Machine Vision and Applications (2006) 16(6): 364–373
DOI 10.1007/s00138-005-0008-9

ORIGINAL PAPER

Xianghua Xie · Majid Mirmehdi · Barry Thomas

Colour tonality inspection using eigenspace features

Received: 25 April 2005 / Accepted: 26 September 2005 / Published online: 25 November 2005

Abstract In industrial quality inspection of colour texture surfaces, such as ceramic tiles or fabrics, it is important to maintain a consistent colour shade or tonality during production. We present a multidimensional histogram method using a novelty detection scheme to inspect the surfaces. The image noise, introduced by the imaging system, is found mainly to affect the chromatic channels. For colour tonality inspection, the difference between images is very subtle and comparison in the noise dominated chromatic channels is error prone. We perform vector-ordered colour smoothing and extract a localised feature vector at each pixel. The resulting histogram represents an encapsulation of local and global information. Principal component analysis (PCA) is performed on this multidimensional feature space of an automatically selected reference image to obtain reliable colour shade features, which results in a reference eigenspace. Then unseen product images are projected onto this eigenspace and compared for tonality defect detection using histogram comparison. The proposed method is compared and evaluated on a data set with groundtruth.

Keywords Colour tonality · Surface inspection · Image noise analysis · Vector directional processing · Multidimensional histogramming

1 Introduction

Automatic visual inspection for texture and colour abnormalities has application on a variety of flat surfaces e.g. wood, steel, wafer, ceramics and other products, and is highly demanded by industry in order to replace the subjective and repetitive process of manual inspection. For example, in ceramic tile production, chromato-textural properties of the final product can be affected by a variety of external factors that are difficult to control, such as colour pigments, humidity and temperature. Thus, online monitoring and feedback control of the whole production line may be necessary. Aside from inspecting *textural* faults, such as cracks, pin holes, undulations, mis-registration and misprint, inspecting *chromatic* defects in terms of overall visual impression is also a significant production quality factor. The variation of colour characteristics from tile to tile is known as the colour tonality or colour shade problem. Any changes in the colour shade, however subtle, will still become significant once the tiles are placed together. Such assessment of tile surfaces for constant colour tonality is one of the key problems in the manufacturing process; it is also tiresome and difficult when inspection is carried out manually. The problem is compounded when the surface of the object is not just plain-coloured, but textured. In short, colour shade irregularities on plain or textured surfaces are regarded as defects and manufacturers have long sought to automate the identification process.

Numerous studies on tile defect detection are available, such as [1–4]. For example, in [1], Kittler et al. presented a method for detecting random texture tile defects consisting of K-means clustering of the image and perceptual merging of clusters in Luv space into a stack of binary images to represent each final cluster. This was followed by morphological analysis to obtain blobs, on which statistical measures were computed to represent the texture. The overall method was computationally very expensive but worked well.

Investigating colour tonality in particular, Baldrich et al. [5] also segmented the tile image into several stacks using a K-means approach. Then statistical measures were used to represent local and global colour information and segment chromatic and shape characteristics of blobs within each stack. However, this was designed for a specific family of grainy tiles and may not be applicable to other types of randomly textured tiles. In [6], Lumbreras et al. used wavelet transforms to assess different colour channels and various decompositions schemes to find appropriate features in order to sort tiles into perceptually homogeneous classes. The feature vectors were classified to the nearest class by using

X. Xie (✉) · M. Mirmehdi · B. Thomas
Department of Computer Science, University of Bristol, Bristol BS8 1UB, England
E-mail: {xie,majid}@cs.bris.ac.uk

Fisher's linear discriminant function. Similar work has been reported in [7], using wavelet analysis in RGB channels. Penaranda et al. [8] computed the first and second histogram moments of each channel of the RGB colour space as colour and texture descriptors to classify tiles according to visual perception. The most relevant work to us on the consistency of product colour tonality has been by Boukouvalas et al., for example in [9, 10]. In [9], the authors presented spatial and temporal constancy correction of the image illumination on the surfaces of uniform colour and two-colour patterned tiles. Later in [10], they proposed a colour histogram based method to automatically grade colour shade for randomly textured tiles by measuring the difference between the RGB histograms of a reference tile and each newly produced tile. By quantising the pixel values to a small number of bins for each band and employing an ordered binary tree, the 3D histograms were efficiently stored and compared. Several measures were investigated to perform the histogram comparison. Normalised cross correlation (NCC) was found to be the most appropriate one as it gave the most consistent performance and also had a bounded range, which allowed the *a priori* definition of thresholds for colour shade.

