

Lattice Estimation from Images of Patterns that Exhibit Translational Symmetry²

Junwei Han¹, Stephen J. McKenna²,

Abstract

The analysis of regular texture images is cast in a model comparison framework. Texel lattice hypotheses are used to define statistical models which are compared in terms of their ability to explain the images. This approach is used to estimate lattice geometry from patterns that exhibit translational symmetry (regular textures). It is also used to determine whether images consist of such regular textures. A method based on this approach is described in which lattice hypotheses are generated using analysis of peaks in the image autocorrelation function, statistical models are based on Gaussian or Gaussian mixture clusters, and model comparison is performed using the marginal likelihood as approximated by the Bayes Information Criterion (BIC). Experiments on public domain images and a commercial textile image archive demonstrate substantially improved accuracy compared to several alternative methods.

Keywords: Texture, Statistical Image Representation, Pattern Recognition

Email addresses: jhan@nwpu.edu.cn (Junwei Han),
stephen@computing.dundee.ac.uk (Stephen J. McKenna)

¹This research was carried out while J. Han was with the School of Computing, University of Dundee, DD1 4HN, UK. He is currently with the School of Automation, Northwestern Polytechnical University, Xian 710072, China.

²S. J. McKenna is with the School of Computing, University of Dundee, DD1 4HN, UK.

²This is an author-created version of an article accepted for publication in the journal *Image and Vision Computing*. The final publication is available at <http://www.journals.elsevier.com/image-and-vision-computing/>

1. Introduction

Regular textures can be modelled as consisting of repeated texture elements, or *texels*. The texels tessellate (tile) the image or, more generally, a surface. Varying illumination, varying physical characteristics of the textured surface, geometric deformations and sensor noise all result in images of such patterns exhibiting approximately regular, as opposed to exactly regular, texture. In this paper, we consider the task of automatically extracting texels from images. In particular we are interested in analysing patterns in printed textiles. Such patterns often exhibit translational symmetry and the visual structure within each texel can itself be complex. Translationally symmetric regular textures can always be generated by a pair of shortest vectors (two linearly independent directions), \mathbf{t}_1 and \mathbf{t}_2 , that define the size, shape and orientation (but not the position) of the texel and the lattice that it generates. The lattice topology is always then quadrilateral. Such textures form one class of wallpaper pattern [1]. We restrict ourselves to consideration of images of planar, approximately regular textures viewed under orthographic projection. While this might at first seem restrictive, this problem is, as will become apparent, far from completely solved. Furthermore, solutions will find application in analysis, retrieval and restoration of textile, wallpaper and tile design images, for example. The aim then, is to automatically decide whether a translationally symmetric regular texture class provides a good model for a texture image, and in cases in which it does, to extract the most predictive texel geometry.

In an earlier conference paper [2], this problem was cast in a statistical model comparison framework; models representing different hypotheses for lattice geometry were compared using the Bayesian Information Criterion (BIC). In this paper we extend this work in several ways. A more extensive discussion of related literature is provided setting the work more clearly in context, and certain aspects of the method have been clarified. False hypotheses are pruned using a constraint on the angle between \mathbf{t}_1 and \mathbf{t}_2 . Classification of texture images as regular or irregular using model comparison is compared empirically to an alternative method [3, 4]. Model comparison is performed using the Akaike Information Criterion (AIC) as well as BIC. New experiments are reported on larger data sets, and results are reported separately for private and public data sets thus enabling future comparisons to be made using the public data. A detailed comparison of lattices obtained using several methods is presented. Experiments involving qualitative observer

evaluations are expanded. The proposed method and implementations of two alternative methods from the literature are compared quantitatively against ground-truth annotations using two quantitative evaluation measures.

The rest of this paper is organized as follows. Section 2 outlines the most closely related previous work. Section 3 presents the model comparison framework. Section 4 describes details of lattice models. Section 5 describes the method used in our experiments for generating lattice hypotheses. Evaluations are described in Section 6 and conclusions are drawn in Section 7.

2. Related work

Previous work proposed for texel and lattice extraction can be grouped broadly into two categories: the local feature-based approach [5, 6, 7, 8, 9, 10, 11, 12, 13] and the global structure-based approach [1, 14, 15, 16, 17, 18]. All texture analysis is necessarily both local and global. The categorisation is in terms of the computational approach: whether it starts by identifying local features and proceeds to analyse global structure, or starts with a global analysis and proceeds by refining estimates of local structure.

The local feature-based approach starts by identifying a number of texel candidates. Matching based on visual similarity between these potential texels and their neighbours is then performed. Successful matching leads to the connection of texels into a lattice structure. The approach iterates until no more new texels are found. Methods vary in the way they initialise texel candidates and in the parametric models used to cope with geometric and photometric variation. Lin et al. [6] asked users to provide an initial texel. Interest points and edges have been used to generate texel candidates automatically [7, 8, 9]. However, Hays et al. [5] pointed out that interest points often fail to find texel locations and instead initialized by combining interest points and normalized cross correlation patches. Affine models have been adopted to deal with local variation among texels [7, 10, 11]. Global projective transformation models have also been used, taking advantage of the spatial arrangement of texels [8, 9]. Hays et al. [5] formulated lattice detection as a texel correspondence problem and performed texel matching based on visual similarity and geometric consistency. Lin et al. [6] proposed a Markov random field (MRF) model with a lattice structure to model global topological relationships among texels and an image observation model able to handle local variations. Park et al. [12, 19] similarly used a lattice MRF formulation, iteratively refining, growing and warping the lattice. Methods

described by Hays et al. [5] and Park et al. [19, 20] addressed the problem of detecting deformed lattice structures in images when those structures did not occupy the entire image. Without needing the assumption of a lattice layout, Zeng et al. [13] defined a point-wise repetition measure and incorporated it into an energy minimization formulation to segment images with repetitive elements.

As indicated in [1], available local feature-based methods can be effective at detecting small texels under significant texture surface deformation and are more suited to such situations. However, these methods require texels that can be identified based on local features (such as corners) and perform matching between individual texels. They often fail to detect texels with complex internal appearance. Fig. 2 shows two examples of such failure. While the methods described in [5] and [12] are able to automatically estimate lattices from near-regular texture images, they are not designed for use with patterns in which the visual structure within each texel is complex, as is the case for many of the images of textile designs in our data (see Fig. 9).

The global structure-based approach [1, 14, 15, 16, 17, 18] tries to extract texels using methods that emphasise the idea of periodic patterns as global processes. Starovoitov et al. [18] used features derived from cooccurrence matrices to extract texels. Gimel'farb et al. [21] described a heuristic scheme to extract *rectangular* texels from *model-based interaction maps*. These maps were constructed from Gibbs random field models and represented partial pairwise pixel interaction energies over large rectangular pixel neighbourhoods. They noted that, more generally, the shape and orientation of texels is in principle governed by certain clusters of such energies. Charalampidis *et al.* [17] used a Fourier transform and made use of peaks corresponding to fundamental frequencies to identify texels. The autocorrelation (AC) function is generally more robust than the Fourier transform for the task of texel extraction especially in cases in which a regular texture image contains only a few texel repetitions [1, 14]. Chetverikov [15] developed a regularity measure by means of finding the maximum over all directions on the AC function. Leu [16] used the several highest peaks in the AC function computed on the gradient field of the image to capture translation vectors. A promising approach was presented by Lin et al. [14] in which salient peaks were identified using Gaussian filters to iteratively smooth the AC function. The Hough transform was then applied to find translation vectors, \mathbf{t}_1 and \mathbf{t}_2 . Liu et al. [1] highlighted the fact that spurious peaks often result in incorrect lattice vectors. Therefore, they proposed a *region of dominance* operator that

