

A Computational Model for Periodic Pattern Perception Based on Frieze and Wallpaper Groups

Yanxi Liu, *Member, IEEE*, Robert T. Collins, *Member, IEEE*, and Yanghai Tsin

Abstract—We present a computational model for periodic pattern perception based on the mathematical theory of crystallographic groups. In each N -dimensional Euclidean space, a finite number of symmetry groups can characterize the structures of an infinite variety of periodic patterns. In 2D space, there are seven *frieze groups* describing monochrome patterns that repeat along one direction and 17 *wallpaper groups* for patterns that repeat along two linearly independent directions to tile the plane. We develop a set of computer algorithms that “understand” a given periodic pattern by automatically finding its underlying lattice, identifying its symmetry group, and extracting its representative motifs. We also extend this computational model for near-periodic patterns using geometric AIC. Applications of such a computational model include pattern indexing, texture synthesis, image compression, and gait analysis.

Index Terms—Periodic pattern, frieze group, wallpaper group, symmetry group, lattice, tiles, motifs, gait analysis.

1 INTRODUCTION

SYMMETRY is a pervasive phenomena in natural and man-made environments. Humans have an innate ability to perceive symmetry of patterns [50], [20], [36], yet it is not obvious how to automate this powerful insight. This paper studies symmetries of periodic patterns in a plane. Periodic and near-periodic patterns can be found in regular textures, indoor and outdoor scenes (e.g., brick walls, tiled surfaces, textiles, windows on buildings, and cars in a parking lot), and intermediate data representations (e.g., spatiotemporal patterns formed from human and animal gaits).

Crystallographic group theory is a mature mathematical theory for analyzing periodic patterns [9], [10]. The theme of this paper is to develop concepts from crystallographic group theory into computer algorithms that can automatically analyze patterns in real images. Mathematically speaking, a *symmetry* of a subset S of Euclidean space R^n is a rigid transformation in R^n that keeps S setwise invariant. The set of all rigid transformations that are symmetries of a pattern has a group structure, and is called the *symmetry group* of the pattern [34], [29]. An essential mathematical fact about periodic patterns is the answer to Hilbert’s 18th question: Despite an infinite variety of instantiations for periodic patterns, a finite set of symmetry groups completely characterizes the possible structural symmetry of any periodic pattern spanning n dimensions [1]. In particular, seven *frieze groups* describe all patterns generated by translation along one dimension [11], [43], 17 *wallpaper groups* describe all patterns generated by two linearly independent translations [39], and 230 *space groups* describe regular crystal patterns generated by three linearly independent translations [10], [6].

Our computational model for periodic pattern perception based on crystallographic groups includes three main components:

1. recovering the underlying translational lattice of the pattern,
2. classifying the symmetry group of the pattern, and
3. identifying a representative motif that perceptually characterizes the pattern.

We do not address automatic detection of a periodic pattern within a larger image; many papers in computer vision have addressed this problem, e.g., [21], [19], [38], [47]. By using group theory to reason about the interrelationships between translation, rotation, reflection, and glide reflection¹ symmetries of a periodic pattern, we gain a much deeper understanding of the patterns than previous algorithms have achieved. The authors of [2] use the mathematical tiling theory to analyze texture, but they do not take advantage of the rich description of reflection and rotation symmetry afforded by symmetry group theory. Much work treats specific types of symmetry in isolation, for example, bilateral symmetry [7], [33], [14], [48], rotation symmetry [31], [23], [52], [35], [46], translation symmetry [15], [21], [19], [38], or reflection and translation symmetries [47]. Glide-reflection symmetry has yet to be dealt with computationally. Although mathematical education tools have existed for some time (e.g., [51]) to synthesize a periodic pattern based on rules derived from an understanding of group theoretic pattern structure, we are the first to offer a computational approach to *analyze* symmetry group structure in images.

In this paper, Section 2 addresses automatic lattice extraction from an image that is known to contain a periodic pattern. Section 3 addresses classification of the symmetry group of a frieze or wallpaper pattern. Section 4 explores the use of geometric AIC to classify noisy, near-periodic patterns into their closest symmetry groups.

1. **Glide-reflection** means symmetries that are composed of a translation (half the size of its minimum translation generator) along the reflection axis followed by a reflection about the axis.

• The authors are with the Robotics Institute, School of Computer Science, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213. E-mail: {yanxi, rcollins, tsin}@cs.cmu.edu.

Manuscript received 25 June 2002; revised 15 July 2003; accepted 8 Aug. 2003.

Recommended for acceptance by R. Nelson.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 116853.

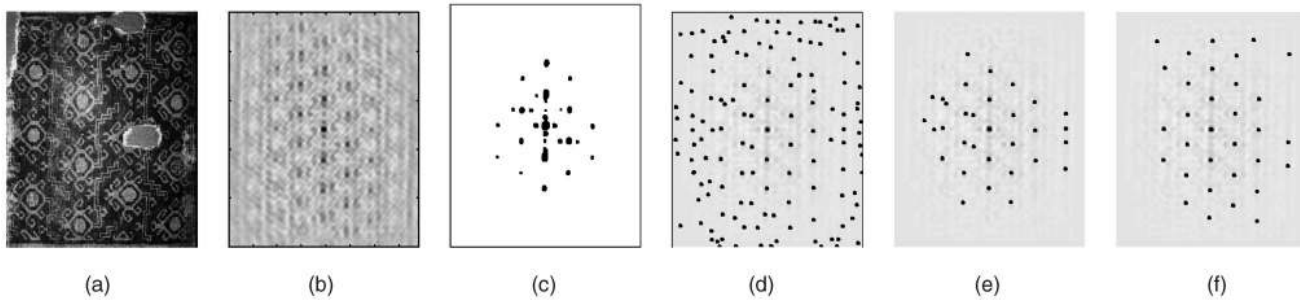


Fig. 1. (a) Original image of a rug. (b) An autocorrelation surface. (c) Peaks found using a global threshold. (d) Peaks extracted using the threshold-free method of [22]. (e) The highest 32 peaks from those return by [22]. (f) The 32 most-dominant peaks found using our approach described in the text.




Fig. 2. Even noise-free images can be hard to process. (a) An incorrect lattice found by the algorithm of Lin et al. [22]. (b) Correct lattice found by our algorithm. (c) Frieze pattern and (d) its 1D autocorrelation response, used to explain how spurious peaks can form (see text for details).

Section 5 shows how knowledge of a pattern's symmetry group can guide extraction of a representative *motif* that exhibits the same local symmetries within a single tile as the global pattern, and conforms well with human perception. Section 6 discusses limitations and potential applications.

2 TRANSLATIONAL LATTICE DETECTION

The key issue in periodic pattern analysis is whether the 2D lattice of the pattern can be correctly extracted. Previous work on lattice detection can be clustered into three approaches. One approach is to extract a sparse set of features and to hypothesize links (translations) between them based on visual similarity or conformance to a parametric model [21], [19], [38]. The benefits of this approach are the ability to detect small regions of a repeating pattern within a larger image, and to group pattern elements despite local surface deformations (such as the folds of a shirt). The drawback is the need for a pattern with distinct corner/edge/contour features. The more traditional image processing approach to detecting global pattern repetition is to use autocorrelation [22], the Fourier transform [37], or periodicity measures defined over cooccurrence matrices [2], [54], [42], [44]. These approaches work for any intensity image, not just ones with strong features. The main drawback is the assumption that a single periodic pattern occupies a large portion of the image, limiting the approach to analysis of patterns that have already been segmented in some other way. For periodic patterns with a low number of repeating cycles (two to three cycles), we have found autocorrelation to be a more appropriate method for quantifying translational symmetry than the Fourier transform, an observation also made in [22]. A third approach, used in structural texture analysis, is based on the idea of a *unit pattern* together with a set of well-defined placement rules. However, the generality and computational tractability of this work is limited in its present form: unit patterns are either regions centered about a local maximum that is bounded on all

sides by local minima [5] or square texture regions with an unspecified window size [32]. We show later in this paper how group theory can be used to automatically specify tile shape, orientation, and placement rules for any periodic pattern.

2.1 The Problem of Peak Detection

The first step in analyzing a periodic pattern is to extract a set of linearly independent vectors that describe the translational symmetry of the pattern. Our approach is to look for peaks in the autocorrelation surface of the pattern.

Fig. 1a shows an image of a rug and Fig. 1b shows its autocorrelation surface.² Although the grid of peaks in Fig. 1b is apparent to the human eye, finding it automatically is very difficult. Simple approaches such as setting a global threshold yield spurious results in Fig. 1c. The trouble is that many legitimate grid peaks have a lower value than some of the spurious peaks. Lin et al. [22] present a threshold-free approach based on performing Gaussian smoothing of the image, followed by selection of local maxima (Fig. 1d). It is hard to see the grid structure interspersed with the spurious points. If we sort peaks by height (correlation score) and take the first 32, we have the answer in Fig. 1e. In comparison, Fig. 1f presents the top 32 peaks resulting from our peak detection algorithm (Section 2.2).

Even relatively noise-free computer-generated patterns, such as Fig. 2, can cause problems for lattice detection algorithms. Fig. 2a shows an incorrect lattice found by the algorithm of [22], together with the correct lattice found by our algorithm (Fig. 2b). Obviously, Lin et al.'s algorithm has picked up smaller peaks between the major ones and, indeed, the pattern's autocorrelation surface has high ridges between the major peaks. Our experience with a variety of periodic patterns indicates that it is common to find spurious peaks of comparable height to desired peaks

2. Note: Autocorrelation generates a correlation surface C_I twice the size of the original image I , i.e., if I is $x \times y$, C_I is $2x \times 2y$.




