

A Computational Model for Repeated Pattern Perception using Frieze and Wallpaper Groups

Yanxi Liu and Robert T. Collins

The Robotics Institute, Carnegie Mellon University
{yanxi, rcollins}@cs.cmu.edu

Abstract

Humans have an innate ability to perceive symmetry, but it is not obvious how to automate this powerful insight. In this paper the mathematical theory of Frieze and wallpaper groups is used to extract visually meaningful building blocks (motifs) from a repeated pattern. A novel peak detection algorithm based on “regions of dominance” is used to automatically detect the underlying translational lattice of a repeated pattern. Following automatic classification of the pattern’s symmetry group, knowledge of the interplay between rotation, reflection, glide-reflection and translation in that group leads to a small set of candidate motifs that exhibit local symmetry consistent with the global symmetry of the entire pattern. Although other work has addressed detection of the translational lattice of a repeated pattern, ours is the first to seek a principled method for determining a representative motif. Experiments show that the resulting pattern motifs conform well with human perception.

1 Introduction

Symmetry is pervasive in both natural and man-made environments. Symmetries of repeated patterns in a plane are of particular interest in computer vision. Repeated patterns can be found in regular textures, indoor and outdoor scenes (brick walls, tiled floors, wallpapers, ceilings, clothes, windows on buildings, cars in a parking lot), or intermediate data representations (periodicity analysis of human and animal gaits in the spatio-temporal domain).

A mature mathematical theory for repeated patterns has been known for over 100 years [2, 3], namely the theory of Frieze and wallpaper groups¹. For monochrome planar repeated patterns, there are seven *Frieze groups* for 2D patterns repeated along one dimension, and seventeen *wallpaper groups* describing patterns extended by two linearly independent translational generators. Despite the infinite variety of repeated patterns, this finite set of symmetry groups

completely characterizes the possible structural symmetry of any repeated pattern. The symmetry group of a repeated pattern is an attractive descriptor in computer vision research due to its independence of scale, absolute color, lighting, density and orientation/position of the pattern.

A computational model of repeated pattern perception can be divided into three parts: generating the underlying translational lattice from the image of a repeated pattern (Section 2); classifying the symmetry group of the repeated pattern (Section 3); and identifying the preferred “motif” of the pattern (Section 4). Our work, initially inspired by [13], appears to be the first to use the theory of Frieze and wallpaper groups for automated analysis of repeated patterns, although there exist flowcharts and computer programs that allow humans to interactively generate and identify repeated patterns for educational purposes [16, 17, 7].

2 Reliable Lattice Construction

The key issue in repeated pattern analysis is whether the 2D lattice of the pattern can be correctly extracted. Previous work on lattice detection can be clustered into two 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 particular parametric model [4, 8, 12]. The benefits of this approach are the ability to detect small regions of repeating pattern within a larger image, and the ability 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 pattern repetition uses autocorrelation [9], the Fourier transform [11], or periodicity measures defined over cooccurrence matrices [1, 18, 14, 15] computed globally over the whole pattern. These approaches work for any intensity image, not just ones with strong features. The main drawback is the assumption that a single repeated pattern occupies a large portion of the image, limiting the approach to analysis of patterns

¹These groups are also called the one and two dimensional Crystallographic groups [5]