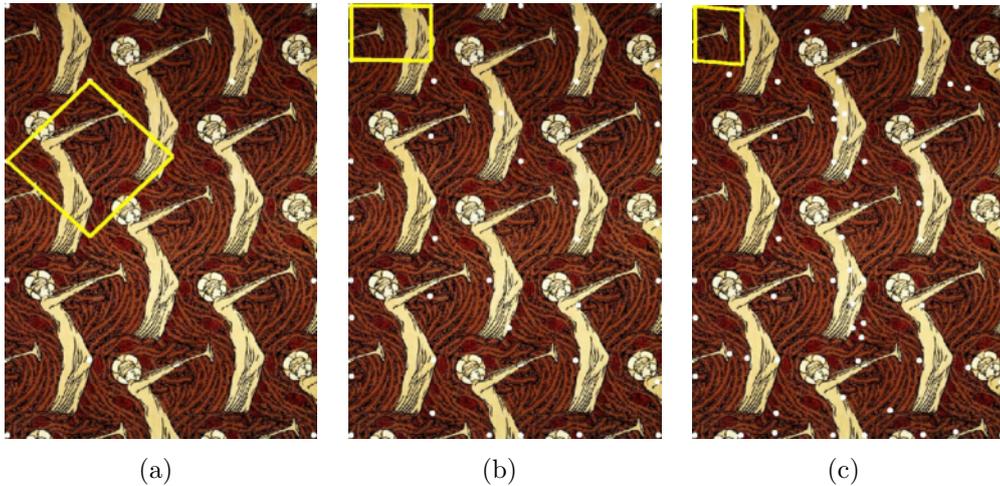


Figure 1: Texels obtained using (a) ten, (b) forty, and (c) seventy dominant peaks in the autocorrelation function. The peak locations are marked with white dots.

takes into account the spatial relationships among peaks to select a list of dominant peaks for translation vectors construction. They observed that a peak with low height located far from any larger neighbours is often more perceptually important than a high peak close to another high peak. However, the problem of how to determine the number of dominant peaks was not addressed. Wu et al. [22] also noted this important limitation. Whilst it is usually relatively easy for a human to select an appropriate subset of peaks, automating this process is difficult. Fig. 1 shows three different texels obtained similarly to Lin et al. [14] from the same image by using different numbers of peaks. The peaks were obtained using the region of dominance method [1]. Whilst using only the first ten peaks can result in success, the method is rather sensitive to this choice. In summary, global structure-based methods are suitable for textures that do not exhibit large geometric deformation and often successfully identify larger texels with more complicated appearances. However, existing methods have free parameters for which a fixed value that works on a wide range of images can often not be found. Methods based on finding peaks in an AC function often yield many unreliable peaks and the number which are reliable can vary dramatically between images. This serious drawback currently makes these methods difficult to apply to large image collections.

There are relatively few papers concerned with discriminating regular

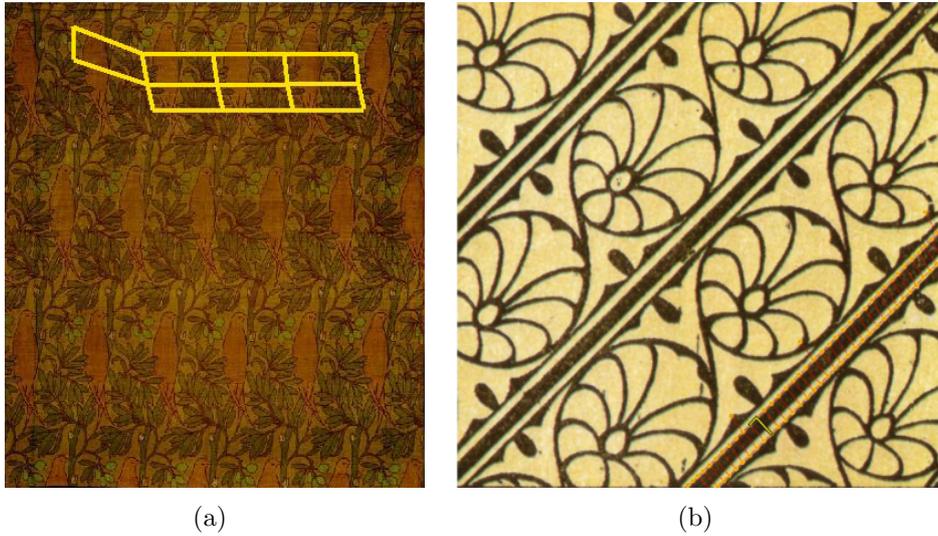


Figure 2: Examples of two local feature-based methods extracting incorrect lattices. (a) Result obtained using an implementation of the method in [5]. (b) Result obtained using an implementation of the method in [12]. A partial, incorrect, lattice is found along the lower diagonal stripe.

texture images with translational symmetry from images not exhibiting this regularity. Lee and Chen proposed a Fourier transform-based measure called Radial Wedge Distribution Variance (RWDV) for this purpose [3, 4]. Others have proposed texture descriptors related to texture regularity without attempting to infer texel geometry. For example, Wu et al. [23] proposed a texture descriptor based on Gabor filters to characterize the level of ”‘structuredness’” in an image. Liu and Picard [24] described the use of the 2D Wold decomposition to obtain features that approximate human perception of ”‘periodicity’”, ”‘directionality’”, and ”‘randomness’”.

3. Model comparison framework

The approach we take is to formulate texel hypotheses as statistical models and then compare these models given the image data. It is not sufficient for a model to be able to fit the data well. The best texel hypothesis under this criterion would be the image itself whereas our purpose is to extract the smallest texture element. Therefore, overfitting must be guarded against by penalising model complexity. Texel hypothesis comparison can be regarded as a model comparison problem for unsupervised statistical modelling of data.

Such a problem can be formulated as Bayesian model comparison which naturally penalises complexity (Occam’s razor).

Let $I = \{x_1, x_2, \dots, x_N\}$ be an image with N pixels. Here, $x_n, 1 \leq n \leq N$ is the intensity of the n^{th} pixel. Let $H \equiv (\mathbf{t}_1, \mathbf{t}_2)$ denote a texel hypothesis for I , H_k the k^{th} in a set of hypotheses, and M_k a statistical model defined based on H_k with parameters θ_k . Texel extraction can be formulated as choosing the most probable texel hypothesis given the image. According to Bayes’ theorem,

$$p(H_k|I) = \frac{p(I|H_k)p(H_k)}{p(I)}. \quad (1)$$

For each H_k , we define a unique model M_k . Hence, two hypotheses can be compared by evaluating the posterior ratio of the corresponding models:

$$\frac{p(M_k|I)}{p(M_l|I)} = \frac{p(M_k)p(I|M_k)}{p(M_l)p(I|M_l)} \quad (2)$$

where

$$p(I|M_k) = \int p(I|\theta_k, M_k)p(\theta_k|M_k)d\theta_k \quad (3)$$

and similarly for $p(I|M_l)$. Here $p(I|\theta_k, M_k)$ is the probability of the image data given the model M_k and its parameters θ_k , and $p(\theta_k|M_k)$ is the prior probability of parameters θ_k given the model M_k . Equal prior probabilities are assigned to each of the hypotheses to be compared. The posterior ratio is then equal to the ratio of the marginal likelihoods. This ratio expresses how well one hypothesis predicts the image compared to another hypothesis [25].

The proposed approach to regular texture analysis involves (i) generation of multiple texel hypotheses, and (ii) comparison of hypotheses based on statistical models. The hypothesis, $H_{\hat{k}}$, with the model that has the largest marginal likelihood is selected:

$$\hat{k} = \arg \max_k \{p(H_k|I)\} \quad (4)$$

This method can be used to classify textures as regular or irregular by comparing the most probable lattice found with a reference hypothesis in which the entire image is a single ‘texel’.

The lattice model should be able to account for both regularity from periodic arrangement and statistical photometric and geometric variability.