Colour histograms have proved their worth as a simple, low level approach in various applications, e.g. [10–12]. They are invariant to translation and rotation, and insensitive to the exact spatial distribution of the colour pixels. These characteristics make them ideal for use in application to colour shade discrimination, irrespective of the texture pattern. The colours on textured (tile) surfaces are usually randomly or pseudo-randomly applied. However, the visual colour impression of the decoration should be consistent from tile to tile. In other words, the amount of ink and the types of inks used for decoration of individual tiles should be very similar in order to produce a consistent colour shade, but the spatial distribution of particular inks is not necessarily fixed from one tile to the next (see Fig. 1). Thus, colour histogram based methods are highly appropriate for colour shade inspection tasks.

Principal component analysis (PCA) [13] has been widely used in data transformation and reconstruction, e.g. image compression, reconstruction and synthesis [14, 15], and also in pattern classification [16, 17]. In [18], the authors derived multivariate texture features from a small local window and applied PCA to generate eigenfilters for a textile defect detection application. Classification was based on the Mahalanobis distance of the filter responses and that of a reference textile image. More recently, Bharati and MacGregor [19] generated multivariate image data from the greylevel image of steel surfaces by shifting the pixel grid in different directions and stacking them together. Then feature vectors were extracted from each pixel location and PCA was applied to find principal components and corresponding eigenvalues. Classification was based on the feature spaces derived from the eigenvalues.

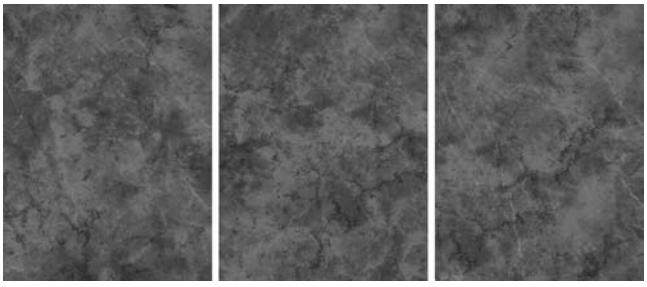

Fig. 1 An example of ceramic tiles with different colour shades. From left: The first two images belong to the same colour shade, the last one is an example of off-shade

In this paper, we present a multidimensional histogram approach to inspect colour tonality defects on randomly textured surfaces. In particular, we are interested in discriminating subtle tonality difference, which conventional methods, such as [9, 10], find difficult to detect. Also, the tonality inspection is treated as a novelty detection problem rather than a classification task, as in [5–8], where both the normal and abnormal samples are pre-defined. However, in reality defects are usually unpredictable. Another motivation is to develop a more accurate approach, particularly for textured designs.

The proposed method combines local colour distribution with global colour distribution to characterise the colour shade properties as part of the histogrammed data. The tiles used were captured by a line-scan camera and manually classified into "standard" and "off-shade" categories by experts. A reference tile image is selected from a small set of good samples using a voting scheme. Initially, a vector directional processing method is used to compute the *local common vector* (LCV) amongst pixels in the RGB space. This is first used to eliminate local noise and smooth the image. Then, a nine element feature vector is computed for each colour pixel in the image composed of the colour pixel itself, its LCV and its local colour variance measured against the local common colour. To minimise the influence of noise, PCA is performed in this 9D feature space. The first few eigenvectors with the largest eigenvalues are selected to form the *reference eigenspace*. The colour features are then projected onto this eigenspace and used to form a multidimensional histogram. By projecting the colour features of an unseen tile onto the same reference eigenspace, a reconstructed image is obtained and histogram distribution comparison can be performed to measure the similarity between the new and the reference tiles.