Fig. 3. Detected lattices for three real-world patterns.

superimposed over the autocorrelation image at twice (or other multiples of) the frequency of the lattice grid structure. An illustration of how this can happen is shown by the frieze pattern in Fig. 2c, displayed next to the 1D autocorrelation response of the pattern when slid along its axis of translation. The correct peaks are A, C, and E. Halfway between actual lattice translations, the large features in the pattern partially match smaller features interspersed between them, causing the spurious peaks B and D to form. Furthermore, these spurious peaks can have higher value than actual peaks located at the periphery of the autocorrelation image (e.g., height of peak B is greater than height of peak in E). These difficulties are exacerbated by complicated patterns in two dimensions.

2.2 Regions of Dominance

It is a nontrivial task to find a proper set of peaks in an autocorrelation surface of a periodic pattern. We have wrestled with the problem of peak finding in many contexts over several years. Our observation is that the absolute height of a peak is not as important as the size of its **region of dominance**, defined as the largest circle centered on the candidate peak such that no higher peaks are contained in the circle. A peak with a low height, but located far from any larger neighbors, is much more perceptually important than a high peak that is close to an even higher one. Referring back to Fig. 2d, the true peak E is lower than the spurious peak B, but is located twice as far away from any higher peak than B is and, thus, dominates a larger region. Revisiting Fig. 1f, the first 32 most-dominant peaks found using our method are well distributed over the whole image, with very few spurious peaks.

A list of peaks in decreasing order of dominance can be computed using a simple, order N^2 algorithm, where N is the number of candidate peaks to be considered. In our

case, these initial peaks P_1, P_2, \dots, P_N are computed using nonmaximum suppression over a sliding $M \times M$ window, where M is a constant 5 in our implementation, but could be chosen based on the scale of the pattern to yield fewer initial candidates. First, sort candidate peaks in descending order of peak height to yield a list Q_1, Q_2, \dots, Q_N . Next, for each Q_j , compute the distance to each $Q_i, 1 \leq i < j$ that comes before it in the list, and denote the minimum distance D_i . Finally, sort the list of peaks again in descending order of D_i , the minimum distance to a higher peak. The peaks are now arranged in decreasing order of dominance.

This approach to peak detection has proven to work well on a diverse set of autocorrelation images. The method generalizes readily to any dimension and is potentially useful in other vision contexts where multiple peaks must be detected within accumulated noisy data, for example, finding modes in intensity or color histograms to perform region segmentation.

2.3 Determine the Shortest, Independent Translation Vectors

Having a set of candidate lattice points extracted as dominant peaks in the autocorrelation surface, the next task is to find the shortest linearly independent translation vectors that generate the lattice. For frieze patterns, this is a single vector; for wallpaper patterns, we need two vectors. Finding lattice vectors is complicated by missing points as well as additional spurious points interspersed with the good data. We use a Hough transform approach similar to [22] to find the two shortest translation vectors that best explain the majority of the point data, but include wallpaper pattern constraints such as requiring that the angle between the two vectors must be between 60 and 90 degrees [39]. Fig. 3 shows detected lattices for some real-world patterns.

TABLE 1
Symmetries of the Seven Frieze Patterns

P_1		Symmetry Group	180° Rotation	Horizontal Reflection	Vertical Reflection	Glide Reflection
P_2		F1	no	no	no	no
P_3		F2	no	no	no	yes
P_4		F3	no	no	yes	no
P_5		F4	yes	no	no	no
P_6		F5	yes	no	yes	yes
P_7		F6	no	yes	no	no
		F7	yes	yes	yes	no




Fig. 4. The unit lattices for the 17 wallpaper groups (courtesy of [39]).

3 SYMMETRY GROUP CLASSIFICATION

In this section, we review the definitions and properties of frieze and wallpaper groups, leading to an algorithm (Section 3.3) for determining the symmetry group of a periodic pattern. A symmetry g of a 2D periodic pattern P is a distance preserving mapping (translation, rotation, reflection, or composition of these) that maps every pixel in the pattern to a pixel of the same gray-value or color such that $g : R^2 \times I \Rightarrow R^2 \times I$, and $g(P) = P$, where I can either be gray values in the range of $[0, 255]$ or triplets of RGB intensity values. Note that a periodic pattern requires the existence of a *nontrivial* translation symmetry, which excludes patterns consisting of identical stripes along the direction of translation. The set of all symmetry transformations of P comprises the pattern's symmetry group.

3.1 Frieze Groups

A *frieze pattern* is a 2D strip in the plane that is periodic along one dimension. Any frieze pattern P is associated with one of seven unique symmetry groups F_i , where $i = 1, \dots, 7$, and $\forall g \in F_i, g(P) = P$. These seven symmetry groups are called the *frieze groups*, and their properties are summarized in iconic and tabular form in Table 1. Without loss of generality, assume the direction of translation symmetry of a frieze pattern is horizontal, the frieze pattern can exhibit five different types of symmetries:

1. horizontal translation,
2. 2-fold rotation (rotation by 180 degrees),
3. horizontal reflection (reflection axis is horizontal),
4. vertical reflection, and
5. horizontal glide-reflection, composed of a half-unit translation horizontally followed by a horizontal reflection.

TABLE 2

Wallpaper Group Classification: The Group Associated with a Wallpaper Pattern Can Be Determined by Checking the Small Number of Symmetries: 180, 120, 90, or 60 Degree Rotation Symmetry, Reflection Symmetry, and Glide-Reflection Symmetry about Axes Parallel to Lattice Unit Parallelogram Boundary Vectors T_1 and T_2 , or Diagonal Vectors D_1 and D_2

	p1	p2	pm	pg	cm	pmm	pmg	pgg	cmm	p4	p4m	p4g	p3	p3m1	p31m	p6	p6m
180		Y				Y	Y	Y	Y	Y	Y	Y				Y	Y
120										Y	Y	Y	Y	Y	Y	Y	Y
90										Y	Y	Y					
60																Y	Y
T_1			Y	Y(g)		Y	Y(g)	Y(g)			Y	Y(g)			Y		Y
T_2						Y	Y	Y(g)			Y	Y(g)			Y		Y
D_1					Y				Y		Y	Y		Y	Y		Y
D_2									Y		Y	Y					Y

"Y" means the symmetry exists for that symmetry group; empty space means no. Y(g) denotes a glide reflection.




Fig. 5. A hand-drawn frieze pattern from [45]. (a) Original image, (b) the median tile, and (c) match scores for tested frieze symmetries.

A frieze pattern can be classified into one of the seven possible frieze groups based on what combination of these five primitive symmetries are present in the pattern (Table 1). Not all possible combinations of symmetries form legitimate symmetry groups, for example, a frieze pattern cannot exhibit both horizontal reflection and glide-reflection symmetries simultaneously [11], [43].

3.2 Wallpaper Groups

A periodic pattern extended in two linearly independent directions to cover the 2D plane is known as a *wallpaper pattern*. The two smallest linearly independent translation vectors T_1 and T_2 in the pattern's symmetry group are generators for the underlying lattice structure of the pattern. The lattice divides the plane into identical parallelogram-shaped subimages, called **lattice units** or **tiles**. The symmetry group of a wallpaper pattern has to be one of the 17 distinct wallpaper groups [39], [41]. Fig. 4 shows a diagram from [39] that depicts precisely the unit lattice shape for each of the 17 wallpaper groups, with the geometric configuration of translation generators, rotation, reflection, and glide-reflection symmetries superimposed for each group. A lattice unit is typically chosen with centers of highest order of rotation at the vertices.

The practical value of understanding the 17 wallpaper groups is that correct pattern classification can be performed after verifying the existence of only a small set of symmetries, specifically four rotations (180, 120, 90, and 60 degrees), and four reflections along axes parallel to either unit lattice parallelogram boundaries T_1 and T_2 or unit lattice diagonals D_1 and D_2 . It is clear from Table 2 that each symmetry group corresponds to a unique sequence of yes/no answers to whether the pattern contains each of these eight types of fundamental symmetry (an additional test may be needed when reflections are present to determine if they are "proper" or "glide" reflections).

3.3 Group Classification Algorithm

We have constructed an algorithm to automatically classify the symmetry group of a given 2D periodic pattern by checking a small set of rotation and reflection symmetries listed in Table 1 for frieze patterns and Table 2 for wallpaper patterns. The algorithm is robust to moderate amounts of pixel noise and outliers. For patterns with large amounts of intensity or geometric distortion, the Geometric-AIC algorithm described in Section 4 will be more effective.

Input: an image dominated by a frieze or a wallpaper pattern.

Output: symmetry group of the input periodic pattern and its median tile.

