

LATTICE EXTRACTION BASED ON SYMMETRY ANALYSIS

Manuel Agustí-Melchor, Jose-Miguel Valiente-González and Ángel Rodas-Jordá

Dept. de Informàtica de Sistemes y Computadores, Universidad Politècnica de Valencia

Camino de Vera s/n, Valencia, Spain

magusti@disca.upv.es, jvalient@disca.upv.es, arodas@disca.upv.es

Keywords: Lattice, grid, periodicity, wallpaper, regular pattern, symmetry, phase analysis.

Abstract: In many computer tasks it is necessary to structurally describe the contents of images for further processing, for example, in regular images produced in industrial processes such as textiles or ceramics. After reviewing the different approaches found in the literature, this work redefines the problem of periodicity in terms of the existence of local symmetries.

Phase symmetry analysis is chosen to obtain these symmetries because of its robustness when dealing with image contrast and noise. Also, the multiresolution nature of the technique offers independence from using fixed thresholds to segment the image. Our adaptation of the original technique, based on lattice constraints, has result in a parameter free algorithm for determining the lattice. It offers a significant increase in computational speed with respect to the original proposal. Given that there is no set of images for assessing this type of techniques, various sets of images have been used, and the results are apresented. A measure to enable the evaluation of results is also introduced, so that each calculated lattice can be tagged with an index regarding its correctness. The experiments show that using this statistic, good results are reported from image collections. Possible applications of the lattice extraction are suggested.

1 INTRODUCTION

In many computer vision tasks a structural description of the contents of images is needed for further processing. This is the case of pattern and tiling images obtained in industrial processes such as textile or ceramics. These images are formed by the combination of motifs that are regularly repeated using geometric transformations to fill the 2D plane without overlaps or gaps. This plane tessellation forms a lattice or regular tiling which is a principal feature of these kind of images. For this reason, lattice extraction is an important concern in tasks such as textile inspection, tile cataloging, or QBIC (*Query By Image Contents*) applications.

We are interested in those wallpaper images that are the most common expression of plane patterns. These patterns are created following a strict set of geometric rules that are described in the “Tiling and Pattern Design Theory” (Horne, 2000). This theory establishes that any planar design must include two translational symmetries as well as other symmetries, such as rotations, reflections or glide reflections. It also establishes that there are only 17 combi-

nations or groups of these isometric transformations, called Wallpaper Groups (Plane Symmetry Groups). Therefore, the structural description of these designs is composed of the translational symmetries (lattice), the plane symmetry group they belong to, and the motif or fundamental domain used to create them.

In this work, we approach the first problem: lattice extraction. The wallpaper images always include the translational symmetry, which means that the image can be reconstructed by repeating the pattern, or the motif, using two linearly independent unaligned displacements. For this reason, lattice extraction can be formulated as a problem of obtaining a grid of discrete points that are related by being separated an integer number of times by the length of the motif in both directions. This grid is described, as is shown in fig. 1, by two direction vectors (four numerical values: two angles and two lengths). The parallelogram formed by these vectors is called a *Fundamental Parallelogram* (FP) and, ideally, the image of this FP is the motif originally used to create the pattern. But, from the point of view of lattice extraction, the location of the starting point of the grid does not matter. Changing this point only changes the content of the

FP, not the grid parameters. Wherever the grid is located; the pattern can be rebuilt by repeating the FP defined at that point.




Figure 1: Grid computed from the detection of a repeated area and characterized by two directional vectors.

For these reasons, lattice extraction is often formulated in terms of finding periodicity in the image, which takes us to the concept of grid previously enunciated. Ideally, in synthetic images, it is possible to compute, with complete accuracy, the repetitiveness. The problem, in the case of real images, is due to the existence of deficiencies that makes it impossible to find an exact match between pixel pairs and, therefore, the need to adjust automatically some tolerances (differently in each image). Typical examples of deficiencies in textile or ceramic images are: noise or perspective distortion introduced in the process of image acquisition; small variations in the industrial manufacturing process, or the variations inherent in hand-made productions.