After presenting our approach in Sect. 2, we demonstrate in Sect. 3 how the reconstructed image shows much less noise in the chromatic channels. Implementation details and comparative results are presented in Sect. 4 and the paper is concluded in Sect. 5.

2 Proposed approach

Normal and abnormal randomly textured colour shades can exhibit only very subtle differences not easily visible to the human eye. Figure 1 shows a particularly difficult example

where the left and centre tiles belong to the same colour shade class and are considered normal samples, while the right tile is an "off-shade" example and should be detected as a defect.

2.1 Noise analysis

The problem of noise, which can be introduced in the imaging system chain and at the printing stage, requires special attention beyond facilitating uniform spatial lighting and temporal consistency during image capture. Noise interference will inevitably enlarge the intra-class variations and make it more difficult to distinguish subtle colour shade differences. The tiles collected in this application were imaged by a 2,048 pixel resolution "Trillium TR-32" RGB colour line-scan camera. The acquired image size varied from 600 pixel × 800 pixel to 1,000 pixel × 1,000 pixel corresponding to the physical size of the tiles.

To examine the nature of the noise in the image acquisition process, we performed PCA directly on the image in RGB colour space. The pixel colours were then projected to the three orthogonal eigenvectors and finally mapped back to the image domain to obtain one image for each eigenchannel. An example of this is shown in Fig. 2 for the leftmost tile in Fig. 1. The first eigenchannel presents the maximum variation in the RGB space, which in most cases is the intensity. The other two orthogonal eigenchannels mainly show the chromatic information. The last eigenchannel is dominated by image noise. The vertical lines are introduced mainly by spatial variation along the line-scan camera's scan line and the horizontal lines are introduced by temporal variations, ambient light leakage and temperature variations.

Clearly, the noise is present in all the image channels, but can dominate in certain chromaticity channels more than others. The noise poses a minor effect on the intensity channel which usually has the largest variation for tile images. Direct comparison in the chromatic channels is likely to be error prone. For colour histogram based methods, each bin has identical weight and the image noise can make the distribution comparison unreliable when colour shade difference is small, as inter-class difference can be demolished. For most tile images, the actual colours only occupy a very limited portion of the RGB space. In other words, the variations in chromaticity are much smaller than those of brightness. However, it can still overwhelm the chromaticity. A variety of smoothing or diffusion methods can be used to explicitly minimise the negative effect of chromatic noise. We found vector directional smoothing [20] to be an effective and robust approach for this purpose. We adopt its underlying principles to compute the LCV, which is later also used as an additional component of our colour feature set to characterise surface shade.

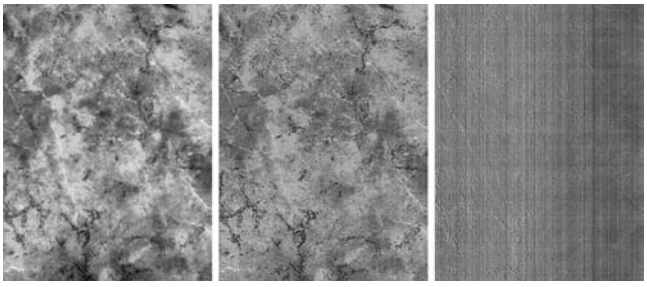

Fig. 2 Image noise analysis showing the three eigenchannels. The noise is highly visible in the third channel. The images have been scaled for visualisation purposes

Other noise can be introduced at the printing stage where subtle temporal inconsistency usually occurs. Further, temporal variations in the imaging system can also contribute to intra-class differences. Thus, more robust colour features, other than RGB alone, are necessary to characterise the colour shade. Vector directional processing can reduce the chromatic noise effect, but on its own it is still not sufficient to distinguish subtle colour shade difference (this will be illustrated in Sect. 4). In this work, we use the LCV as the local salient chromatic feature, along with local statistics and pixel RGB colours themselves as an overall colour shade feature vector. Then, PCA is performed on these feature vectors to obtain reliable chromatic features for colour shade inspection.

Next, we detail the procedure for performing vector directional processing and obtaining the LCV.

