Symmetry Sensitivities of Derivative-of-Gaussian Filters

Lewis D. Griffin and Martin Lillholm

Abstract—We consider the measurement of image structure using linear filters, in particular derivative-of-Gaussian (DtG) filters, which are an important model of V1 simple cells and widely used in computer vision, and whether such measurements can determine local image symmetry. We show that even a single linear filter can be sensitive to a symmetry, in the sense that specific responses of the filter can rule it out. We state and prove a necessary and sufficient, readily computable, criterion for filter symmetry-sensitivity. We use it to show that the six filters in a second order DtG family have patterns of joint sensitivity which are distinct for 12 different classes of symmetry. This rich symmetry-sensitivity adds to the properties that make DtG filters well-suited for probing local image structure, and provides a set of landmark responses suitable to be the foundation of a nonarbitrary system of feature categories.

Index Terms—Group theory, invariance, pattern analysis

✦

---

1 INTRODUCTION

WE continue an investigation of a model of an early stage of front-end vision, when local measurements of an image are performed by computation of inner products between an image and derivative-of-Gaussian (DtG) filters [1]. See Fig. 2 for an overview of this model. Interest in the model arises from 1) the biological relevance of DtGs, 2) the elegant formal properties of families of DtGs, and 3) the effectiveness of DtGs in computer vision. The biological relevance of DtGs is that they provide a model [1], [2], [3], [4], [5], [6], [7], [8] of the measurement of image structure performed by the linear aspects of V1 simple cell neurons of the mammalian primary visual cortex. The formal properties of measurement with DtGs are linked to the several different interpretations that can be put on such measurement. These interpretations are: as computing the Hermite Transform, a local analogue of the Fourier Transform [9], [10], or as implementing a generalization of differential geometry to noninfinitesimal local structure [5], [6], [11], [12], or as, in a single operation, rescaling the image and computing a regular Taylor series [13], [14]. The effectiveness for computer vision comes from the Cartesian separability of DtGs, their good joint-localization in space and frequency, the adequacy of a small basis set tuned to only a few orientations, and the closure of the family with respect to composition by convolution [15].

Measurement with DtGs is a satisfactory solution for linear measurement of *quantitative* local image structure, but researchers in computational vision have long found that it is useful to go beyond quantitative to some notion of *qualitative* structure [16], such as categorizing image locations as being edges, bars, or junctions. Similar ideas can be found in the biological literature (see [17]). The central puzzle of this "filters-to-features" [4] process is how to establish a partitioning of filter response space into categories corresponding to different features. What makes it a puzzle is that linear filter response space has, apart from the origin, no special landmarks or subspaces picked out: There seem to be no clues as to how to carve it up.

Various sources of additional structure that could drive the partitioning of filter response space have been suggested and explored (reviewed in [17]). At one extreme is the use of natural image statistics [18], [19], [20], [21], [22]; intermediate is the use of tractable idealized models of the statistics of natural images, for example, as Brownian noise [23], [24]; at the other extreme are purely geometric arguments. It is arguable that no approaches have yet been purely geometric. The furthest in this direction is an approach [25], [26], [27], [28] based on the assumption that image intensity is bounded above and below.

This paper is concerned with defining additional structure on filter response space purely through geometric considerations, with no assumptions being made about statistics or restrictions on the images measured. Our approach is to consider in what circumstances filter responses are informative about the local symmetry of the measured image. For generality, our results when possible will concern linear filters in general, but we also specialize to DtG filters and show that a rich structure of landmarks does arise. We leave it to a future paper to show that these landmarks can be an effective basis for a partitioning into feature categories.

1.1 Paper Organization

The paper is written to be self-contained and is organized as follows: Section 2 concerns symmetry in general and of images in particular. Section 3 describes the detection of image symmetry using general linear filters. Section 4 reviews the DtG family of linear filters. Section 5 describes the detection of image symmetry using DtGs. Finally, Section 6 is about summary and conclusion.

1.2 Conventions Used

If a variable name is introduced as an unconstrained example of a mathematical type (e.g., $q : \mathbb{R} \to \mathbb{R}$, a scalar function of one variable), it remains so throughout the paper.

- *The authors are with the Department of Computer Science, University College London, Malet Place Engineering Building, London WC1E 6BT, UK. E-mail: {L.Griffin, M.Lillholm}@cs.ucl.ac.uk.*

If $c \in \mathbb{R}$ is a scalar, $q \equiv c$ indicates that q takes the same constant value c for all arguments. We use the symbol $\circ$ for the composition of transformations and their application to scalars, functions, and images; the placeholder $\_$ is used as an unnamed variable; $|\_|$ is used for the signed determinant of an isometry (-1 or $+1$). The abbreviations D1, P1, T1, and C1 will be used to refer to numbered definitions, propositions, theorems, and corollaries. Italics are used for emphasis, especially of a first use.

2 SYMMETRY

In this section, we review material relating to symmetry. We start with a discussion of the role of transformations in understanding structures and how the transformation perspective gives rise to the distinction between intrinsic and extrinsic properties and to the possibility of symmetrical structures. We will argue that symmetry type is a particularly fundamental intrinsic property. In the next sections, we move on from the symmetry of structures in general to the symmetry of images in particular. We summarize previous results cataloguing image symmetry types.

2.1 The Transformation Perspective

Generally speaking, a structure embedded in a space can usefully be analyzed from the perspective of the space group (C), a group of transformations of the containing space [29]. The choice of space group makes all the difference between one geometry and another [30], so we assume that the choice has been made and is fixed. This transformation perspective has two consequents—intrinsic versus extrinsic properties, and symmetrical structures—which we consider in turn.

Intrinsic properties are those that are unchanged by application of any transformation of the space group; the rest are *extrinsic properties*. This dichotomy is familiar for planar figures. Let the space group be the planar isometries (i.e., $C = E(2)$), which are the distance-preserving transformations. What we commonly call the *shape* of a figure is its intrinsic properties, while its *location, orientation, and handedness* are extrinsic properties. The intrinsic/extrinsic dichotomy is of philosophical importance [31], and cognitively, it may play a role in carving the world into distinct objects [32].

Symmetrical structures are defined in terms of *automorphisms*, which are those transformations of the space group that cause the structure to be repositioned so that it is indistinguishable from the original. The automorphisms of a structure are closed under composition and inversion and so form a mathematical group ($Aut[X] = \{t \mid t \in C, t \circ X = X\}$). When this *automorphism group* contains more than just the identity transformation, the structure is said to be symmetrical.

What then of the automorphism group, is that an intrinsic property of a structure? No. Translating a circle, for example, changes the automorphism group as rotations about a different center are needed. The automorphism group has an invariant though, its *symmetry type*, which is simply an equivalence ($\cong$) class of symmetry groups that can be made identical by conjugation, which is pre and postcomposition by a transformation and its inverse, i.e., $\forall U, V \subseteq C \ U \cong V \Longleftrightarrow \exists w \in C \ U = w \circ V \circ w^{-1}$.

Thus, are the two consequents of the transformation perspective—intrinsic versus extrinsic properties, and symmetrical structures—reunited in the concept of the symmetry type of a structure. We claim that symmetry type is a particular fundamental intrinsic property. We note that:

- It is defined without reference to any external measure.
- It is a property of all structures.
- Its definition requires only a notion of indistinguishability.

Contrast this with the "radius" (for example) of a circle. "Radius" is an intrinsic property of a circle as it is unaffected by an isometry, but it is considerably more complex than symmetry type. In particular, observations 1-3 are not true of "radius."

2.2 Types of Image Symmetry

The precise split between intrinsic and extrinsic, and the variety of possible symmetry types both depend on the class of transformations considered [30]. Consider, for example, allowing isotropic scaling of planar figures in addition to the rigid transformations considered above: With their inclusion, all circles are the same intrinsic shape, without them they are not; with their inclusion, there are symmetrical spirals, without them there are not.

