \sum_{1 \leq u+v \leq N} c_{uv} \frac{\left(2^{-\frac{1}{2}} \sigma \right)^{u+v}}{u!v!} H_u \left(\frac{x}{\sigma\sqrt{2}} \right) H_v \left(\frac{y}{\sigma\sqrt{2}} \right). \quad \square$$

Having determined $P_{\vec{j}}$, the element of the metamery class of \vec{j} that minimizes the Scale Space norm, we can now compute that its Scale Space norm is

$$\|P_{\vec{j}}\|_{\sigma} = \left(\sum_{1 \leq u+v \leq N} \frac{(\sigma^{u+v} c_{uv})^2}{u!v!} \right)^{\frac{1}{2}}.$$

This minimal Scale Space norm is what we define the norm of the jet to be, i.e., $\|\vec{j}\| := \|P_{\vec{j}}\|_{\sigma}$. The norm induces a metric in the standard way, i.e., $d(\vec{j}, \vec{k}) = \|\vec{j} - \vec{k}\|$.

So, to conclude this section by being specific, in the case of the second order jet, the metamery class norm-minimizer is the polynomial

$$P_{\vec{j}}(x, y) := \left(c_{00} - \frac{1}{2} \sigma^2 (c_{20} + c_{02}) \right) + (c_{10}x + c_{01}y) + \frac{1}{2} (c_{20}x^2 + 2c_{11}xy + c_{02}y^2)$$

and the Gaussian-windowed, mean-centred L^2 -norm of this polynomial and, thus, by definition the norm of the second order jet is

$$\|\vec{j}\| := \left(\sigma^2 (c_{10}^2 + c_{01}^2) + \frac{1}{2} \sigma^4 (c_{20}^2 + 2c_{11}^2 + c_{02}^2) \right)^{\frac{1}{2}}.$$

This norm satisfies \mathbf{R}_1 - \mathbf{R}_3 as required and by its derivation measures "the magnitude of the structure measured by the jet" as desired.

2.2 A Reparameterization of the Jet

Our aim is to discover the structure of jet space modulo the transformation group T . This is difficult when using the default coordinate system, which is based on individual derivative measurements, as each of the coordinates are affected by some of the transformations. What is needed is a reparameterization of jet space such that only some of the parameters are affected by the transformations while the rest remain invariant. There are many reparameterizations that thus qualify, but not all make the computations that we wish to perform using it easy. We give below a reparameterization that does make these computations tractable. In the discussion (Section 5), we will explain how the major results of the paper are independent of the reparameterization used.

The roots of the reparameterization that we have devised can be seen in the following rearrangement of the expression for the second order jet space norm:

$$\|\vec{j}\| = \left(\sigma^2 (c_{10}^2 + c_{01}^2) + \frac{1}{4} \sigma^4 (c_{20} + c_{02})^2 + \frac{1}{4} \sigma^4 \left((c_{20} - c_{02})^2 + 4c_{11}^2 \right) \right)^{\frac{1}{2}}.$$

The first bracketed component measures the first order structure, which is necessarily anisotropic, the second component measures the isotropic part of the second order structure, and the third the anisotropic component of second order structure. The magnitudes of each component are invariant to rotation and reflection of the image domain, so parameters built from these components will likewise be invariant. In particular, we will use one parameter (b) that measures the balance between the first and third terms, and one (l) that measures the balance of the second term relative to the sum of the first and third. The third (a) measures the angle between the two anisotropic components (i.e., the first and third above). Again, we note that these are not the only choices possible but they are the only combinations with which we have been able to complete the calculations we present later. This is discussed further in Section 5.

We will write $\vec{k} = (z \ n \ \theta \ m \ l \ b \ a)^T$ for a vector composed of the new parameters and $\Omega = \mathbb{R} \times \mathbb{R}^+ \times (-\pi, \pi] \times \{-1, 1\} \times [-\pi/2, \pi/2] \times [0, \pi/2] \times [0, \pi/2]$ for the domain of \vec{k} . The mapping (\vec{r}) from the derivative representation \vec{j} to the reparameterization is given by

$$\vec{k} = \begin{pmatrix} z \\ n \\ \theta \\ m \\ l \\ b \\ a \end{pmatrix} = \vec{r}(\vec{j}) = \begin{pmatrix} c_{00} \\ \|\vec{j}\| \\ \tan^{-1} c_{01}/c_{10} \\ \operatorname{sgn}^+ \left(\tan^{-1} \left(2 \frac{(c_{01}^2 - c_{10}^2)c_{11} + c_{10}c_{01}(c_{20} - c_{02})}{(c_{01}^2 - c_{10}^2)(c_{02} - c_{20}) + 4c_{10}c_{01}c_{11}} \right) \right) \\ \tan^{-1} \left(\frac{\sigma(c_{20} + c_{02})}{\sqrt{4(c_{10}^2 + c_{01}^2) + \sigma^2((c_{20} - c_{02})^2 + 4c_{11}^2)}} \right) \\ \tan^{-1} \left(\sigma \sqrt{\frac{(c_{20} - c_{02})^2 + 4c_{11}^2}{4(c_{10}^2 + c_{01}^2)}} \right) \\ \frac{1}{2} \left| \tan^{-1} \left(2 \frac{(c_{01}^2 - c_{10}^2)c_{11} + c_{10}c_{01}(c_{20} - c_{02})}{(c_{01}^2 - c_{10}^2)(c_{02} - c_{20}) + 4c_{10}c_{01}c_{11}} \right) \right| \end{pmatrix}.$$

The values of the parameters on a typical natural image are shown in Fig. 3. As the figure shows, they are not particularly intuitive. Parameter $z \in \mathbb{R}$ is just a renaming of the 0th order term from the jet. Parameter $n \in \mathbb{R}^+$ is the jet space norm defined in Section 2.2. Parameter $\theta \in (-\pi, \pi]$ gives the direction of the gradient vector. Parameter $m \in \{-1, 1\}$ specifies whether the acute angle from the gradient vector to the orientation of most positive second derivative is clockwise or counter clockwise. Note that the function sgn^+ , which give the value 1 for an argument of 0 but is otherwise like a regular sgn function, is used to ensure that $m \in \{-1, 1\}$. Parameter $l \in [-\pi/2, \pi/2]$ measures the balance between the

