

Basic Image Features (BIFs) Arising from Approximate Symmetry Type

Lewis D. Griffin¹, Martin Lillholm¹, Mike Crosier¹, and Justus van Sande²

¹ Computer Science, University College London, London WC1E 6BT, UK

² Biomedical Engineering, Eindhoven University of Technology, The Netherlands
l.griffin@cs.ucl.ac.uk

Abstract. We consider detection of local image symmetry using linear filters. We prove a simple criterion for determining if a filter is sensitive to a group of symmetries. We show that derivative-of-Gaussian (DtG) filters are excellent at detecting local image symmetry. Building on this, we propose a very simple algorithm that, based on the responses of a bank of six DtG filters, classifies each location of an image into one of seven Basic Image Features (BIFs). This effectively and efficiently realizes Marr's proposal for an image primal sketch. We summarize results on the use of BIFs for texture classification, object category detection, and pixel classification.

Keywords: Gaussian Derivatives, Hermite Transform, Group Theory.

1 Introduction

Previous schemes for detection of image symmetry are fairly complex [1-6]; requiring, for example, comparison of the outputs of filters at *multiple* positions. Herein we show that symmetries may be detected by *single* linear filters. Building on this we present a simple algorithm that computes a Marr-type primal sketch [7] by categorizing local image structure according to its approximate symmetry.

The paper is organized as follows. In section 2 we present results on image symmetries. In 3 we show how to test whether a linear filter is sensitive to a symmetry. In 4 we review image measurement with derivative-of-Gaussian (DtG) filters. In 5 we consider the symmetry-sensitivity of these DtG filters. In 6 we show how this sensitivity gives rise to a system of Basic Image Features (BIFs). In 7 we summarize results on using BIFs for texture categorization, object category detection and pixel classification. In 8 we conclude. Sections 2-5 are a distillation of work published, in press and under review in fuller form elsewhere [8-14]; 6 is new; parts of 7 have been presented or are under review in fuller form elsewhere [9, 11].

2 Image Symmetries

Symmetry of a structure (X) is always relative to some class of admissible transformations. A structure is said to have a symmetry when a non-trivial group of

admissible transformations, known as the automorphism group, each leave it indistinguishable from the original. This is denoted by $Aut[X] := \{t \mid t \circ X = X\}$.

Considering images, an obvious class of transformations are the spatial isometries; and the possible symmetries, relative to this class, have long been catalogued [15-17]. A broader class of transformations, where each spatial isometry is combined with a permutation of a finite set of image ‘colour’ values, has also been considered. These allow the symmetries of, for example, Escher’s ‘Reptiles’ to be expressed [18]. The gamut of possible ‘colour symmetries’ has been fully determined [19, 20].

We propose that the class of ‘image isometries’, defined as a spatial isometry combined with an intensity isometry, is appropriate for images. We write an image isometry as $g = (i, s)$, where $i: \mathbb{R} \rightarrow \mathbb{R}$ is an intensity isometry, and $s: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a spatial isometry. Such an image isometry is applied to an image $I: \mathbb{R}^2 \rightarrow \mathbb{R}$ according to $g \circ I = i \circ I(s \circ _) = i(I(s(_)))$. Choosing a class of transformations is tantamount to choosing a geometry [21], and the geometry that corresponds to the class of image isometries has previously been considered for images [22] and much earlier, abstractly, as one of a larger class of possible geometries [23].

We have employed a method for determining the possible automorphism groups of images, relative to the class of image isometries. The method relies on two results. First, that the projection of a group of image isometries onto their spatial or intensity components in both cases makes a group. Second, that (except for a special case) the intensity projection group must be isomorphic to a factor group of the spatial projection group [8]. Using the method, we have determined the possible automorphism groups of 2-D images, except for cases that contain discrete periodic translations. A summary of these possible symmetries, together with our notational system is shown in fig. 1. The symmetries include: familiar ones, such as reflectional ($J_{2,1}$), reflect-and-negate ($J_{6,1}$), and Yin-Yang type ($J_{7,2}$); simple but often ignored ones, such as variation in one direction only (J_3); simple but novel, such as continuous translate-and-increment in one direction, plus a line of reflection parallel to that direction (J_{11}); and some wholly novel, such as continuous translate-and-increment in one direction, plus a continuous line of centres of Yin-Yang type symmetry (J_{12}).