In this paper, the structures of concern are *images*, which we characterize as scalar functions of the euclidean plane. For images, as with planar figures, the gamut of possible symmetry types depends on the class of transformations considered. If the transformation class is $E(2)$, then the symmetries of a football pitch can be fully expressed, but those of Escher's lithograph "Reptiles" (where the plane is tiled by interlocking black, white, and red lizard shapes [33]) cannot. To deal with patterns such as "Reptiles," a broader class of "colored isometries" ($E(2) \times S_n$, where S_n is one of the symmetric groups), where each spatial isometry is combined with a tonal transformation that permutes a finite set of image "color" values, can be considered and the possible symmetries can be determined relative to this class.

In this paper, the transformations that we will consider are, like colored isometries, a pairing of a spatial isometry with a tonal transformation; but rather than a permutation of a discrete label set, we follow [34] and use isometrical transformations of the image intensity. These intensity isometries are of two possible types only: increment/decrements ($v \mapsto c + v$) and negations about a value ($v \mapsto 2c - v$). We call a pairing of a spatial and an intensity isometry an *image isometry*, and will consider the group of them $E(2) \times E(1)$.

Let $i \in E(1)$ be a 1D (intensity) isometry. The following trivial propositions about such isometries will be useful later:

Proposition 1 (Derivative of a 1D isometry and its inverse). *The derivative of i and i^{-1} are necessarily constant-valued, equal, and of unit absolute value. They will be denoted by $|\_|$, so $i' \equiv (i^{-1})' \equiv |i| = |i^{-1}| \in \{-1, 1\}$.*

Proposition 2 (1D isometry commutes with affine combination). $i \circ \sum_r z_r v_r = \sum_r z_r \ i \circ v_r$, *where* $z_r, v_r \in \mathbb{R}$ *and* $\sum_r z_r = 1$.

Let $s \in E(2)$ be a 2D (spatial) isometry. We denote its application to an image $I : \mathbb{R}^2 \to \mathbb{R}$ by $s \circ I = I(s^{-1} \circ \_)$. We write an image isometry as $g = (i, s)$ and denote its application to an image by $g \circ I = i \circ I(s^{-1} \circ \_)$. In Section 4,
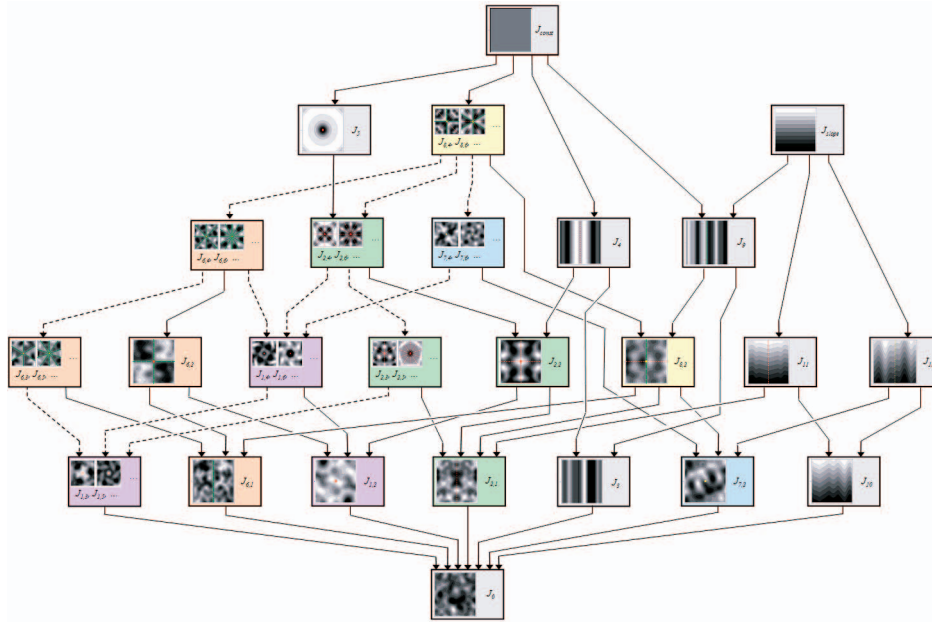
Fig. 1. The different types of image symmetry, relative to the class of image isometries, but cases without periodic translation only. The links show supertype-subtype relationships from above to below. Colored panels indicate the membership of a countably infinite family of that symmetry type, e.g., all $J_{2,n}$ panels are pale green. Dashed links indicate that the type/subtype relationships only hold between certain orders of the groups at either end of the link. Type/subtype relationship can also hold within a node when the orders permit it; for example, $J_{1,3}$ is a subtype of $J_{1,9}$. The contour plots within the panels show an example image with the symmetry type. These were generated by applying the symmetrization procedure described in Section 4 to random images. Overlaid on these plots are lines of reflection (red), lines of negating-reflection (green), centers of rotation (red), centers of negating-rotation (yellow), and lines of centers of negating-rotation (yellow). These lines and centers are not shown when they are too numerous, e.g., lines of reflection for J_3.

we will have the need of the following easily proven proposition about image isometries:

Proposition 3 (First variation of an image isometry).

$$\frac{\delta(g \circ I)}{\delta I}\delta I = |i|s \circ \delta I.$$

Each type of planar isometry—identity, reflection, translation, rotations, and glide reflection—can be paired with each type of intensity isometry—identity, increment, negation—to make an image isometry. Although all pairings are valid image isometries, some (e.g., the pairing of a spatial identity transformation with an intensity increment) can never be automorphisms. We call the pairing of a reflection and a negation an *antireflection*, and a rotation and a negation an *antirotation*.

Having chosen a class of transformations, the possible image symmetries, relative to the class, can be catalogued. This was achieved in the first half of the 20th century for the planar isometries ($E(2)$) [35], [36], [37], and in the second half for the colored isometries ($E(2) \times S_n$) [38], [39]. The catalogue of possible image symmetries, relative to the class of image isometries ($E(2) \times E(1)$), is partially complete [40], [41]—the complete list of symmetries *other than those involving periodic translations* has been provably determined and is shown in Fig. 1 and listed in Table 1. The figure and table indicate a notational system for the symmetry types established in [41]. Individual types are denoted by J_n or J_{label}; families of types are denoted by $J_{n,m}$, where n indicates the family and m the order. The lattice of Fig. 1 indicates inclusion relations between symmetry types.

In Section 6, when we consider the sensitivity of families of DtG filters to the image symmetries shown in Fig. 1, we will have cause to more finely subdivide the symmetry types. The subdivision will be based on the relation between the transformations of a group and the point in the plane, where the filters are centered. We will distinguish between symmetry types that are centered (indicated by c superscript), aligned (by a or $a-$ superscript), and in general position (by g). The possible subtypes are shown in Table 1.

2.3 Detection of Image Symmetry

The detection and classification of image symmetry has been studied in the context of Computer [42], [43], [44], [45], [46], [47] and Biological Vision [48], [49], [50], [51], [52], [53], [54], [55]. As discussed in Section 3.2, there are many types of symmetries potentially relevant to image understanding, but often detection of the reflectional type only has been considered; less often considered are skew-reflectional (reflect and shear) [56], antireflectional [54], and periodic-translational [47], [51]. Researchers in Biological Vision generally focus on global symmetry, which, having an established definition, makes it suitable for stimulus characterization. Researchers in Computational Vision, recognizing that global symmetry is a rather rare circumstance in natural images, often focus on advancing definitions of local symmetry and algorithms for its detection [42], [43], [57], [58], [59], [60], [61]. For biologists and psychologists, the interest in symmetry arises mostly from informal observation of sensitivity, from the suggestion of it as a Gestalt grouping principle, from its potential for use in detecting biological structures such as faces and flowers [62], and for its role in mate selection [63]. For computer vision engineers, symmetry is thought to be useful for image description in general, and object recognition in particular [64].