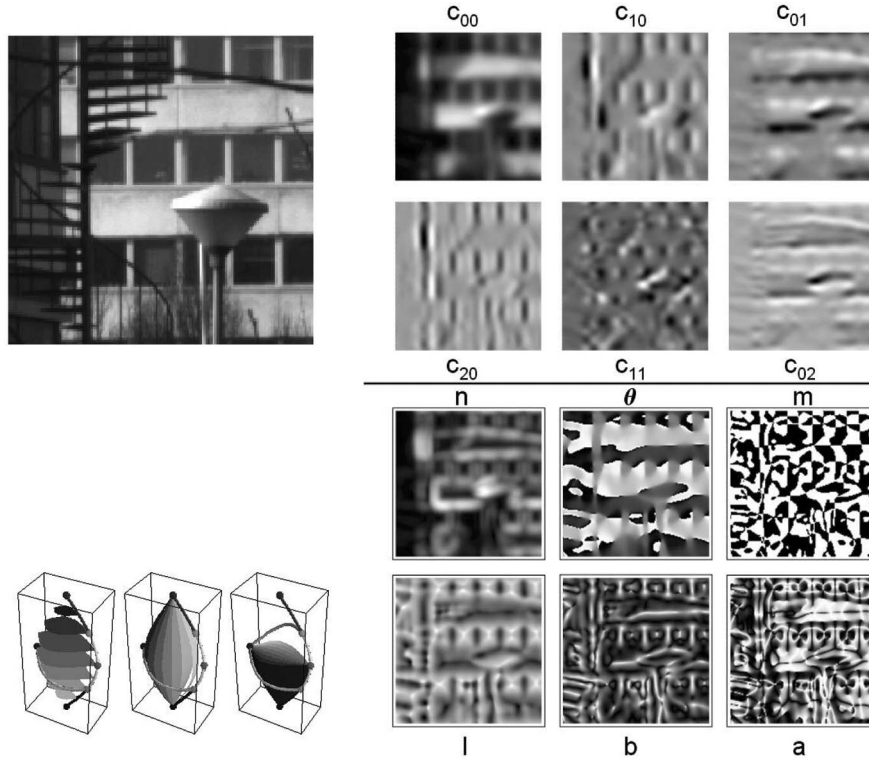


Fig. 3. The natural image patch on the left is 145×145 pixels. The panels to the right are all calculated at a scale of $\sigma = 4$. At the top, the components of the second order jet are shown. At the bottom, the parameters of the reparameterization of jet space are shown, other than z which is the same as c_{00} . Computation of Gaussian derivatives was performed in Mathematica using the Front-End Vision add-on [37]. The three plots at the lower left are not discussed until Section 5—they show, left-to-right, surfaces of constant l , b , and a over the solid.

isotropic and anisotropic parts of the local structure. Parameter $b \in [0, \pi/2]$ measures the balance between the first and second order contributions to anisotropic local structure. Parameter $a \in [0, \pi/2]$ measures the angle between first and second order anisotropic components.

The mapping in the other direction, from Ω the domain of the new parameters, to derivative values is given by the function \vec{q}

$$\vec{q}(\vec{k}) = \begin{pmatrix} c_{00} \\ c_{10} \\ c_{01} \\ c_{20} \\ c_{11} \\ c_{02} \end{pmatrix} = \vec{j} = \begin{pmatrix} z \\ \sigma^{-1}n \cos l \cos b \cos \theta \\ \sigma^{-1}n \cos l \cos b \sin \theta \\ \sigma^{-2}n(\sin l + \cos l \sin b \cos(2(\theta - ma))) \\ \sigma^{-2}n \cos l \sin b \sin(2(\theta - ma)) \\ \sigma^{-2}n(\sin l - \cos l \sin b \cos(2(\theta - ma))) \end{pmatrix}$$

\vec{q} is one-to-one in the interior of Ω , but many-to-one along parts of the boundary. In detail,

1. if $n = 0$, then θ, m, l, b, a are irrelevant,
2. if $|l| = \pi/2$, then m, b, a are irrelevant,
3. if $b = 0$, then m, a are irrelevant, and
4. if $a \in \{0, \pi/2\}$, then m is irrelevant.

In the remainder, when a parameter is irrelevant this will be denoted with an underscore (-).

2.3 The Orbit Structure of Jet Space

If we take an arbitrary point in jet space, and apply group T to it, we obtain a subspace of jet space known as the orbit of

the point. Because the group is closed under composition of transformations, each orbit is the orbit of any point within it, not just the point that was chosen originally. Orbits do not intersect, so the complete system of orbits stratifies jet space. In the problem at hand, there are seven kinds of orbit which we now list:

- O₁.** These are generated by the most generic type of jet, i.e., those for which $n > 0 \wedge |l| < \pi/2 \wedge b \in (0, \pi/2) \wedge a \in (0, \pi/2)$. Such jets give rise to 3D orbits because they, being fully generic, are nontrivially changed by the three transformations $T_2, T_4,$ and T_5 . The orbits are the image of $(\mathbb{R} \ \mathbb{R}^+ \ (-\pi, \pi] \ \{-1, 1\} \ l \ b \ a)^T$ under the mapping \vec{q} . They consist of a disjoint pair of individually connected 3D submanifolds of jet space—one component from $m = -1$ and the other from $m = 1$. There is a 3D family of this type of orbit indexed by the parameters $l, b,$ and a .
- O₂.** These are generated by points for which $n > 0 \wedge |l| < \pi/2 \wedge b \in (0, \pi/2) \wedge a = 0$. The orbits of such points are images of $(\mathbb{R} \ \mathbb{R}^+ \ [-\pi, \pi] \ - \ l \ b \ 0)^T$ under the mapping \vec{q} . These images each consist of a connected 3D submanifold of jet space. There is a 2D family of this type of orbit indexed by the parameters l, b .
- O₃.** Like **O₂** but with $a = \pi/2$.
- O₄.** These are generated by points for which $n > 0 \wedge b = 0 \wedge |l| < \pi/2$. The orbits of such points are the image of $(\mathbb{R} \ \mathbb{R}^+ \ [-\pi, \pi] \ - \ l \ 0 \ -)^T$ under the mapping \vec{q} . These images consist of a single connected 3D submanifold

of jet space. There is a 1D family of such orbits indexed by the parameter l .

O₅. Like **O₄** but with $b = \pi/2$.

O₆. These are generated by points for which $n > 0 \wedge |l| = \pi/2$. There are exactly two orbits of this type: The images of $(\mathbb{R} \ \mathbb{R}^+ \ - \ - \ \pi/2 \ - \ -)^T$ and $(\mathbb{R} \ \mathbb{R}^+ \ - \ - \ -\pi/2 \ - \ -)^T$ under the mapping \vec{q} . Whereas orbits of type **O₁**-**O₆** are 3D submanifolds of jet space, both orbits of this type are connected 2D submanifolds of jet space. The lower dimensionality is because the jets are unaffected by rotation (**T₂**).

O₇. This type of orbit is generated by points for which $n = 0$. There is exactly one of these orbits: The image of $(\mathbb{R} \ 0 \ - \ - \ - \ -)^T$ under the mapping \vec{q} . This image is a connected 1D submanifold of jet space.