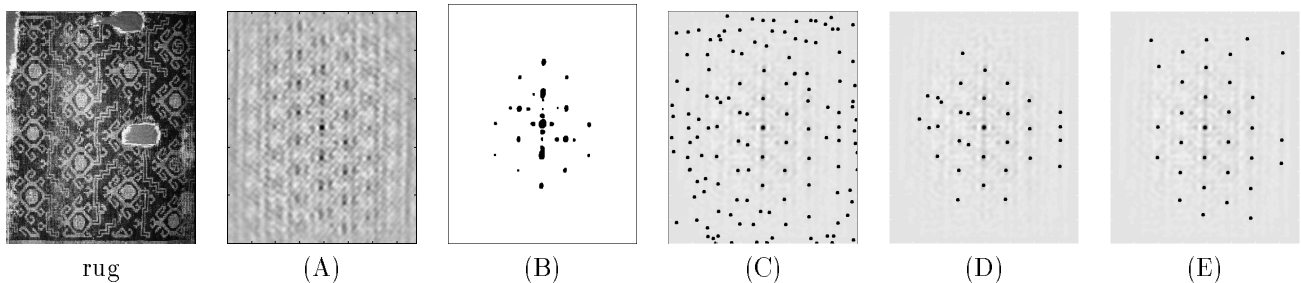


Figure 1: An oriental rug image and A) its autocorrelation surface, B) peaks found using a global threshold, C) peaks extracted using the threshold-free method of Lin et.al, D) the highest 32 peaks from those return by Lin et.al, E) the 32 most-dominant peaks found using our approach described in the text.

that have already been segmented in some other way. For the isolated patterns considered in this paper, autocorrelation proves to be an appropriate method for quantifying translational symmetry.

2.1 The Problem of Peak Detection

Figure 1(A) shows an autocorrelation surface for the rug image on the left. Although the grid of peaks is apparent to the human eye, finding it automatically is very difficult. Simple approaches such as setting a global threshold yield spurious results (Figure 1B). The trouble is that many legitimate grid peaks have a lower value than some of the spurious peaks. Lin et.al. [9] present a threshold-free approach based on finding an optimal Gaussian smoothing of the image, followed by selection of local maxima. Their method yields 135 local maxima peaks on this image, shown in Figure 1C. It is hard to see the grid structure interspersed with the spurious points. If we sort peaks by height and take the first 32, we have the answer in Figure 1D. In comparison, Figure 1E presents the first 32 peaks found by our peak detection algorithm (Section 2.2).

Even noise-free computer-generated patterns can cause problems for lattice detection algorithms. Figure 2(A) shows an incorrect lattice found by the algorithm of Lin et.al., together with the correct lattice found by our algorithm (Figure 2B). Obviously, their algorithm has picked up smaller peaks between the major ones, and indeed, the pattern’s autocorrelation surface has high furrows between the major peaks. Our experience with a variety of repeated patterns indicates that it is common to find spurious peaks superimposed over the autocorrelation image at twice (or other multiples of) the frequency of the lattice. An illustration of why this happens is shown by the Frieze pattern in Figure 3, displayed next to the 1D autocorrelation response computed by sliding the pattern

strip along its axis of translation. The correct peaks are A, C and E; the main spurious peaks are B and D, although there are smaller ones as well. Halfway between actual lattice translations, the large features in the pattern partially match smaller features interspersed between them, causing spurious peaks to form. Furthermore, these spurious peaks can have higher value than actual peaks located at the periphery of the autocorrelation image (e.g. $\text{height}(\text{B}) > \text{height}(\text{E})$). These difficulties are exacerbated in two dimensions.

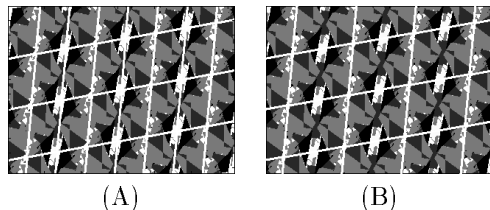


Figure 2: Even noise-free images can be hard to process. (A) An incorrect lattice found by the algorithm of Lin et.al. (B) Correct lattice found by our algorithm.

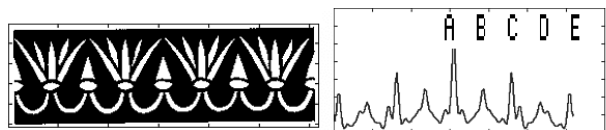


Figure 3: Frieze group and its associated 1D autocorrelation response.

2.2 Regions of Dominance

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 overshadowed by an even higher one. Referring back to Figure 3, 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 Figure 1E, the first 32 most-dominant peaks found using our method are well distributed over the whole image, with very few spurious peaks included (there are some at the periphery).