TABLE 1
Summary of the Possible Nonperiodic Image Symmetries

Group	Subtypes	Description of Symmetry
J_0		The only automorphism is the identity transformation.
$J_{1,2}, J_{1,3}, \ldots$	c,g	A finite number of rotations about a common centre.
$J_{2,1}, J_{2,2}, \ldots$	2,1: a,g 2,2+: c,a,g	A finite number of mirrors, intersecting in a point.
J_3		Unchanging in one direction.
J_4	a,g	Unchanging in one direction, with a line of reflection in the direction.
J_5	a,g	All possible rotations about a common centre.
$J_{6,1}, J_{6,2}, \ldots$	6,1: a,g; 6,2+: c,a,g	A finite number of anti-mirrors, intersecting in a point.
$J_{7,2}, J_{7,4}, \ldots$	c,g	A finite number of anti-rotations and rotations about a common centre.
$J_{8,2}, J_{8,4}, \ldots$	c,a,a-,g	A finite number of anti-mirrors and mirrors, intersecting in a point.
J_9	a,g	Unchanging in one direction, with a line of anti-reflection in the direction.
J_{10}		The image can be continuously translated in one direction, if the translation is matched by a proportionate intensity increment.
J_{11}	a,g	As J_{10} plus a line of reflection in the translation direction.
J_{12}	a,g	As J_{10} plus a line, in the translation direction, of centres of 180° anti-rotations.
J_{slope}		The symmetries of a non-flat plane.
J_{const}		The symmetries of a constant function.

The subtypes column is blank when only the general subtype is possible.

The methods and models proposed for symmetry detection and classification can be fairly complex. Some require first processing the image to create an edge map [61], [65], others are based on nonlinear combination of the outputs of *multiple* linear filters positioned at distinct image locations (e.g., either side of a line of reflection) [66], and yet others detect and match point features [67], [68].

Our approach to image symmetry, presented in the remainder of the paper, is focused on local symmetry. We recognize, as others have done before, that whereas extended symmetry in images is a strong but rare image cue, local symmetry is actually commonplace; and it is not just the local reflectional symmetry to which the Medial Axis Transform [69] is attuned, but a range of local symmetry types. The results that we will present concern when a symmetry can be ruled *out* on the basis of the response of linear filters. Possibly, this is useful for symmetry detection in the conventional sense, but our aim is rather to discover what natural landmarks exist in filter response space that could serve as a basis for solving the filters-to-features problem. The symmetries that we shall consider for detection are those defined by groups of images isometries. Detection of other types of symmetries would also be of interest.

3 LINEAR FILTERS AND IMAGE SYMMETRY

Linear filters are typically the first stage in Computer Vision systems and are a good approximation to the first stage of cortical analysis in mammalian vision systems (see [70] for a

discussion of the limits of this approximation). We cannot conceive of a simpler mode of image measurement. In this section, we consider whether measurement with linear filters has any sensitivity to the symmetry group of an image; in Section 6, we will extend our consideration to the sensitivity to symmetry type of a particular family of linear filters, the DtGs.

On cursory analysis, linear filters can detect image symmetry but it seems to require multiple filters at distinct locations to do so. For example, a necessary condition for a particular reflectional symmetry would be that some filter F gives the same response as F', the reflected filter in the reflected position. The cursory analysis is wrong though, and multiple filters are *not* necessary for symmetry detection. Consider, for instance, the two-pixel-support $\boxed{+1}\boxed{-1}$ filter, such as used in finite-difference schemes. If this filter is positioned so that it straddles a candidate line of reflection, then a necessary criterion for the symmetry is that this single filter gives a zero response—we will call this its *critical response*.

Note that the above criterion—zero response by the $\boxed{+1}\boxed{-1}$ filter—is a necessary condition for the reflectional symmetry to be present, but is *not* sufficient. This is a general limitation that exists for all filter-based symmetry detection. It is inevitable given either the linearity or the locality of the measurements. We do not believe that this limitation undermines the correctness or significance of our analysis, but it should not be ignored. We characterize the limitation as follows: the responses of appropriate, localized, linear

filters can tell you when a symmetry is present *so far as the filters are concerned*.

We can generalize from the above with the definition: A (nonzero) filter is said to be sensitive to a symmetry if it gives the same (critical) response to all images that have the symmetry.

To formalize this, we use an inner product notation $\langle F|I\rangle = \int_{\vec{x}\in\mathbb{R}^2} F(\vec{x})I(\vec{x})$ to denote the measurement of an image by a filter $F:\mathbb{R}^2 \to \mathbb{R}$. Let $K \subseteq E(2) \times E(1)$ be a group of image isometries. We can then say:

Definition 1 (Sensitivity of a filter F to a group of image isometries K). $F \not\equiv 0$ *is sensitive to* $K \Longleftrightarrow \exists f \in \mathbb{R} \ \forall I \ Aut[I] \supseteq K \Rightarrow \langle F \mid I \rangle = f.$

The definition of sensitivity, although unambiguous, does not allow us readily to determine whether a given filter is sensitive to a given symmetry—the difficulty lies in ascertaining the truth of the universal quantification across images. We have solved this by finding a provably equivalent criterion without any universal quantification. Our proof makes use of two types of operations on 2D scalar functions: i-symmetrization, which we typically apply to images and f-symmetrization that we typically apply to filters. In *i-symmetrization*, denoted and defined by $I^K := |K|^{-1}\sum_{g\in K} g \circ I$, we separately apply to an image each transformation of a group of image isometries and take the mean of the results. In *f-symmetrization*, denoted and defined by $F^{(K)} := |K|^{-1}\sum_{(i,s)\in K} |i|s \circ F$, we separately apply to a filter each of the spatial parts of a group of image isometries, weight the transformed filters by the determinant of the corresponding intensity isometry, and take the mean of the results. For both types of symmetrization, the summation over the group becomes an integration if the group is continuous. We now state and prove three propositions concerning i-symmetrization, followed by a proposition about f-symmetrization, followed by the main theorem of the paper.

Proposition 4 (i-symmetrization makes symmetric images). $Aut[I^K] \supseteq K$ *(or, equivalently, $\forall g \in K \ g \circ I^K = I^K$).*

Notes. The example symmetric images of Fig. 1 were generated by i-symmetrizing noise images.

Proof. Let $h \in K$. Then, making use of P2:

$$h \circ I^K = h \circ \left(|K|^{-1}\sum_{g\in K} g \circ I\right) = |K|^{-1}\sum_{g\in K} h \circ g \circ I$$
$$= |K|^{-1}\sum_{g\in K}(h \circ g) \circ I = |K|^{-1}\sum_{g\in K} g \circ I = I^K.$$

Hence, $Aut[I^K] \supseteq K$ as required.   □

Proposition 5 (i-symmetrization has no effect on symmetric images).

$$Aut[I] \supseteq K \quad \Rightarrow \quad I^K = I.$$

Proof. Assume that $Aut[I] \supseteq K$ and so $\forall g \in K \ g \circ I = I$. Plug this into the definition of i-symmetrization:

$$I^K = |K|^{-1}\sum_{g\in K} g \circ I = |K|^{-1}\sum_{g\in K} I = I.$$

  □

Proposition 6 (Variation of i-symmetrization is f-symmetrization).

$$\frac{\delta I^K}{\delta I}\delta I = (\delta I)^{(K)}.$$

Proof. Trivial given P3.   □

Proposition 7 (f-symmetrization makes quasisymmetric filters).

$$\forall (i,s) \in K \quad |i|s \circ F^{(K)} = F^{(K)}.$$

Notes. Compare with P4 to see a difference in effect between i and f-symmetrizations.

Proof.

$$|i|s \circ F^{(K)} = |i|s \circ \left(|K|^{-1}\sum_{(j,t)\in K}|j|t \circ F\right)$$
$$= |K|^{-1}\sum_{(j,t)\in K}|i||j|s \circ t \circ F$$
$$= |K|^{-1}\sum_{(j,t)\in K}|i \circ j|(s \circ t) \circ F$$
$$= |K|^{-1}\sum_{(j,t)\in K}|j|t \circ F = F^{(K)}.$$

  □

Theorem 1 (Sensitivity of a filter (F) to a group (K) of image isometries). *The following are true or false together:*

- F *is sensitive to* K,
- $F^{(K)} \equiv 0$,
- $\langle F \mid F^{(K)} \rangle = 0$.