The space of orbits is known as an orbifold, which is a mathematical structure like a manifold but allowed to have a few corners, cone points and punctures where the local topology is unlike \mathbb{R}^n ([38, chapter 13]). The manifold part of the orbifold is composed from the **O₁**, **O₂**, and **O₃** points, the rest is nonmanifold like. Informally, the orbifold has the following structure: Take a cross-product of three open intervals (**O₁**); onto each of a pair of opposite faces glue the product of two open intervals (**O₂** and **O₃**); for each of another pair of opposite faces collapse the face down into a line, bringing the **O₂** and **O₃** components together, and glue to an open interval (**O₄** and **O₅**); collapse the remaining two faces down into points (bringing along the **O₂**-**O₅** components) and glue single points (**O₆**). Finally, there is a single point (**O₇**) that is adjacent to every other point of the orbifold. This topological structure is easier to grasp once a metric structure has been imposed as we will do in the next section.

2.4 Deriving the Metric Tensor

We can use the metric on jet space to induce a metric on the orbifold. The obvious way to do this is to define the distance between two orbits as their closest distance in jet space. However, this approach fails to define a metric of use or interest. This is because transformation **T₅** ensures that all orbits contain jets with norms arbitrarily close to zero and so the distance between any pair of orbits according to this approach would be zero. To circumvent this, we instead focus on the spherical subspace of jet space that consists of jets with unit norm.

The unit-norm spherical subspace intersects all orbits apart from the single **O₇** orbit, which corresponds to vanishing first and second order structure. Setting aside this exceptional point, we can then define the distance between two orbits as the length of the shortest path, on the unit-norm sphere, that starts on one orbit and ends on the other. In the remainder of this section, we use this definition to determine the metric tensor of the orbifold. This metric tensor describes the distances between infinitesimally close orbits. In theory, the geodesic structure of the orbifold could be determined from the metric tensor and, so, a general formula for the distance between arbitrary orbits could be determined, but we do not do this in this paper.

The norm we derived in Section 2.1 gives rise to an inner product structure on the jet space. Let $\vec{a} = (a_{00} \ a_{10} \ a_{01} \ a_{20} \ a_{11} \ a_{02})^T$ and $\vec{b} = (b_{00} \ b_{10} \ b_{01} \ b_{20} \ b_{11} \ b_{02})^T$

be two jets expressed in the derivative coordinate system. The inner product between the two jets is

$$\vec{a} \cdot \vec{b} := \sigma^2(a_{10} b_{10} + a_{01} b_{01}) + \frac{1}{2}\sigma^4(a_{20} b_{20} + 2a_{11} b_{11} + a_{02} b_{02}).$$

Suppose now that both jets are on the unit-norm sphere, i.e., $\|\vec{a}\| = \|\vec{b}\| = 1$. The distance-squared between them, measured along the surface of the unit norm sphere, is thus $d^2(\vec{a}, \vec{b}) = (\cos^{-1}(\vec{a} \cdot \vec{b}))^2$. If the two jets are sufficiently close, then this distance-squared is approximately $d^2(\vec{a}, \vec{b}) \approx 2(1 - \vec{a} \cdot \vec{b})$; this will be used below to calculate the distance between a jet and a perturbation of it.

Consider a pair of nearby orbits of the general **O₁** type, indexed by the parameters $\langle l, b, a \rangle$ and $\langle l + \lambda, b + \beta, a + \alpha \rangle$. The intersections of these orbits with the unit-norm sphere are the sets $\vec{q}(\mathbb{R}, 1, (-\pi, \pi], \{-1, 1\}, l, b, a)$ and $\vec{q}(\mathbb{R}, 1, (-\pi, \pi], \{-1, 1\}, l + \lambda, b + \beta, a + \alpha)$, respectively. Consider the point $\vec{f}(l, b, a) = \vec{q}(0, 1, 0, 1, l, b, a)$ from the first intersection and the family of points $\vec{g}(l, b, a; \lambda, \beta, \alpha; \theta) = \vec{q}(0, 1, \theta, 1, l + \lambda, b + \beta, a + \alpha)$ from the second. The distance-squared, along the surface of the unit-norm sphere, between \vec{f} and \vec{g} is denoted: $h(l, b, a; \lambda, \beta, \alpha; \theta) = d^2(\vec{f}(l, b, a), \vec{g}(l, b, a; \lambda, \beta, \alpha; \theta))$.

Now, we define functions 1) $\omega(l, b, a; \lambda, \beta, \alpha)$ that selects the value of θ that minimizes h (i.e., $\frac{\partial h}{\partial \theta}|_{\theta=\omega} = 0$, $\frac{\partial^2 h}{\partial \theta^2}|_{\theta=\omega} > 0$) and 2) t that gives the value of h at this minimum, i.e., $t(l, b, a; \lambda, \beta, \alpha) = h(l, b, a; \lambda, \beta, \alpha; \omega(l, b, a; \lambda, \beta, \alpha))$.

To determine the metric tensor of the orbifold, we calculate the second derivatives of t with respect to λ, β, α , evaluate these at $\lambda = \beta = \alpha = 0$, and simplify. We have used the symbolic mathematics functionality of Mathematica to do this (see electronic supplement), making use of

1. all the definitions given above,
2. the definition of \vec{q} ,
3. the distance-squared approximation for nearby jets,
4. that by definition $\omega(l, b, a; 0, 0, 0) = 0$, and
5. derivatives of ω obtained using the implicit function theorem.

At the end of the simplification one obtains

$$\begin{aligned} \frac{\partial^2 t}{\partial \lambda^2} \Big|_{\lambda=\beta=\alpha=0} &= 2, & \frac{\partial^2 t}{\partial \beta^2} \Big|_{\lambda=\beta=\alpha=0} &= 2 \cos^2 l, \\ \frac{\partial^2 t}{\partial \alpha^2} \Big|_{\lambda=\beta=\alpha=0} &= \frac{4 \cos^2 l \sin^2 2b}{5 - 3 \cos 2b} \\ \frac{\partial^2 t}{\partial \lambda \partial \beta} \Big|_{\lambda=\beta=\alpha=0} &= \frac{\partial^2 t}{\partial \lambda \partial \alpha} \Big|_{\lambda=\beta=\alpha=0} = \frac{\partial^2 t}{\partial \beta \partial \alpha} \Big|_{\lambda=\beta=\alpha=0} = 0. \end{aligned}$$

So, the metric tensor, with respect to the indexing of orbits by $\langle l, b, a \rangle$, is

$$g = \text{diag} \left(1, \cos^2 l, \frac{2 \cos^2 l \sin^2 2b}{5 - 3 \cos 2b} \right).$$