Algorithm:

- 1. Find the pattern lattice:** Compute the underlying pattern lattice using the algorithm described in Section 2.
- 2. Estimate a median tile and noise model:** Cut out a set of tile-shaped regions from the input pattern using the overlaid lattice. Choose one of these as a reference tile and register it with all other tiles in the set using a sum of squared difference (SSD) measure. This yields a set of corresponding intensity measurements for each pixel in the tile. A median tile is obtained by assigning each pixel the median value of its corresponding intensity measurements. Also, the pixel noise level is estimated by computing the standard deviation σ of the residuals between all pixels and their corresponding median tile values.
- 3. Test symmetries:** Test for the existence of rotation and reflection symmetries in the pattern. Four symmetries are tested for frieze groups (Table 2). The steps taken for each symmetry are:
 - (a) Apply the symmetry to the original image (e.g., rotate by 180 degrees) to obtain a transformed image I' .
 - (b) Correlate the median tile with image I' at all lags to get a rough registration map.
 - (c) Start at the point with highest correlation value and register the median tile with I' by finding a small translation that minimizes SSD registration error.
 - (d) At the position of best registration, compute the trimmed normalized residual error

$$d = \sum_{k=1}^{N'} e_k = \sum_{k=1}^{N'} \frac{(m_k - i_k)^2}{\sigma^2},$$

where N is the total number of pixels in the tile, $N' = (1 - b)N$ is a smaller number of pixels as determined by b , the *trim rate*, which is the percentage of pixels to be discarded from the end of the e_k error sequence when e_k is sorted in ascending order (the higher on the queue the noisier the pixels), m_k, i_k are corresponding pixel intensity values of the median tile and image I' , respectively, and σ is the standard deviation of the pixel noise model.

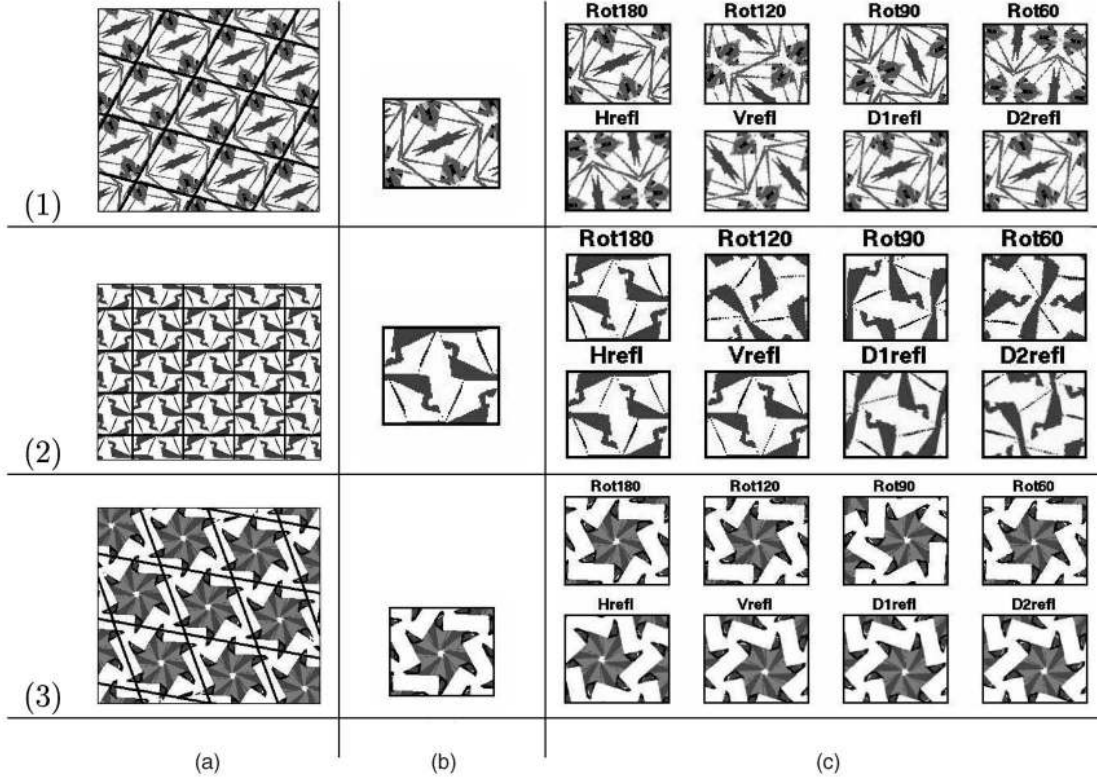


Fig. 6. Three sample patterns. (a) Original image overlaid with detected lattice. (b) Median tile. (c) Best matched positions of the median tile on the transformed images.

- (e) Repeat the computation of trimmed normalized residual errors d_i at neighboring lattice points, and keep the error value $d_{med} = \text{median}\{d_i\}$ which is the median among all computed errors. The idea behind this step is to guard against accepting accidental good alignments between the median tile and transformed image as evidence of the existence of symmetry. A proper symmetry must also preserve the original image's lattice structure.
- (f) If we assume that pixel values are corrupted by independent Gaussian noise with mean 0 and standard deviation σ , then d_{med} should obey a $\chi^2(N')$ distribution with N' degrees of freedom. Evaluate whether the tested symmetry exists by comparing d_{med} to a threshold t_0 where $\int_0^{t_0} \chi_{N'}^2(x) dx = 0.99$. The symmetry is said to exist if $d_{med} < t_0$, otherwise, not. Note that $t_0 \approx N'$, for large N' , and thus an approximate test for symmetry is whether $d_{med}/N' < 1.0$.
- (g) When a reflection symmetry is found to exist, decide if it is a glide reflection or not by examining the offset of the location of best registration between the median tile and the transformed image I' to the location where the center of the original reference tile is mapped to in I' . This offset should be roughly an integer multiple of one of the lattice vectors if we have a proper reflection, otherwise, it falls roughly halfway between integer multiples, and is labeled as a glide reflection.
- 4. Classify the symmetry group:** Validate the symmetry test results against the symmetries listed in Table 1 or Table 2 to classify the symmetry group of the pattern.

3.4 Symmetry Group Classification Examples

The examples in this section serve to illustrate the symmetry group classification algorithms for frieze and wallpaper patterns.

Example 1: A Hand-Drawn Pattern. Fig. 5a shows a frieze pattern scanned in from [45], and Fig. 5b shows the estimated median tile. The original pattern was hand-drawn and contains many small geometric imperfections, for example, missing lines on the second "vase" from the left, and missing internal ear markings on the rightmost face. Computation of a median tile and the use of a trim rate term was designed in our algorithm to compensate for these types of errors. All of the results in this section were computed using a trim rate of 0.1. Fig. 5c contains a table of match scores for testing different types of frieze symmetries. The values in the table are reported as d_{med}/N' , which was motivated earlier as being an approximate χ^2 test. A table element less than 1.0 (marked in bold font) indicates the presence of the symmetry. Based on these results and Table 1, the pattern's symmetry group is classified as F3.

Example 2: Synthetic Patterns. Fig. 6 shows three sample wallpaper patterns adapted from a set of synthetic wallpaper images from [12]. We have successfully processed all the 17 wallpaper groups from the wallpaper patterns on this site [26]. Fig. 6a shows the detected lattice for each pattern, overlaid on the pattern. The computed median tile is shown in Fig. 6b. For ease of representation we always use a rectangular median tile, corresponding to the bounding box of the actual tile shape. To test for the presence of a rotation or reflection symmetry, this median tile is correlated with rotated and reflected versions of the

TABLE 3
The Matching Scores of the Eight Symmetry Tests (Table 2) for the Three Sample Patterns in Fig. 6

Example	rot180	rot120	rot90	rot60	T1 refl	T2 refl	D1 refl	D2 refl	Group
(1)	0.00	6.04	6.04	6.50	4.69	4.79	0.25	0.31	<i>cm</i>
(2)	0.00	11.03	13.76	12.18	0.00 (g)	0.00 (g)	12.16	13.21	<i>pgg</i>
(3)	0.00	0.00	8.77	0.00	5.51	5.49	5.49	5.42	<i>p6</i>

original pattern. The positions of highest correlation for each of the four rotations and four reflections that need to be tested (Table 2) are shown in Fig. 6c.

We can see by inspection that pattern (1) has 180 degree rotation and reflection about the tile diagonals. This observation is echoed by the numeric values returned by the algorithm, shown in Table 3, thus the the pattern is classified as having wallpaper symmetry group *cm*. Pattern (2) also has both 180 degree rotation symmetry and two reflection symmetries. However, the reflections in this case are about the unit lattice vectors rather than the lattice diagonals, and further processing determines that these are glide reflections rather than proper reflections. Referring to Table 2, the pattern is determined to have the *pgg* wallpaper symmetry group. Pattern (3) in Fig. 6 is found to have 180, 120, and 60 degree rotation symmetries, and no reflection symmetry. Referring to Table 2, the pattern is determined to have the *p6* wallpaper symmetry group.

Example 3: Real-World Patterns. Figs. 7, 8, and 9 demonstrate various types of wallpaper patterns captured from the real environment around us. These include patterns of rugs, metal gates, honeycomb, windows, tiled floors, tiled walls, and chrome surfaces. Due to length limitations, we present details on only one real-world pattern, a photograph of a damaged oriental rug (Fig. 10). The table of symmetry test scores for this example is shown below

rot180	rot120	rot90	rot60	T1 refl	T2 refl	D1 refl	D2 refl	Group
0.57	1.34	1.15	1.25	1.48	1.16	0.46	0.57	<i>cm</i>

From these results, the algorithm concludes that the pattern has 180 degree rotation symmetry and reflection symmetry about the lattice diagonals. Referring to Table 2, the pattern is determined to have *cm* wallpaper symmetry group. Note that this real-world classification example is not as clear-cut as earlier synthetic pattern examples, where the existence of symmetries is indicated by a value close to zero. The scores for both symmetries and nonsymmetries of the real patterns may lie closer to the threshold 1.0 due to the high amount of noise in the image.

Testing symmetry using a median tile becomes necessary for treating real-world periodic patterns when the pattern intensity is very noisy. Since the median tile is a better description of the underlying pattern, we could alternatively “regenerate” the original image using the median tile and use this new image to test for the presence of rotation and reflection symmetries. A generalization of this alternative approach is presented in the next section.