Proof. Define the functional $k(I) = \langle F \mid I^K \rangle$. Compute the variation of k with respect to I, and simplify the result making use of P1 and P6.

$$\frac{\delta k(I)}{\delta I}\delta I = \langle F \mid (\delta I)^{(K)} \rangle = \left\langle F \middle| |K|^{-1}\sum_{(i,s)\in K}|i|s \circ \delta I \right\rangle$$
$$= |K|^{-1}\sum_{(i,s)\in K}\langle F||i|s \circ \delta I\rangle$$
$$= |K|^{-1}\sum_{(i,s)\in K}\langle |i^{-1}|s^{-1} \circ F|\delta I\rangle$$
$$= \left\langle |K|^{-1}\sum_{(i,s)\in K}|i|s \circ F|\delta I\right\rangle$$
$$= \langle F^{(K)}|\delta I\rangle.$$

1. F is sensitive to $K \Rightarrow F^{(K)} \equiv 0$.

   Assume the premise. Then, since by P4 $Aut[I^K] \supseteq K$, it follows that $k(I)$ is constant, and so the variation of $k(I)$ with respect to I must be 0. For this to be true for all δI, it must be that $F^{(K)} \equiv 0$ as required to prove.

2. $F^{(K)} \equiv 0 \Rightarrow F$ is sensitive to K.

   Assume that $F^{(K)} \equiv 0$; therefore $\frac{\delta}{\delta I} k(I) = 0$; therefore $k(I)$ is constant. Define $f := k(\_)$ to be that constant value. Then, since by P5 $Aut[I] \supseteq K \Rightarrow I = I^K$, it follows that $Aut[I] \supseteq K \Rightarrow \langle F \mid I \rangle = \langle F \mid I^K \rangle = k(I) = f$.

3. $F^{(K)} \equiv 0 \Rightarrow \langle F \mid F^{(K)} \rangle = 0$.

   Trivial.

4. $\langle F \mid F^{(K)} \rangle = 0 \Rightarrow F^{(K)} \equiv 0$.

   Assume the premise, then

   $$\forall (i, s) \in K \langle |i| s \circ F || i | s \circ F^{(K)} \rangle = 0.$$

   Then P7 allows simplification to

   $$\forall (i, s) \in K \langle |i| s \circ F | F^{(K)} \rangle = 0.$$

   Continuing:

   $$\forall (i, s) \in K \langle |i| s \circ F | F^{(K)} \rangle = 0$$
   $$\Rightarrow \left\langle |K|^{-1} \sum_{(i,s) \in J} |i| s \circ F \middle| F^{(K)} \right\rangle = 0$$
   $$\Rightarrow \left\langle F^{(K)} \middle| F^{(K)} \right\rangle = 0$$
   $$\Rightarrow F^{(K)} \equiv 0.$$

   Together, steps 1-4 imply the theorem. □

3.1 Notation for Sets of Filters

In the following sections, we will typically deal with sets of filters and their span. It is convenient for presentation to assemble a set of filters $(F_1, \ldots, F_n)$ into a vector, denoted by $\vec{F} := (F_1 \ldots F_n)^{\mathrm{T}}$. We will then use a dot product notation to denote the filter in the span of $\vec{F}$ picked out by a weight vector $\vec{w} = (w_1 \ldots w_n)^T$ as follows: $\vec{w} \cdot \vec{F} := \sum_{i=1}^{n} w_i F_i$. We will also use the same inner product notation we use for measurement by a single filter to denote the measurement by a family of filters, i.e., $\langle \vec{F} \mid I \rangle := (\langle F_1 \mid I \rangle \ldots \langle F_n \mid I \rangle)^{\mathrm{T}}$. We extend the notation for i and f-symmetrizations to vectors of filters in the obvious way, e.g., $\vec{F}^{(K)} = (F_1^{(K)} \ldots F_n^{(K)})^{\mathrm{T}}$. Finally, we define an outer product between sets of filters as follows:

$$\vec{A} \otimes \vec{B} := \begin{pmatrix} \langle A_1 \mid B_1 \rangle & \ldots & \langle A_1 \mid B_n \rangle \\ \vdots & \ddots & \vdots \\ \langle A_n \mid B_1 \rangle & \ldots & \langle A_n \mid B_n \rangle \end{pmatrix}.$$

4 DTG FILTERS, JETS, AND JET SPACE

As this paper is self-contained, in this section, we review the definitions and properties of DtGs and the measurement of image structure using them.
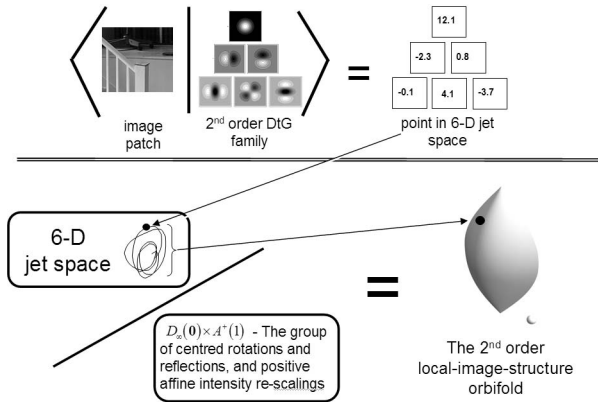

Fig. 2. Schematically, (top-left) the measurement of image structure using a bank of DtG filters up to second order and the resulting jet (top-right). The orbit in jet space that would result if the patch was rotated, brightened, etc., is illustrated by a squiggle (bottom-left). Factoring the jet space by a group collapses each orbit to a single point of the orbifold, a different point for each orbit. A mildly distorting embedding of the orbifold into 3D space is shown at bottom-right.

4.1 Families of DtG Filters

Gaussian Derivative (DtG) filters of scale $\sigma > 0$ are defined in 1D by

$$G(x) := G_0(x) := (2\pi\sigma^2)^{-\frac{1}{2}} e^{-\frac{x^2}{2\sigma^2}},$$
$$G_n(x) := \frac{d^n}{dx^n} G(x) = \left(\frac{-1}{\sigma\sqrt{2}} \right)^n H_n \left(\frac{x}{\sigma\sqrt{2}} \right) G(x), \quad n \in \mathbb{Z}^+,$$

where H_n is the nth Hermite polynomial; and in 2D by

$$G_{mn}(x, y) := G_m(x) \, G_n(y), \quad m, n \in \mathbb{Z}^+,$$
$$G(x, y) := G_{00}(x, y).$$

They are used as a general purpose method to probe an image location (which, for simplicity, we assume is at the origin $\mathbf{0}$) by computation of inner products:

$$j_{mn} := \langle G_{mn} | I \rangle.$$

Typically, one measures with a family of DtG filters up to some order. In this paper, we are concerned with the 2nd order family $\{G_{mn} \mid 0 \leq m + n \leq 2\}$; henceforth *the* DtG family. See Fig. 2. We will write $\vec{G} := (G_{00} \ldots G_{02})^{\mathrm{T}}$ for a vector of these filters. The vector of DtG responses $\vec{j} = (j_{00} \ldots j_{02})^{\mathrm{T}}$ is called a jet and is said to be an element of jet space. We will write $\vec{j}(I) := \langle \vec{G} \mid I \rangle$ for the jet arising from measuring the image I.

The suitability of DtGs as the front-end of an uncommitted vision system arises from the symmetries that individual filters and families possess [8]. First, among these symmetries is a scale symmetry, which manifests in a change of size, but not of shape, when a DtG is blurred with a Gaussian kernel [8]. Second is that the linear span of a family of DtGs is rotationally symmetric [5], a property which is typically called steerability [71]. Steerability allows, for example, a first-order derivative DtG filter, at any desired orientation, to be constructed as a linear sum of any two first-order derivative DtG filters in distinct directions. For example, the sum of the horizontal and vertical first-order derivative DtG filters shown in the middle row of the triangle of filters in the

top-left of Fig. 2 is precisely equal to a DtG filter, which is a first-order derivative in a direction at 45 degrees to the horizontal and vertical. It is worth noting that, in practice, there is no need first to construct a desired filter in this manner, and then to apply it to the image. Instead, because of linearity, one needs only to apply the full basis set of filters, and then linearly combine the filter outputs to compute the responses that would have been obtained had the desired filter actually been synthesized and applied.