Or, written as a line element

$$ds^2 = dl^2 + \cos^2 l \left(db^2 + \frac{2 \sin^2 2b}{5 - 3 \cos 2b} da^2 \right).$$




Fig. 4. Shows the second-order local-image-structure solid, which is the embedding of the orbifold into \mathbb{R}^3 using the mapping $\vec{\chi}$. The four surface-rendered panels show different views of the solid. The density plots show examples of local structure, with the connectors showing where they map to in the solid; the local structure at bottom-right, without a connector, corresponds to the center of the solid. The shape of the solid is roughly that of a lemon, flattened perpendicular to the main axis of symmetry to produce two sharp edges. Notice though that the two sharp edges are not the same shape (right-most view). The sharp vertices of the solid at its top and bottom correspond to umbilic points. The sharp colored edge marks the locus of local forms that have vanishing first order structure. The locus runs from the maximum umbilic point (blue) at top, through to a ridge shape (orange), a balanced saddle [3] (brown), a rut (green), to the minimum umbilic point (blue) at bottom. The purple point on the other sharp edge corresponds to vanishing second order structure. The pink arc, which travels over the smooth faces of the solid, marks local structure which is effectively 1D. The arc is on the surface of the solid, because the surface corresponds to local structure where the angular first and second order components of the local structure are at 0° or 90° , which is a necessary condition for local structure to be 1D. The plots were rendered using Mathematica and an add-on package for the tubular curves.

The argument above determined the metric tensor for an orbit of the most general type \mathbf{O}_1 , i.e., $\langle l, b, a \rangle \in (-\pi/2, \pi/2) \times (0, \pi/2) \times (0, \pi/2)$. If we assume (as seems reasonable) that the metric tensor is continuous over the entire orbifold, we can extend the applicability of the derived equations for the metric tensor to orbits of \mathbf{O}_2 - \mathbf{O}_6 , i.e., $\langle l, b, a \rangle \in [-\pi/2, \pi/2] \times [0, \pi/2] \times [0, \pi/2]$. The definition does not extend to the single orbit of type \mathbf{O}_7 , as this orbit is disconnected from the rest of the orbifold.

3 AN EMBEDDING INTO EUCLIDEAN SPACE

In Section 2, we developed a coordinate system $\langle l, b, a \rangle$ for the orbifold and determined the metric tensor of the space relative to that coordinate system. To further understand the intrinsic shape of the orbifold, we can calculate the scalar curvature of the space. This quantity is computed from the metric tensor but yields a value which is independent of the coordinate system (like the Gaussian curvature of surfaces) [39]. We compute (using Mathematica and a package devoted to general relativity calculations) that

$$\kappa_{scalar} = 6 + \frac{96 \sec^2 l}{(5 - 3 \cos 2b)^2}.$$

We observe that the scalar curvature varies in value dependent on l and b . The minimum of the curvature is $\kappa_{scalar} = \frac{15}{2}$ at $l = 0, b = \pi/2$, while the maximum is $+\infty$ at $|l| = \pi/2$.

The pattern of scalar curvature shows that the orbifold does not have the intrinsic geometry of some region of euclidean 3-space. This is inconvenient as it prevents the simplest approaches to visualizing its shape. Instead, we are

forced to find an embedding [40] of the orbifold into euclidean 3-space. Fortunately, since the $\langle l, b, a \rangle$ coordinate system led to a diagonal metric tensor, finding an embedding is not difficult. For instance,

$$\vec{\chi} \left(\begin{pmatrix} l \\ b \\ a \end{pmatrix} \right) = \begin{pmatrix} (b - \frac{\pi}{4}) \cos l \\ (a - \frac{\pi}{4}) \sqrt{2} \frac{\cos l \sin 2b}{\sqrt{5-3 \cos 2b}} \\ -l \end{pmatrix}_{xyz}.$$

We show this embedding of the orbifold into \mathbb{R}^3 in Fig. 4, which shows that the embedded orbifold is a convex solid, shaped like a flattened lemon. We refer to it as the second-order local-image-structure solid (or solid for short). The interior of the solid corresponds to orbits of type \mathbf{O}_1 , the front and back surfaces to \mathbf{O}_2 and \mathbf{O}_3 , the sharp edge with the purple point to \mathbf{O}_4 , the other sharp edge to \mathbf{O}_5 , and the sharp vertices to \mathbf{O}_6 .

The embedding is not isometric as this is not possible, but it has been constructed so that it preserves intrinsic volumes. To demonstrate this, one shows the equality of 1) the intrinsic volume of the infinitesimal subvolume of the orbifold picked out by $[l, l + \delta l] \times [b, b + \delta b] \times [a, a + \delta a]$ and 2) the euclidean volume of the embedding of this subvolume. The intrinsic volume is computed from the metric tensor by $\sqrt{|g|} \delta l \delta b \delta a = \sqrt{2} \cos^2 l \frac{\sin 2b}{\sqrt{5-3 \cos 2b}} \delta l \delta b \delta a$. The embedded volume is given by the scalar triple product $\frac{\partial \vec{\chi}}{\partial l} \cdot (\frac{\partial \vec{\chi}}{\partial b} \times \frac{\partial \vec{\chi}}{\partial a}) \delta l \delta b \delta a$ which yields the same value.

To quantify the distortion of the embedding, we consider the embedding of infinitesimal spheres in the orbifold. Since the embedding is volume-preserving but not isometric, these project into ellipsoids of the same volume as the spheres. The

distortion of the embedding can be well quantified by considering the eccentricity (largest over smallest diameter) of the ellipsoids. Eccentricities are never degenerate and range from 1 at the center of the solid up to 5.4 at the two vertices. Using numerical integration, we have computed that the median eccentricity is 1.3 and the mean is 1.46.

4 DISTRIBUTION OF LOCAL STRUCTURES IN IMAGES

We expect that there are many ways in which the orbifold and the solid can be used to understand local image structure. As an example, we consider the density of different second order local forms in natural images and in noise images. The types of noise we consider are those generated using a fixed Fourier amplitude spectrum but a random phase spectrum. In particular, we consider blue noise which has an amplitude spectrum proportional to frequency ($\|\mathcal{F}[N_b](\omega)\| \propto |\omega|$), white noise which has constant amplitude spectrum ($\|\mathcal{F}[N_w](\omega)\| = c > 0$), and pink noise which has an amplitude spectrum inversely proportional to frequency ($\|\mathcal{F}[N_p](\omega)\| = |\omega|^{-1}$). Blue noise has found uses in generating dithering patterns for half-tone printing [41], [42]. White noise is also known as Gaussian noise, as it can be produced by generating an independent random Gaussian variable at each domain location. Pink noise is also known as one-over-f noise, and has been studied because its amplitude spectra is typical of natural images [43].