A list of peaks in decreasing order of dominance can be computed using a simple 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 non-maximum suppression over a sliding $M \times M$ window, where $M = 5$ 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 ($N \cdot \text{LOGN}$) 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 precedes it in the list, and denote the minimum distance D_i . This is an order N^2 process. 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 in noisy data, for example finding peaks in intensity or color histograms to perform region segmentation.

2.3 Determine the Shortest, Independent Translation Vectors

Having a set of candidate lattice points, our goal is to find two linearly independent translation vectors that generate the lattice. This task is complicated by missing points as well as additional spurious points. A Hough transform approach adapted from [9] is used to find the two shortest translation vectors that best explain the majority of the point data. From the study of wallpaper groups, we know that the angle between the two generating vectors must be between 60 and 90 degrees. Figure 4 shows detected lattices for two real-world patterns.

3 Symmetry Group Classification

A 2D *repeated* or *periodic* pattern has the following property: there exists a finite region bounded by two linearly independent translations, which when acted

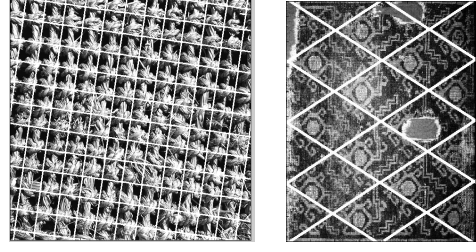


Figure 4: Detected lattices for two real-world patterns.

on by the group generated by the translations produces simultaneously a covering (no gaps) and a packing (no overlaps) of the original image [13, 3]. The smallest such bounded region is called a *unit of the pattern* or *lattice unit*, since the translational orbit of any single point on the plane is a lattice. A *symmetry* of a subset S of Euclidean space is an isometry that keeps S setwise invariant. All symmetries of S form the *symmetry group* of S under composition. It has been proven that there are only seven *Frieze groups* (Figure 5) for 2D patterns repeated along one

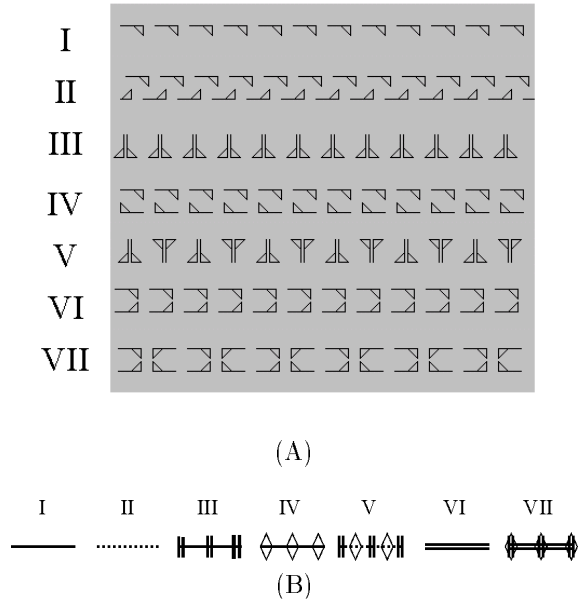


Figure 5: (A) The seven Frieze patterns. (B) Their unit lattices, with specific symmetries indicated: \diamond meaning 2-fold rotation, \parallel meaning reflection, and dotted lines meaning glide-reflections.

dimension, and seventeen *wallpaper groups* (Figure 6) describing patterns extended by two linearly indepen-