4.2 The Intrinsic Component of Jets

The symmetries of a DtG filter family give it the potential to measure intrinsic image structure; but additional work needs to be done to make it explicit, as raw filter responses are still a mixture of intrinsic and extrinsic aspects of image structure. A representation that isolates the intrinsic component of DtG measurements (up to second order) has been developed, and will now be summarized so that it can be used in the remainder of the paper; see [1] for a fuller treatment. The representation (see Fig. 2) works by factoring out of the 6D second order jet space the changes due to the group (denoted by $D_{\infty}(\mathbf{0}) \times A^+(1)$, i.e., the cross product of the infinite Dihedral group with the 1D special Affine group) generated by rotations about the measurement point, reflection in lines through the measurement point, and strictly positive affine rescalings of the intensity (i.e., $v \mapsto \alpha + \beta v, \ \beta > 0$). We regard these changes as extrinsic as they typically arise from changes in the relationship between imager and scene, rather than from scene changes. We do not claim this as the final word in what group best models extrinsic image changes.

The factoring proceeds by considering the foliation of the jet space into orbits, each corresponding to the range of jets obtained as a particular image patch is intensity scaled, reflected, and rotated. A new space—an orbifold—is constructed that has one point for each orbit. Like manifolds, orbifolds are locally like $\mathbb{R}^n$ but they can have boundaries, creases, and corners. The topology of an orbifold arises from the neighborhood relations among the orbits in the factored space. In the case of factoring the 6D jet space, the induced orbifold has a 0D component that corresponds to the orbit, which consists of the jet space origin only, and a 3D component, each point of which corresponds to a 3D orbit in the jet space. When we refer to the orbifold without qualification, we will mean the 3D component; we will be explicit when we refer to the 0D component.

A coordinate system (l, b, a) for the orbifold is given by Griffin [1]:

$$l = \tan^{-1}\left(\frac{c_{20} + c_{02}}{\sqrt{4(c_{10}^2 + c_{01}^2) + ((c_{20} - c_{02})^2 + 4 c_{11}^2)}}\right) \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right],$$

$$b = \tan^{-1}\left(\frac{1}{2}\sqrt{\frac{(c_{20} - c_{02})^2 + 4 c_{11}^2}{c_{10}^2 + c_{01}^2}}\right) \in \left[0, \frac{\pi}{2}\right],$$

$$a = \frac{1}{2}\left|\tan^{-1}\left(2\frac{(c_{01}^2 - c_{10}^2)c_{11} + c_{10} c_{01}(c_{20} - c_{02})}{(c_{01}^2 - c_{10}^2)(c_{02} - c_{20}) + 4c_{10} c_{01} c_{11}}\right)\right| \in \left[0, \frac{\pi}{2}\right],$$

where c_{ij} are scale-normalized derivatives defined by $c_{ij} := \sigma^{i+j} j_{ij}$. Roughly speaking, the l-coordinate expresses

a normalized Laplacean value, and the a and b-coordinates express, respectively, the relative magnitudes and angle between the first and second order structures. In Section 6, and the Appendix, which can be found on the Computer Society Digital Library at http://doi.ieeecomputersociety.org/10.1109/TPAMI.2009.91, we will sometimes find it convenient to use the *shape index* (s) descriptor of pure second order structure [72], [73], [74], which can be computed using $\tan s = \csc b \tan l$.

In [1], we argue for the following norm on jet space, $\|\vec{c}\|_{\sigma} = ((c_{10}^2 + c_{01}^2) + \frac{1}{2}(c_{20}^2 + 2c_{11}^2 + c_{02}^2))^{\frac{1}{2}}$, and from it we induce a metric on the orbifold that can be expressed as the line $dL^2 = dl^2 + \cos^2 l(db^2 + \frac{2\sin^2 2b}{5 - 3\cos 2b} da^2)$. We defined an embedding

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} (b - \frac{\pi}{4}) \cos l \\ (a - \frac{\pi}{4})\sqrt{2} \frac{\cos l \sin 2b}{\sqrt{5 - 3\cos 2b}} \\ -l \end{pmatrix}$$

of the orbifold into euclidean space. Using the line element, we were able to show that this embedding is only mildly distorting, so its use for visualization (Figs. 2, 3, and 4) is not too misleading.

5 DtG FILTERS AND IMAGE SYMMETRY

In Section 2.2, we summarized the possible symmetries of images. In Section 3, we established a link between these symmetries and linear image filters via an easily applicable test of the sensitivity of any given filter to any given symmetry group. In Section 4, we described the DtG family of filters. We now determine the symmetry sensitivity of the DtG family.

The symmetry sensitivity of the DtG family is fully determined by considering the sensitivities of individual linear filters only; there is no need to consider the sensitivity of nonlinear combinations of filter responses. This is because the possible jets (measured as always at the origin) of images with a particular symmetry fill an affine subspace of jet space (follows from P3) and an affine subspace is fully determined by the span of directions in its complementary orthogonal subspace and where the subspace intersects these directions. Each direction corresponds to a filter in the DtG family, which is sensitive to the symmetry, and the value of the intersect corresponds to the filter's critical response.

We can go beyond determining the sensitivity to (mere) symmetry groups, and determine the sensitivity to symmetry types. To do this, we will make use of the jet space factorization of Section 4.2. In the Appendix, which can be found on the Computer Society Digital Library at http://doi.ieeecomputersociety.org/10.1109/TPAMI.2009.91, we provide further details on the method: In this Section, we only describe and interpret the results.

In presenting our results, we will make use of the concept of a *sensitivity submanifold* (SS). This is the affine subspace of jet space in which the jets measured at the origin arising from images with a particular symmetry must lie. We will also use SS to refer to the (topological closure of the) image of the jet space SS when projected into the
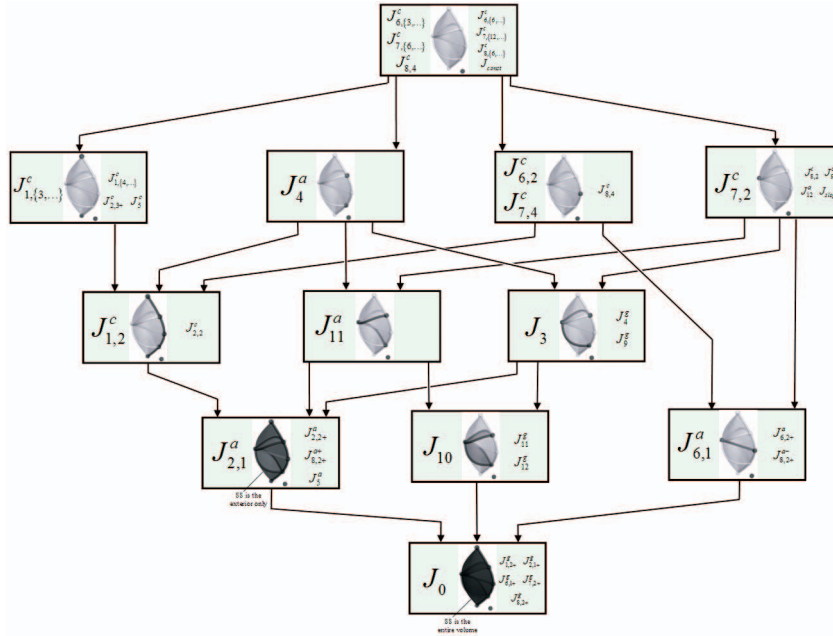
Fig. 3. The sensitivity submanifolds (SSs) of different symmetry types are shown in red. The different possible SSs are arranged in a lattice induced by inclusion relations. The symmetry-type labels correspond to those used in Fig. 1 and [41]. Where necessary, superscripts are used to indicate the spatial relationship between jet measurement point and the symmetry type. The variable c superscript indicates that the measurement point is at a center of rotation; a indicates that it is "aligned" on exactly one line of reflection or antireflection, but not centered ($a+$ or $a-$ is used if necessary to indicate the type of the reflection); g indicates that it is in general position, neither centered nor aligned.

orbifold—context will make it clear which sense is intended. Working with the closure allows for simpler presentation without loss of important detail.

5.1 The Symmetry-Sensitivity Lattice

Fig. 3 shows the complete results on symmetry-type sensitivity of the DtG family. Periodic symmetries are unmentioned in the figure because, as proved in the Appendix, which can be found on the Computer Society Digital Library at http://doi.ieeecomputersociety.org/10.1109/TPAMI.2009.91, DtGs are never sensitive to these.