4.1 Noise

From Longuet-Higgin's results [28], we know that for noise defined by a fixed power spectrum and random phases the distribution of local forms across jet space will be a multidimensional normal distribution ($P(\vec{j})d\vec{j} = |2\pi\Xi|^{-\frac{1}{2}} e^{-\frac{1}{2}\vec{j}^T\Xi^{-1}\vec{j}}d\vec{j}$), with a covariance matrix (Ξ) depending on the power spectrum. We calculate (see electronic supplement) that the covariance matrices for blue (Ξ_b), white (Ξ_w), and pink (Ξ_p) noise are

$$\Xi_b = \frac{1}{16\pi\sigma^8} \begin{pmatrix} 4\sigma^2 & 0 & 0 & 0 & 0 \\ 0 & 4\sigma^2 & 0 & 0 & 0 \\ 0 & 0 & 9 & 0 & 3 \\ 0 & 0 & 0 & 3 & 0 \\ 0 & 0 & 3 & 0 & 9 \end{pmatrix},$$

$$\Xi_w = \frac{1}{16\pi\sigma^6} \begin{pmatrix} 2\sigma^2 & 0 & 0 & 0 & 0 \\ 0 & 2\sigma^2 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 3 \end{pmatrix},$$

$$\Xi_p = \frac{1}{32\pi\sigma^4} \begin{pmatrix} 4\sigma^2 & 0 & 0 & 0 & 0 \\ 0 & 4\sigma^2 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 3 \end{pmatrix}.$$

The first step in deriving the density over the orbifold is to change variables from \vec{j} to (n, θ, l, b, a) . This is done using the mapping \vec{r} (Section 2.2) and ignoring the z and m parameters. The determinant of the Jacobian of this

map is $2n^4\sigma^{-8}\cos^3l\sin 2b$. The parameters n and θ can then be integrated out. Finally, the densities need to be corrected so that they are relative to the metric of the orbifold rather than relative to $\langle l, b, a \rangle$. This is done by dividing by $\sqrt{|g|} = \sqrt{2}\cos^2l\frac{\sin 2b}{\sqrt{5-3\cos 2b}}$. The resulting densities are

$$P_b(l, b, a) = \frac{576 \cos l \sqrt{5 - 3 \cos 2b}}{\pi(11 - 2 \cos 2b \cos^2 l + 3 \cos 2l)^{\frac{3}{2}}},$$

$$P_w(l, b, a) = \frac{48 \cos l \sqrt{10 - 6 \cos 2b}}{\pi(5 - \cos 2b + 2 \cos 2l \sin^2 b)^{\frac{3}{2}}},$$

$$P_p(l, b, a) = \frac{192 \cos l \sqrt{5 - 3 \cos 2b}}{\pi(9 - 6 \cos 2b \cos^2 l + \cos 2l)^{\frac{3}{2}}}.$$

The patterns of these densities arrayed across the local structure solid are shown in Fig. 5. Since the orbifold embedding that defines the solid is volume-preserving, the structure of the density (i.e., the iso-density contours and critical points) as displayed on the solid is a faithful reproduction of the structure of the density across the orbifold. From the figure, the general trend is clear—there is a reduction in mostly second order local forms and an increase in mostly first order local forms as one moves from blue to white to pink noise. This reflects the visual differences between the different types of noise.

4.2 Natural Images

To study the distribution of structures in natural images, we compute a histogram over the structure solid. But first, to confirm that our methods for doing this are correct, we computed such histograms for the three types of noise and compared them to the analytically-derived results described in Section 4.1.

Noise images were generated, derivatives were computed at a scale of $\sigma = 4$, and the jets at 3×10^7 random points were used to populate histograms over the shape solid. The histograms of the solid used a system of cubic bins each of size 0.05^3 (in xyz space). A matrix of $32 \times 64 \times 22$ such bins completely contains the solid. Of these 45,056 bins, 24 percent are completely filled and 11 percent are partially filled. The volume of partially-filled bins was computed by numerical integration so that counts could correctly be turned into densities. The resulting histograms are shown in Fig. 5 (bottom row). Visually, they are an excellent match to the analytical densities shown above them. To quantify these matches, we have computed that 1) the Jensen-Shannon divergences [44] between the empirical and ideal distributions are approximately 0.006 for all three types of noise, while the angles between the square-rooted empirical and ideal densities treated as vectors (related to the Bhattacharyya distance [45]) are all approximately 5° .

Having validated our technique for histogram computation we now apply it to natural images. For this purpose, we used 20 images from a calibrated natural image collection [46]. As we have described previously [35], only images without saturation and blur were used and the images were perturbed to remove quantization plateaus. Jets were extracted as per the noise images, but with the additional precaution of not extracting jets within 7σ of the image border. The total number of samples collected was 3×10^7 . We also repeated this exercise at the coarser scale $\sigma = 8$.




Fig. 5. The figure shows from the analysis of the three types of noise studied in this section. The columns concern, from left-to-right, blue, white, and pink noise. The top row shows examples of the noise. The central row shows a cross-section through the second order local structure solid, with the analytically-derived density of local forms for that type of noise. These cross-sections are rotated relative to Fig. 4 so that the colored edge (pure second order structure) is along the top. The bottom row shows empirical densities based on 3×10^7 samples.




Fig. 6. The figure shows iso-density contours of the histogram of local forms for natural images and pink noise. The left panel of each pair shows the iso-density surface that surrounds 50 percent of the total histogram density; the right panels 10 percent of the density.

We found that the distributions at the two scales examined were effectively identical (Jensen-Shannon divergence = 0.004, Bhattacharyya angle = 4°) confirming previous findings [19], [47], [36], [35], [32], [48], [49], [50] of approximate scale-invariance in natural images. By examining cross-sections through the histograms we discovered that there was a diffuse clustering of density around the locus of effectively 1D local forms. This is illustrated in Fig. 6. To quantify the degree of the clustering, we offer the following: 50 percent of the density is in 20 percent of the solid close to the locally 1D forms and 10 percent of the density is in 1.7 percent of the solid close to the locally 1D forms. For comparison, we looked at pink noise (which has a similar power spectrum to natural images [50]) using the same analysis. For pink noise, we find slightly less clustering, and around the pure first order form rather than the locally 1D forms (see Fig. 6). The figures for pink noise are: 50 percent of the density is in 25 percent of the solid and 10 percent of the density is in 3.5 percent of the solid.

5 DISCUSSION

In this section, we discuss the various choices we have made in the derivation of the orbifold and solid and consider whether our results are independent of them (Sections 5.1, 5.2, and 5.3), the zero-scale limit (Section 5.4), and previous work on the distribution of local forms in natural images (Section 5.5).