2.2 Vector directional median and LCV

Following the work in [20], as shown in Fig. 3, a colour is represented as a vector in the 3D RGB space with the three primaries, R, G and B, defined by the axes. The triangular plane connecting the three primaries in the RGB cube is known as the *Maxwell triangle*. The intersection point, marked as "×" in Fig. 3, of a colour vector with the Maxwell triangle gives an indication of the chromaticity of the colour, i.e. its hue and saturation, in terms of the distance of the point from the vertices of the triangle [21]. As the position of the intersection point only depends on the direction of the

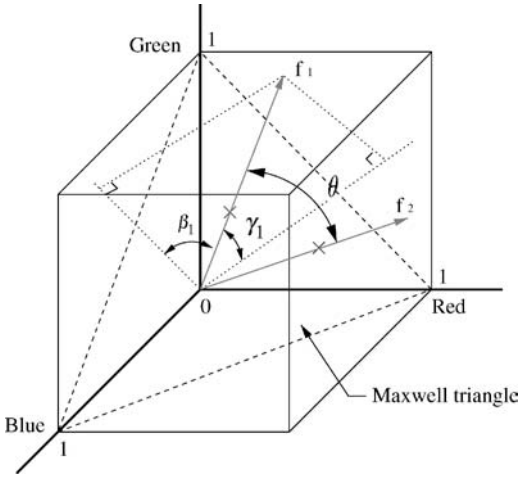

Fig. 3 Perspective representation of the RGB colour cube

colour vector and not the magnitude, the direction then represents the chromaticity. The angle between any two colour vectors, e.g. between f_1 and f_2 in Fig. 3, represents the chromaticity difference between them. So, the directional median of the set of vectors $f_1, f_2, \ldots, f_n$ within a window on the image can be considered as the vector that minimises the sum of the angles with all the other vectors in the set. The median is insensitive to extremes; as the vector direction (chromaticity) determines the colour perception, the noise due to the imaging system can be approximately suppressed using this median vector.

Let $f(x) : \mathcal{R}^2 \rightarrow \mathcal{R}^m$ be the image, a map from a continuous plane to the continuous space $\mathcal{R}^m$. For a colour image, $m = 3$. A window $W \in \mathcal{R}^m$ with a finite number of pixels is implied in calculating the directional median. The pixels in the processing window W are denoted as $\{g_i, i = 1, 2, \ldots, n\}$. The element $f(g_i)$, hereafter referred to as f_i for convenience, is an m-dimensional vector in the space of $\mathcal{R}^m$. Thus, the vectors in W define the input set $\{f_i, i = 1, 2, \ldots, n\}$. Let α_i be the sum of the angles between the vector f_i and each of the vectors in the set. Then,

$$\alpha_i = \sum_{j=1}^{n} \mathcal{A}(f_i, f_j), \quad i = 1, 2, \ldots, n \tag{1}$$

where $0 \leq \mathcal{A}(f_i, f_j) \leq \pi/2$ specifies the angle between vectors f_i and f_j in a colour image. Generally, $0 \leq \mathcal{A}(f_i, f_j) \leq \pi$. In the case of colour images, $0 \leq \mathcal{A}(f_i, f_j) \leq \pi/2$. Then, the ascending order of all the αs gives

$$\alpha_{(1)} \leq \alpha_{(2)} \leq \cdots \leq \alpha_{(k)} \leq \cdots \leq \alpha_{(n)} \tag{2}$$

The corresponding order of the vectors in the set is given by

$$f^{(1)} \leq f^{(2)} \leq \cdots \leq f^{(k)} \leq \cdots \leq f^{(n)} \tag{3}$$