Consider a particular symmetry type: for example, the centered subtype of $J_{2,2}$ that we denote by $J_{2,2}^c$. As shown in Fig. 1 and [41], this symmetry type consists of invariance to perpendicular mirrors intersecting at the measurement point. Consider a particular group, denoted by K, of this type: for example, the one with horizontal and vertical mirror lines. There will be various filters in the span of the DtG family that are sensitive to K-symmetry, as can be determined using the test of T1. These filters will always give the same response if the image has K-symmetry, or, to put it the other way, there are responses that these filters will not give if the image has K-symmetry; hence, there are restrictions on the possible jets of K-symmetric images.

The restrictions on the jets of K-symmetric images can be mapped into the orbifold. It turns out that there are restrictions there too: The orbifold components of the jets of K-symmetric images must lie on the right-hand crease of the orbifold. Thus, the right-hand crease is the SS of the group K. Moreover, since the orbifold is, by construction, "blind" to orientation, the right-hand crease is the SS of any group of type $J_{2,2}^c$; hence, we say that it is the SS of the $J_{2,2}^c$ symmetry type. This is indicated by the leftmost box of Fig. 3, middle row.

The orbifold is not blind to the position of the measurement point relative to the symmetry. So, for example, the SS for noncentered $J_{2,2}$ groups is different. If the measurement point is on exactly one of the mirrors, we say that the symmetry type is "aligned" and denote it by $J_{2,2}^a$. The SS of $J_{2,2}^a$ is the exterior of the orbifold—this is indicated by the leftmost box of the second row (from the bottom) in Fig. 3. If the measurement point is on neither of the mirrors, we say that the symmetry is in "general" position and denote it by $J_{2,2}^g$—the SS of this is the entirety of the orbifold.

Looking across the different boxes of Fig. 3, one can see that the SS is assembled from a limited collection of components. Some of these components are aligned in an obvious manner to the geometry of the orbifold: its exterior surface ($|l| \leq \frac{\pi}{2} \wedge b \in [0, \frac{\pi}{2}] \wedge a \in \{0, \frac{\pi}{2}\}$), one of the two creases on the exterior ($|l| \leq \frac{\pi}{2} \wedge b = \frac{\pi}{2}$), as shown in the $J_{1,2}^c$ panel, and the pair of apical points at the ends of the creases ($|l| = \frac{\pi}{2}$) as shown in the J_5 pane. Some are less obviously aligned: a pair of leaf-shaped (more precisely lanceolate) internal cross sections ($|s| = \frac{\pi}{4} \wedge (b, a) \in [0, \frac{\pi}{2}]^2$, where s is the shape index defined in Section 4.2), shown in the J_{10} panel; two curves, each made of one boundary from each of the two leaf-shaped cross sections ($b \in [0, \frac{\pi}{2}] \wedge (s, a) \in \{(-\frac{\pi}{4}, 0), (\frac{\pi}{4}, \frac{\pi}{2})\}$ and $b \in [0, \frac{\pi}{2}] \wedge (s, a) \in \{(-\frac{\pi}{4}, \frac{\pi}{2}), (\frac{\pi}{4}, 0)\}$), shown in the J_{11}^a and J_3 panels; and a rod across the interior of the orbifold ($l = 0 \wedge b \in [0, \frac{\pi}{2}] \wedge a = \frac{\pi}{4}$), shown in the $J_{6,1}^a$ panel.

The lattice structure of Fig. 3 is based on inclusion relations between the various SSs. For example, the SS for J_3 is linked to and below that of J_{slope} because $SS(J_3) \supset SS(J_{slope})$. Observe that this relation is similar to the relation $J_3 \subset J_{slope}$. This is an example of a general rule expressed in the following proposition:
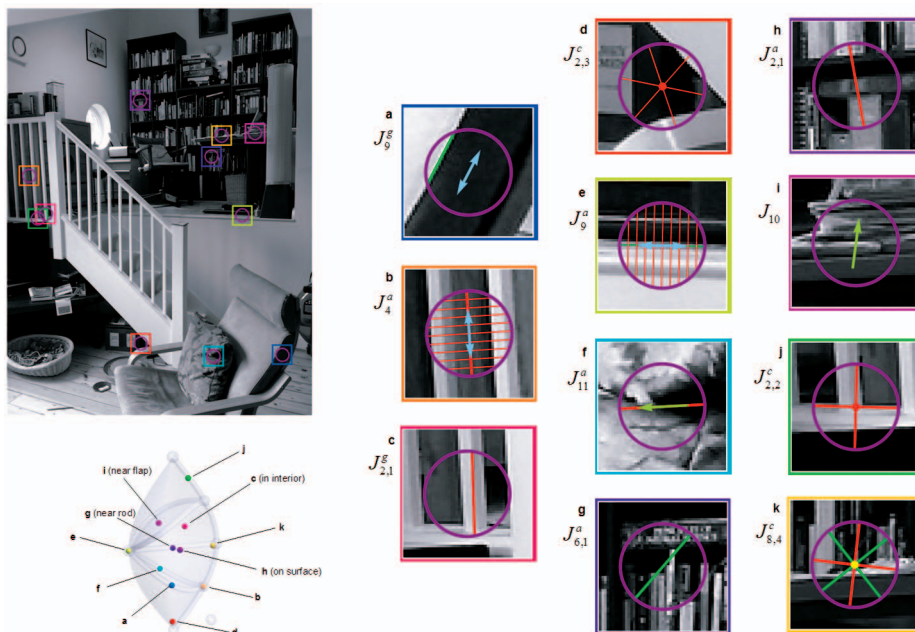
Fig. 4. Eleven sample patches (right) taken from an example gray-level image (left). Patch border color indicates where in the image the patch came from. Patches are labeled on their left with the symmetry that they are judged approximately to have. The orbifold locations of the jet measured at each patch center, at a scale indicated by the purple circles, are indicated by the small spheres in the bottom-left diagram. Patches are ordered (a)-(k) on the basis of the degree of approximation (locally) of the symmetry written to their left; alphabetically, earlier patches have the symmetry more clearly.

Proposition 8 (Symmetry subtype relation implies SS subset relation). *If A, B are symmetry types, then $A \subseteq B \Rightarrow SS(A) \supseteq SS(B)$.*

Proof. Assume the premise. Hence, any B-symmetric image is A-symmetric. Hence, the jets that can arise from measuring B-symmetric images are a subset of the jets that can arise from measuring A-symmetric images. This nesting is trivially preserved under projection into the orbifold.　　　　　　　　　　　　　　□

P8 is strong enough for us to characterize the lattice of Fig. 3 as "compatible" with that of Fig. 1, but its limitations should be borne in mind. In particular, the implied subset relations between SS can be improper (for example, $J_{1,2}^c \subset J_{2,2}^c$ but $SS(J_{1,2}^c) = SS(J_{2,2}^c)$) and the implication of P8 is not from right-to-left (for example, $SS(J_{6,2}^c) = SS(J_{7,4}^c)$ but $J_{6,2}^c \not\subset J_{7,4}^c \land J_{7,4}^c \not\subset J_{6,2}^c$).

5.2　Examples on an Image

In this section, we aim to encourage informal understanding of what it means for *an image to have a symmetry only in so far as a DtG filter family has assessed*. The caveat means that the symmetry is only asserted to manifest *locally* and *partially*.

The "local" part of the caveat arises because a family of DtG filters (effectively) probes only a restricted "local" part of the image. The idea of a symmetry holding only over a *local* region is a variant of the idea of a symmetry holding over only a *limited* region, such as occurs in all real-world tiling patterns [75]. The "partial" part of the caveat arises because assessment of the symmetry is incomplete, even locally. This is because of the much greater number of degrees of freedom that an image has locally than the

number of filters used to probe those degrees of freedom. We are familiar with many examples of partial local symmetry, for instance, the translational symmetry of a row of soldiers. Since the row is finite, the symmetry is only local. Since the soldiers agree in their personhood, dress, and posture, but disagree in other aspects of their appearance, the symmetry is only partial.