5.1 The Norm

The orbifold metric is induced by the jet space norm. Since no prior literature exists on the choice of jet space norm, in order to proceed, we were forced to argue for a particular choice in this paper. The choice of jet space norm is likely to have implications for other issues in DtG type analyses, so we are nervous of having made an error on this fundamental issue. We urge the reader to evaluate the argument for the choice of norm in its own right and not simply to accept it because we have gone on to make use of it. We do note two features of our choice that recommend its use. First, is that its expression $\|\vec{j}\|_\sigma^N = (\sum_{1 \leq u+v \leq N} \frac{(\sigma^{u+v} c_{uv})^2}{u!v!})^{\frac{1}{2}}$ is pleasingly simple. Second is that, in the limit as the jet order goes to infinity, the jet space

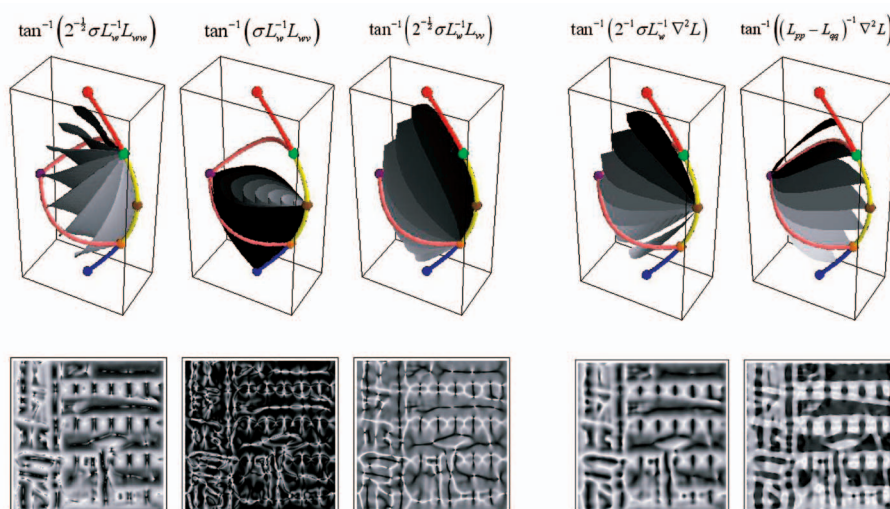


Fig. 7. Shows (top) iso-value surfaces of various differential invariants. From left-to-right: edge phase, flowline curvature, isophote curvature, Laplacean phase, and shape index. The bottom panels show the invariants calculated for the image shown in Fig. 3.

norm tends to the windowed variance of the image which is convenient as the statistics of windowed variance can readily be determined [51].

5.2 The Reparameterization

In Section 2.2, we chose a reparameterization of jet space. The choice was made with two aims in mind: 1) that three components of the new coordinate system should be invariant to the transformation group and 2) that calculation of the metric tensor relative to these three coordinates (Section 2.3) should be tractable. The reparameterization chosen was sufficient to meet both these aims, but we make no claim that it is uniquely so. In particular, there are several obvious possibilities for a reparameterization based on differential invariants more well-known than l , b , and a . In Fig. 7, we show the coordinates that would have arisen had we used some of these possibilities. This figure should be looked at in conjunction with Fig. 3 in which equivalent diagrams are shown for the parameters l , b , and a .

The three left-most panels of Fig. 7 show differential invariants based on a gauge coordinate system aligned with the first order structure; where the w -direction is along the gradient and the v -direction is normal to that [52]. These invariants have attracted attention because two of them (isophote and flowline curvature) have an intuitive interpretation in terms of the geometry of the blurred image [53], [54], [55]. The right-most panel shows the foliation of the solid by surfaces of constant shape index—Koenderink and van Doorn's well-motivated differential invariant for describing pure second order structure [56]. A natural partner for the shape index is the Laplacean phase, shown in the fourth panel; this has an interpretation in terms of the local slope of the iso-surface in scale space [57].

There is no reason why the coordinates shown in Fig. 7 should not be used for the reparameterization, we simply report our experience that we were unable to simplify the maths when we did so. It is important to understand however that the important aspects of our results do not depend on the coordinate system used. While the expression for the metric

tensor is relative to a particular coordinate system (as they always are), the following results are not coordinate system relative:

1. varying scalar curvature,
2. volume-preservation of the embedding,
3. degree of distortion of the embedding, and
4. the distributions of structures in noise and natural images.

5.3 The Embedding

In Section 3, we presented an embedding of the orbifold into euclidean 3-space. The embedding respects the intrinsic geometry of the orbifold in that it is volume-preserving. This is a useful property as it means that when we display densities across the solid (as in Figs. 5 and 6) the pattern of iso-density surfaces and critical points is the same as in the orbifold itself. The embedding does however distort the metric of the orbifold and since we make no claim that it does so minimally, must be regarded as ad hoc and potentially misleading. However, our results on the densities of local forms in natural and noise images have shown that the parts of the solid where distortion is greatest (the top and bottom vertices) are where the density is always low. This is like maps of the globe constructed so that the greatest metrical distortions occur where the population and geography is sparse.

5.4 The Zero-Scale Limit

It is of interest to consider the zero-scale limit of the jet and orbifold. In the zero-limit, the elements of the jet are the standard infinitesimal derivatives of the image since $\lim_{\sigma \rightarrow 0} C_{\sigma}^{(u,v)} = \delta^{(u,v)}$. Of course, such infinitesimal derivatives may not exist or be infinite (which is one of the motivations for working with nonzero scale derivatives in the first place as they do always exist) [12], but even if we assume that the underlying zero-scale image is smooth enough (e.g., Morse) so that they do exist and are finite, then the results we obtain are a little surprising. Consider the expression for the jet space norm

$$\|\vec{j}\| := \sigma(c_{10}^2 + c_{01}^2 + \sigma^2(c_{20}^2 + 2c_{11}^2 + c_{02}^2))^{\frac{1}{2}}.$$

In the limit, as scale goes to zero, the jet values c_{uv} (which depend on the scale of measurement) will tend to the standard derivatives of the underlying image, which by assumption exist and are finite. Therefore, in the zero-scale limit, the jet norm will be zero because of the leading σ factor. This makes sense as the jet space norm is defined to approximate windowed variance. We can prevent the norm going to zero by instead considering $\sigma^{-1} \|\vec{j}\|$, but then in the limit the norm is sensitive to the first order derivatives only. As a consequence of this, if we follow through the analysis to derive the distribution of structures over the solid for infinitesimal derivatives we find that almost all the weight of the distribution will be in a delta spike at the pure first order location (i.e., the purple point in Fig. 4), while the almost-negligible remainder of the distribution is arranged along the pure second order structure locus (the multicolored arc in Fig. 4). This shows that infinitesimal-scale local structure does not exist in mixed first/second order form as it does for nonzero scale structure which accords well with the analysis of Morse functions reviewed in Section 1.3.