4 MODEL SELECTION USING GEOMETRIC AIC

Complications of symmetry group classification for periodic patterns arise from two main sources: 1) real-world patterns

can be very noisy, thus departing from ideal frieze or wallpaper patterns, and 2) symmetry groups have a hierarchical relationship among themselves. This second observation refers to the fact that symmetry groups are *not* disjoint, mutually exclusive classes—some symmetry groups are subgroups of others. Kanatani points out in [17] that each symmetry group should be given a fair chance to be selected, otherwise, a classification algorithm faced with ambiguous data always favors the most general class. In this section, we address the first issue by defining a distance measure between a given pattern and a family of perfect frieze (wallpaper) patterns. We address the second issue by using geometric AIC for symmetry group model selection. The result is an alternative symmetry group classification algorithm that is more effective at handling near-periodic real-world patterns. We use frieze patterns for the purposes of illustration, but a similar approach can be applied to wallpaper groups [27].

4.1 Symmetry Group Distance and Geometric-AIC (G-AIC)

Given a near-periodic frieze pattern P with t tiles, we define a set of *perfect frieze tiles* $\{P_n\}$, $n = 1..7$ for P , as follows:

1. For $t > 1$ and $n = 1$, P_n is the pixel-wise average of all the tiles in P .
2. For $t = 1$ and $n > 1$, $P_n = \frac{(F_n(P)+P)}{2}$, where $F_n(P)$ is the pattern obtained by applying the set of symmetry operations in F_n to P .
3. For $t > 1$ and $n > 1$, P_n is the pixel-wise average of each P_n obtained above.

Fig. 11 shows how a tile P is transformed into a perfect tile P_n for each of the seven frieze groups.

Now, we define a **symmetry group distance** (SD) of a near-frieze pattern P to each of its perfect frieze patterns $\{P_n\}$ as

$$SD_n(P) = \min_{p_i \in P, q_i \in P_n} \left\{ \sum_{i=1}^{tN} \left(\frac{p_i - q_i}{s_i} \right)^2 \right\}, \quad (1)$$

where N is the number of pixels in a tile (smallest 2D repeating region), t is the number of tiles being studied, p_i and q_i are intensity values of corresponding pixels of near-frieze pattern P and perfect frieze pattern P_n , respectively, and s_i is the standard deviation of the frieze pattern at pixel i . SD_n thus represents the minimum SSD between a near-frieze pattern P and any frieze pattern in $\{P_n\}$. For independent Gaussian noise, the distance SD_n has a χ^2 distribution with tN degrees of freedom. It can be proven that the perfect frieze tile P_n defined above has the minimal distance to P among all frieze tiles with symmetry group F_n [18]. Our definition of frieze pattern symmetry group distance is analogous to that of Zabrodsky et al. [53]




Fig. 7. Real-world patterns (left column) processed by our algorithm with overlaid detected lattices (right column). The symmetry group of each pattern is classified as: (a) Rug has group $p1$. (b) Rug has group cm . (c) and (d) are both metal gates with symmetry group $cm\bar{m}$. Note the background clutter visible through the gaps of the metal gates.

for polygons. The difference is that we are dealing with pattern intensity variations, while the authors of [53] compute vertex locations of polygons.

The frieze symmetry groups form a hierarchical structure (Fig. 12a) where frieze group F_1 is a subgroup of all the other groups, and so on. Frieze groups F_5 and F_7 are the two least general symmetry groups for frieze patterns. If no care is taken, a symmetry group classification algorithm based on raw symmetry group distance scores will always favor a more general class, say F_1 , over a more special class, say F_5 and F_7 . To address this problem, we adopt the

concept of Geometric-AIC (G-AIC) proposed by Kanatani [16], [17]. The *degrees of freedom* (DOF) of a frieze pattern depends on how the intensity (or color) of each pixel on a tile is constrained. For frieze patterns with translation symmetry only, there is no constraint on the value of any of the N pixels on a tile, thus the DOF is N . On the other hand, pixels on a P_3 pattern have an additional vertical reflection symmetry constraint to satisfy and, thus, half of the pixel intensities need to be the same as the other half. The DOF of a P_3 pattern is therefore $N/2$. Fig. 12a shows the frieze

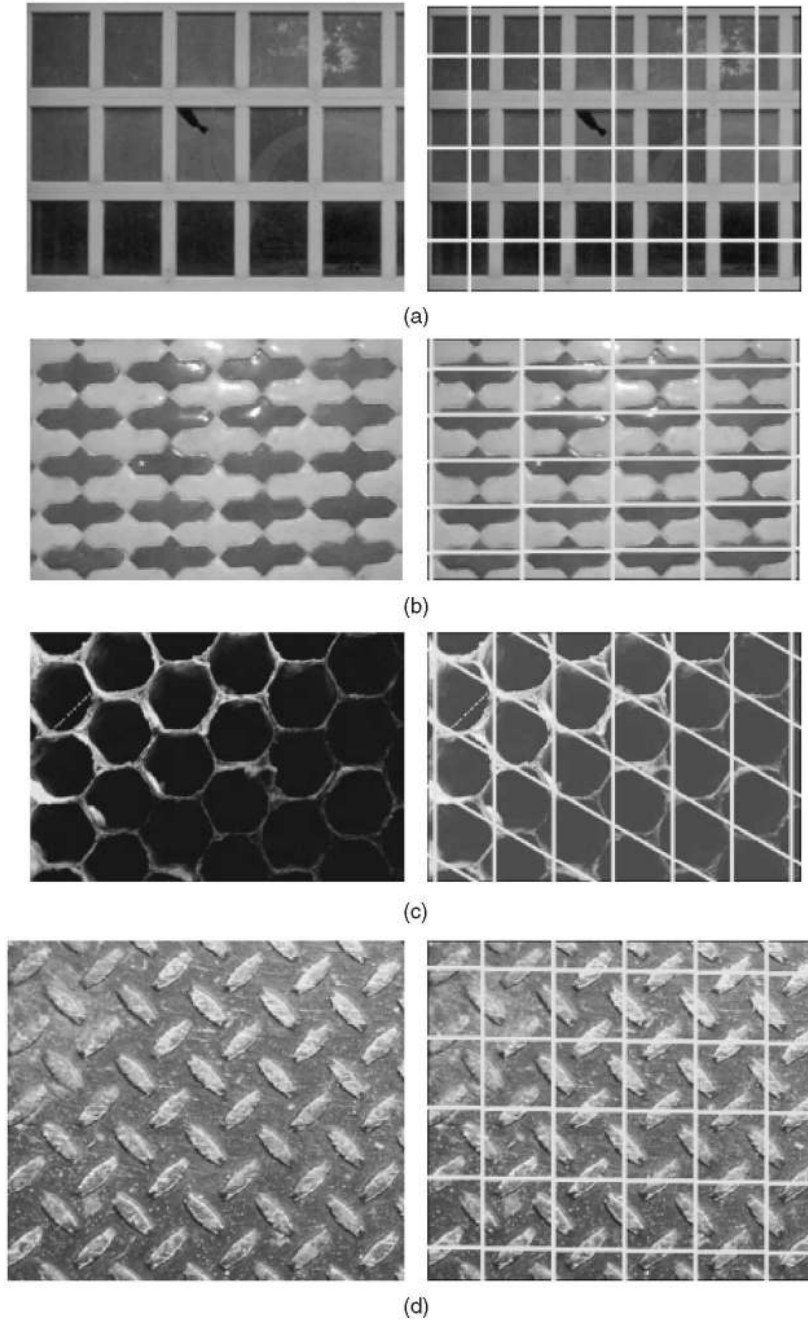


Fig. 8. More real-world patterns. Both (a) windows and (b) tiles have symmetry group $p6m$. (c) Honeycomb has one of the most complicated symmetry groups, $p6m$ (lower right corner in Fig. 4). (d) This nonskid metal surface has symmetry group $p4g$. Note pattern imperfections caused by lighting and surface defects in all of these examples.

group hierarchy and Fig. 12b illustrates the DOFs and the pixel constraints for each of the seven frieze groups.

G-AIC states that group F_m is preferred over F_n for a near-frieze pattern P if

$$\frac{SD_m(P)}{SD_n(P)} < 1 + \frac{2(d_n - d_m)}{r(tN) - d_n}, \quad (2)$$

where d_m and d_n are the degrees of freedom for frieze patterns of F_m and F_n , respectively, and r is the codimension. Since the data space (the intensity space) is dimension one, and our model space (point in multidimensional

intensity space) dimension is 0, the codimension $r = 1 - 0 = 1$.

In summary, P has frieze group F_m rather than F_n if

$$\frac{SD_m(P)}{SD_n(P)} < \frac{t}{t-1}, \quad \text{for } m = 2, 3, 4, 6 \text{ and } n = 1, t > 1 \quad (3)$$

$$\frac{SD_m(P)}{SD_n(P)} < \frac{2t}{2t-1}, \quad \text{for } m = 5, 7 \text{ and } n = 2, 3, 4, 6, t \geq 1 \quad (4)$$

$$\frac{SD_m(P)}{SD_n(P)} < \frac{2t+1}{2t-2}, \quad \text{for } m = 5, 7 \text{ and } n = 1, t > 1. \quad (5)$$

We combine these concepts of symmetry group distance and geometric AIC into the following new algorithm for




Fig. 9. More real-world patterns: (a) is a rug with group pm . (b) and (c) have slight affine and projective distortions in their cloth and rug patterns, respectively, with symmetry group pmm . (d) Chrome, (e) metal gate, and (f) tiles all have symmetry group cm . (g) Tiled wall has symmetry group $p4m$ (another complicated symmetry group, see Fig. 4).

frieze group classification (a similar algorithm can be constructed for wallpaper groups).