3 Sensitivity of Linear Filters to Symmetries

Detection of a symmetry seems to require multiple measurements, but this is incorrect. Consider a +1/-1 filter, such as used in finite-difference schemes. When positioned so that it straddles a putative line of reflection, a necessary criterion for the symmetry is that the filter gives a 0 response. We generalize this: a filter F is *sensitive to a symmetry* K if it gives the same response to all images that have the symmetry (i.e. $\exists f \in \mathbb{R} \quad Aut[I] \supseteq K \Rightarrow \langle F|I \rangle = f$). This definition is impractical because it requires assessment across all images. However, we have found a necessary and sufficient test that requires only a single integral to be computed. We present this below in Theorem 1, after introducing some notation.




Fig. 1. The group/subgroup lattice of the possible image symmetries, excluding those with discrete periodic translation

We use an inner product notation ($\langle F|I \rangle := \int_{\bar{x} \in \mathbb{R}^2} F(\bar{x})I(\bar{x})$) to denote the measurement of an image $I: \mathbb{R}^2 \rightarrow \mathbb{R}$ by a filter $F: \mathbb{R}^2 \rightarrow \mathbb{R}$; and we define an operator $F^{(K)} := \sum_{(i,s) \in K} |i\rangle s \circ F$ which, roughly speaking, ‘smears’ a filter by a group.

Theorem: Symmetry-Sensitivity Test for Filters

F is sensitive to K if and only if $\langle F|F^{(K)} \rangle = 0$

Proof

A formal proof will be published elsewhere [14]. Intuitively the truth of the theorem can be understood as follows. The signal that a filter ‘sees’ best is a copy of itself. Of all the symmetric signals, a symmetrised version of the filter should be the most easily seen. If the filter cannot see a symmetrised version of itself, then it is insensitive to the symmetry.

4 Gaussian Derivative Filters

Gaussian Derivative (DtG) filters are defined in 1-D by

$$G_\sigma(x) := (2\pi\sigma^2)^{-\frac{1}{2}} e^{-\frac{x^2}{2\sigma^2}}, \quad G_\sigma^{(n)}(x) := \frac{d^n}{dx^n} G_\sigma(x) = \left(\frac{-1}{\sigma\sqrt{2}}\right)^n H_n\left(\frac{x}{\sigma\sqrt{2}}\right) G_\sigma(x)$$

where H_n is the n th Hermite polynomial; and in 2-D by

$$G_\sigma^{(m,n)}(x, y) := G_\sigma^{(m)}(x) G_\sigma^{(n)}(y).$$

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

$$j_{mm} = \langle G_\sigma^{(m,n)}|I \rangle.$$

Typically, one measures with a family of DtG filters up to some order e.g. the 2nd order family $\{G_\sigma^{(m,n)} \mid 0 \leq m+n \leq 2\}$. Scale-normalized filter responses $c_{pq} := \sigma^{p+q} j_{pq}$ make later equations simpler.

The suitability of DtGs as the front-end of an uncommitted computational vision system arises from the symmetries that individual filters and families possess [24]. First amongst these is a scale symmetry, which manifests as a change of size, but not of shape, when a DtG is rescaled by blurring with a Gaussian kernel. Second is that the linear span of a family of DtGs is rotationally symmetric.

The responses of a bank of DtG filters entangle intrinsic and extrinsic aspects of image structure. For example, an in-plane rotation of the image about the measurement point causes the DtG responses to change. A representation that

disentangles these aspects for measurement up to 2nd order has been developed [13]. The representation works by factoring out of the 6-D 2nd order DtG response space the changes due to the group of image isometries that fix the measurement point and do not invert the intensity axis which we denote $D_{\infty}(\mathbf{0}) \times A^+(1)$.