5.5 Previous Work on the Distribution of Local Forms

We can compare our results on the distribution of local structures to previous similar results. Longuet-Higgins [28] showed that the joint distribution of derivatives of noise defined by its power spectrum is a multinormal distribution, but gave no results on the distribution after extrinsic aspects of the structure had been factored out. Koenderink and van Doorn [2] gave analytical and numerical results for the distribution of pure second order structures in noise and natural images after the same group that we have considered had been factored out. They noted that all the classes of image he looked at showed an excess of structure with close to parabolic curvature. The distributions they determined are marginal distributions of the ones we have presented. We have shown that our results are precisely compatible with theirs (see the Appendix, which can be found at <http://computer.org/tpami/archives.htm>). In the case of natural images, we can interpret their observation of an excess of parabolically curved points as a consequence of the more general finding of an excess of locally-1D points. Pedersen et al. [29], [31], [19], [30], has numerically computed the distribution of local structures defined up to third order, modulo intensity scaling and translation (but not rotation or reflection), for noise and natural images. Because the distribution they obtained was across a 7D space, they were not able to visualize it, but they did show for natural images that there was a significant clustering of density around the jets for blurred, displaced straight step edges. This is fully consistent with our finding of an excess of locally-1D structures in natural images.

6 CONCLUDING REMARKS

The results we have presented are specific to the Gaussian Derivative model of local visual processing. As we hope to have shown, one of the advantages of this model is the

simplicity with which sophisticated results can be obtained, and the breadth of results that can be accommodated within this single framework. We claim that this is because, uniquely among such models, the Gaussian derivative framework is a principled generalization of the differential calculus which is surely the acme of local analysis.

ACKNOWLEDGMENTS

EPSRC-funded project "Basic Image Features" EP/D030978/1.

REFERENCES

- [1] J.G. Daugman, "Uncertainty Relation for Resolution in Space, Spatial-Frequency, and Orientation Optimized by Two-Dimensional Visual Cortical Filters," *J. Optical Soc. Am. A-Optics Image Science and Vision*, vol. 2, no. 7, pp. 1160-1169, 1985.
- [2] J.J. Koenderink and A.J. van Doorn, "Local Structure of Gaussian Texture," *IEICE Trans. Information & Systems*, vol. 86, no. 7, pp. 1165-1171, 2003.
- [3] L.D. Griffin and A.C.F. Colchester, "Superficial and Deep-Structure in Linear Diffusion Scale-Space- Isophotes, Critical-Points and Separatrices," *Image and Vision Computing*, vol. 13, no. 7, pp. 543-557, 1995.
- [4] S. Wolfram, *The Mathematica Book*, fourth ed. Cambridge Univ. Press, 1999.
- [5] J.J. Koenderink, "The Structure of Images," *Biological Cybernetics*, vol. 50, no. 5, pp. 363-370, 1984.
- [6] *Scale-Space Theory in Computer Vision*. B.M. ter Haar Romeny, LNCS, vol. 1252. 1997. Springer-Verlag.
- [7] "Scale-Space and Morphology in Computer Vision," *Lecture Notes in Computer Science*, M. Kerckhove, ed., 2001.
- [8] "Scale Space and PDE Methods in Computer Vision," *Lecture Notes in Computer Science*, R. Kimmel, N. Sochen, and J. Weickert, eds., 2006.
- [9] "Scale-Space Theories in Computer Vision," *Lecture Notes in Computer Science*, M. Nielsen, ed., 1999.
- [10] "Scale Space Methods in Computer Vision," *Lecture Notes in Computer Science*, L.D. Griffin and M. Lillholm, eds., 2003.
- [11] T. Lindeberg, *Scale-Space Theory in Computer Vision*. Kluwer Academic, 1994.
- [12] L.M.J. Florack, *Image Structure*. Springer, 1997.
- [13] J.D. Victor and B.W. Knight, "Simultaneously Band and Space Limited Functions in Two Dimensions, and Receptive Fields of Visual Neurons," *Perspectives and Problems in Nonlinear Science*, E. Kaplan, J. Marsden, and K. Sreenivasan, eds., pp. 370-420, Springer-Verlag, 2002.
- [14] M. Abramowitz and I.A. Stegun, *Handbook of Mathematical Functions*. Dover Publications, 1964.
- [15] D.L. Ringach, "Spatial Structure and Symmetry of Simple-Cell Receptive Fields in Macaque Primary Visual Cortex," *J. Neurophysiology*, vol. 88, pp. 455-463, 2002.
- [16] L.M.J. Florack, B.M. ter Haar Romeny, M.A. Viergever, and J.J. Koenderink, "The Gaussian Scale-Space Paradigm and the Multi-scale Local Jet," *Int'l J. Computer Vision*, vol. 18, pp. 61-75, 1996.
- [17] M. Lillholm and K.S. Pedersen, "Jet Based Feature Classification," *Proc. 17th Int'l Conf. Pattern Recognition*, vol. 2, pp. 787-790, 2004.
- [18] T. Iijima, "A System of Fundamental Functions in an Abstract Figure Space," *Systems, Computers, Controls*, vol. 2, no. 6, pp. 96-103, 1971.
- [19] K.S. Pedersen and A.B. Lee, "Toward a Full Probability Model of Edges in Natural Images," *Proc. Seventh European Conf. Computer Vision*, pp. 328-342, 2002.
- [20] L.D. Griffin and M. Lillholm, "Hypotheses for Image Features, Icons and Textons," *Int'l J. Computer Vision*, vol. 70, no. 1, Oct. 2006.
- [21] J.J. Koenderink and A.J. van Doorn, "Image Processing Done Right," *Lecture Notes in Computer Science*, A. Heyden, G. Sparr, M. Nielsen, and P. Johansen, eds., pp. 158-172, 2002.
- [22] J.C. Pinoli, "The Logarithmic Image Processing Model: Connections with Human Brightness Perception and Contrast Estimators," *J. Math. Imaging and Vision*, vol. 7, no. 4, pp. 341-358, 1997.
- [23] L.D. Griffin, "Descriptions of Image Structure," PhD thesis, Univ. of London, 1995.