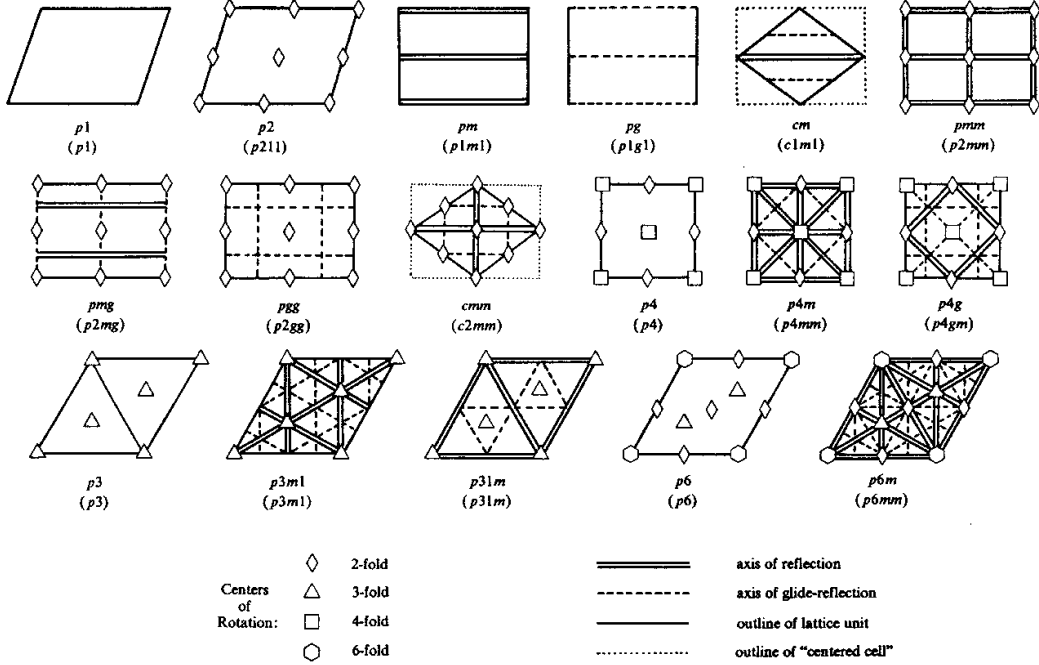


Figure 6: The generating regions for the 17 Wallpaper groups

dent translational generators [13, 3]. Mathematically, wallpaper groups are defined only for infinite patterns that cover the whole plane. In practice, we analyze a repeated pattern P of a finite area, and use the phrase “symmetry group G of P ” to mean that G is the symmetry group of the infinite repeated pattern that has P as a finite patch.

Figure 5A shows the seven possible symmetry group configurations for a horizontal Frieze pattern. Figure 5B shows the corresponding lattice units and their respective rotation, reflection and glide-reflection symmetries². Due to limited space, the rest of this paper focuses on wallpaper patterns.

Figure 6 depicts unit lattices for the 17 distinct wallpaper groups (from [13]). Each unit is characterized in terms of its translation generators, rotation, reflection and glide-reflection symmetries. The two linearly independent translations of minimum length are the two basic generators of each group, and they construct a lattice for the group. Even though the variety of pattern instantiations is endless, the underlying relationship between translation, rotation, reflection and glide-reflection in any 2D repeated pattern **must** conform to one of these seventeen cases.

Since a symmetry of a 2D repeated pattern has to

²Here *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.

map the lattice associated with the pattern onto itself, i.e. mapping centers of rotation to new centers of rotation having the same order, the only possible rotation symmetries are 2, 3, 4, 6-fold rotations. This restriction is often referred to as the *crystallographic restriction*. Furthermore, reflection axes can only be oriented parallel, diagonal, or perpendicular to the lattice translation vectors. Under these constraints, there are only five possible lattice shapes: (1) parallelogram, (2) rectangular, (3) rhombic, (4) square and (5) hexagonal. Each lattice unit is a parallelogram. Rectangular units have angles of 90°. Rhombic units have equal-length edges. Square and hexagonal lattices are special cases of rectangle and rhombic, respectively.

We have constructed an algorithm that can automatically classify which symmetry group a 2D repeated pattern under Euclidean transformations belongs to. 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 rotation and/or reflection symmetries, and without knowing the offset of the underlying lattice. Table 1 lists the eight symmetries checked in the classification algorithm. It is clear that each group corresponds to a unique sequence of values listed in Table 1, and is mutually exclusive from each other. The determination of a specific rotation or reflection or glide-reflection symmetry is performed by applying the symmetry to be tested to the entire pattern, then

Table 1: Wallpaper group classification: numbers 2,3,4 or 6 denote n-fold rotational symmetry, Tx (or Dx) denotes reflectional symmetry about one of the translation (or diagonal) vectors of the unit lattice. “Y” means the symmetry exists for that symmetry group; empty space means no. Y(g) denotes a glide reflection.