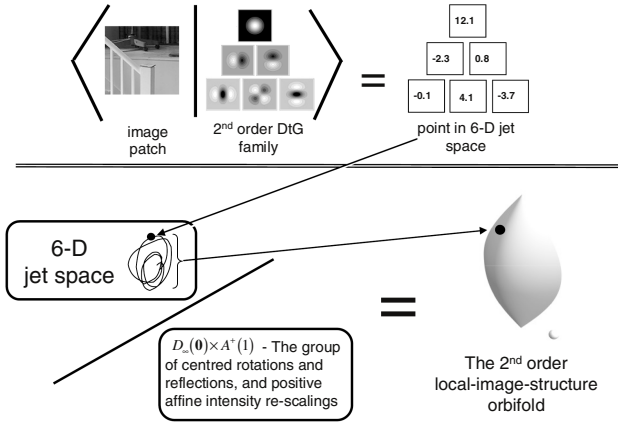


Fig. 2. The top part illustrates schematically the probing of an image patch by a bank of DtG filters resulting in a point in jet space; the bottom, the factoring of the jet space by a group of transformations resulting in the local-image-structure orbifold

The result is an orbifold – a type of manifold with boundaries, creases and corners allowed – consisting of a 3-D and a 0-D component (figure 2). The intrinsic aspect of a 6-tuple of filter responses corresponds to a particular location in the orbifold, and is invariant to rotating the image about the measurement point, reflecting it in a line through the measurement point, or affinely scaling the intensity. When the responses of the 1st and 2nd order DtG filters are all zero, the intrinsic aspect is the 0-D part of the orbifold; all other responses map to the 3-D component. A coordinate system $(l, b, a) \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right] \times \left[0, \frac{\pi}{2}\right] \times \left[0, \frac{\pi}{2}\right]$ for the 3-D component is given by [13]:

$$l = \arctan\left(\sqrt{4(c_{10}^2 + c_{01}^2) + (c_{20} - c_{02})^2 + 4c_{11}^2}, c_{20} + c_{02}\right)$$

$$b = \arctan\left(2\sqrt{c_{10}^2 + c_{01}^2}, \sqrt{(c_{20} - c_{02})^2 + 4c_{11}^2}\right)$$

$$a = \frac{1}{2} \left| \arctan\left(\left((c_{01}^2 - c_{10}^2)(c_{02} - c_{20}) + 4c_{10}c_{01}c_{11}, 2\left((c_{01}^2 - c_{10}^2)c_{11} + c_{10}c_{01}(c_{20} - c_{02})\right)\right)\right) \right|$$

The orbifold has been equipped with a metric, induced by one on the filter response space, which expressed as a line element in the lba -system is $ds^2 = dl^2 + \cos^2 l \left(db^2 + da^2 2(5 - 3\cos 2b)^{-1} \sin^2 2b \right)$. The orbifold is intrinsically curved, but it can be embedded into Euclidean 3-space with only mild distortion.

5 Symmetry-Sensitivity of DtG Filters

Using the elements of sections 2-4, we can determine which DtG filters are sensitive to which symmetries. We consider not just canonical filter forms (e.g. an x -derivative) but any linear combination of filters in the 2nd order filter family. This allows us to determine the symmetry-sensitivity of the entire filter family, independent of the particular basis filters used. For example, while the x -derivative filter is sensitive to a reflectional symmetry with a vertical mirror line through the measurement point, the x - and y -derivatives together are sensitive to any reflectional symmetry in a line through the measurement point.




Fig. 3. The sensitivity-submanifolds (SS) of different symmetry types are shown in red. The different possible SS are arranged in a lattice induced by inclusion relations. The symmetry type labels correspond to those used in figure 1. Superscripts indicate the spatial relationship between the symmetry and the origin: a c indicates origin-centred rotation; an $a+$ that the origin is contained in a line of reflection, but is not a centre of rotation; similarly for $a-$ and anti-reflections; a g indicates general position, neither centred nor aligned. All symmetries labelled in a box have the indicated SS; those on the left are minimal.

The filter family sensitivities can be projected into the orbifold to determine where the intrinsic component of the jet responses must lie whenever the image has any of a class of symmetries equivalent by conjugation with an element of $D_{\infty}(\mathbf{0}) \times A^+(1)$. We

call the restricted set of possible responses the sensitivity-submanifold (SS). For example the SS is the orbifold exterior ($a = 0 \wedge a = \pi/2$) for reflectional symmetry in a mirror through the measurement point. The results are summarized in fig 3.