- [24] J.J. Koenderink and A.J. van Doorn, "The Structure of Relief," *Advances in Imaging and Electron Physics*, vol. 103, pp. 65-150, 1998.
- [25] B.M. ter Haar Romeny, L.M.J. Florack, A.H. Salden, and M.A. Viergever, "Higher-Order Differential Structure of Images," *Image and Vision Computing*, vol. 12, no. 6, pp. 317-325, 1994.
- [26] L.M.J. Florack, B.M. ter Haar Romeny, J.J. Koenderink, and M.A. Viergever, "General Intensity Transformations and Differential Invariants," *J. Math. Imaging and Vision*, vol. 4, no. 2, pp. 171-187, 1994.
- [27] E.P. Simoncelli, "Statistical Modeling of Photographic Images," *Handbook of Image and Video Processing*, A. Bovik, ed. Academic Press, 2005.
- [28] M.S. Longuet-Higgins, "Statistical Analysis of a Random Moving Surface," *Philosophical Trans. Royal Soc. London. A.*, vol. 249, pp. 321-364, 1956.
- [29] K.S. Pedersen, "Properties of Brownian Image Models in Scale-Space," *Proc. Fourth Scale-Space Conf.*, pp. 281-296, 2003.
- [30] K.S. Pedersen, "Statistics of Natural Image Geometry," PhD, thesis, Univ. of Copenhagen, 2003.
- [31] A.B. Lee, K.S. Pedersen, and D. Mumford, "The Nonlinear Statistics of High-Contrast Patches in Natural Images," *Int'l J. Computer Vision*, vol. 54, nos. 1-2, pp. 83-103, 2003.
- [32] D. Geman and A. Koloydenko, "Invariant Statistics and Coding of Natural Microimages," *Proc. IEEE Workshop Statistical and Computational Theories of Vision*, 1999.
- [33] E. Tagliati and L.D. Griffin, "Features in Scale Space: Progress on the 2D second Order Jet," *Lecture Notes in Computer Science*, M. Kerckhove, ed., pp. 51-62, 2001.
- [34] J. Van de Weijer and R. van den Boomgaard, "Least Squares and Robust Estimation of Local Image Structure," *Int'l J. Computer Vision*, vol. 64, no. 2-3, pp. 143-155, 2005.
- [35] L.D. Griffin, "Feature Classes for 1D, 2nd Order Structure Arise from the Maximum Likelihood Statistics of Natural Images," *Network: Computation in Neural Systems*, vol. 16, nos. 2/3, pp. 301-320, 2005.
- [36] L.D. Griffin, M. Lillholm, and M. Nielsen, "Natural Image Profiles Are Most Likely to be Step Edges," *Vision Research*, vol. 44, no. 4, pp. 407-421, 2004.
- [37] B.M. ter Haar Romeny, *Front-End Vision and Multi-Scale Image Analysis*. Kluwer, 2003.
- [38] W.P. Thurston, *The Geometry and Topology of Three-Manifolds*, Electronic Version 1.1 ed., Math. Sciences Research Inst., 2002.
- [39] C.W. Misner, K.S. Thorne, and J.A. Wheeler, *Gravitation*. WH Freeman, 1973.
- [40] C. Doran and A. Lasenby, *Geometric Algebra for Physicists*. Cambridge Univ. Press, 2003.
- [41] R.A. Ulichney, "Dithering with Blue Noise," *Proc. IEEE*, vol. 76, no. 1, pp. 56-79, 1988.
- [42] R.A. Ulichney, *Digital Halftoning*. MIT Press, 1987.
- [43] D.J. Field, "Wavelets, Vision and the Statistics of Natural Scenes," *Philosophical Trans. Royal Soc. London Series A-Math. Physical and Eng. Sciences*, vol. 357, no. 1760, pp. 2527-2542, 1999.
- [44] J. Lin, "Divergences Measures Based on the Shannon Entropy," *IEEE Trans. Informatin Theory*, vol. 37, pp. 145-151, 1991.
- [45] T. Kailath, "The Divergence and Bhattacharyya Distance Measures in Signal Selection," *IEEE Trans. Comm.*, vol. 15, no. 1, pp. 52-60, 1967.
- [46] J.H. van Hateren and A. van der Schaaf, "Independent Component Filters of Natural Images Compared with Simple Cells in Primary Visual Cortex," *Proc. Royal Soc. London Series B-Biological Sciences*, vol. 265, no. 1394, pp. 359-366, 1998.
- [47] E. Kretzmer, "Statistics of Television Signals," *Bell System Technical J.*, vol. 31, pp. 751-763, 1952.
- [48] Y. Gousseau, "Morphological Statistics of Natural Images," *Proc. SPIE, Wavelet Applications in Signal and Image Processing*, vol. 4119, pp. 208-214, 2000.
- [49] D.J. Field and N. Brady, "Visual Sensitivity, Blur and the Sources of Variability in the Amplitude Spectra of Natural Scenes," *Vision Research*, vol. 37, no. 23, pp. 3367-3383, 1997.
- [50] D.L. Ruderman, "Origins of Scaling in Natural Images," *Vision Research*, vol. 37, no. 23, pp. 3385-3398, 1997.
- [51] R.A. Frazor and W.S. Geisler, "Local Luminance and Contrast in Natural Images," *Vision Research*, vol. 46, pp. 1585-1598, 2006.
- [52] L.D. Griffin, "Mean, Median and Mode Filtering of Images," *Proc. Royal Soc.: Series A*, vol. 456, pp. 2995-3004, 2000.
- [53] L.M.J. Florack, B.M.T. Romeny, J.J. Koenderink, and M.A. Viergever, "Scale and the Differential Structure of Images," *Image and Vision Computing*, vol. 10, no. 6, pp. 376-388, 1992.
- [54] J.J. Koenderink and A.J. van Doorn, "Representation of Local Geometry in the Visual-System," *Biological Cybernetics*, vol. 55, no. 6, pp. 367-375, 1987.
- [55] J.J. Koenderink and W. Richards, "Two-Dimensional Curvature Operators," *J. Optical Soc.Am.*, vol. 5, no. 7, pp. 1136-1141, 1988.
- [56] J.J. Koenderink and A.J. van Doorn, "Surface Shape and Curvature Scales," *Image and Vision Computing*, vol. 10, no. 8, pp. 557-565, 1992.
- [57] L.D. Griffin, A.C.F. Colchester, S.A. Roell, and C. Studholme, "Hierarchical Segmentation Satisfying Constraints," *Proc. British Machine Vision Conf.*, E. Hancock, ed., pp. 135-144, 1994.



Lewis D. Griffin received the BA (Hons.) degree in mathematics and philosophy in 1988 from Oxford University, United Kingdom. He researched for the PhD degree while he was a research assistant in the Department of Neurology, Guys Hospital, London. In 1995, he was awarded the PhD degree from the University of London for a thesis ("Descriptions of Image Structure") in the area of computational vision. In 1997, following postdoctoral positions at INRIA Sophia-Antipolis, France and University of Surrey, United Kingdom, he was appointed lecturer in vision sciences at Aston University, United Kingdom. In 2000, he moved to Imaging Sciences, Medical School, King's College London. In 2005, he moved to the Computer Science Department, University College London where he is at present. His research interests are: the computational theory of image structure, visual psychophysics, color vision, and machine learning.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.