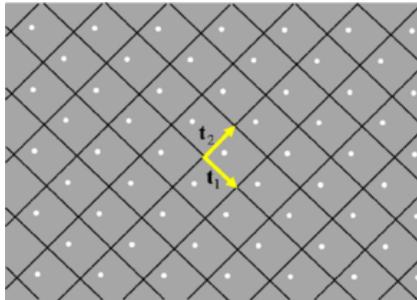


Figure 3: An example of cluster allocation according to a texel hypothesis, $H \equiv (\mathbf{t}_1, \mathbf{t}_2)$. The value of $f(n, H)$ is the same for each of the highlighted pixels. There are Q pixels in each parallelogram.

Let us first suppose a regular texture image I with N pixels x_1, x_2, \dots, x_N , and a hypothesis H with Q pixels per texel. Based on H , each pixel of the image is assigned to one of Q positions on the texel according to the lattice structure. Thus, the N pixels are partitioned into Q disjoint sets, or clusters. If we choose to assume that the N pixels are independent given the model, we have,

$$p(I|M) = \prod_{n=1}^N p(x_n|M) = \prod_{q=1}^Q \prod_{n:f(n,H)=q} p(x_n|M) \quad (5)$$

where f is a function that maps n and a lattice hypothesis to an element of the set $\{1, \dots, Q\}$. f is used to consider all pixels in the image that have the same location relative to the texel within which they are situated. Fig. 3 illustrates the pixels assigned to one of the Q clusters.

The integrals in Equation (3) can only be computed analytically in certain cases such as exponential likelihoods with conjugate priors. Otherwise, approximations can be obtained using sampling methods, for example. While it would be interesting to explore these alternatives in future work, this paper uses the Bayes Information Criterion (BIC) as a readily computable approximation. BIC approximates the marginal likelihood integral via Laplace's method and the reader is referred to the papers by Schwarz [26] and Raftery [27] for full details of its derivation. Given a maximum likelihood parameter estimate, $\hat{\theta}$, we have

$$\log p(I|M) \approx \log p(I|\hat{\theta}, M) + \log p(\hat{\theta}) + \frac{d}{2} \log 2\pi \quad (6)$$

$$-\frac{d}{2} \log N - \frac{1}{2} \log |\mathbf{i}| + O(N^{-1/2})$$

where d is the number of parameters and \mathbf{i} is the expected Fisher information matrix for one observation. The subscript k has been dropped here. The term $\log p(I|\hat{\theta}, M)$ is of order $O(N)$, $(d/2) \log N$ is of order $O(\log N)$, and the remaining terms are of order $O(1)$ or less. The log marginal likelihood can be approximated by removing all terms of order $O(1)$ or less. The BIC for the model is then

$$BIC(M) = -\log p(I|\hat{\theta}, M) + (d/2) \log N \approx -\log p(I|M) \quad (7)$$

The first term can be interpreted as an error of fit to the data while the second term penalises model complexity.

Using the BIC approximation, hypothesis $H_{\hat{k}}$ is selected where,

$$\hat{k} = \arg \min_k \{BIC(M_k)\} \quad (8)$$

A widely used alternative to BIC for model selection is Akaike's information criterion (AIC) [28]. According to that method, models with low AIC values are preferred, where,

$$AIC(M) = -\log p(I|\hat{\theta}, M) + d. \quad (9)$$

The second term penalises model complexity. When $N > e^2$, as in the models considered here, the second term in BIC is greater than d . Therefore, AIC penalises model complexity less than BIC. The use AIC in place of BIC is evaluated for comparison in Section 6.

Our model, M_{irr} , for the irregular texture hypothesis is that the 'texel' is the entire image. If a texture is regular, a model with a texel smaller than the entire image should better predict the image. If no such texel can be found that better predicts the image, the irregular hypothesis should be preferred. Therefore, if the irregular hypothesis results in a higher BIC value, the image is classified as regular. Otherwise, it is classified as irregular, i.e.

$$\begin{aligned} BIC(M_{irr}) &\leq BIC(M_{\hat{k}}) && \text{Irregular texture} \\ BIC(M_{irr}) &> BIC(M_{\hat{k}}) && \text{Regular texture} \end{aligned} \quad (10)$$

where $M_{\hat{k}}$ is the best lattice hypothesis selected by Equation (8).

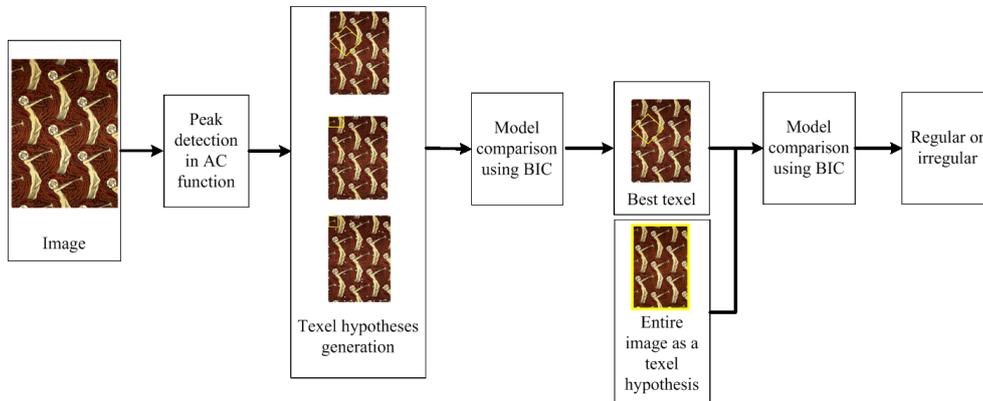


Figure 4: System overview. Model comparison is used for identifying the best texel hypothesis and for comparing it with the entire image as a texel hypothesis. Hypotheses were generated using analysis of the autocorrelation function.

Fig. 4 shows a system overview including the use of model comparison for texel estimation. We will now describe some particular choices for lattice models and lattice hypothesis generation.

4. Lattice models

Each of the Q clusters was modelled as Gaussian, a modelling choice motivated by exploratory data analysis which showed Gaussian-like distributions, as well as by parsimony and computational simplicity. Experimental results (see Section 6) suggest that this was a reasonable choice, at least for the data used here. However, when data sets contain many spatially localised disruptions (due to physical surface damage, local photometric variations such as specular highlights, local geometric distortions, or small occluding objects for example), a more heavy-tailed distribution that penalises resulting outliers less heavily can be used. For example, a cluster can be modelled as a mixture of two Gaussians with the same mean but different variances, (σ_1^2, σ_2^2) , and a mixing weight, π_1 that places greater weight on the low variance Gaussian.

Modelling each of the Q clusters as Gaussian with fixed variance gives:

$$BIC(M) = -\log p(I|\hat{\theta}, M) + (Q/2) \log N \quad (11)$$

$$= -\sum_{q=1}^Q \sum_{n:f(n,H)=q} \log p(x_n|\hat{\mu}_q, \sigma^2) + (Q/2) \log N \quad (12)$$

$$= \frac{1}{2\sigma^2} \sum_{q=1}^Q \sum_{n:f(n,H)=q} (x_n - \hat{\mu}_q)^2 + (Q/2) \log N + C_1 \quad (13)$$

where C_1 is a constant that depends on σ^2 , and $\hat{\mu}_q$ is a maximum likelihood estimate of the mean of the q^{th} cluster. Alternatively, modelling each cluster as a mixture of two Gaussians gives:

$$BIC(M) = -\log p(I|\hat{\theta}, M) + (Q/2) \log N \quad (14)$$

$$= -\sum_{q=1}^Q \sum_{n:f(n,H)=q} \log p(x_n|\hat{\mu}_q, \sigma_1^2, \sigma_2^2, \pi_1) + (Q/2) \log N \quad (15)$$

$$= -\sum_{q=1}^Q \sum_{n:f(n,H)=q} \log\left(\frac{\pi_1}{\sigma_1} \exp\frac{-(x_n - \hat{\mu}_q)^2}{2\sigma_1^2} + \frac{1 - \pi_1}{\sigma_2} \exp\frac{-(x_n - \hat{\mu}_q)^2}{2\sigma_2^2}\right) + (Q/2) \log N + C_2 \quad (16)$$

where C_2 is a constant. In each case, there is one free parameter per cluster, i.e., $d = Q$.