	p1	p2	pm	pg	cm	pmm	pmg	pgg	cmm	p4	p4m	p4g	p3	p3m1	p31m	p6	p6m
2		Y				Y	Y	Y	Y	Y	Y	Y				Y	Y
3													Y	Y	Y	Y	Y
4										Y	Y	Y					
6																Y	Y
T1			Y	Y(g)		Y	Y(g)	Y(g)			Y	Y(g)			Y		Y
T2						Y	Y	Y(g)			Y	Y(g)			Y		Y
D1					Y				Y		Y	Y		Y	Y		Y
D2									Y		Y	Y					Y

checking the similarity between the original and transformed images. The derivation, experiments, and an important extension of this classification algorithm to handle repeated patterns under affine and perspective distortions, are described in [10].

4 Extracting Representative Motifs

Although other work has addressed detection of the translational lattice of a repeated pattern, ours is the first to seek a principled method for determining a representative motif. The issue here is that consideration of translational symmetry alone fixes the size, shape and orientation of the lattice, but leaves open the question of where the lattice is located in the image. Any offset of the lattice carves the pattern into a set of identical tiles, but these tiles typically appear nonintuitive to a human observer and cannot be considered a good specification of the pattern’s motif (Figure 7).

Choosing a good motif should help one see, from a single tile, what the pattern “looks like”. This is related to the problem of figure-ground separation. We would like to find a motif that is centered on the foreground “figure”, rather than having the figure be split into two or more pieces on opposite edges of the tile. From work in perceptual grouping, it is known that the human perceptual system often has a preference for symmetric figures. 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 pattern fragments on them is maximized.

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 point fixed by the largest stabilizer subgroup. Candidate motifs can then be de-

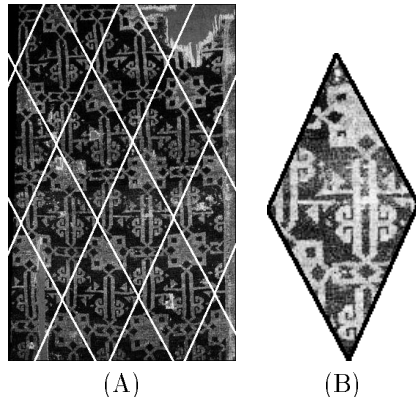


Figure 7: A and B shows an automatically extracted lattice and the tile that it implies. The tile is not a good representation of the pattern motif.

termined 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 in the symmetry group.

4.1 Oriental Rug Image

To make this idea more concrete, we walk through a detailed example. Figure 7 shows an automatically extracted lattice, and the tile that it carves out. The pattern fragment on the tile appears nonintuitive to a human observer, even though it is geometrically correct. If the goal is to tile the plane, any parallelogram of the same size and shape is an equally good lattice unit. However, from a perception point of view, some parallelograms are much better descriptors of the underlying symmetry of the overall pattern than others.

Based on an analysis of symmetry, it is found that this pattern belongs to the wallpaper group cm . Referring to Figure 6, we see that the highest order of rotation in cm is two-fold (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 (and no others) that can be extracted from this pattern, each centered on a representative of one of the distinct orbits of the two-fold rotation centers (Figure 8). Although each has the same degree of rotational 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.

Aside from motif selection, knowledge of the lattice structure of a repeated pattern allows us to determine which pixels in an image should look the same. Taking the median of corresponding pixels across the multiple tiles of the rug image, for example, creates a “median tile” with noise and irregularities filtered out. Figure 9 compares the original worn rug with a virtual rug generated from the median tile.

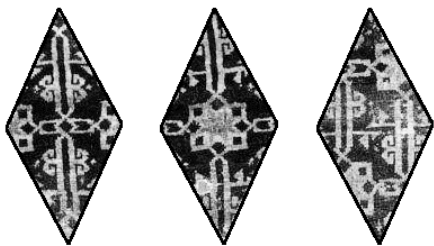


Figure 8: Most-symmetric motifs found for the oriental rug image.