The first term in (3) minimises the sum of the angles with all the other vectors within the set and is considered as the directional median. Meanwhile, the first k terms in (3) constitute a subset of colour vectors which have generally the same direction. In other words, they are similar in chromaticity, but they can be quite different in brightness, i.e. magnitude. However, if they are also similar in brightness, we need to choose the vector closest to $f^{(1)}$. By considering the first k terms $f^{(i)}, i = 1, 2, \ldots, k$ (i.e. the ones with closest directionality to the median), we define a simple measure so that the difference between any pair of vectors in the set is computed as

$$|\Delta\lambda_{(f^{(i)}, f^{(j)})}| + \min(\lambda_{f^{(i)}}, \lambda_{f^{(j)}})\mathcal{A}(f^{(i)}, f^{(j)}) \tag{4}$$

where λ denotes the magnitude of a vector. Thus, when the magnitudes vary significantly in the subset $f^{(i)}, i = 1, 2, \ldots, k$, the first term of (4) will dominate and the vector of median brightness will be selected. On the other hand, if the vectors in the subset have similar brightness, the second term of (4) will help select the LCV vector as the vector that has the least sum of chromaticity differences to other vectors. However, for computational efficiency, we select the LCV from the first k terms as the one that possesses the median brightness attribute with approximately similar accuracy. The value of k was empirically chosen as $\frac{W}{2}$. Alternatively, an adaptive method [22] can be used to determine its value. Thus, the LCV is computed in a running local window to smooth the image. An example of this is shown in Fig. 4. The LCV will also then be used as a component of the colour feature vector applied for shade comparison as detailed in the next section.

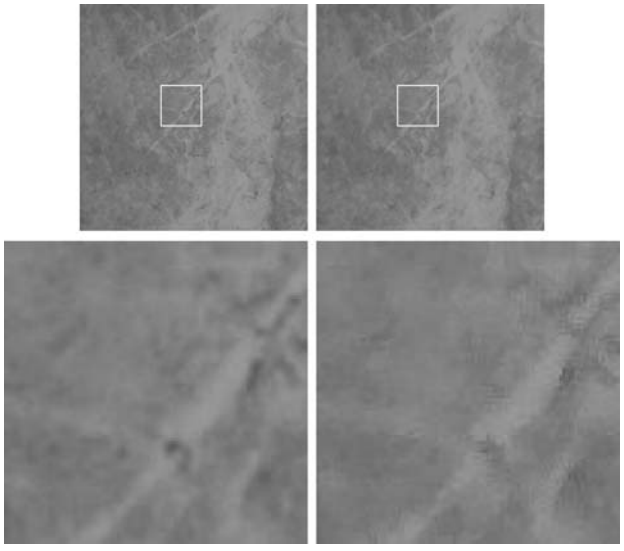

Fig. 4 Vector directional smoothing using a 5×5 window: (*top row*) original tile image and enlarged detail marked in white; (*bottom row*) smoothed image again with the enlarged detail showing the effect of smoothing

2.3 Distribution comparison in eigenspace

Comparing global colour distributions between a reference tile and an unseen tile alone is not always enough, as subtle variations may be absorbed in the colour histograms. The evaluation of local colour distribution becomes a necessity.

2.3.1 Setting up the reference

A reference tile is selected using a simple voting scheme (details in Sect. 4). For any pixel g_i with its colour vector f_i, its brightness is represented by the magnitude λ_i and its direction (chromaticity) is determined by the two angles β_i and γ_i (that it makes with two of the axes in the RGB cube as shown in Fig. 3). Thus, we form a nine element colour feature vector

$$x_i = \left[\lambda_i, \beta_i, \gamma_i, \lambda_i^0, \beta_i^0, \gamma_i^0, \sigma_i^\lambda, \sigma_i^\beta, \sigma_i^\gamma\right]^{\mathrm{T}} \tag{5}$$

comprising the magnitude and directions of the colour pixel itself, its LCV denoted as $(\lambda^0, \beta^0, \gamma^0)$ and the variances

of the local colours in brightness (σ_i^{λ}) and chromaticity ($\sigma_i^{\beta}, \sigma_i^{\gamma}$) measured against the LCV. The basis for this representation is that it is designed to encapsulate the local information at pixel level, which in turn can be used in a global multidimensional histogramming framework for the entire image.