5. Lattice hypothesis generation

In principle there is an infinite number of lattice hypotheses. However, probability density will be highly concentrated at multiple peaks in the hypothesis space. The posterior distribution can therefore be well represented by only considering a, typically small, number of hypotheses at these peaks. In the maximum a posteriori setting adopted here, the approach taken is to identify multiple hypotheses in a data-driven manner and then compare these hypotheses using BIC. The approach is general in that any algorithms that generate a variety of reasonable hypotheses can be used.

Peaks in AC functions are associated with texture periodicity but automatically deciding which peaks can characterize the arrangement of texels is problematic and has not been properly addressed in the literature [1, 14, 15, 16]. In particular, changing the number of peaks considered can result in different lattice hypotheses. Since the total number of peaks is limited we can only obtain a limited number of hypotheses.

Given a grey-scale image $I(x, y)$, $1 \leq x \leq W$, $1 \leq y \leq L$ where W and L are image width and height respectively, its AC function can be computed as follows:

$$AC(x, y) = \frac{\sum_{i=1}^W \sum_{j=1}^L I(i, j) I(i + x, j + y)}{\sum_{i=1}^W \sum_{j=1}^L I^2(i, j)} \quad (17)$$

Applying the fast Fourier transform (FFT) to calculate the AC function is a more efficient alternative.

$$AC(x, y) = F^{-1}[F[I(x, y)]^* F[I(x, y)]] \quad (18)$$

where F and F^{-1} denote FFT and inverse FFT.

This paper combines aspects of the methods of Lin et al. [14] and Liu et al. [1] to generate lattice hypotheses. First, we apply Gaussian filters to iteratively smooth the AC function. Then, salient peaks obtained from the first stage are ranked according to their dominance. The most highly ranked peaks are selected as input for lattice hypothesis construction. Specifically, the generalized Hough transform was used to find pairs of displacement vectors from the peaks similarly to Lin et al. [14]. The number of peaks in the rank-ordered list to use was varied in order to generate multiple hypotheses. Typically a few tens of the generated hypotheses will be distinct. However, according to wallpaper pattern theory, only hypotheses where the angle between \mathbf{t}_1 and \mathbf{t}_2 is between 60 and 90 degrees are valid translation vectors [1]. Therefore, false hypotheses are removed using this constraint. This step reduces computation time and improves the accuracy of texel extraction.

6. Evaluation

There are a few systematic evaluations of lattice extraction algorithms in the literature [29, 30, 19]. Chen et al. [29] reported an evaluation of translational symmetry in terms of the number of texels but not in terms of their shape. Park et al. [19] described an evaluation method for detection of deformed lattice regions in photographs of real-world scenes. Ground-truth lattice structures were annotated using an interactive editing tool. Detection rate was defined as the number of valid texels detected divided by the number of ground-truth texels. They described a semi-automatic method for establishing correspondence between detected texels and ground-truth texels. A detected texel did not need to be the one with the shortest vectors t_1 and t_2 . In order to be valid it had to have an area between 50% and 150% of the corresponding texel in the ground-truth. However, the aim of our method and thus the evaluation methods we use differ from that of Park et al. [19]. We

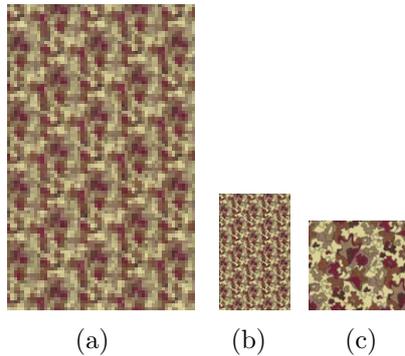


Figure 5: Thumbnails can be generated by cropping a rectangular sub-image that bounds a region a little larger than a texel, in this case $(1.5\mathbf{t}_1, 1.5\mathbf{t}_2)$. (a) original image (b) thumbnail generated by sub-sampling only (c) thumbnail generated by automatic cropping. (b) and (c) are shown here so as to have equal areas.

are not concerned with segmenting lattice regions in unsegmented scenes so do not report detection rates. Rather, we aim to find the correct lattice with the shortest vectors. We use quantitative measures that directly compare the size and shape of estimated texels with annotated ground-truth.

The textures in the images considered are never exactly regular; texels vary in appearance and there can also be small geometric distortions. For this reason, the ground-truth lattice for an image is not uniquely determined; human observers asked to annotate it will exhibit some variation. Furthermore, it can be difficult for a human observer to select the minimal lattice; the mathematically most correct lattice is not always the most perceptually obvious one. In some applications, such as automatic thumbnail generation or clustering based on lattice structure for image browsing [31], what matters to users is that they perceive the lattice to be good. (Fig. 5 shows an example of thumbnail generation in which automatic texel-based cropping results in a thumbnail that conveys very different information about content than simply resampling the image). For these reasons, in addition to quantitative measures comparing to annotated ground truth, we report the results of experiments in which observers qualitatively assessed estimated texels.

A dataset of 200 regular texture images was used for evaluation, comprising 147 images of printed textiles from a commercial archive (courtesy of Liberty Art Fabrics) and 53 images taken from three public domain databases (the Wikipedia Wallpaper Groups page, a Corel database, and the PSU texture database [1]). These images ranged in size from 300×225 pixels to

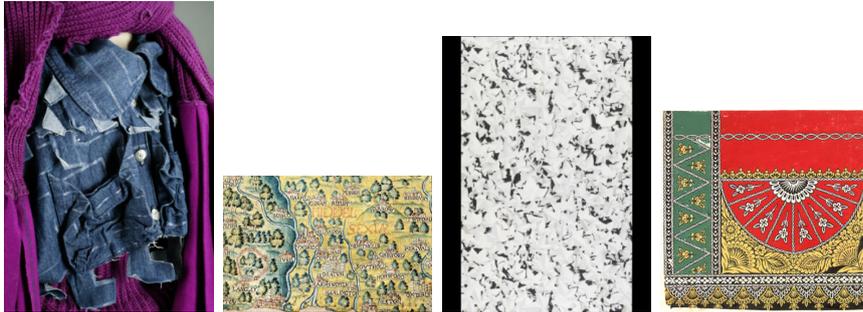


Figure 6: Examples of images to be classified as not having a *regular* texture.

2648 × 1372 pixels. The number of texels along one direction in an image ranged from 2 to a few tens. The criterion for inclusion was that each image was judged qualitatively to contain a translationally symmetric regular texture. This data set includes images that are challenging because of appearance variations between texels, large texels with intricate appearance, small geometric deformations, occluding labels, and stains, wear and tear in some of the textile images. Ground-truth texel data were obtained by overlaying lattice annotations on regular images under user control. This was performed by a researcher who had an understanding of the geometry of wallpaper patterns (the first author). These ground-truth data are available from the authors. In addition, a set of 200 images that were not regular texture images was selected at random from a museum art database and from the same commercial textile archive used earlier. Examples of these images are shown in Figure 6.

When the proposed method used Gaussians to model clusters, the only free parameter was the variance, σ^2 . A suitable value for σ^2 was estimated from a set of 20 training images as follows. Many texel hypotheses were automatically generated using different numbers of AC peaks and a user then selected from them the best translation vectors, $\mathbf{t}_1, \mathbf{t}_2$. Pixels were allocated to clusters according to the resulting lattice and a maximum likelihood estimation of σ^2 was computed. The result was $\sigma^2 = 264$. Since this semi-automatic method might not be using precise texel estimates, it might overestimate the variance compared to that which would be obtained using optimal lattices. Therefore, further values for σ^2 (100, 144 and 196) were also used for evaluation in order to test the sensitivity of the method. In any particular experiment, σ^2 was fixed for all test images. The method was

also evaluated using a Gaussian mixture to model each cluster, with free parameters set to $\sigma_1^2 = 60$, $\sigma_2^2 = 800$, and $\pi_1 = 0.9$.