4.2 Computer-Generated Images

A sample of some of the seventeen wallpaper groups and their extracted canonical motifs is shown in Fig-

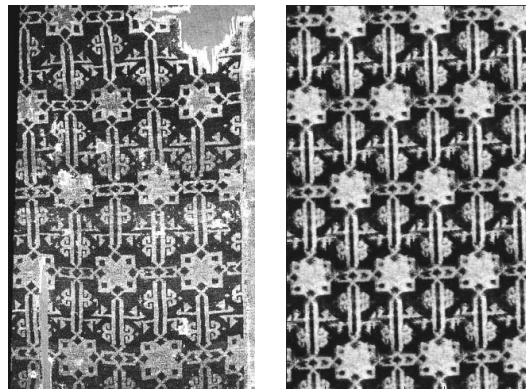


Figure 9: Real oriental rug and a perfectly symmetric virtual rug formed by translating the median tile.

ure 10. Each of the wallpaper samples was downloaded from a web site [6]. Generation of representative motifs proceeds as described in the previous example, by centering the motif on distinct centers of the highest-order rotation. There are some 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. However, even in these cases we are exploring how *approximate symmetries* in the pattern can be used to fix the unconstrained 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. Wallpaper groups with order 3 and order 6 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 extracted as hexagonal tiles.

5 Conclusion

We propose a computational model for repeated pattern perception based on the mathematical theory of crystallographic groups, in particular, the wallpaper groups. This mature mathematical theory provides guidance for analyzing and classifying repeated patterns, and for extracting the pattern’s visually meaningful building blocks, namely motifs.

This computational model has been implemented and tested on several synthetic and real-world repeated patterns. Detection of the underlying trans-

lational lattice is based on a novel scheme for detecting peaks in an autocorrelation image. The relevance of a peak is determined by the size of its “region of dominance”, the largest circle centered on the peak, and that contains no higher peaks. This approach is the best we have seen for ignoring spurious peaks in the autocorrelation surface caused by partial matches among pattern elements.

Once the pattern’s symmetry group is known, the relationship between translation, rotation, reflection and glide-reflection in the pattern is completely determined. Furthermore, it is then possible to analyze how the lattice should be placed in order to carve out tiles that exhibit maximal local symmetry. We hypothesize that symmetric tiles form good candidates for the human-perceived motif of the pattern. More importantly, understanding of the symmetry group of the pattern enables a principled enumeration of all symmetric tiles.

One scheme for symmetry-based motif selection was explored – centering tiles on distinct fixed points of the pattern’s highest-order rotation – and was found to produce tiles that are indeed plausible candidates for the motif of the pattern. There are other symmetry-based 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 reflections axes intersect, would also be locally symmetric. Inclusion of these reflection intersection points together with centers of rotations would lead to an exhaustive enumeration of symmetric tiles, and potentially a rank ordering based on the degree of rotational and reflectional 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. All of the factors of perceptual organization, including proximity, alignment, convexity, form and orientation, will have a role to play in a comprehensive theory of repeated pattern perception.