6 Symmetry-Based Basic Image Features (BIFs)

We have used the symmetry sensitivities of the DtG filters as a starting point in defining a set of Basic Image Features (BIFs) that realize Marr’s idea of a primal sketch of image structure, in a computationally simple scheme. We do not claim that the scheme is derived as rigorously as the results on symmetry sensitivity.

Our scheme works by considering the orbifold projection of jets, and classifying them according to the SS that they are closest to i.e. we define a Voronoi cell partitioning of the orbifold with the SS as cell centres. We find that this works best when only seven 0-D SS (the first and second rows of figure 3) are used, though we cannot justify this beyond that it produces nice results. The resulting orbifold partitioning is shown in the top-left of figure 4.




Fig. 4. Top left: the partitioning of the orbifold into BIF categories. Bottom left: BIFs calculated across a range of scales for a simple image of a figure ‘8’; in each cube scale increases right-to-left. Lower cubes sectioned for visualisation. Right: an example complex greyscale image, with BIFs calculated at one particular scale.

The orbifold distance to the six of these SS that lie in the 3-D component of the orbifold are simple to compute; for example, the distance to the $J_{7,2}^c$ SS is $\tan^{-1}\left(\frac{1}{2}c_{20}^2 + c_{11}^2 + \frac{1}{2}c_{02}^2\right) / \left(c_{10}^2 + c_{01}^2\right)$. To find which distance is shortest it is

computationally equivalent but simpler to find which of six quantities is maximum. The distance to the seventh SS, which corresponds to the origin of jet space where all the 1st and 2nd derivative filters have zero response, is not well-defined. We incorporate it into our scheme by using a multiple of the 0th order jet response. The full resulting scheme for computing BIFs is as follows.

i) compute scale-normalized DtG filter responses as described in section 7.

ii) compute $\lambda := \frac{1}{2}(c_{20} + c_{02})$ and $\gamma := \sqrt{\frac{1}{4}(c_{20} - c_{02})^2 + c_{11}^2}$

iii) classify according to which is the largest of

$$M = \left\{ \varepsilon \cdot c_{00}, \sqrt{c_{10}^2 + c_{01}^2}, \lambda, -\lambda, (\gamma + \lambda)/\sqrt{2}, (\gamma - \lambda)/\sqrt{2}, \gamma \right\}.$$

In our scheme the only free parameters, that have to be tuned to the application are the filter scale σ and ε which controls the amount of image classified as flat; a setting of $\varepsilon = 0.05$ is an effective default. For display purposes we find the following colour scheme effective: if $\varepsilon \cdot c_{00}$ is the largest of M then colour the pixel pink; if

$\sqrt{c_{10}^2 + c_{01}^2}$ is largest colour it grey; then black, white, blue, yellow and green.

7 Example Applications Using BIFs

We summarize results on using BIFs for texture, object and pixel classification.

7.1 Texture Classification

Textures are often classified based on a representation of them by a histogram over a texton vocabulary [25-29]. Textons are categorical patch classifications [25, 30]. To define the texton vocabulary, a space of patch descriptions is typically Voronoi partitioned into on-the-order-of 1000 texton categories, usually around centres found by k-means clustering of the responses from many images. Textures are then classified by nearest-neighbour matching of histograms.

We have investigated the classification performance of an approach in which images are labelled using spatial complexes of BIFs instead of Voronoi cells in a local description vector space. Our approach is (i) simpler because we have eliminated the clustering step needed to produce a dictionary of features, and (ii) faster because we assign image patches to histogram bins without having to use a high-dimensional nearest-neighbour computation.

We call the spatial complexes of BIFs that we use analogously to textons, Basic Image Patterns (BIPs). The type of BIP that we have found effective for texture description is a scale-template of the BIFs at the same location but at four, octave-separated scales. Unlike *spatial*-template BIPs, these *scale*-templates retain the rotation invariance of BIFs, which has been shown [30] to be advantageous in texture classification tasks. For textures, we do not use the pink/flat BIF category, so four scales produces a $6^4=1296$ bin histogram representation, which seems to capture the right trade-off between specificity and generality (see figure 5).




Fig. 5. Left: An image from the CURET 'polyester' texture class. Centre: BIFs computed at four octave-separated scales, stacked to form an array of 'column-BIPs'. Right: Occurrence histogram of column-BIPs from every position in the image".