Input: an image dominated by a near-frieze pattern P

Output: the frieze group corresponding to the pattern, and a “perfect” tile that conforms exactly to that frieze group and is closest to the original pattern in a least squares sense.

Algorithm:

1. **Find the pattern lattice:** Compute the underlying pattern lattice using the algorithm described in Section 2.
2. **Construct perfect tiles:** Cut out a set of tile-shaped regions from the input pattern using the overlaid lattice. Choose one of these as a reference tile and register it with all other tiles in the set using the SSD measure. This yields a set of corresponding intensity measurements for each pixel in the tile. For each frieze group




Fig. 10. (a) Original image overlayed with detected lattice. (b) Median tile. (c) Positions of best match of median tile with transformed images.




Fig. 11. A perfect frieze tile can be constructed from a given tile P1 for each of the seven frieze groups.




Fig. 12. (a) The subgroup relationship among the seven frieze symmetry groups ($F_1 \dots F_7$ in Table 1), where $F_i \rightarrow F_j$ means F_i is a subgroup of F_j . Left column of (a) indicates the degrees of freedom in each level of the symmetry group hierarchy. (b) Determining the degrees of freedom of frieze patterns by counting symmetry constraints on the intensity value of a pixel. The figure shows representative sets of pixels within a tile (here, (a) through (g) indicate frieze patterns P1 through P7) that have to have the same intensity value by the symmetry constraints. The dotted lines are horizontal and vertical midlines of the tile.

$F_n, n = 1, \dots, F_7$, construct a perfect frieze tile P_n that is closest to the original pattern P as measured by the SD_n symmetry distance.

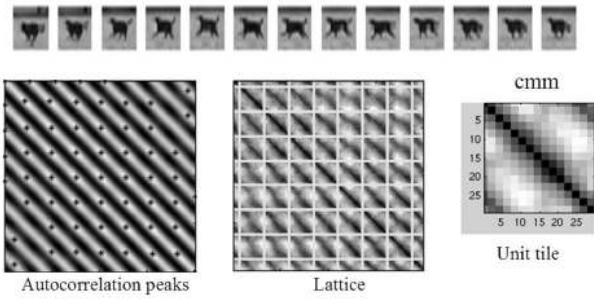
3. **Compute distances to perfect tiles:** Compute SD_n between the given pattern tile P and each perfect tile P_n in group F_n (1).
4. **Classify the symmetry group:** Find the perfect tile P_{min} that has minimum distance to P , and locate the corresponding symmetry group F_{min} in the subgroup hierarchy (Fig. 12a). If F_{min} has the least degrees of freedom in the group hierarchy (in frieze group case these groups are F5 or F7), then we classify P 's symmetry group as F_{min} ; otherwise, we compute distances ratios as defined in (3), (4), and (5) between two groups that have a subgroup relationship and choose the preferred symmetry group accordingly.

4.2 Gait Analysis Using Spatiotemporal Frieze Patterns

We illustrate this algorithm by analyzing human and animal gaits, which are known to be approximately periodic. Unlike previous examples in this paper which used patterns that were constructed to contain a particular type of symmetry, the gait patterns here are generated by human or animal motions that are not guaranteed to be strictly symmetric.


We look at two spatiotemporal image representations of gait. The first is a wallpaper pattern composed of the correlation scores between all pairs of video frames i, j from a gait sequence. Fig. 13 shows that the gait patterns of a running dog and a walking human have different symmetry groups, cm and $p4m$, respectively. This result reinforces that human gait is more bilaterally symmetric than a dog's running gait and, thus, symmetry of the gait patterns may be useful for classification when shape cues are hard to extract.

Symmetry of A Running Dog



rot180	rot120	rot90	rot60	flipT1	flipT2	flipD1	flipD2
0.513	1.301	1.284	1.241	1.311	1.311	0.613	0.558

Symmetry of A Walking Human



rot180	rot120	rot90	rot60	flipT1	flipT2	flipD1	flipD2
0.484	1.110	0.924	1.143	0.567	0.892	0.891	0.835

Fig. 13. Human walking gait is more symmetrical than a dog’s gait pattern (dog sequence courtesy of [3]).




Fig. 14. (a) Spatio-temporal gait representations are generated by projecting the body silhouette along its columns and rows, then stacking these 1D projections over time to form frieze-like patterns f_C, f_R repeating along the time dimension. (b) Observation #1: f_C of a humanoid avatar exhibits a wide variety of symmetric structures when viewed from different orientations. (c) Frieze pattern extracted from a 30-second long human walking sequence. Observation #2: The symmetry group variations of f_C of a human echo those of the avatar when viewed from different orientations.

The second type of spatiotemporal representation of gaits is a frieze pattern. Fig. 14a illustrates how a frieze pattern is generated from a time sequence of projections of 2D silhouettes along the X or Y axes. Note that each frieze pattern shown in Fig. 14a is a 1D (along the time axis) near-periodic pattern in a 2D space: time versus X -axis for column projection f_C and time versus Y -axis for row-projection f_R . Fig. 14b shows frieze pattern variation of f_C for a walking avatar viewed from different orientations. Fig. 14c shows the trajectory and f_C frieze pattern of a 30-second human walking sequence, in which one can observe the variations of frieze pattern f_C with the change of the viewing angles.

One hypothesis is that the f_C frieze patterns of different subjects viewed from the same direction share the same symmetry group. To test this hypothesis, we have computed frieze patterns for human gaits in the CMU MoBo database [8], containing motion sequences of 25 subjects walking on a treadmill. The result of symmetry group classification from side-view frieze patterns is dominated by symmetry group $F3$ (16 of 25), followed by symmetry group $F7$ (7 of 25) and symmetry group $F5$ (2 of 25). Some sample patterns are shown in Fig. 15. The walking avatar also has symmetry group $F3$ from this viewing angle. From the frieze group hierarchy (Fig. 12) we see that $F3$ is a subgroup of $F7$ and $F5$, thus they belong to the same “family” of groups. From Fig. 14a, we observe that a lower side view has symmetry group $F3$, while the higher side view has $F7$. Since the camera

is in a fixed location, the variation of symmetry groups between $F3$ and $F7$ of the 23 (out of 25) subjects can be caused by the height of each subject (viewing elevation angle changes). The only two female subjects (shorter subjects) have the $F7$ group. The two individuals with $F5$ symmetry




Fig. 15. Sample of frieze tiles and their symmetry groups computed from side views of the 25 walking subjects from CMU MoBo database [8]. The result of symmetry group classification from side-view frieze patterns is dominated by symmetry group $F3$ (16 of 25), followed by symmetry group $F7$ (7 of 25) and symmetry group $F5$ (2 of 25).




Fig. 16. (a) and (b) show an automatically extracted lattice and the tile that it implies. The tile is not a good representation of the pattern motif. (c) and (d) show the lattice position in terms of one of the three most-symmetric motifs found for the oriental rug image. The latter was generated automatically by an algorithm that analyzes pattern symmetry based on knowledge of the 17 wallpaper groups.

group present a walking style with an additional glide-reflection symmetry (Table 1) than those with $F3$. Here, the glide-reflection symmetry in a side view gait pattern can be caused by the equal amount of forward and backward body swings within one gait cycle.

This frieze pattern-based gait analysis idea was explored further in [27], where gait angle estimation is carried out using f_C frieze patterns while human identification from gaits observed from similar orientations is carried out using f_R . Classification rates range from 81 percent to 100 percent when training and testing on three different types of human gaits: slow walk, fast walk, and walk holding a ball.

5 REPRESENTATIVE MOTIFS OF PERIODIC PATTERNS

When translational symmetry is determined for a periodic pattern, it fixes the size, shape, and orientation of the unit lattice, but leaves open the question of where the unit lattice is located in the image. Figs. 16a and 16b shows an automatically extracted lattice, and the tile that it carves out, on an image of an oriental rug. The pattern fragment on the tile appears nonintuitive to a human observer, even though it is mathematically correct—if the goal is to tile the plane, any parallelogram of the same size and shape produces an equally good lattice unit. However, from a perception point of view, some parallelograms produce tiles that are better descriptors of the underlying symmetry of the overall pattern than others. For example, if the whole pattern has reflection symmetries, we would like the lattice unit in isolation to also have reflection symmetry. Our contribution in this section is to show how a small set of tiles can be chosen, in a principled way, such that the symmetry of the local tile pattern is maximized.

We define a *motif* as a representative tile of a periodic pattern. If we entertain the idea that the most representative motif is the one that is most symmetrical, one plausible strategy for generating motifs is to align the motif center with the center of the highest-order rotation in the pattern. This is the lattice point with the largest stabilizer group. Candidate motifs can then be determined systematically by

enumerating each distinct center point of the highest-order rotation. Two rotation centers are *distinct* if they lie in different *orbits* of the symmetry group, that is, if one cannot be mapped into the other by applying any translation, rotation, reflection, or glide-reflection symmetry within the symmetry group. This strategy for motif selection is explored in the examples below.