The method was implemented using Matlab on a 2.4GHz, 3.5GB PC. As an indication of computation time, when processing a 550×745 pixel image it took 0.5s to compute the autocorrelation function (using Matlab `fft` and `ifft2` functions), a further 3s to filter and obtain salient peaks, and 2s to obtain hypotheses using the Hough transform. The complexity of model comparison is then linear in the number of pixels and the number of hypotheses compared. Computing a BIC value for a hypothesis took 90ms.

6.1. Classification

The method was first used to classify texture images as regular or irregular (Equation (10)). It was compared with the method of Lee and Chen [3, 4] which used RWDV to discriminate regular texture from irregular images, and assumed that distributions of RWDV's for the regular and irregular texture images were Gaussian. Classification was then performed using Bayes' rule.

The set of 200 regular texture images was split into two sets of 100 images each. One was used as training and the other as testing. Similarly, half of the 200 irregular texture images were used for training and the other half for testing. The experiment was then repeated after switching the training and test sets and the test error was estimated by averaging these two tests. The parameters of the Gaussian distributions in the method of Lee and Chen were estimated from the training set. Plots of false positive rates against false negative rates were obtained by varying the class priors in the method of Lee and Chen ($p(\text{regular}) \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$), and by varying the value of σ^2 in the proposed method using a Gaussian model ($\sigma^2 \in \{100, 121, 144, 169, 196, 225, 256\}$). Figure 7 shows these plots. The proposed method has greater accuracy. The equal error rate was approximately 0.38 for the method of Lee and Chen. The equal error rate for the proposed method was 0.24.

6.2. Lattice Geometry Estimation

The proposed method for lattice estimation was compared with two related algorithms from the literature [1, 14]. Liu et al. [1] did not specify how to determine the number of peaks in the AC function. Therefore, results are reported here using three different values for the number of peaks, namely

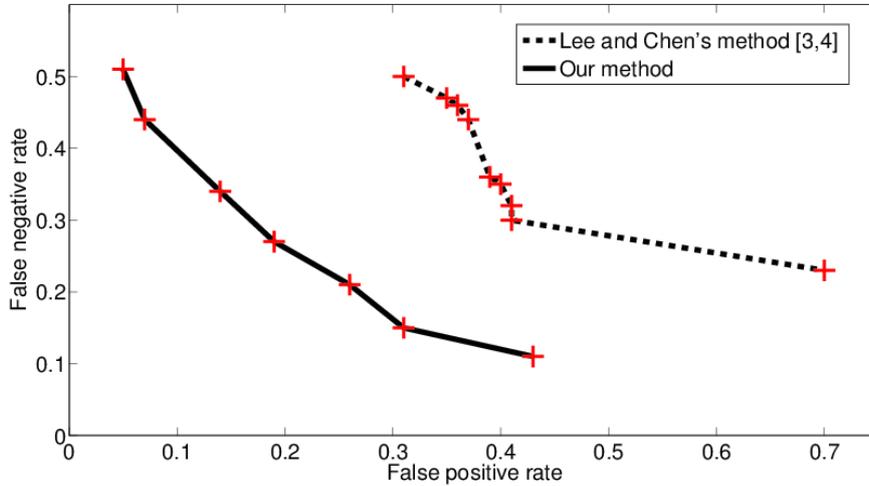


Figure 7: Classification of texture as regular or irregular: false positive rate versus false negative rate.

10, 40, and 70. ⁴

There are other methods which, although related to the proposed approach, are not well-suited for direct empirical comparison with it [5, 12, 19]. They are essentially local-feature based methods that can be effective at detecting small texels under significant texture surface deformation. They require texels that can be identified based on local features (such as corners) and perform matching between individual texels. They often fail to detect texels with complex internal appearance and large size. Although these methods are able to automatically estimate lattices from near-regular texture images, they are not designed for use with patterns in which the visual structure within each texel is complex, as is the case for many of the images analysed here. Furthermore, such methods [5, 12, 19] express the texel geometry as varying with location and have been evaluated primarily in terms of texel *detection* rates. In contrast, the evaluation presented in this paper focuses on the *accuracy* of estimated lattices since for images of textile designs and wallpaper images such as those analysed here, a single pair, $(\mathbf{t}_1, \mathbf{t}_2)$, is sufficient to characterise texel geometry. The proposed approach

⁴Our implementation made use of J. Hays' code implementing the *region of dominance* operator and constructed translation vectors based on pseudocode given in [14].

is more closely aligned in terms of motivation, assumptions and evaluations with the work of Liu *et al.* [1] and Lin *et al.* [14]. Therefore, direct comparison was performed with the methods described in those papers.

6.2.1. Qualitative evaluation

Two volunteers (one male and one female) jointly scored and rank-ordered the algorithms. Initially they assessed the results independently and then in cases of disagreement they were asked to reach agreement through discussion. Disagreement happened in less than 5% of cases. They were able to reach agreement in all cases. The observers were shown lattices overlaid on images and were asked to label each lattice as obviously correct (OC), obviously incorrect (OI), or neutral. They were to assign OC if the lattice was exactly the same or very close to what they expected, OI if the result was far from their expectations, and *neutral* otherwise. Presentation of results to observers was randomised to hide from them which algorithms produced which results. Table 1 summarises the results. In order to facilitate future comparisons, values in parentheses show results using only the 53 images from public domain databases as test images. The angle constraint mentioned at the end of Section 5 enabled 39% of hypotheses to be rejected as false without the need for BIC evaluation. It resulted in a small (3%) increase in OC results. The proposed algorithm resulted in approximately 80% of test images being rated OC ($\sigma^2 = 100$, BIC). The use of AIC in place of BIC resulted in a small decrease in this percentage.

Fig. 8 shows some example texel lattices obtained when different algorithms were applied to the same images. Fig. 9 shows some further examples of lattices obtained using the proposed method. In Fig. 9, examples (a)-(g) were judged to be correct (OC). Examples (a) and (b) show clear appearance variations between texels due to staining and wear. The texels in many examples are intricate e.g., (d)-(g). Example (h) was judged as *neutral*, a result that does seem reasonable at a glance. However, close inspection reveals that a more correct tessellation would merge four neighbouring texels into one. Recovering this preferred tessellation automatically is problematic because only one complete texel is present. Results shown in (i), (j), (k) and (l) were judged to be incorrect (OI). Inspection of the image in (i) reveals significant foreshortening. Much of the texture in (j) is due to the cloth rather than the repetitive nature of the woven design. This example also exhibits geometric deformation. The tessellation in (k) looks reasonable at first glance but closer inspection reveals that doubling the width and height of

the texel is a better solution. The tessellation in (1) is also reasonable but was judged incorrect because every second motif is mirrored and there is some misalignment. Example (1) has a border that includes a colour calibration pattern. Note that several examples had only a few texel repetitions within the image.

In the above experiment, the pair of observers performed a joint decision making process. A separate experiment was performed to assess the extent to which independent observers would agree when they did not have the opportunity for communication with each other. Eight observers (3 females, 5 males) independently evaluated the results of the proposed algorithm (with $\sigma^2 = 100$ as this parameter setting gave the best of the results reported in Table 1). They were asked to label each lattice as OC, OI, or neutral. In 90% of cases a majority of the observers agreed, i.e., at least five of them gave the same answer. We assigned judgements of OC the value 1, judgements of neutral the value of 0, and judgements of OI the value -1. Given values from eight observers, the standard deviation, σ_{obs} , of these values was calculated. For example, if all observers agreed then $\sigma_{obs} = 0$, if half the observers judged a texel to be OC and the other half judged it as neutral then $\sigma_{obs} = 0.25$, and if half the observers judged a texel to be OC and the other half judged it to be OI then $\sigma_{obs} = 1$. On average for the 200 test images, σ_{obs} was 0.26.