Several techniques have been employed to approach the 2D periodicity problem, as described in detail in (O'Mara, 2002). The most common method is autocorrelation. It is based in using part of the image (typically one half) as a mask and performing the autocorrelation of this mask with the whole image. The peaks, or local maxima, of the autocorrelation map can be extracted and the vectors that define their unions used to propose a grid. Fig. 2 shows two examples of autocorrelation maps on grey level images and the resulting grid superimposed on the image from the point (0,0). As may be noted, the second case is wrong because the same threshold values are used in both images and the correct tessellation is not found. This example illustrates a common case of autocorrelation failure. The threshold selection is an image dependent problem, whereby images with low dynamic range produce flat autocorrelation maps and it is therefore difficult to find significant local maxima. In this work, we approach the problem in a different way. Instead of using the grey level values to find the repetitiveness between image points, we use a symmetries space whose values depend on the local symmetries in the neighborhood of each image point. Because wallpaper images have strong inherent symmetries, the values in this new space peak clearly and are less dependent on the image contrast and noise.




Figure 2: Autocorrelation results on grey level images. Left: autocorrelation peaks. Right: grid obtained from connecting a fixed number of peaks, as a ratio of the maximum.

The symmetries space is obtained through the analysis of phase discrepancy (or congruency) between local frequency components at each image point. These frequency components are computed at several orientations and scales using wavelet filters. This phase analysis enables the automatic detection of these locally oriented symmetries and, therefore, computes a high value of entropy, or symmetry, when it estimates the orientation and scale corresponding to the lengths or distances of appearing symmetries. If more local symmetries appear at different orientations with the same periodicity value, then confidence in that value increases. In short, this process obtains the periodicity of the signal from the identification of two local symmetries that repetitively appear in this new space.

2 PHASE ANALYSIS

The chosen method of phase analysis was developed by Kovesei (Kovesei, 1997). He established that image features such as step edges, lines, roof edges and match bands, all give rise to points where the Fourier components of the image, computed at those points, are maximally in phase. Oppositely, highly symmetric image points will have minimally in phase frequency components. Therefore, he introduces two dimensionless quantities, called Phase Congruency (PC) and Symmetry (Sym), which provide absolute measurements of the significance of these feature points. The use of phase congruency for marking fea-

tures has significant advantages over gradient based methods. It is invariant to changes in image brightness and contrast. Moreover it does not require any prior recognition or segmentation of objects, thus enabling the use of universal thresholds values that can be applied over a wide range of images.

Kovesi proposes a method to obtain these features, based on the determination of the energy at a local level. Kovesi's method uses a multiresolution approach based on a pair of symmetric/antisymmetric logGabor wavelet filters (M^e, M^o), from which an accumulated response for a given scale s and orientation o is obtained using the eq. 1.

$$[ef_{s,o}(x), of_{s,o}(x)] \leftarrow [I(x) * M_{s,o}^e, I(x)M_{s,o}^o] \quad (1)$$

The values $ef_{s,o}(x)$ (even) and $of_{s,o}(x)$ (odd) can be thought of as the real and imaginary parts of complex valued frequency components. These values enable the content of a grey image to be rewritten as the combination of a phase component (A_s) and an orientation component (ϕ) computed as follows:

$$\begin{aligned} A_{s,o}(x) &= \sqrt{ef_{s,o}^2(x) + of_{s,o}^2(x)} \\ \phi_{s,o}(x) &= \text{atan2}(ef_{s,o}(x) + of_{s,o}(x)) \end{aligned} \quad (2)$$

The *Symmetry* measurement is computed through the eq. 3. It provides a value of local symmetry combining the phase and orientation components obtained from filter responses over multiple scales and orientations, $Sym(x)$, as:

$$\frac{\sum_{o=1}^O \sum_{s=1}^S |A_{s,o}(x)| [|\cos(\phi_o(x))| - |\sin(\phi_o(x))|] - T_o}{\sum_{o=1}^O \sum_{s=1}^S A_{s,o}(x) + \epsilon} \quad (3)$$

being $\bar{\phi}(x)$ the overall mean phase angle, T_o is a noise compensation term that represents the expected response from noise in the image (it is derived from the computed values of $A_o = \sum_{s=1}^S A_{s,o}$), The ϵ is a small constant to avoid ill-conditioned situations caused by the very small values of the amplitude value $A_{s,o}$. Eq. 3 can be rearranged as follow:

$$Sym(x) = \frac{\sum_{o=1}^O \sum_{s=1}^S [|ef_{s,o}(x)| - |of_{s,o}(x)|] - T_o}{\sum_{o=1}^O \sum_{s=1}^S A_{s,o}(x) + \epsilon} \quad (4)$$

meaning that the symmetry value is the difference between the absolute responses of even-symmetric and odd-symmetric wavelet filters. This measure varies between $[-1..+1]$ and is almost linear with phase deviation. It must be noted that this model looks for image transitions from white to black and in the reverse direction. It can also be instructed to look for just one transition by changing the term that involves the combination of filter response. This is made by using a *polarity* parameter. Using the formulation for only one

transition (positive polarity) provides a wider range of symmetry values, which facilitates the search for similar values. Positive polarity is assumed in the rest of the paper and the term $-[ef_{s,o}(x) - of_{s,o}(x)]$ is used in the eq. 3. Thus, the *Sym* value ranges from $[0..+1]$.