Although partial symmetry can be found in commonplace examples, it is not always easy to perceive. To help in understanding the partial local symmetries that DtG filter families determine, we have prepared Fig. 4, which shows 11 example patches extracted from a single gray-level image. The patches were selected by finding image locations where it is unambiguous which is the smallest (tightest) SS close to the orbifold position of the jet. In selecting the 11 example patches, some low-contrast patches were rejected, but otherwise the patches were not hand-selected. Jets were calculated at a scale of $\sigma = 7$. The patches are 49^2 pixels. The overlaid circles, of radius $\sqrt{3 + \sqrt{6}}\sigma \approx 16$ pixels, show the fuzzy region (defined by the outermost inflexion of the filters) that the DtG family roughly probes.

Each patch has been labeled with the symmetry of those corresponding to the tightest nearby SS that, by eyeball measure, it best possesses. An overlay of mirror lines, translation directions, rotation centers, etc., also indicates the symmetry. Again by eye, we have ordered the patches depending on the quality of the symmetry as assessed across the aperture probed by the DtGs. Thus, patch (a) has an excellent J_9^g symmetry, patch (e) has approximate J_5^g symmetry, and patch (i) has rough $J_{2,1}^a$ symmetry.

The message of Fig. 4 is that having a jet compatible with some symmetry guarantees only that the patch, at full resolution, will roughly have the symmetry. The stronger result, already clearly shown by the theorems in this paper, but confirmed by Fig. 4, is that if the image locally has a symmetry that the DtGs are sensitive to, then the orbifold location of the jet will correctly reflect that. We thus stress that what we have mapped out for DtGs are merely sensitivities; we make no claim about their usefulness for symmetry detection in the normally understood sense. Also, we observe that there can be symmetries strongly present in an image (e.g., the periodic symmetry of the banisters in Fig. 4) which do not give rise to strongly constrained jets. This is because the second order DtG family is not sensitive to this symmetry, at least not in the way in which we have defined sensitivity.

6 SUMMARY AND CONCLUSION

We started by describing the filters-to-features problem: the problem of partitioning into qualitatively distinct feature categories, the joint-response space of a collection of linear filters. We explained that the difficulty is that filter response space is a landmark-free vector space, so there is no clue how to carve it up. We reviewed previous approaches to defining structure in filter response space for the purpose of guiding the partitioning. We described how these approaches varied in their reliance on natural image statistics versus geometrical considerations, but that no previous method had been truly purely geometrical. We proposed that consideration of what filter responses told one about the symmetry of the measured image had the potential to define the needed additional structure, in a purely geometrical manner. We stated and proved a theorem that gave an easy test for whether a filter was informative about the presence of a symmetry. We used the test to explore the symmetry sensitivity of DtG filters, a class of filters that we introduced as an important model of biological vision and a useful tool in computer vision. We found that DtG filters showed excellent sensitivity to different symmetries, and a rich network of landmarks in their response space was induced, each corresponding to responses that are reliably obtained whenever the image has particular symmetries.

Our stated goal, defining additional structure on DtG response space, was therefore reached. It remains for future work to determine whether this structure provides a suitable basis for a useful system of feature categories. Optimistically, we recall our argument that symmetry type is not merely an intrinsic property of the image, but a particularly fundamental one; pessimistically, we object that the concern of visual systems is to see through the image to the scene beyond, and we lack an argument that ties image symmetries to scene symmetries, even in a noisy manner.

ACKNOWLEDGMENTS

REFERENCES

[1] L.D. Griffin, "The Second Order Local-Image-Structure Solid," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 29, no. 8, pp. 1355-1366, Aug. 2007.

[2] R.A. Young and R.M. Lesperance, "The Gaussian Derivative Model for Spatial-Temporal Vision: II. Cortical Data," *Spatial Vision,* vol. 14, nos. 3/4, pp. 321-389, 2001.

[3] J. Weickert, S. Ishikawa, and A. Imiya, "On the History of Gaussian Scale-Space Axiomatics," *Gaussian Scale-Space Theory,* pp. 45-59, Kluwer, 1997.

[4] M.A. Georgeson, "From Filters to Features—Location, Orientation, Contrast and Blur," *Higher-Order Processing in the Visual System,* pp. 147-165, Wiley, 1994.

[5] 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.

[6] J.J. Koenderink, "Operational Significance of Receptive-Field Assemblies," *Biological Cybernetics,* vol. 58, no. 3, pp. 163-171, 1988.

[7] J.J. Koenderink and A.J. van Doorn, "Receptive-Field Families," *Biological Cybernetics,* vol. 63, no. 4, pp. 291-297, 1990.

[8] J.J. Koenderink and A.J. van Doorn, "Generic Neighborhood Operators," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 14, no. 6, pp. 597-605, June 1992.

[9] L. Debnath, "On Hermite Transforms," *Mathematicki Vesnik,* vol. 1, no. 16, pp. 285-292, 1964.

[10] L. Debnath, *Integral Transforms and Their Applications.* CRC Press, 1995.

[11] L.M.J. Florack et al., "Scale and the Differential Structure of Images," *Image and Vision Computing,* vol. 10, no. 6, pp. 376-388, 1992.

[12] L.D. Griffin, "Histograms of Infinitesimal Neighbourhoods," *Scale Space and Morphology in Computer Vision,* M. Kerckhove, ed., pp. 326-334, Springer, 2001.

[13] T. Lindeberg and B.M. ter Haar Romeny, "Linear Scale-Space Theory I: Basic Theory," *Geometry-Driven Diffusion in Computer Vision,* B.M. ter Haar Romeny, ed., pp. 1-38, Kluwer, 1994.

[14] B.M.T. Romeny et al., "Higher-Order Differential Structure of Images," *Image and Vision Computing,* vol. 12, no. 6, pp. 317-325, 1994.

[15] B.M. ter Haar Romeny, *Front-End Vision and Multi-Scale Image Analysis.* Kluwer, 2003.

[16] D. Marr, *Vision.* W.H. Freeman & Co., 1982.

[17] L.D. Griffin and M. Lillholm, "Hypotheses for Image Features, Icons and Textons," *Int'l J. Computer Vision,* vol. 71, no. 3, pp. 213-230, 2006.

[18] X.W. Liu and D.L. Wang, "A Spectral Histogram Model for Texton Modeling and Texture Discrimination," *Vision Research,* vol. 42, no. 23, pp. 2617-2634, 2002.

[19] S.-C. Zhu et al., "What Are Textons?" *Int'l J. Computer Vision,* vol. 62, no. 1, pp. 121-143, 2005.

[20] M. Varma and A. Zisserman, "A Statistical Approach to Texture Classification from Single Images," *Int'l J. Computer Vision,* vol. 62, nos. 1/2, pp. 61-81, 2005.

[21] L.D. Griffin, "Feature Classes for 1D, Second 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.

[22] 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.

[23] L.D. Griffin, "Local Image Structure, Metamerism, Norms, and Natural Image Statistics," *Perception,* vol. 31, no. 3, pp. 377-377, 2002.

[24] 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, Springer, 2001.

[25] J.J. Koenderink and A.J. van Doorn, "Receptive Field Assembly Specificity," *J. Visual Comm. and Image Representation,* vol. 3, no. 1, pp. 1-12, 1992.

[26] J.J. Koenderink, "What Is a Feature?" *J. Intelligent Systems,* vol. 3, no. 1, pp. 49-82, 1993.

[27] J.J. Koenderink and A.J. van Doorn, "Metamerism in Complete Sets of Image Operators," *Advances in Image Understanding,* pp. 113-129, Wiley-IEEE CS Press, 1996.

[28] J.J. Koenderink and A.J. van Doorn, "Local Image Operators and Iconic Structure," *Algebraic Frames for the Perception-Action Cycle,* G. Sommer and J.J. Koenderink, eds., pp. 66-93, Springer, 1997.

[29] E. Cassirer, "The Concept og Group and the Theory of Perception," *Philosophy and Phenomenological Research,* vol. 5, no. 1, pp. 1-36, 1944.

[30] F. Klein, "A Comparative Review of Recent Researches in Geometry," *Bull. New York Math. Soc.,* vol. 2, pp. 215-249, translated by M.W. Haskell, 1892.