References

- [1] R.W. Connors and C.A. Harlow. Toward a structural textural analyzer based on statistical methods. *CGIP*, 12(3):224–256, March 1980.
- [2] E.S. Fedorov. The elements of the study of figures. [Russian]. In *Zapiski Imperatorskogo S. Peterburgskogo Mineralogicheskogo Obshchestva*, volume 2(21), pages 1–289, 1885.
- [3] B. Grünbaum and G.C. Shephard. *Tilings and Patterns*. W.H. Freeman and Company, New York, 1987.
- [4] L. Hamey and T. Kanade. Computer analysis of regular repetitive textures. In *Image Understanding Workshop 1989*, pages 1076–1088, 1989.
- [5] N.F.M. Henry and K. Lonsdale, editors. *International Tables for X-ray Crystallography, Volume 1, Symmetry Groups*. The Kynoch Press, England, 1969. The International Union of Crystallography.
- [6] D. Joyce. Wallpaper groups (plane symmetry groups). aleph0.clarku.edu/~djoyce/wallpaper/wall1.html.
- [7] Kali. Programs that generate planar crystallographic patterns. <http://www.geom.umn.edu/apps/kali/>.
- [8] T. Leung and J. Malik. Detecting, localizing and grouping repeated scene elements. In *ECCV LNCS 1064, vol 1*, pages 546–555, 1996.
- [9] H. Lin, L. Wang, and S. Yang. Extracting periodicity of a regular texture based on autocorrelation functions. *Pattern Recognition Letters*, 18:433–443, 1997.
- [10] 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 Institute, Carnegie Mellon University, Pittsburgh, PA, 1998.
- [11] T. Matsuyama, S. Miura, and M. Nagao. A structural analysis of natural textures by fourier transformation. *CVGIP*, 24(3):347–362, December 1983.
- [12] F. Schaffalitzky and A. Zisserman. Geometric grouping of repeated elements within images. In Forsyth, Di Gesu, Mundy, and Cipolla, editors, *Shape, Contour and Grouping in Computer Vision*, LNCS. Springer-Verlag, 1999.
- [13] D. Schattschneider. The plane symmetry groups: Their recognition and notation. *American Mathematical Monthly*, 85:439–450, 1978.
- [14] K. Selkainaho, J. Parkkinen, and E. Oja. Comparison of x2 and k statistics in finding signal and picture periodicity. In *ICPR*, pages 1221–1224, 1988.
- [15] V.V. Starovoitov, S.Y. Jeong, and R.H. Park. Texture periodicity detection: Features, properties, and comparisons. *IEEE SMC-A*, 28(6):839–848, 1998.
- [16] D.K. Washburn and D.W. Crowe. *Symmetries of Culture: Theory and Practice of Plane Pattern Analysis*. University of Washington Press, 1991.
- [17] S. Whitesides, K. Abadjian, P. Hum, and N. Rodjanapiches. Computer-aided instruction for a theorem in geometry. In *Proc MINI and Microcomputers and their application*, pages 115–118. ACA Press, 1985.
- [18] S.W. Zucker and D. Terzopoulos. Finding structure in co-occurrence matrices for texture analysis. *CGIP*, 12(3):286–308, March 1980.