Figure 3: Kovesi's phase analysis with *polarity* = +1: a) original image, b) Symmetry values (equalized), and c) partial 3D view of b).

Fig. 3 shows the results obtained applying eq. 4 over a grey level image (the luminance of the original RGB image), with the default parameters proposed by Kovesi. The motivation for making the transformation can be seen in figure 3-c, where a detailed view of the left part of the symmetry image is shown. It can be seen that the transformation retains the repetitiveness and the spatial relationships of the original image points, but it is now represented as peaks distributed along the image in places where there is a high local symmetry.

The phase analysis also produces, as a result, a second component: the orientation component. This reflects the orientation angle, for each image point, for which the maximum symmetry value was reached. We return to this component later.

3 LATTICE EXTRACTION ALGORITHM

In this work we propose an methodology based on the transformation from the bitmap space to a symmetry domain where this search can be easily performed. It is based on a three-stage algorithm to detect and extract the grid of a regular image as follows: first, transform from the original RGB values of image do-

main into a symmetry domain (this is achieved by using Kovesei's symmetry model applied to a grey level version of the original image); second, segmentation, i.e. look for similar symmetry locations as identifiers of equivalent local pixel contents; and third, detect object unions by obtaining the direction vectors that define the lattice by the studying the unions between the identified observations.

The first step is made through the use of the phase analysis to obtain symmetry values. The phase analysis is ran over a grey scale version (I : the luminance component) of the original RGB image. The analysis is repeated over a negative version of the image ($\sim I$) to be able to cope with an initially unknown distribution of original pixel values, dark backgrounds, and clear motifs, or viceversa. This results in two symmetry images $SymD$ (direct) and $SymI$ (inverse), together with their corresponding orientation images.

Some decisions about the parameters of this transformation must be taken. Firstly, the number of levels of the multiresolution analysis depend on the image resolution, and goes from 1 to $\log_2(\max(M,N))$, $M \times N$ being the image size. Secondly, the number of orientations (O) is initially chosen as 180 (the complete range of directions on the plane $[0..\pi]$) by using increments of 1 degree. Rotational symmetries lower than 30 degrees are not allowed in wallpaper designs, so a 7 degree step in orientations can be used with sufficient accuracy - as the Nyquist theorem shows.

Finally, as the logGabor wavelet weighted by a spread function is the filter used, the parameters of both have to be chosen. For the wavelets, the cut-off frequency value is computed from a minimum: the minWaveLength value chosen as its related period can be as much as the half the largest dimension of the image, that is, to detect a minimum periodicity of 2 in the image. For the spread that controls the sharpness of the directional selectivity of the filters (as shown in fig. 4), the orientation is used together with a constant ($dThetaOnSigma = 1.7$) value that ensures an even coverage of the 2-D frequency spectrum. The combined response of the filters for a constant value of frequency and spread value, but different values of orientation can be seen in fig. 4.

At this point, the two versions (*direct* and *inverse*) of the transformed domain of representation must be examined and/or combined to obtain a binarized intermediate image - from which the last step can be applied. This segmentation and direction determination needs a little more consideration if they are to be achieved without imposing predetermined thresholds. The subsections below discuss this topic.




Figure 4: From the left: filter appearance for a fixed value of frequency, orientation of 0 and spread values of 0.5, and 1.0, and examples of filter directionality, for a spread of 1,77, a constant frequency value and orientations of 60, and 180 degrees.

3.1 Segmentation Step: Looking for Objects

At this point, the two versions (*direct* and *inverse*) of the transformed domain of representation must be examined and/or combined to obtain a binarized intermediate image - from which the last step can be applied.

This stage consist in analysing each one of the symmetry images by looking for points that exhibit the same maximum symmetry value. At least, three are needed to establish two directions and, with them, the resulting grid. But because we are working only with the maximum values of the symmetry image, the influence of noise is very marked. A more robust strategy is possible by taking into account the neighbourhood points of each local maximum because they correspond to nearby points with high symmetry values.