A further experiment was performed to compare the proposed method to the two other methods. For each image, lattice results from our algorithm using Gaussians, our algorithm using Gaussian mixtures, the algorithm of Liu et al. [1], and the algorithm of Lin et al. [14] were shown on the screen simultaneously. The two subjects rank ordered those four results. Algorithms shared the same rank if they yielded equally good results. For example, if three of the algorithms gave good lattices of equal quality and the fourth algorithm gave a poor lattice then three algorithms shared rank 1 and the other algorithm was assigned rank 4. If two of the algorithms gave good lattices of equal quality and the other two gave poor lattices of equal quality, then the former shared rank 1 and the latter shared rank 3. Table 2 summarizes the rankings. Values in parentheses show the results obtained when only the 53 images from public domain databases were used. For the Gaussian model, we set $\sigma^2 = 264$ (which yielded the worst accuracy of the variance values tried in Table 1). For the algorithm of Liu et al. [1], we set the number of dominant peaks to 40, which achieved the best performance of the values tried. The lattices produced by the proposed method with Gaussian mixture model were ranked first or equal first for 81.5% of the images.

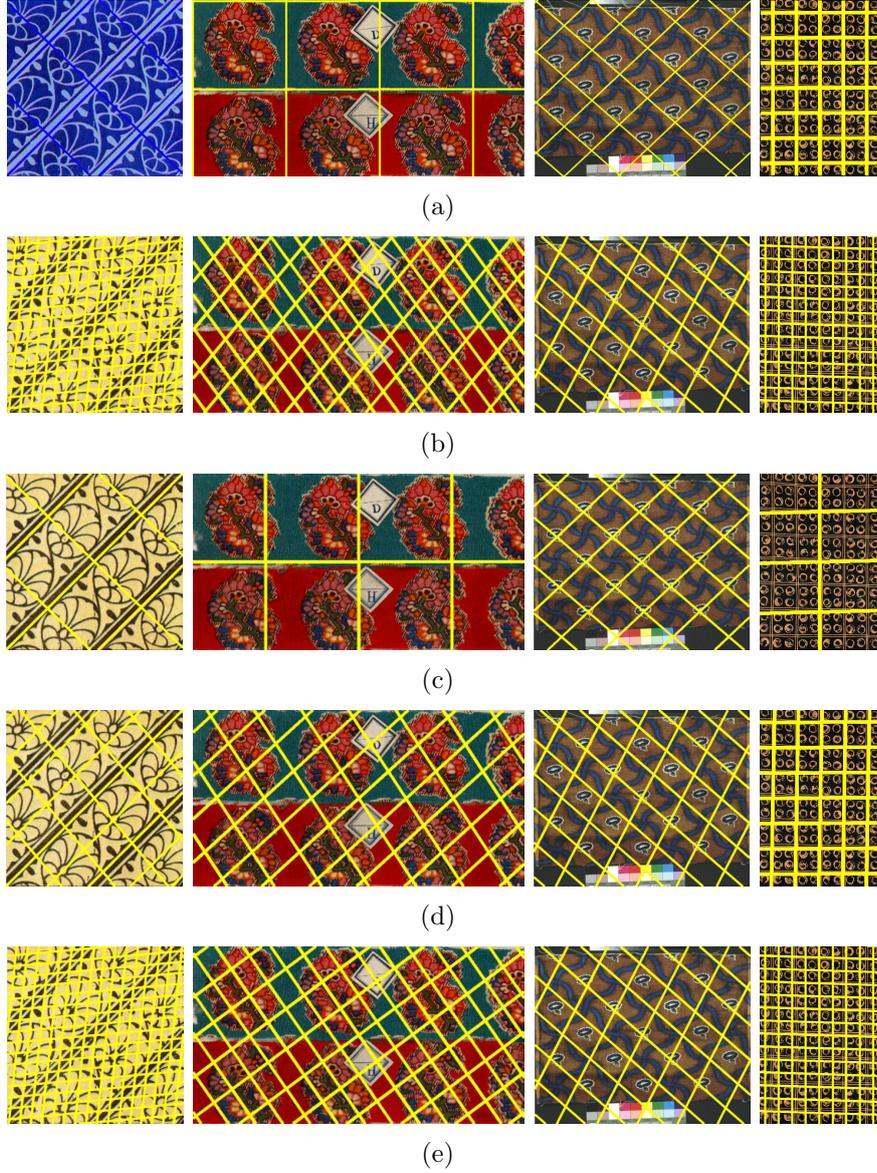


Figure 8: Comparison of texel lattices obtained by different methods for some example images. (a) Proposed method. (b) Lin et al. [14]. (c) Liu et al. [1] (10 peaks). (d) Liu et al. [1] (40 peaks). (e) Liu et al. [1] (70 peaks).

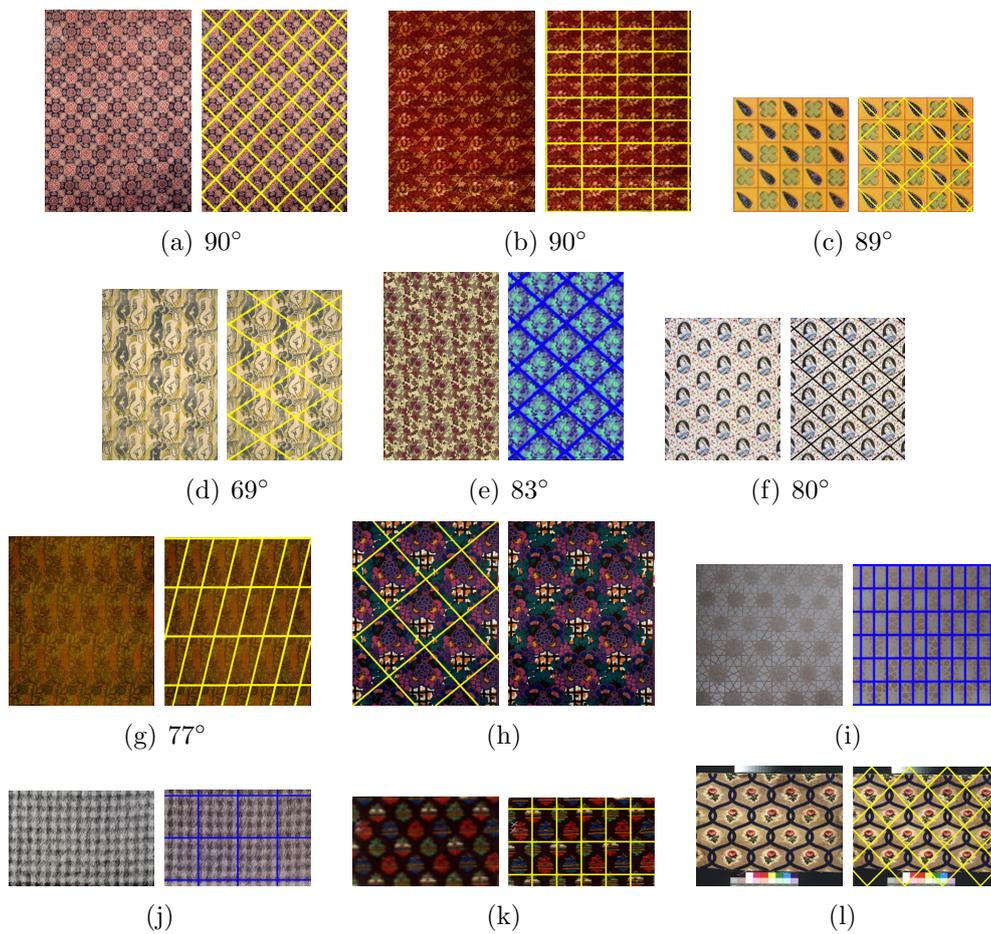


Figure 9: Results from the proposed algorithm. Note that the lattice translations shown here are arbitrary (due to the translational symmetry of the regular texture class being considered). (a)-(g) OC results showing angle between \mathbf{t}_1 and \mathbf{t}_2 . (h) A neutral result. (i)-(l) OI results.