Let $X = \{x_i \in \mathcal{R}^p, i = 1, 2, \ldots, q\}$ be a set of q p-dimensional vectors, which are feature vectors derived from the colour tile image. Let w and h specify the dimensions of the image. Then $q = w \times h$ and p is the dimension of the feature space. Let $\bar{x} = \frac{1}{q}\sum_{x \in X} x$ be the mean vector of X. The feature matrix X is then mean-centred by deducting the mean vector $\bar{x}$. Next, PCA is performed to obtain the eigenvectors (principal axes) denoted by $e_i \in \mathcal{R}^p$.

Singular value decomposition (SVD) [23] can be used to obtain these principal components. The matrix of eigenvectors are then given as $E = [e_1, e_2, \ldots, e_p] \in \mathcal{R}^{p \times p}$. The columns of E are arranged in descending order corresponding to the eigenvalues $\omega_i, i = 1, 2, \ldots, p$. Only j, $j < p$, eigenvectors with the largest eigenvalues are needed to represent X to a sufficient degree of accuracy determined by a simple threshold T, $T = \sum_{i=1}^{j} \omega_i / \sum_{i=1}^{p} \omega_i$, with corresponding eigenvectors, $E_j = [e_1, e_2, \ldots, e_j]$. The threshold T is usually empirically selected so that statistically significant features are retained. Thus, the resulting number of eigenvectors j is different from one type of texture to another. We refer to the subset thresholded with T as the reference eigenspace $\Phi_{\bar{x},E_j}$, where our colour features are well represented and surfaces with the desired shade should have a similar distribution. Characteristics not included in $\Phi_{\bar{x},E_j}$ are small in variation and likely to be redundant noise. Colour feature comparison is now possible to be performed in this eigenspace for unseen tiles. The reference set-up is completed by projecting the original feature matrix X onto the reference eigenspace $\Phi_{\bar{x},E_j}$:

$$X' = \overrightarrow{PCA}(X, \Phi_{\bar{x},E_j}) = E_j^T (X - \bar{x} J_{1,q}) \tag{6}$$

where $J_{1,q}$ is a $1 \times p$ unit matrix consisting of all 1s, resulting in $X' = \{x_i^{'} \in \mathcal{R}^j, i = 1, 2, \ldots, q\}$. Also, we can reconstruct the tile image from this eigenspace through backward projection of a matrix $X' \in \Phi_{\bar{x},E_j}$ onto the original feature space:

$$\hat{X} = \overleftarrow{PCA}(X', \Phi_{\bar{x},E_j}) = E_j X' + \bar{x} J_{1,q} \tag{7}$$

2.3.2 Verifying new surfaces

For a novel tile image, the same feature extraction procedure is performed to obtain the colour feature matrix Y. However, Y is then projected onto the reference eigenspace $\Phi_{\bar{x},E_j}$, resulting in Y', i.e. $Y' = \overrightarrow{PCA}(Y, \Phi_{\bar{x},E_j})$. Note PCA is not performed on Y. This projection provides a mapping of the new tile in the reference eigenspace where defects will be expected to stand out. Finally, multidimensional histogram comparison is performed to measure the similarity between X' and Y' in the reference eigenspace $\Phi_{\bar{x},E_j}$. In [10], Boukouvalas et al. found that for comparing distributions of such kinds the NCC measure performs best as it is bounded in the range $[-1..1]$ and easily finds partitioning which assigns only data with acceptable correlation to the same class. For pairs of quantities (s_i, r_i), $i = 1, 2, \ldots, n$, then,

$$NCC = \frac{\Sigma_i (s_i - \bar{s})(r_i - \bar{r})}{\sqrt{\Sigma_i (s_i - \bar{s})^2} \sqrt{\Sigma_i (r_i - \bar{r})^2}} \tag{8}$$

where $\bar{s}$ and $\bar{r}$ are the respective means. The NCC represents an indication of what residuals are to be expected if the data are fitted to a straight line using least squares. When a correlation is known to be significant, NCC lies in a pre-defined range and the partition threshold is easy to choose. Direct multidimensional histogram comparison is computationally expensive, however for tile images, the data usually only occupies a small portion of the feature space. Thus, only those bins containing data are stored in a binary tree structure. Unlike [10], we found it unnecessary to quantise the histogram, thus preserving further accuracy.