To do this we perform a thresholding of the symmetry images with a iteratively decreasing threshold value ($TSym$) that is a percentage of the absolute maximum value. In each iteration, two binary images ($objectsD, objectI$) are made showing "objects", in the sense of areas, around the local maxima that appear repeatedly and show certain shape regularity. A simple labelling and comparison step enable us to extract these objects and obtaining the more frequently repeated objects. A comparison measure between objects is taken as the euclidean distance between the objects bounding boxes equalized and centered on the mass center. As we do not know how much similar the objects are, we use a threshold ($TSimilitude$) that is a percentage of the normalized euclidean distance. This value is iteratively decreased until a minimum number of similar objects is found. By putting this all together we have the first version of the segmentation algorithm as:

```
function [objs,nObjs]<-extractObjs(symD,symI)
  TSym <- 100%;
  while (nObjs < 4) and (TSym > 0)
    objectsD <- threshold( SymD, TSym )
    objectsI <- threshold( SymI, TSym )
    TSimilitude <- 100%;
    while (nObjs < 4) and (TSimilitude > 0)
```

```

objs, nObjs<-count & compare (objectsD,
                             objectsI, TSimilitude)
if (nObjs < 4)
  TSimilitude <- TSimilitude - 1;
endif
endWhile
if (nObjs < 4)
  TSym <- TSym - 1
endif
endWhile
endFunction

```

The algorithm progress until a minimum number of objects (three) are found, or until the T_{Sym} range has been exhausted. For simplicity, error checking and convergence tests have not been included. Segmentation in the symmetry domain was still difficult because of the existence of noise and strange values obtained on the image border.

As an alternative, we studied the feasibility of the orientation component that can be derived from the phase analysis. We found that, in general, high values in the symmetry domain were the result of the accumulative contribution of filter responses over a wide range of orientations. Furthermore, if we only look in one orientation, a directly binarized version of the orientation image can be obtained. In this image, a value of 1 indicates that there was an “appreciable” filter response in that direction, and a zero if not. The decision about the appreciable filter response is based on comparison with an estimation of the noise in each image. These values are computed from the symmetry value of the minimum analysis scale.




Figure 5: Segmentation stage can be conducted over (a) symmetry (equalized) or (b) orientation components.

We found that this binary orientation image reflects the same objects as those found around the local maxima in the symmetry space. Fig. 5 shows this idea. In this work, we propose the use of this binary orientation component because it is easier, than the symmetry component, as thresholds do not have to be established to segment objects. We propose a second version of the segmentation algorithm as follows:

```

function [objs, nObjs]<-extractObjs(
                             orientacionD, orientacionI )
objectsD <- orientacionD
objectsI <- orientacionI
TSimilitude <- 100%;

```

```

while (nObjs < 4) and (TSimilitude > 0)
objs, nObjs<-count & compare (objectsD,
                             objectsI, TSimilitude)
if (nObjs < 4)
  TSimilitude <- TSimilitude - 1;
endif
endWhile
endFunction

```

3.2 Object Union Step

This point is devoted to the techniques that can be used to establish the correct connection of the points detected by the previous phase. They are focused on the use of Hough and computational geometry techniques. This is the same with any of the techniques set forth in the introduction to solve the problem of detecting periodicity.

Our proposal is to solve this stage in a three-step filtering function. Firstly, to be consistent with the definition of periodicity of a signal, we will restrict the list to the connections that are “observable” between recurrences of the same object in the image. Two elements are considered equal (fig. 6-a)) if they are similar in their characterization of “locally symmetrical”, i.e, the form of objects. This is the set of possible directions.

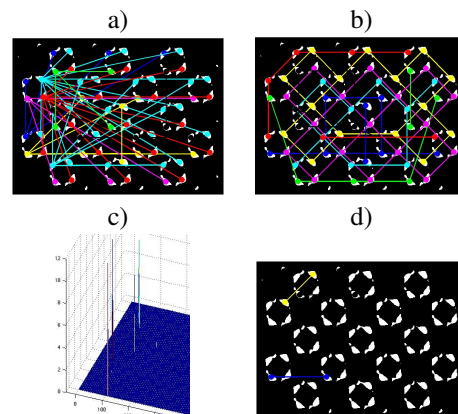


Figure 6: Stages of the object union: (a) all possible connections, (b) removing by grid restrictions, (c) voting results (number of votes for each orientation and displacement pair), and (d) chosen directions.