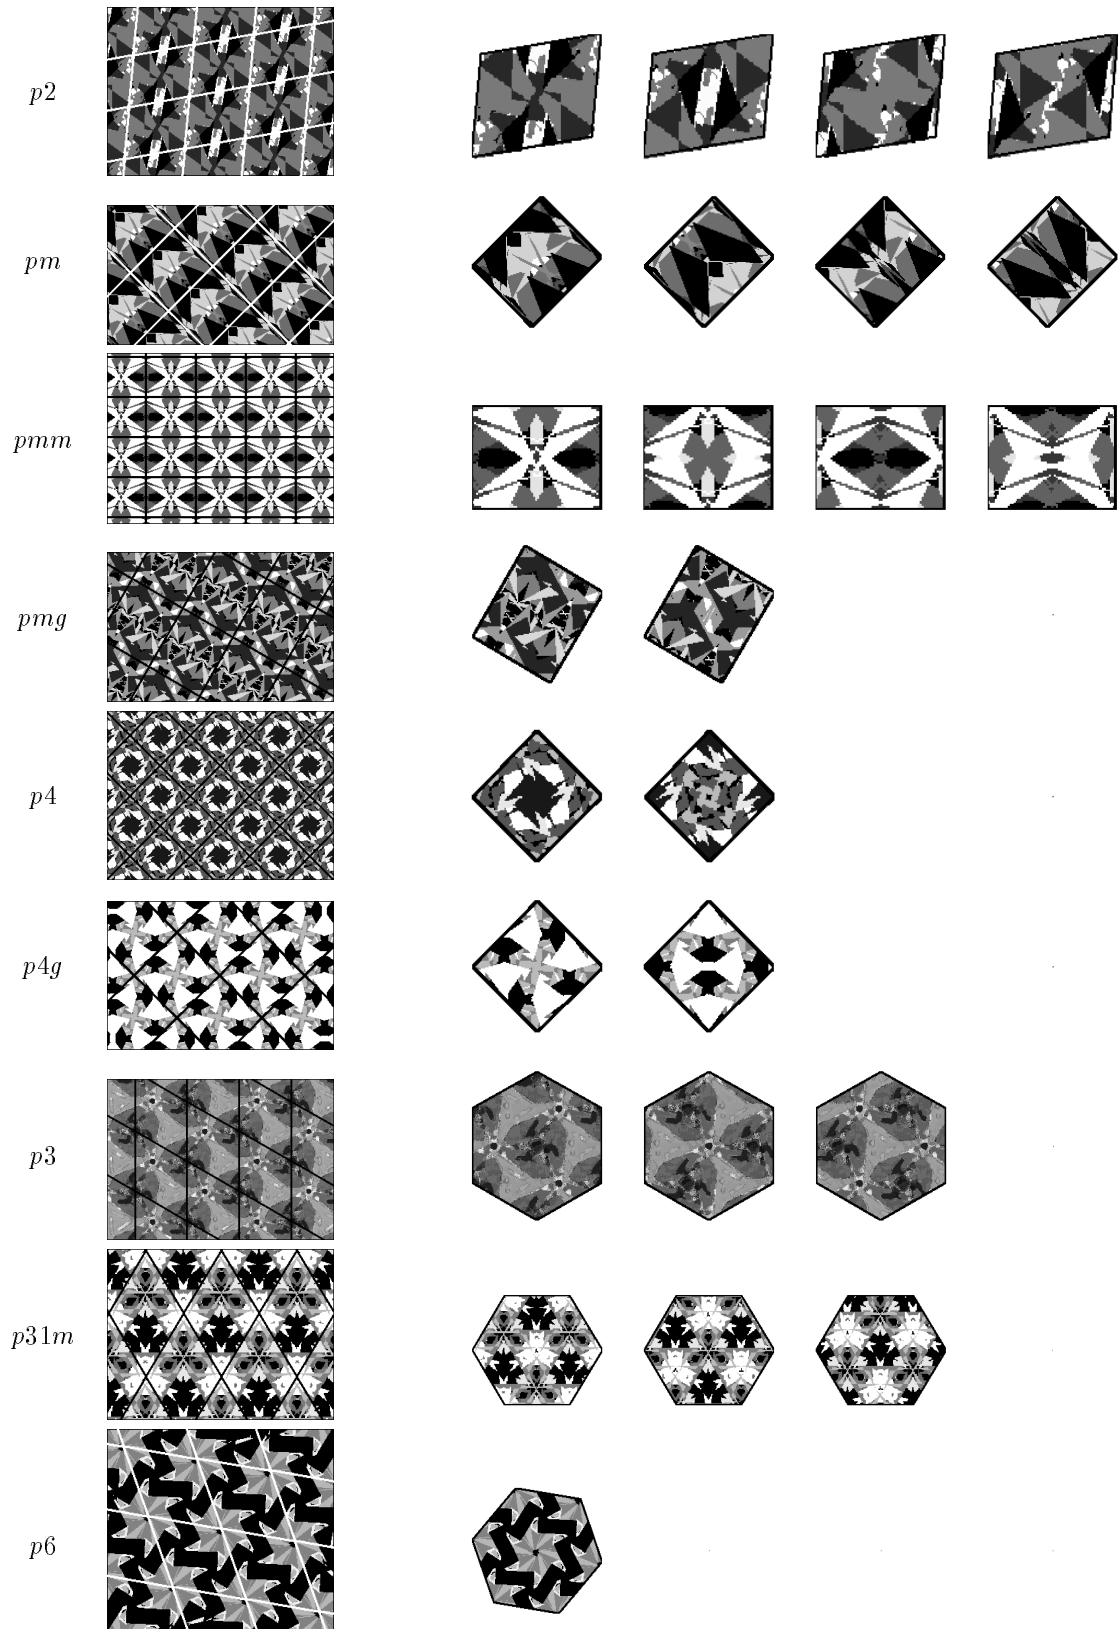


Figure 10: Automatically detected lattices and motifs for examples of some of the seventeen wallpaper groups.