[31] D. Lewis, "Extrinsic Properties," *Philosophical Studies,* vol. 44, pp. 111-112, 1983.

[32] S.Y. Auyang, *Mind in Everyday Life and Cognitive Sciences.* MIT Press, 2000.

[33] D. Schattschneider, *MC Escher: Visions of Symmetry.* Plenum Press, 1990.

[34] J.J. Koenderink and A.J. van Doorn, "Image Processing Done Right," *Proc. European Conf. Computer Vision,* 2002.

[35] L. Bieberbach, "Über die Bewegungsgruppen der Euklidischen Raume I," *Mathematische Annalen,* vol. 70, p. 297, 1911.

[36] J.H. Conway et al., "On Three-Dimensional Space Groups," *Contributions to Algebra and Geometry,* vol. 42, no. 2, pp. 475-507, 2001.

[37] B. Grünbaum and G.C. Shephard, *Tilings and Patterns.* W.H. Freeman & Co., 1987.

[38] W.T. Holser, "Classification of Symmetry Groups," *Acta Crystallographica,* vol. 14, pp. 1236-1242, 1961.

[39] A.A. Loeb, *Color and Symmetry.* Robert E. Krieger, 1978.

[40] L.D. Griffin, "Symmetries of 1D Images," *J. Math. Imaging and Vision,* vol. 31, nos. 2/3, pp. 157-164, 2008.

[41] L.D. Griffin, "Symmetries of 2D Images: Cases without Periodic Translation," to be published in *J. Math. Imaging and Vision.*

[42] Y. Bonneh, D. Reisfeld, and Y. Yeshurun, "Quantification of Local Symmetry—Application to Texture-Discrimination," *Spatial Vision,* vol. 8, no. 4, pp. 515-530, 1994.

[43] M. Mellor and M. Brady, "A New Technique for Local Symmetry Estimation," *Proc. Conf. Scale Space and PDE Methods in Computer Vision,* pp. 38-49, 2005.

[44] R. Scognamillo et al., "A Feature-Based Model of Symmetry Detection," *Proc. Royal Soc. of London Series B—Biological Sciences,* vol. 270, no. 1525, pp. 1727-1733, 2003.

[45] M. Xu and D. Pycock, "Estimating True Symmetry in Scale-Space," *Proc. 1998 IEEE Int'l Conf. Systems, Man, and Cybernetics,* pp. 4620-4625, 1998.

[46] S.C. Zhu, "Embedding Gestalt Laws in Markov Random Fields," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 21, no. 11, pp. 1170-1187, Nov. 1999.

[47] Y.X. Liu, R.T. Collins, and Y.H. Tsin, "A Computational Model for Periodic Pattern Perception Based on Frieze and Wallpaper Groups," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 26, no. 3, pp. 354-371, Mar. 2004.

[48] S. Oka et al., "VEPs Elicited by Local Correlations and Global Symmetry: Characteristics and Interactions," *Vision Research,* vol. 47, no. 16, pp. 2212-2222, 2007.

[49] S. Sally and R. Gurnsey, "Symmetry Detection across the Visual Field," *Spatial Vision,* vol. 14, no. 2, pp. 217-234, 2001.

[50] S.C. Dakin and A.M. Herbert, "The Spatial Region of Integration for Visual Symmetry Detection," *Proc. Royal Soc. of London Series B—Biological Sciences,* vol. 265, no. 1397, pp. 659-664, 1998.

[51] G.C. Baylis and J. Driver, "Perception of Symmetry and Repetition within and across Visual Shapes: Part-Descriptions and Object-Based Attention," *Visual Cognition,* vol. 8, no. 2, pp. 163-196, 2001.

[52] D.M. Levi and J. Saarinen, "Perception of Mirror Symmetry in Amblyopic Vision," *Vision Research,* vol. 44, no. 21, pp. 2475-2482, 2004.

[53] S.J.M. Rainville and F.A.A. Kingdom, "The Functional Role of Oriented Spatial Filters in the Perception of Mirror Symmetry—Psychophysics and Modeling," *Vision Research,* vol. 40, no. 19, pp. 2621-2644, 2000.

[54] S. Mancini, S.L. Sally, and R. Gurnsey, "Detection of Symmetry and Anti-Symmetry," *Vision Research,* vol. 45, no. 16, pp. 2145-2160, 2005.

[55] S.C. Dakin and R.J. Watt, "Detection of Bilateral Symmetry Using Spatial Filters," *Spatial Vision,* vol. 8, no. 4, pp. 393-413, 1994.

[56] T.-J. Cham and R. Cipolla, "Symmetry Detection through Local Skewed Symmetries," *Image and Vision Computing,* vol. 13, no. 5, pp. 439-450, 1995.

[57] G. Marola and M. Schiaffino, "Using Local Symmetry for Detecting 2-Dimensional Objects," *Proc. Int'l Conf. Signal Processing,* pp. 699-702, 1990.

[58] M. Milgram, R. Belaroussi, and L. Prevost, "Multi-Stage Combination of Geometric and Colorimetric Detectors for Eyes Localization," *Proc. Conf. Image Analysis and Processing,* pp. 1010-1017, 2005.

[59] O. Hansen and J. Bigun, "Local Symmetry Modeling in Multidimensional Images," *Pattern Recognition Letters,* vol. 13, no. 4, pp. 253-262, 1992.

[60] K.G. Cho and S.M. Dunn, "Hierarchical Local Symmetries," *Pattern Recognition Letters,* vol. 12, no. 6, pp. 343-347, 1991.

[61] M. Brady and H. Asada, "Smoothed Local Symmetries and Their Implementation," *Int'l J. Robotics Research,* vol. 3, no. 3, pp. 36-61, 1984.

[62] V. Braitenberg, *Vehicles. Experiments in Synthetic Psychology.* MIT Press, 1984.

[63] D. Perrett, "Symmetry and Human Facial Attractiveness," *Evolution and Human Behavior,* vol. 20, no. 5, pp. 295-307, 1999.

[64] M. Park et al., "Performance Evaluation of State-of-the-Art Discrete Symmetry Detection Algorithms," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* pp. 1-8, 2008.

[65] Y. Liu et al., "Digital Papercutting," *Proc. ACM SIGGRAPH,* 2005.

[66] S.M. Pizer et al., "Object Shape before Boundary Shape: Scale-Space Medial Axes," *J. Math. Imaging and Vision,* vol. 4, pp. 303-313, 1994.

[67] G. Loy and J.O. Eklundh, "Detecting Symmetry and Symmetric Constellations of Features," *Proc. European Conf. Computer Vision,* pp. 508-521, 2006.

[68] V. Prasad and L. Davis, "Detecting Rotational Symmetries," *Proc. Int'l Conf. Computer Vision,* pp. 346-352, 2005.

[69] H. Blum, "Biological Shape and Visual Science (Part I)," *J. Theoretical Biology,* vol. 38, pp. 205-287, 1973.

[70] B.A. Olshausen and D.J. Field, "How Close Are We to Understanding V1?" *Neural Computation,* vol. 17, no. 8, pp. 1665-1699, 2005.

[71] W.T. Freeman and E.H. Adelson, "The Design and Use of Steerable Filters," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 13, no. 9, pp. 891-906, Sept. 1991.

[72] 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.

[73] J.J. Koenderink and A.J. van Doorn, "Local Structure of Gaussian Texture," *IEICE Trans. Information and Systems,* vol. IE86D, no. 7, pp. 1165-1171, 2003.

[74] M. Lillholm and L.D. Griffin, "Statistics and Category Systems for the Shape Index Descriptor of Local Image," *Image and Vision Computing,* 2008.

[75] A. Weinstein, "Groupoids: Unifying Internal and External Symmetry," *Notices of the Am. Math. Soc.,* vol. 43, no. 7, pp. 744-752, 1996.

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

Martin Lillholm received the BSc degree in computer science and mathematics from the University of Copenhagen, and the MSc degree in 2002 and the PhD degree in computational vision in 2004 from the IT University of Copenhagen. He worked two years as an assistant professor at the IT University before taking up his current employment as a research fellow in the Department of Computer Science at the University College London. He is currently working on visual alphabets in the context of image description and object recognition.

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