Secondly, we reduce this connection list by using the spatial constraints imposed by grid definition. Accordingly, a path between any of the grid points is composed of a minimum number of minimum length local connections. As described in (Mount, 2001), this is the area of computational geometry. We use the *Relative Neighbourhood Graph* (RNG) algorithm over the established connections to obtain structurally logical connections, because of its ability (Toussaint,

1980), to extract a perceptually meaningful structure from a set of points. Finally, in a last step we use a voting scheme for the selected connections to choose the two most repeated (fig. 6-c)). These two connections define the grid parameter (fig. 6-d)) and they define the grid geometry (fig. 6-d)) as shown onf the next algorithm:

```
function [grid] <- vDirectors(objects, nObjs)
  directionList <- tentativeGrid( objects )
  directionList<-reduceConections( directionList )
  v1,v2 <- vote( directionList )
  grid <- (v1, v2)
endFunction
```

4 EXPERIMENTS

Once the lattice geometry has been extracted, a final step is needed in order to evaluate the correctness of the method. To achive this, correct databases together with a comparison method must be established. “Wallpaper Groups” (Joyce, 1997), “Wikipedia Wallpaper Groups page” (Wikipedia, 2007), “Basic Tilings” (Savard, 2001), and “PSU Near-Regular Texture Database” (Lee, 2001) contain some image collections closely related to the ones in our field of application. An example of results obtained with the first of these databases is shown in fig. 7. We have complemented the set by adding new images resulting from the application of a set of rigid geometric transformations and different noise models in order to provide for wide range in the samples available.




Figure 7: Results for some images (from left to right): original image, established directions, and *FP* from image center.

With respect to the comparison criteria, the lack of uniformity and correct labeling of databases not including grid geometry, makes expert intervention sometimes necessary to evaluate the quality of the result. Nevertheless, an effort to define automatic criteria based only on image parameters must be carried out.

A first approximation is to define the error in terms of the euclidean distance between the original and a reconstructed image replicating a selected FP obtained using grid parameters. Thus we define, for example, the *Mean FP (MNFP)* and *Median FP (MDFP)* as the parallelogram obtained, respectively, through the mean or median of image pixels related to the grid geometry extracted. The choice of the representative parallelogram is important for the final result of the reconstruction process, so it is not a good statistical. Nevertheless, this method has several problems: the selection of a representative FP is needed, influencing the final statistical results; the image comparison process not only evaluates the correctness of grid geometry but also evaluates the image content; and the lack of measure normalization.

To solve these problems we propose to calculate a measure of *Grid Adjust Error (GAE)* to quantify the error between the grid geometry computed and the true image geometry. The formulation is done using a geometrical criteria intended to minimise errors related to pixel distribution over the image, as characterized by its variance S^2 , which can be expressed as:

$$S^2 = \frac{1}{n} \sum_{i=1}^r n_i S_i^2 + \frac{1}{n} \sum_{i=1}^r n_i (\bar{x}_i - \bar{x})^2 \quad (5)$$

Where the image has been divided in r sets of size n_i , mean \bar{x}_i and variance S_i , being n the size of the image, $n = n_1 + n_2 + \dots + n_r$, and \bar{x} is the mean.

We use an image partition based on the computed grid geometry and define the *Variance FP (VFP)* in a similar way to *(MNFP)* and *(NDFP)* but using pixel variance instead of mean or median. Thus the terms in eq. 5 are instanced as $S_i = VFP_i$ and $\bar{x}_i = MNFP_i$. In such a way, when the geometry of the lattice is correctly determined, the first term of equation tends to zero and the second term to image variance. When the computed geometry is incorrect, the opposite effect occurs. Both terms in eq. 5 range from 0 to S^2 .

Since we need to reduce image distribution effects, finally we propose the normalized measure of *GAE* as follow:

$$GAE = \frac{\frac{1}{n} \sum_{i=1}^r n_i VFP_i}{S^2} \quad 0 \leq GAE \leq 1 \quad (6)$$

Experiments with a number of 100 random lattices were attempted on the wallpaper set, with the aim of showing the statistical value tendency for wrong lattices. The obtained variance values are between 553,28 and 9804,8, the statistic produced a mean value of between 0,945 and 1,019, while the standard deviation of the values is between 0,008 and 0,119. The observed mean values over one are inherent in the process of quantization: a result of the maths errors on

Table 1: Lattice evaluation for some images in Wallpaper database.