Table 1: Comparison of lattice estimation algorithms on regular texture images. Accuracy is defined as the number of OC results divided by the total number of test images. (OC denotes "obviously correct" and OI denotes "obviously incorrect"). Test results are reported here for 200 regular texture images. Values in parentheses are for the 53 of these images that are in the public domain.

Algorithm variant	# OC	# Neutral	# OI	Accuracy
Gaussian ($\sigma^2 = 100$)	159 (43)	24 (5)	17 (5)	0.795 (0.811)
Gaussian ($\sigma^2 = 144$)	158 (42)	19 (5)	23 (6)	0.790 (0.792)
Gaussian ($\sigma^2 = 196$)	155 (41)	20 (5)	25 (7)	0.775 (0.774)
Gaussian ($\sigma^2 = 264$)	151 (38)	17 (3)	32 (12)	0.755 (0.717)
Gaussian mixture	153 (36)	14 (2)	33 (15)	0.765 (0.679)
AIC Gaussian ($\sigma^2 = 100$)	154 (42)	25 (5)	21 (6)	0.770 (0.792)
Liu et al. [1] (10 peaks)	80 (21)	7 (2)	113 (30)	0.400 (0.396)
Liu et al. [1] (40 peaks)	97 (20)	9 (3)	94 (30)	0.485 (0.377)
Liu et al. [1] (70 peaks)	60 (13)	5 (1)	135 (39)	0.300 (0.245)
Lin et al. [14]	53 (10)	12 (1)	135 (42)	0.265 (0.189)

Table 2: Comparison based on ranking the results of different algorithms. Entries indicate the numbers of images for which each algorithm was ranked first, second, third or fourth.

Algorithm	# Rank 1	# Rank 2	# Rank 3	# Rank 4
Gaussian, $\sigma^2 = 264$	161 (40)	19 (7)	18 (5)	2 (1)
Gaussian mixture	163 (42)	19 (9)	17 (2)	1 (0)
Liu et al. [1] (40 peaks)	100 (24)	10 (3)	56 (22)	34 (4)
Lin et al. [14]	40 (7)	7 (1)	44 (11)	109 (34)

6.2.2. Evaluation by comparison to ground-truth

Two measures were used to perform geometric comparison of estimated texels and ground-truth texels. The first measure was:

$$g = \frac{(\|\hat{\mathbf{t}}_1\| - \|\mathbf{t}_1\|)^2}{\|\mathbf{t}_1\|} + \frac{(\|\hat{\mathbf{t}}_2\| - \|\mathbf{t}_2\|)^2}{\|\mathbf{t}_2\|} + \frac{(\|\hat{\mathbf{t}}_1 + \hat{\mathbf{t}}_2\| - \|\mathbf{t}_1 + \mathbf{t}_2\|)^2}{\|\mathbf{t}_1 + \mathbf{t}_2\|} + \frac{(\|\hat{\mathbf{t}}_1 - \hat{\mathbf{t}}_2\| - \|\mathbf{t}_1 - \mathbf{t}_2\|)^2}{\|\mathbf{t}_1 - \mathbf{t}_2\|} \quad (19)$$

where the ground-truth texel is $(\mathbf{t}_1, \mathbf{t}_2)$ and the estimated texel is $(\hat{\mathbf{t}}_1, \hat{\mathbf{t}}_2)$. This is based on a related measure called the g-score which was used by Liu et al. [32] to quantify texture regularity. We will refer to the measure in Equation (19) as the g-score in this paper. A potential disadvantage of this measure is that it is invariant to rotation in the image plane.

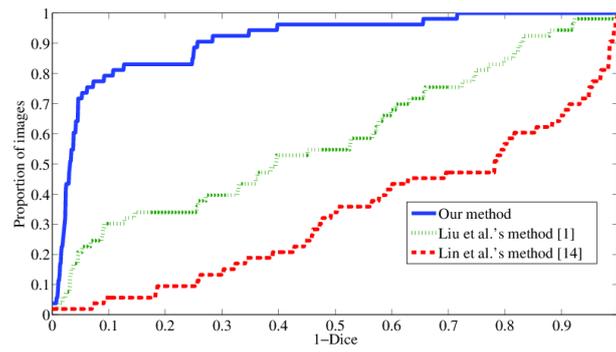
The second was the Dice coefficient [33]. This coefficient was obtained by rendering the estimated and ground-truth parallelograms with their centroids coincident (since their relative translation must be ignored) and then dividing the area of their intersection by their average area. This coefficient is one when the texels are identical and is reduced towards zero by differences in size, shape or rotation of the texels.

Fig. 10 shows the results obtained using the proposed method and the methods of Liu et al. [1] and Lin et al. [14]. Plotted are the proportions of images for which one minus the Dice coefficient (respectively g-score) did not exceed the value on the abscissa.

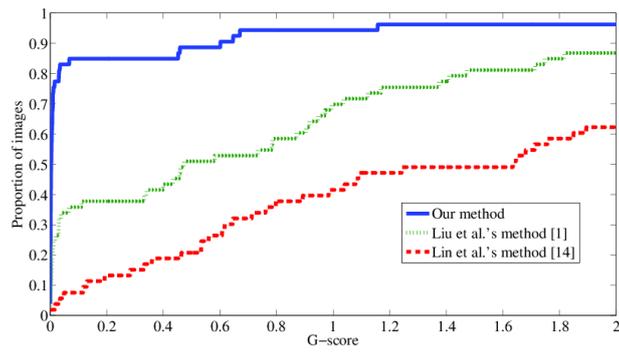
The sensitivity of the proposed method to zero-mean, additive Gaussian noise was investigated by measuring accuracy with 10%, 20%, and 30% noise. Fig. 11 shows the resulting curves; accuracy degrades gracefully as noise increases.

7. Discussion and Conclusions

It seems clear that the method proposed in this paper has superior accuracy to the methods with which it was compared. This was the case on both the commercial and public databases. In particular, the number of lattice extractions not judged to be clearly correct by observers was only a third of that obtained by the closest competing method. Even with parameter settings which disadvantage it relative to the other methods, the proposed



(a)



(b)

Figure 10: Proportions of images for which error did not exceed the value on the abscissa. Error was defined as (a) $1 - \text{Dice}$ coefficient, and (b) g-score.

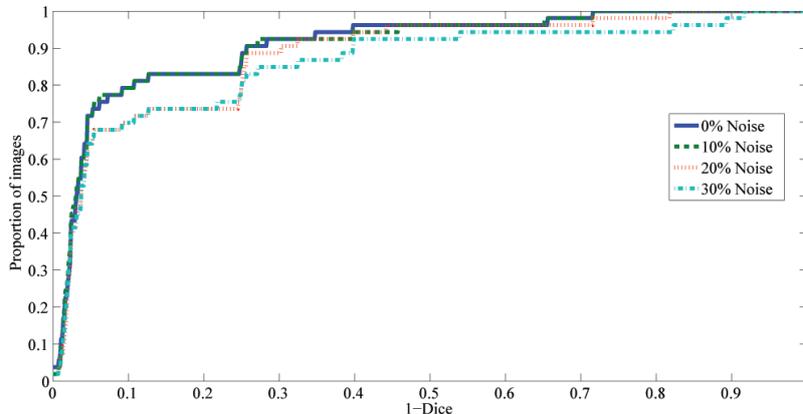


Figure 11: The effect of image noise on the uppermost curve in Fig. 10(a).

method gave markedly better lattice estimates (Table 2). Furthermore, the value of σ^2 had little effect on the results (see Table 1).