5.1 Oriental Rug Example

Refer back to the oriental rug image in Fig. 16. Based on an analysis of symmetry groups, it is found that this pattern belongs to the wallpaper group $cm\bar{m}$. Referring to Fig. 4, we see that the highest order of rotation in $cm\bar{m}$ is 180 degrees. Each placement of the lattice such that a tile is centered on one of these points of two-fold rotation will lead to a symmetrical candidate motif. Nine centers of rotation are shown in the diagram, but they are not all distinct. Consider one of the vertices of the parallelogram; by application of lattice translations, each vertex can be mapped to any other vertex, so they all lie in the same orbit. Consider now the four rotation centers on the diagonal edges of the parallelogram. Through a combination of glide reflections, and two-fold rotation about the center of the tile, these can all be mapped to each other and, thus, form a second orbit. Finally, the rotation center at the center of the tile lies in a third orbit. Therefore, there are three distinct candidate motifs that can be extracted from this pattern, each centered on a representative of one of the distinct orbits of the two-fold rotation centers (Fig. 16d). Although each has the same degree of rotation symmetry, the first two are more “preferred” than the third. This is because their centers also lie on the intersection of two axes of mirror reflection symmetry, while the center of the third pattern lies on the intersection of two glide reflection axes. Glide reflections (a combination of reflection and translation) are harder for humans to identify in a pattern than mirror reflections [40], [13].

5.2 Synthetic Pattern Examples

Generation of representative motifs proceeds similarly for other symmetry groups, by centering the motif on distinct centers of the highest-order rotation. There are some

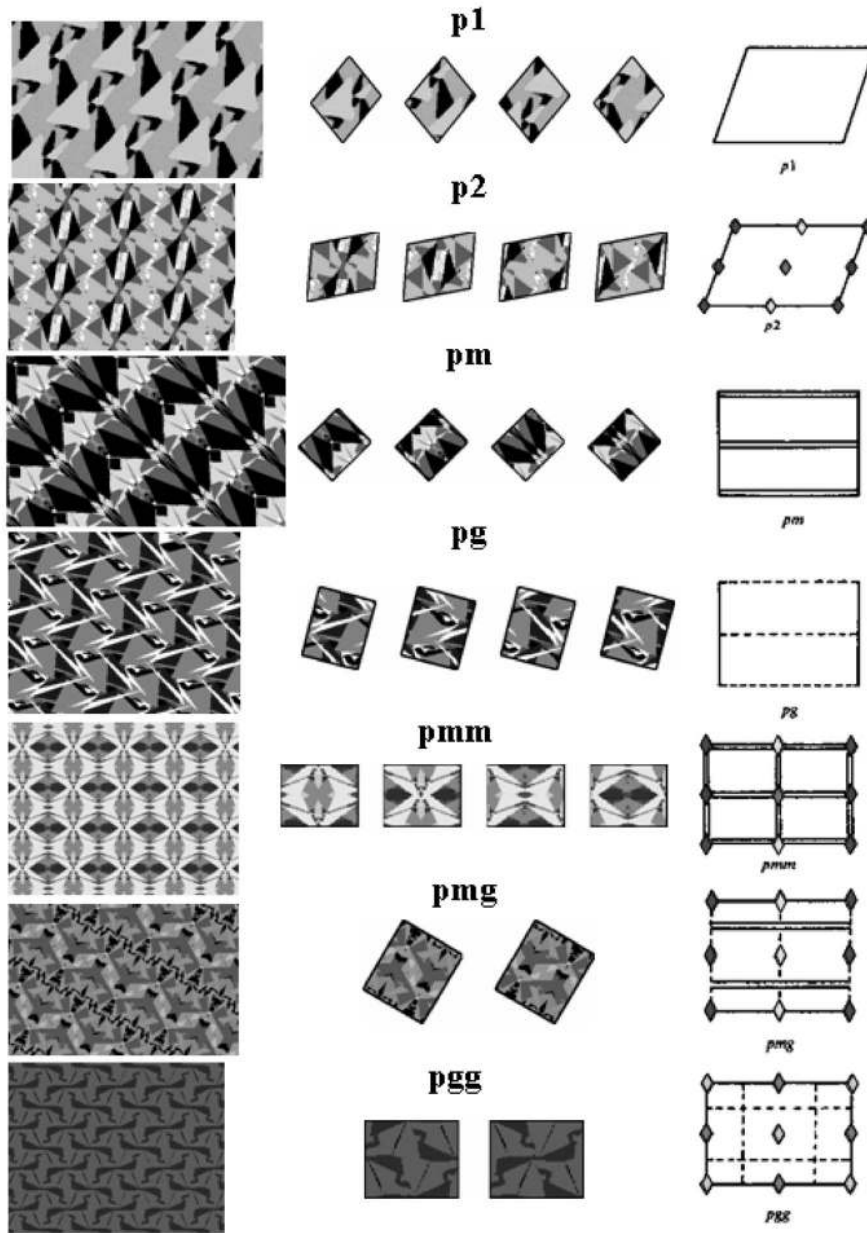


Fig. 17. Motifs centered on the highest order of rotation symmetry. For symmetry groups without rotation centers ($p1$, pm , pg), we use *approximate symmetries* to locate approximate rotation centers.

exceptions, however. Three groups (pm , pg , and cm) have no rotation symmetry, but have parallel reflection axes that fix the lattice offset in one direction (by centering the lattice on the reflection axis). Furthermore, the group $p1$ has no rotation or reflection symmetries and, therefore, there are no symmetry-based constraints on the lattice offset. In these cases, we are exploring how *approximate symmetries* in the pattern can be used to fix the unconstrained lattice offset. Specifically, the failed rotation symmetries are examined to find which rotations and rotation centers produce a transformed pattern that is most similar to the original pattern, as measured by correlation score. This almost-symmetry heuristic often leads to plausible motifs. Wallpaper groups with 120 degree and 60 degree rotation ($p3$, $p3m1$, $p31m$, $p6$, and $p6m$) have so-called “hexagonal lattices” formed by regrouping the vertices of adjacent parallelogram tiles. For these groups, candidate motifs are also extracted as hexagonal tiles.

A sample of the 17 wallpaper groups and their extracted canonical motifs is shown in Figs. 17, 18, and 19.

6 DISCUSSION

We propose a computational model for periodic pattern perception based on the mathematical theory of crystallographic groups, in particular, the theory of frieze and wallpaper groups. This mature mathematical theory provides guidance for developing computer algorithms to extract underlying lattices, analyze and classify periodic patterns, and capture perceptually meaningful motifs. Being able to understand the structural details of ALL possible 2D periodic patterns by a computer has both intellectual and practical significance [24], [25], [27], [28], [30] as well as computational challenges. We have shown the feasibility of using computer




Fig. 18. Further examples of motifs centered on the highest order of rotation symmetry. For symmetry groups without rotation centers (cm), we use *approximate symmetries* to locate approximate rotation centers.

algorithms for periodic pattern analysis of hand-drawn, synthetic, spatiotemporal, and real world images.

We now discuss the limitations of our current implementation and potential applications of a computational model for periodic pattern perception.

6.1 Limitations

One of the limitations of our work is the focus on intensity variations in patterns as the sole measure for determining the presence of symmetry. This approach presumes the patterns are only corrupted by Gaussian noise in the pixel intensity values. Geometric variations and distortions are not explicitly handled, although a trim rate was introduced to mitigate the effects of unmodeled geometric distortion and intensity outliers. Furthermore, correlation as a similarity measure requires a good distribution of color or grey values in the image. For example, images of brick walls are not handled well by correlation, since the distinguishing features of the pattern, the grout lines, make up only a small fraction of the pixels in the image and, indeed, their effects are overwhelmed by the large expanses of uniform-colored brick material. We

observe through introspection that the human visual system takes a much more feature-based approach to determining the symmetry of a pattern. The correspondence of geometric visual elements such as points of high boundary curvature, or small compact regions in the image, play a much larger role in evaluating pattern symmetry than a strict sum of squared differences comparison of image intensity values. A feature correspondence approach also would be potentially more tolerant of geometric pattern distortion.

A second limitation of the work described in this paper is the restriction to pattern classification under the scaled Euclidean transformation group. This implies that analysis of a pattern on a planar surface in the world must be performed on images taken by a camera oriented perpendicular to the surface. In recent work, we have been developing methods for classification of periodic patterns viewed from an off-axis viewing direction. For planar periodic patterns with small spatial extent with respect to surface distance from the camera, we model the pattern projection onto the image as an affine transformation, and the resulting projection of symmetry as a skewed symmetry