Our method has been tested on the CURET texture dataset [31]. As reported in [9], the simple column-BIP representation and nearest-neighbour matching using the Bhattacharyya distance correctly classifies $98.2 \pm 0.1\%$ of the remaining 49 images per class, which is at least as good as other methods using nearest-neighbour classifiers. Extending this method by using a multi-scale histogram comparison [9] results in an improvement to $98.6 \pm 0.1\%$ on CURET, which is comparable to methods [27] using SVMs for classification; and produces what are, to the best of our knowledge, the best reported results [9] on the more challenging UIUCTex [32] and KTH-TIPS [27] datasets, which include variations in scale.

7.2 Object Categorization

Texton approaches have also been shown to be useful for object categorization [28, 33]. Similar to texture, the 'standard' approach is to partition a patch descriptor space, such as that used by SIFT [34, 35], into on-the-order-of 1000 categories (visual words) and then to describe each image to be analyzed by what visual words it contains, and to use machine learning techniques to determine a classifier that can predict the category of object based on such descriptions.

We have conducted preliminary experiments to assess whether visual words built from BIFs could be used rather than SIFT-space categories. As with texture this would be simpler and faster. For our initial experiments, we have labelled pixels according to their BIF type and, inspired by SIFT, with an orientation, quantized at the $\pi/4$ level. The orientation depends on the BIF type: grey BIFs have one of eight possible orientations based on 1st order structure; yellow, green and blue BIFs have one of four possible orientations based on 2nd order structure; black, white and pink BIFs are unoriented. Thus we have twenty-three possible orientation-augmented BIF (oBIF) labels. oBIFs are a natural 2nd order generalization of the gradient orientation alphabet typically used in SIFT [34, 35]. See figure 6 for an example image and calculated oBIFs.

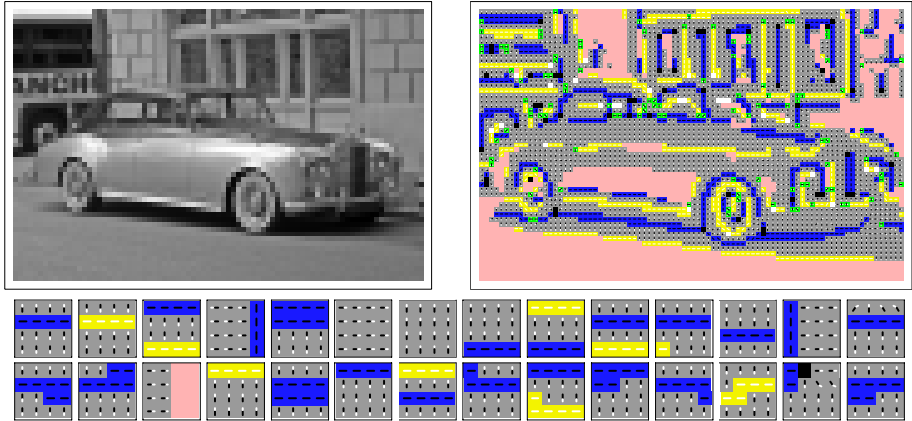


Fig. 6. The top row (left) shows an image from the PASCAL challenge, labeled with direction-augmented BIFs at right. On the bottom are shown the 4x4 template BIPs whose occurrence in an image most informatively signals the presence of a car.

We have tested three different types of visual word, which when built from BIFs or oBIFs we call Basic Image Patterns (BIPs); two based on geometrical partitioning of patch space and one based on more standard data-driven quantization. Each BIP system has been used with simple un-optimised-of-the-shelf classifiers and applied to the 20-class PASCAL VOC 2008 [33] object recognition challenge dataset. Our score in figure 7 is based on a late fusion of the three schemes and is mid-field: above other first-time entrants and below well-optimised veteran entries. Using the PASCAL VOC 2008 [33] dataset, examples of 4x4 template BIPs whose presence in images is approximately independent, and which are maximally informative for the ‘car’ category are shown in figure 6.




Fig. 7. Results for the PASCAL VOC 2008 challenge. Each bar in the chart is a challenge entry - our result is highlighted.