Most errors made by BIC were due to overfitting, e.g., over half of those BIC results (Gaussian, $\sigma^2 = 100$) rated OI were due to overfitting. Replacing BIC with AIC exacerbates this problem since AIC penalises model complexity less and thus tends to overfit more. This is why the AIC result in Table 1 (row 6) is a little worse than the corresponding BIC result (row 1).

The method proposed is fully automatic and can be used to both classify images as regularly or irregularly textured and to estimate lattice geometry from those textures that are regular. Adopting the theoretical framework of statistical model selection has yielded accuracy clearly superior to other methods with which it was compared. Nevertheless, the error rates could be usefully lowered further by future work.

The framework can easily embed other existing work. Alternative methods for generating hypotheses could be explored in the context of the approach. Further work is needed to explore the relative merits of alternative statistical texture models. This should enable better performance on images of damaged textiles, for example. The ratio of the two BIC values used in (10) might provide a useful measure of the extent to which a texture is regular (in the sense studied in this paper). Validation of this hypothesis is left for future work. BIC can sometimes agree poorly with the marginal likelihood and it would be worth exploring approximations based on sampling methods, for example. Finally, it should be possible to extend the approach to

handle varied viewing geometry, and to the analysis of near-regular textures on deformed 3D surfaces by allowing relative deformation between texels.

Acknowledgements

The authors are grateful to Ruixuan Wang and Annette Ward for helpful discussions. They would like to thank J. Hays and M. Park for providing their source code. This research was supported by the UK Technology Strategy Board grant “FABRIC: Fashion and Apparel Browsing for Inspirational Content” in collaboration with Liberty Art Fabrics, System Simulation Ltd., and the Victoria & Albert Museum.

References

- [1] Y. Liu, R. T. Collins, Y. Tsin, A computational model for periodic pattern perception based on frieze and wallpaper groups, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26 (3) (2004) 354–371.
- [2] J. Han, S. J. McKenna, R. Wang, Regular texture analysis as statistical model selection, in: *European Conference on Computer Vision*, Marseille, France, 2008, pp. 242–255.
- [3] K. L. Lee, L. H. Chen, A new method for coarse classification of textures and class weight estimation for texture retrieval, *Pattern Recognition and Image Analysis* 12 (4) (2002) 400–410.
- [4] K.-L. Lee, L.-H. Chen, An efficient computation method for the texture browsing descriptor of MPEG-7, *Image and Computing Vision* 23 (2005) 479–489.
- [5] J. Hays, M. Leordeanu, A. Efros, Y. Liu, Discovering texture regularity as a higher-order correspondence problem, in: *European Conference on Computer Vision*, Graz, Austria, 2006, pp. 533–535.
- [6] W. Lin, Y. Liu, A lattice-based MRF model for dynamic near-regular texture tracking, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29 (2007) 777–792.

- [7] T. Leung, J. Malik, Detecting, localizing and grouping repeated scene elements from an image, in: European Conference on Computer Vision, Cambridge, UK, 1996, pp. 546–555.
- [8] T. Tuytelaars, A. Turina, L. Gool, Noncombinational detection of regular repetitions under perspective skew, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25 (2003) 418–432.
- [9] F. Schaffalitzky, A. Zisserman, Geometric grouping of repeated elements within images, in: *Lecture Notes In Computer Science, Shape, Contour and Grouping in Computer Vision*, Springer, 1999, pp. 165–181.
- [10] D. A. Forsyth, Shape from texture without boundaries, in: European Conference on Computer Vision, Copenhagen, Denmark, 2002, pp. 225–239.
- [11] A. Lobay, D. A. Forsyth, Recovering shape and irradiance maps from rich dense texton fields, in: *Computer Vision and Pattern Recognition*, 2004, pp. 400–406.
- [12] M. Park, R. T. Collins, Y. Liu, Deformed lattice discovery via efficient mean-shift belief propagation, in: *European Conference on Computer Vision (ECCV)*, Marseille, France, 2008, pp. 474–485.
- [13] G. Zeng, L. van Gool, Multi-label image segmentation via point-wise repetition, in: *Computer Vision and Pattern Recognition (CVPR)*, 2008, pp. 1–8.
- [14] H. Lin, L. Wang, S. Yang, Extracting periodicity of a regular texture based on autocorrelation functions, *Pattern Recognition Letters* 18 (1997) 433–443.
- [15] D. Chetverikov, Pattern regularity as a visual key, *Image and Vision Computing* 18 (12) (2000) 975–985.
- [16] J. Leu, On indexing the periodicity of image textures, *Image and Vision Computing* 19 (13) (2001) 987–1000.
- [17] D. Charalampidis, Texture synthesis: Textons revisited, *IEEE Transactions on Image Processing* 15 (3) (2006) 777–787.

- [18] V. Starovoitov, S. Y. Jeong, R. Park, Texture periodicity detection: features, properties, and comparisons, *IEEE Trans. on Systems, Man, and Cybernetics-A* 28 (1998) 839–849.
- [19] M. Park, K. Brocklehurst, R. T. Collins, Y. Liu, Deformed lattice detection in real-world images using mean-shift belief propagation, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31 (10) (2009) 1804–1816.
- [20] M. Park, K. Brocklehurst, R. T. Collins, Y. Liu, Translation-symmetry-based perceptual grouping with applications to urban scenes, in: *Asian Conference on Computer Vision (ACCV)*, 2010, pp. 329–342.
- [21] G. Gimelfarb, L. Yu, D. Zhou, Geometric structure and randomness in texture analysis and synthesis, in: *Geometry, Morphology, and Computational Imaging*, Springer, 2003, pp. 137–142.
- [22] F. Wu, C. Zhang, J. He, An evolutionary system for near-regular texture synthesis, *Pattern Recognition* 40 (8) (2007) 2271–2282.
- [23] P. Wu, B. S. Manjunath, S. D. Newsam, H. D. Shin, A texture descriptor for image retrieval and browsing, in: *IEEE Workshop on Content-Based Access of Image and Video Libraries*, 1999, pp. 3–7.
- [24] F. Liu, R. W. Picard, Periodicity, directionality, and randomness: World features for image modeling and retrieval, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18 (7) (1996) 722–733.
- [25] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms*, Cambridge, 2003.
- [26] G. Schwarz, Estimating the dimensions of a model, *Annals and Statistics* 6 (1978) 461–464.
- [27] A. E. Raftery, Bayesian model selection in social research, *Sociological Methodology* 25 (1995) 111–163.
- [28] H. Akaike, A new look at the statistical model identification, *IEEE Transactions on Automatic Control* 19 (6) (1974) 716–723.

- [29] P.-C. Chen, J. H. Hays, S. Lee, M. Park, Y. Liu, A quantitative evaluation of symmetry detection algorithms, Tech. Rep. CMU-RI-TR-07-36, Robotics Institute, Carnegie Mellon University (September 2007).
- [30] M. Park, S. Lee, P. Chen, S. Kashyap, A. A. Butt, Y. Liu, Performance evaluation of state-of-the-art discrete symmetry detection algorithms, in: *Computer Vision and Pattern Recognition Conference (CVPR)*, 2008.
- [31] J. Han, S. J. McKenna, Classifying and comparing regular textures for retrieval using texel geometry, in: *International Conference on Computer Vision Theory and Applications*, Lisboa, Portugal, 2009, pp. 347–354.
- [32] Y. Liu, W.-C. Lin, J. Hays, Near-regular texture analysis and manipulation, in: *ACM SIGGRAPH*, 2004, pp. 368–376.
- [33] L. R. Dice, Measures of the amount of ecologic association between species, *J. Ecol.* 26 (1945) 297–302.