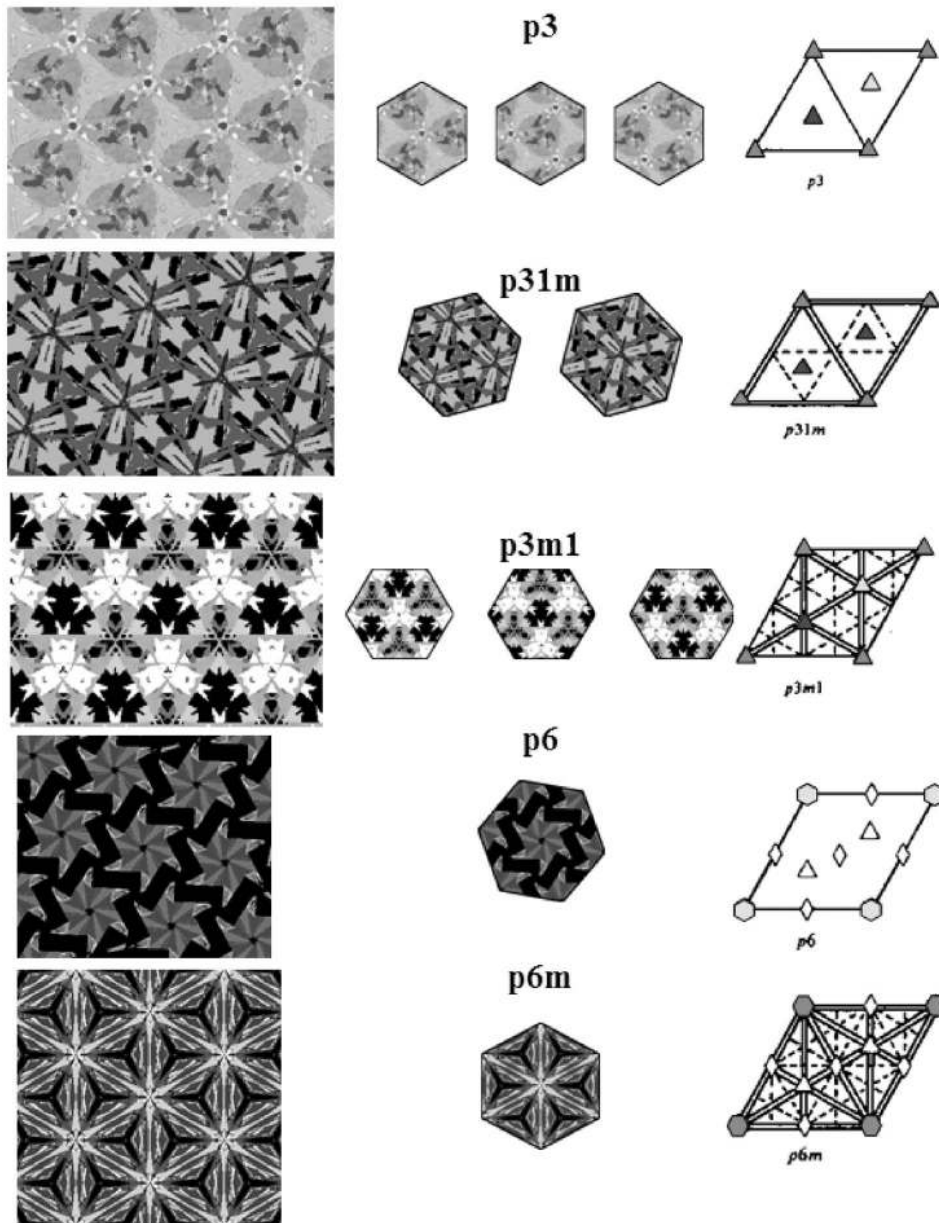


Fig. 19. Further examples of motifs centered on the highest order of rotation symmetry.

group [25]. We are also developing methods to deal with patterns on surfaces close to the camera, where perspective distortion effects become nonnegligible. It is expected that symmetry group analysis under affine and perspective imaging models will lead to a deeper understanding of the appearance of patterned surfaces in the world around us. Local deformations of approximate periodic patterns are also considered in our recent work on texture synthesis of near-regular textures [30], [28].

One scheme for symmetry-based motif selection was explored, namely centering tiles on distinct fixed points of the highest-order rotation of the pattern. This approach was found to produce tiles that are plausible candidates for the motif of the pattern. There are other schemes that could be employed, however. For example, tiles centered on points where a reflection axis intersects a glide reflection axis, or where two glide reflection axes intersect, would also be locally symmetric. Inclusion of reflection axis intersections

together with centers of rotations would lead to an exhaustive enumeration of symmetric tiles, and potentially a rank ordering based on the degree of rotation and reflection symmetry each tile exhibits. Symmetry is clearly just one piece of the perception puzzle. Although our approach can effectively enumerate good candidates for the pattern motif, and the human-perceived motif is typically among these candidates, choosing precisely *which* candidate is preferred by human perception is an open problem.

6.2 Practical Values of Understanding Periodic Symmetry

The practical significance of having a computational means to recognize and understand periodicity in a scene pattern is multifold.

Image indexing: An infinite number of periodic patterns can be indexed into a finite number of symmetry group classes regardless of pattern orientation, complexity, size, absolute

color, and global illumination conditions. Therefore, symmetry group indexing of periodic patterns provides a stable framework that has the potential to speed up the search in a large wallpaper pattern database up to 17 times.

Near-regular texture synthesis: An open problem in texture synthesis work is how to decide the shape and size of a sampling window on the original texture [4]. Because of this, structural fidelity for near-regular textures cannot be guaranteed using state-of-the-art texture synthesis algorithms. We show that a proper analysis of texture regularity in terms of symmetry groups and lattices, combined with texture synthesis, can faithfully reproduce the regularity as well as the randomness in a near-regular texture sample [30].

Image compression: Symmetry implies redundancy. A correct recognition of these redundancies can lead to effective compression algorithms for periodic or near-periodic patterns.

Pattern comparison: The computed pattern motifs, due to their uniqueness, provide a good basis for periodic pattern comparison. We have used this idea to compare spatio-temporal human gait patterns in [27].

Understanding human perception: Our work on motif selection suggests concrete hypotheses about human perceptual organization of periodic patterns, which can be systematically tested.

6.3 Real-World Symmetries

The varieties of real-world periodic symmetries are surprisingly large, even in our daily environment, as shown in Figs. 7, 8, and 9. Schattschneider [39] has observed that 16 of the 17 wallpaper groups can be exemplified by Chinese window structures. Washburn and Crowe [50] illustrated most of the 17 distinct patterns in African textile and pottery. It has also been shown that different subsets of the 17 wallpaper groups are preferred by different cultures [49].

Computationally, the difficulty in automatic analysis of real world symmetries is the fact that real-world patterns may be very noisy. The computed symmetries of a distorted periodic pattern may not correspond to any set of the required symmetries of a frieze or wallpaper group (Tables 1 and 2). In this paper, we deal with this situation by introducing a new concept: defining a minimum distance between a noisy pattern and its corresponding perfect pattern in each of the potential symmetry groups, and using geometric AIC to decide, fairly and unambiguously, on the best fit among the symmetry group hierarchy. An alternative approach to deal with this inconsistency would be using the vector value of symmetry matching scores of a pattern, such that there will be 2^4 (Table 1) or 2^8 (Table 2) different potential classes for near-periodic frieze or wallpaper patterns, respectively. Regardless of the method, there is a clear benefit to using symmetry groups as anchor points to study the high dimensional space of near-periodic, noisy, real-world patterns.

ACKNOWLEDGMENTS

This work is supported in part by the US National Science Foundation grants IIS-0099597, IIS-0208965, and ONR N00014-00-1-0915.

REFERENCES

- [1] L. Bieberbach, "Über die Bewegungsgruppen der n-Dimensionalen Euklidischen Räume Mit Einem Endlichen Fundamentalbereich," *Göttinger Nachrichten*, pp. 75-84, 1910.
- [2] R.W. Conners and C.A. Harlow, "Toward a Structural Textural Analyzer Based on Statistical Methods," *Proc. Int'l Conf. Color in Graphics and Image Processing*, vol. 12, no. 3, pp. 224-256, Mar. 1980.
- [3] R. Cutler and L. Davis, "Robust Real-Time Periodic Motion Detection, Analysis and Applications," *IEEE Trans. Pattern Analysis and Machine Intelligence*, special issue on video surveillance, vol. 22, no. 8, Aug. 2000.
- [4] A.A. Efros and W.T. Freeman, "Image Quilting for Texture Synthesis and Transfer," *Proc. SIGGRAPH*, pp. 35-42, 2001.
- [5] R. Enrich and J.P. Foith, "A View of Texture Topology and Texture Description," *Proc. Conf. Computer Graphics Image Processing*, vol. 8, pp. 174-202, 1978.
- [6] J.A. Gallian, *Contemporary Abstract Algebra*. Lexington, Mass.: D.C. Heath and Company, 1986.
- [7] A.D. Gross and T.E. Boulton, "Analyzing Skewed Symmetries," *Int'l J. Computer Vision*, vol. 13, no. 1, pp. 91-111, Sept. 1994.
- [8] R. Gross and J. Shi, "The CMU Motion of Body (MoBo) Database," Technical Report CMU-RI-TR-01-18, Robotics Inst., Carnegie Mellon Univ., 2001.
- [9] B. Grünbaum and G.C. Shephard, *Tilings and Patterns*. New York: W.H. Freeman and Company, 1987.
- [10] *Int'l Tables for X-Ray Crystallography*, vol. 1, Symmetry Groups, N.F.M. Henry and K. Lonsdale, eds., The Kynoch Press, England, The Int'l Union of Crystallography, 1969.
- [11] O. Jones, *The Grammar of Ornament*. New York: Van Nostrand Reinhold, 1972.
- [12] D. Joyce, Wallpaper Groups (Plane Symmetry Groups), aleph0.clarku.edu/~djoyce/wallpaper/wall1.html, 2003.
- [13] J. Kahn and D. Foster, "Horizontal-Vertical Structure in the Visual Comparison of Rigidly Transformed Patterns," *J. Experimental Psychology: Human Perception and Performance*, vol. 12, pp. 422-433, 1986.
- [14] T. Kanade, "Recovery of the 3-Dimensional Shape of an Object from a Single View," *Artificial Intelligence*, vol. 17, pp. 75-116, 1981.
- [15] T. Kanade and J.R. Kender, "Mapping Image Properties into Shape Constraints: Skewed Symmetry, Affine-Transformable Patterns, and the Shape-from-Texture Paradigm," *Readings in Computer Vision: Issues, Problems, Principles, and Paradigms*, M.A. Fischler and O. Firschein, eds. pp. 257-267, Kaufmann, 1987.
- [16] K. Kanatani, *Statistical Optimization for Geometric Computation: Theory and Practice (Machine Intelligence and Pattern Recognition)*, vol. 18, North-Holland, 1996.
- [17] K. Kanatani, "Symmetry as a Continuous Feature," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 19, no. 3, pp. 246-247, Mar. 1997.
- [18] M. Kazhdan, B. Chazelle, D. Dobkin, A. Finkelstein, and T. Funkhouser, "A Reflective Symmetry Descriptor," *Proc. Seventh European Conf. Computer Vision*, pp. 642-656, May 2002.
- [19] T. Leung and J. Malik, "Detecting, Localizing and Grouping Repeated Scene Elements," *Proc. European Conf. Computer Vision*, vol. 1, pp. 546-555, 1996.
- [20] M. Leyton, *Symmetry, Causality, Mind*. Cambridge, Mass.: MIT Press, 1992.
- [21] L.G.C. Hamey and T. Kanade, "Computer Analysis of Regular Repetitive Textures," *Proc. Image Understanding Workshop*, pp. 1076-1088, 1989.
- [22] H.-C. Lin, L.-L. Wang, and S.-N. Yang, "Extracting Periodicity of a Regular Texture Based on Autocorrelation Functions," *Pattern Recognition Letters*, vol. 18, pp. 433-443, 1997.
- [23] J. Liu, J. Mundy, D. Forsyth, A. Zisserman, and C. Rothwell, "Efficient Recognition of Rotationally Symmetric Surfaces and Straight Homogeneous Generalized Cylinders," *Proc. Conf. Computer Vision and Pattern Recognition*, pp. 123-128, June 1993.
- [24] Y. Liu and R.T. Collins, "A Computational Model for Repeated Pattern Perception Using Frieze and Wallpaper Groups," *Proc. Computer Vision and Pattern Recognition Conf.*, pp. 537-544, http://www.ri.cmu.edu/pubs/pub_3302.html, June 2000.
- [25] Y. Liu and R.T. Collins, "Skewed Symmetry Groups," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 872-879, http://www.ri.cmu.edu/pubs/pub_3815.html, Dec. 2001.