7.3 Pixel Classification

Many image problems involve inferring one of a small class of labels for each pixel of an image. For complex images with unpredictable global structure, most approaches balance the likelihood of the labels, given the local image structure, and the likelihood of the local arrangement of inferred labels. In both cases the likelihoods are computed on the basis of statistics learnt from groundtruth-labelled training data.

We have experimented with the use of BIFs in the computation of label likelihoods given the image i.e. ignoring the likelihood of arrangements of inferred labels. For our experiments we have used 2-D Electron Microscopy images of neuronal grey matter tissue, stained to enhance neuronal membranes. We trained on four images with hand-drawn groundtruth data, indicating the position of membranes, and evaluated on a further four images. We use a k -Nearest Neighbour (k -NN) approach to classification. NN classification starts by compiling a list of descriptors of all the patches in the training data, together with the groundtruth label of the pixel in the patch centre. The classifier is used by extracting a patch around each pixel to be classified, forming a description of it, comparing the description to each the compiled descriptions, finding the k which are most similar, and assigning the pixel being analyzed with the label associated with the majority of the k .

We evaluated a baseline solution based on pixel values. The distance between two patch descriptions is simply the Euclidean distance between their blurred pixel values, minimized over allowing one patch to be rotated or reflected into eight configurations. We jointly optimize blur, patch size and k . The best settings that we find are: no blur, 7×7 patches, and $k=14$. At these settings membrane-labelled pixels overlap (intersection divided by union) with the groundtruth by 48%.

Our solution uses a patch of BIF labels as a patch descriptor. The distance between two descriptors is simply the number of pixels where the label does not agree. As in the baseline, we minimize the distance over one of the patches being rotated or reflected. We jointly optimize the scale (σ) at which the BIFs are computed, the parameter ε which controls the amount of the flat BIF class, patch size and k . The best settings that we find are $\sigma=1.2$, $\varepsilon=0.15$, 9×9 patches, and $k=10$. At these settings we achieve an overlap of 55%. See figure 8.




Fig. 8. Typical results of our pixel-classification system

So, using BIF- rather than greylevel-description raises the score from 48% to 55%. Computation is also faster because the k NN lookup dominates the cost of computing patch descriptions, and with BIFs the distances that need to be computed are of a Hamming rather than Euclidean type.

8 Conclusions

We have derived a scheme for classifying image structure into one of seven BIF types based on the outputs of a bank of six DtG filters. Applied to an entire image, the output realizes Marr's notion of an image primal sketch. Presented results show that BIF description is simple, fast and effective for texture, object and pixel classification.

The BIF system was derived by considering the sensitivity of DtG filters to image symmetry. Although the final algorithm is pleasingly simple, there are some weak points in the derivation of the BIFs from symmetry sensitivities. Specifically, why are only the 0-D SS considered, how exactly does orbifold distance correspond to degree of failure of symmetry, why should least-approximate local symmetry be an effective feature label? We hope that the foundation of symmetry-sensitivity of DtGs can eventually answer all of these questions in a scheme where arbitrary choice has been eliminated. Such a scheme will be extendable to higher-order filter families (where appeal to visual evidence and past practice are less effective), for which a richer alphabet of feature labels is to be expected. We predict that such a richer alphabet will give more effective solutions in the application areas that we have reviewed.

References

1. Liu, Y.X., Collins, R.T., Tsin, Y.H.: A computational model for periodic pattern perception based on frieze and wallpaper groups. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26(3), 354–371 (2004)
2. Scognamillo, R., et al.: A feature-based model of symmetry detection. *Proceedings of the Royal Society of London Series B-Biological Sciences* 270(1525), 1727–1733 (2003)
3. Mellor, M., Brady, M.: A new technique for local symmetry estimation. In: Kimmel, R., Sochen, N.A., Weickert, J. (eds.) *Scale-Space 2005*. LNCS, vol. 3459, pp. 38–49. Springer, Heidelberg (2005)
4. Bonneh, Y., Reissfeld, D., Yeshurun, Y.: Quantification of local symmetry - application to texture-discrimination. *Spatial Vision* 8(4), 515–530 (1994)
5. Mancini, S., Sally, S.L., Gurnsey, R.: Detection of symmetry and anti-symmetry. *Vision Research* 45(16), 2145–2160 (2005)
6. Baylis, G.C., Driver, J.: Perception of symmetry and repetition within and across visual shapes: Part-descriptions and object-based attention. *Visual Cognition* 8(2), 163–196 (2001)
7. Marr, D.: *Vision*. W H Freeman & co., New York (1982)
8. Griffin, L.D.: Symmetries of 1-D Images. *Journal of Mathematical Imaging and Vision* 31(2-3), 157–164 (2008)
9. Crosier, M., Griffin, L.D.: Texture classification with a dictionary of basic image features. In: *CVPR 2008*. IEEE, Los Alamitos (2008)
10. Lillholm, M., Griffin, L.D.: Statistics and category systems for the shape index descriptor of local image. *Image and Vision Computing* (in press) (2008)
11. Lillholm, M., Griffin, L.D.: Novel image feature alphabets for object recognition. In: *ICPR 2008* (2008)
12. Griffin, L.D.: Symmetries of 2D images: cases without periodic translation. *Journal of Mathematical Imaging and Vision* (in press)