I	S^2	\overline{VFP}_{ok}	GAE_{ok}	GAE_{nok} [m, σ]
1	6070,0	529,7	0,09	[0,98, 0,08]
3	7564,4	1722,9	0,23	[0,94, 0,16]
4	6574,3	1581,7	0,24	[0,99, 0,04]
7	2538,3	619,3	0,24	[1,01, 0,02]
8	553,3	137,3	0,25	[1,00, 0,03]
9	7052,3	2258,6	0,32	[1,01, 0,07]
10	3218,9	946,7	0,29	[0,98, 0,19]
13	1115,6	380,5	0,34	[0,95, 0,17]
17	5511,6	2300,4	0,42	[0,95, 0,15]

rounding the values of the computed lattice geometry to obtain pixel coordinates of the image.

Results for the “Wallpaper Groups” database are shown in table 1. This table shows, from left to right, image numbers and variance: the highness of this value means that there is a greater dispersion on pixel values; in regular images (as is the case) this means there is noise. The next columns are the weighted means of VFP and, the GAE values for the correct grid extracted for the current image. Finally, a statistical distribution of GAE appears when using grids extracted from the rest of the images, and this means wrong lattice geometry values. From the last two columns, as it was observed in the other collections, the values of the statistical near to 0 identify a correct lattice, meanwhile when this value is very near to 1 (the greatest deviation is 0,19) can be concluded that the lattice used does not offer a good reconstruction, and it does not explain the image content.

5 CONCLUSIONS AND FUTURE WORK

This work has shown an approach for extracting lattice in regular image based on phase analysis. To achieve this goal we have revised the literature, and observed the existing problems. The correspondence between a formulation of lattice based on periodicity versus symmetry, has led us to examine a multiresolution wavelet transform implementation. The study of this, together with the justification of how to calculate the parameters of the proposed technique, has resulted in an algorithm that produces a process that can auto-adjust itself to the circumstances of each image. A set of experiments was conducted on this algorithm, and the results in the determination of the lattice have been applied to different sets of images.

In addition, as products of this work, an statistic

is proposed to evaluate the correctness of the algorithm. A methodology is also proposed that covers all the variations an image may undergo, typically in the context proposed for implementation. As lines of future work we plan to explore: (a) the computation of the rest of isometries that help us choose to which *Plane Symmetry Group* an image belongs; (b) to obtain a statistical confirmation of the behaviour of the GAE values for geometry values nearby to those of the correct lattice; and (c) extract all the information from the symmetry component, starting from the idea that this aspect leads us to detect directly the axes of reflection symmetry present in the image.

As a practical demonstration of the use of the algorithms submitted, we are using this technique of detecting repetitive patterns for real situations such as: (a) obtaining a structural representation or syntactic of the image that can be used for CBIR, and (b) rebuilding, or recovering, the contents of an image, that has been spoiled by noise, because stains, holes, breaks, and other defects that involve the breakage or incompleteness of the image.

ACKNOWLEDGEMENTS

This work is supported in part by project VISTAC (DPI2007-66596-C02-01).

REFERENCES

- Horne, C. E. (2000). *Geometric symmetry in patterns and tilings*. Woodhead Publishing. ISBN 1 85573 492 3, Abington Hall (England).
- Joyce, D. E. (1997). Wallpaper groups (plane symmetry groups). <http://www.clarku.edu/~djoyce/>.
- Kovesi, P. (1997). Symmetry and asymmetry from local phase. In *AI'97, Tenth Australian Joint Conference on Artificial Intelligence*. pp 185-190.
- Lee, S. (2001). Psu near-regular texture database. <http://vivid.cse.psu.edu/texturedb/gallery/>.
- Mount, D. (2001). *Computational geometry: Proximity and Location, cap. 63*. CRC Press, 0-8493-8597-0.
- O'Mara, D. (2002). *Automated facial metrology, Chapter 4: Symmetry detection and measurement*. PhD thesis, University of Western Australia.
- Savard, J. J. G. (2001). Basic tilings: The 17 wallpaper groups. <http://www.quadibloc.com/math/tilint.htm>.
- Toussaint, G. T. (1980). The relative neighbourhood graph of a finite planar set. *Pattern Recognition*, vol. 12, pp. 261-268.
- Wikipedia (2007). Wallpaper group — wikipedia, the free encyclopedia. <http://www.wikipedia.org/>.