- [26] Y. Liu and R.T. Collins, "Frieze and Wallpaper Symmetry Groups Classification under Affine and Perspective Distortion," Technical Report CMU-RI-TR-98-37, The Robotics Inst., Carnegie Mellon Univ., Pittsburgh, Penn., 1998.
- [27] Y. Liu, R.T. Collins, and Y. Tsin, "Gait Sequence Analysis Using Frieze Patterns," *Proc. Seventh European Conf. Computer Vision*, 2002.
- [28] Y. Liu and W.C. Lin, "Deformable Texture: The Irregular-Regular-Irregular Cycle," *Proc. Third Int'l Workshop Texture Analysis and Synthesis (Texture 2003)*, Oct. 2003.
- [29] Y. Liu and R. Popplestone, "A Group Theoretical Formalization of Surface Contact," *Int'l J. Robotics Research*, vol. 13, no. 2, pp. 148-161, Apr. 1994.
- [30] Y. Liu, Y. Tsin, and W.C. Lin, "The Promise and Perils of Near-Regular Texture," *Int'l J. Computer Vision*, accepted for publication, 2003.
- [31] G. Loy and A. Zelinsky, "A Fast Radial Symmetry Transform for Detecting Points of Interest," *Proc. European Conf. Computer Vision*, p. 358, 2002.
- [32] S.Y. Lu and K.S. Fu, "A Syntactic Approach to Texture Analysis," *Computer Graphics Image Processing*, vol. 7, pp. 303-330, Nov. 1978.
- [33] G. Marola, "On the Detection of the Axes of Symmetry of Symmetric and Almost Symmetric Planar Images," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, no. 1, pp. 104-108, Jan. 1989.
- [34] W. Miller Jr., *Symmetry Groups and Their Applications*. New York: Academic Press, 1972.
- [35] S.C. Pei and L.G. Liou, "Automatic Symmetry Determination and Normalization for Rotationally Symmetric 2D Shapes and 3D Solid Objects," *Pattern Recognition*, vol. 27, pp. 1193-1208, 1994.
- [36] A.R. Rao and G.L. Lohse, "Identifying High Level Features of Texture Perception," *CVGIP: Image Processing*, vol. 55, pp. 218-233, 1993.
- [37] J. Russ, *The Image Processing Handbook*, second ed., CRC Press, 1995.
- [38] F. Schaffalitzky and A. Zisserman, "Geometric Grouping of Repeated Elements within Images," *Shape, Contour, and Grouping in Computer Vision*, D.A. Forsyth, V. Di Gesu, J.L. Mundy, and R. Cipolla, eds., Springer-Verlag, 1999.
- [39] D. Schattschneider, "The Plane Symmetry Groups: Their Recognition and Notation," *Am. Math. Monthly*, vol. 85, pp. 439-450, 1978.
- [40] D. Schattschneider, personal communications, 1999.
- [41] R.L.E. Schwarzenberger, "The 17 Plane Symmetry Groups," *Math. Gazette*, vol. 58, pp. 123-131, 1974.
- [42] K. Selkainaho, J. Parkkinen, and E. Oja, "Comparison of χ^2 and k Statistics in Finding Signal and Picture Periodicity," *Proc. Int'l Conf. Pattern Recognition*, pp. 1221-1224, 1988.
- [43] A.V. Shubnikov and V.A. Koptsik, *Symmetry in Science and Art*. D. Harker, ed., New York: Plenum Press, 1974.
- [44] V.V. Starovoitov, S.Y. Jeong, and R.H. Park, "Texture Periodicity Detection: Features, Properties, and Comparisons," *IEEE Trans. Systems, Man, and Cybernetics, Part A*, vol. 28, no. 6, pp. 839-848, 1998.
- [45] P. Stevens, *Handbook of Regular Patterns*. Cambridge, Mass.: MIT Press, 1987.
- [46] C.M. Sun, "Fast Recovery of Rotational Symmetry Parameters Using Gradient Orientation," *Optical Eng.*, vol. 36, no. 4, pp. 1073-1077, Apr. 1997.
- [47] T. Tuytelaars, A. Turina, and L. Van Gool, "Noncombinatorial Detection of Regular Repetitions under Perspective Skew," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, no. 4, pp. 418-432, Apr. 2003.
- [48] L. Van Gool, T. Moons, and M. Proesmans, "Mirror and Point Symmetry under Perspective Skewing," *Proc. Conf. Computer Vision and Pattern Recognition*, pp. 285-292, 1996.
- [49] D.K. Washburn, "Perceptual Anthropology: The Cultural Salience of Symmetry," *Am. Anthropologist*, vol. 101, no. 3, pp. 547-562, 1999.
- [50] D.K. Washburn and D.W. Crowe, *Symmetries of Culture: Theory and Practice of Plane Pattern Analysis*. Univ. of Washington Press, 1991.
- [51] J. Weeks, Programs that Can Automatically Generate 2D Planar Crystallographic Patterns, Kali, <http://www.geom.umn.edu/apps/Kali/>, 1995.
- [52] R.K.K. Yip, W.C.Y. Lam, P.K.S. Tam, and D.N.K. Leung, "A Hough Transform Technique for the Detection of Rotational Symmetry," *Pattern Recognition Letters*, vol. 15, no. 9, pp. 919-928, Sept. 1994.
- [53] H. Zabrodsky, S. Peleg, and D. Avnir, "Symmetry as a Continuous Feature," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 17, no. 12, pp. 1154-1165, Dec. 1995.
- [54] S.W. Zucker and D. Terzopoulos, "Finding Structure in Cooccurrence Matrices for Texture Analysis," *Proc. Int'l Conf. Color in Graphics and Image Processing*, vol. 12, no. 3, pp. 286-308, Mar. 1980.



Yanxi Liu received the PhD degree in computer science from the University of Massachusetts (Amherst), where she studied group theory applications in robotics. Her postdoctoral training was performed in INRIA Grenoble, France, and DIMACS (a US National Science Foundation Research and Education Center for Discrete Mathematics and Theoretical Computer Science). She is a faculty member affiliated with both the Robotics Institute (RI) and the Center for Automated Learning and Discovery (CALD) of the School of Computer Science, Carnegie Mellon University. Her current research interests are centered on developing a computational framework for near-regular pattern analysis and synthesis, and seeking discriminative subspaces in large image sets or videos, with an emphasis on medical applications. She has been serving as a reviewer for all major journals in computer vision and pattern recognition, and on NIH study groups and panels. She is a member of IEEE and the IEEE Computer Society.



Robert T. Collins received the PhD degree in computer science in 1993 from the University of Massachusetts at Amherst, for work on recovering scene models using stochastic projective geometry. He is currently a member of the Research Faculty at the Robotics Institute of Carnegie Mellon University, where his research interests include video scene understanding, automated surveillance, human activity modeling, and real-time tracking. He was coeditor the August 2000 special issue of *IEEE Transactions on Pattern Analysis and Machine Intelligence* on the topic of video surveillance. In 1999, he was area chair for the computational geometry section of the IEEE Computer Vision and Pattern Recognition Conference. He routinely serves on the program committee for IEEE workshops on visual surveillance, performance evaluation of tracking and surveillance, and detection and recognition of events in video. He has been a reviewer for all of the major computer vision journals and conferences, and serves on US National Science Foundation review panels in the area of computer vision. He is a member of the IEEE and the IEEE Computer Society.



Yanghai Tsin received the BEng and MS degrees from the Tsinghua University and the Chinese Academy of Sciences, respectively. He is a PhD student at Carnegie Mellon University. His research interest is in computer vision and pattern recognition.

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