13. Griffin, L.D.: The 2nd order local-image-structure solid. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29(8), 1355–1366 (2007)
14. Griffin, L.D., Lillholm, M.: Symmetry-sensitivity of derivative of gaussian filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (in press)
15. Bieberbach, L.: Über die bewegungsgruppen der euklidischen raume I. *Mathematische Annalen* 70, 297 (1911)
16. Conway, J.H., et al.: On three-dimensional space groups. *Contributions to Algebra and Geometry* 42(2), 475–507 (2001)
17. Grünbaum, B., Shephard, G.C.: *Tilings and Patterns*. WH Freeman & co., New York (1987)
18. Schattschneider, D.: *MC Escher. Visions of Symmetry*. Plenum Press (1990)
19. Holser, W.T.: Classification of symmetry groups. *Acta Crystallographica* 14, 1236–1242 (1961)
20. Loeb, A.A.: *Color and Symmetry*. Robert E. Krieger (1978)
21. Klein, F.: A comparative review of recent researches in geometry (trans. by MW Haskell). *Bulletin of the New York Mathematical Society* 2, 215–249 (1892)
22. Koenderink, J.J., van Doorn, A.J.: Image processing done right. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) *ECCV 2002*. LNCS, vol. 2350, pp. 158–172. Springer, Heidelberg (2002)
23. Cayley, A.: Sixth memoir upon the quantics. *Philosophical Transactions of the Royal Society* 149, 61–70 (1859)
24. Koenderink, J.J., van Doorn, A.J.: Generic Neighborhood Operators. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14(6), 597–605 (1992)
25. Varma, M., Zisserman, A.: Texture classification: are filter banks necessary? In: *CVPR 2003*. IEEE, Los Alamitos (2003)
26. Varma, M., Zisserman, A.: A statistical approach to texture classification from single images. *International Journal of Computer Vision* 62(1), 61–81 (2005)
27. Hayman, E., et al.: On the significance of real-world conditions for material classification. In: Pajdla, T., Matas, J(G.) (eds.) *ECCV 2004*. LNCS, vol. 3024, pp. 253–266. Springer, Heidelberg (2004)
28. Zhang, J., et al.: Local features and kernels for classification of texture and object categories: a comprehensive study. In: *CVPR 2006* (2006)
29. Perronnin, F., et al.: Adapted vocabularies for generic visual categorization. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) *ECCV 2006*. LNCS, vol. 3954, pp. 464–475. Springer, Heidelberg (2006)
30. Varma, M., Zisserman, A.: *Unifying Statistical Texture Classification Frameworks*. *Image and Vision Computing* (in press) (2005)
31. Cula, O.G., Dana, K.J.: Compact representation of bidirectional texture functions. In: *CVPR 2001*. IEEE, Los Alamitos (2001)
32. Lazebnik, S.C., Schmid, C., Ponce, J.: A spare texture representation using local affine regions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27(8), 1265–1278 (2005)
33. Csurka, G., et al.: Visual categorization with a bag of keypoints. In: *ECCV 2004*, pp. 1–22 (2004)
34. Lowe, D.G.: Towards a computational model for object recognition in IT cortex. In: *Biologically Motivated Computer Vision, Proceeding*, pp. 20–31 (2000)
35. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60(2), 91–110 (2004)