3 Method summary and comments

The proposed method is now summarised with the aid of Fig. 5. The whole procedure comprises two stages: a training stage and a testing stage. The goal of the training stage is to select the best representative of all the training images whose corresponding eigenspace and multidimensional histogram will be treated as the reference data in the testing stage. The inspection starts with the selection of a reference tile using a voting scheme. First, a small number of good samples are each treated as a reference tile and compared with each other. Vector directional processing (as detailed in Sect. 2.2) is used to compute the LCV. Then a nine element feature vector is extracted for each pixel, followed by PCA. Those eigenvectors with the largest eigenvalues (thresholded using T) are used to form the reference eigenspace. Colour features are then projected onto this eigenspace as in (6), which results in a multidimensional histogram for each tile image. NCC is performed to examine the similarity. Thus, for each tile image, we have a similarity measure, quantified by NCC, against each of the other tile images in the training set. The one with the least sum of NCC measures will be treated as the "golden sample". In a practical scenario, this stage would be carried out on the first few good tiles manually selected on the production line. It also allows a threshold on the NCC value to be chosen to classify normal/abnormal colour shades. A simple approach is to set the threshold using the NCC value range for the reference tile in the training stage. Let R_{NCC} denote this range for the reference tile. Then the threshold is selected as $T_{NCC} = 1 - \gamma R_{NCC}$. Here, we empirically choose $\gamma = 1.2$. Alternatively, if a relatively larger number of training samples are available, the threshold can be $T_{NCC} = \mu_{NCC} - \gamma \sigma_{NCC}$ where $(\mu_{NCC}, \sigma_{NCC})$

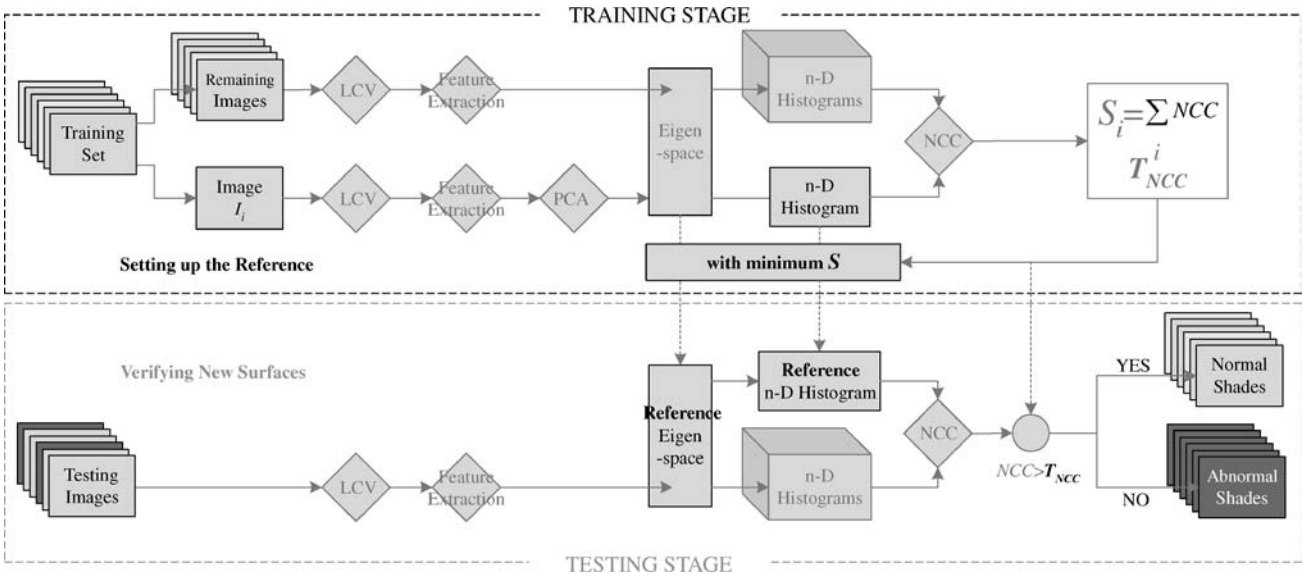

Fig. 5 Flow chart of proposed method

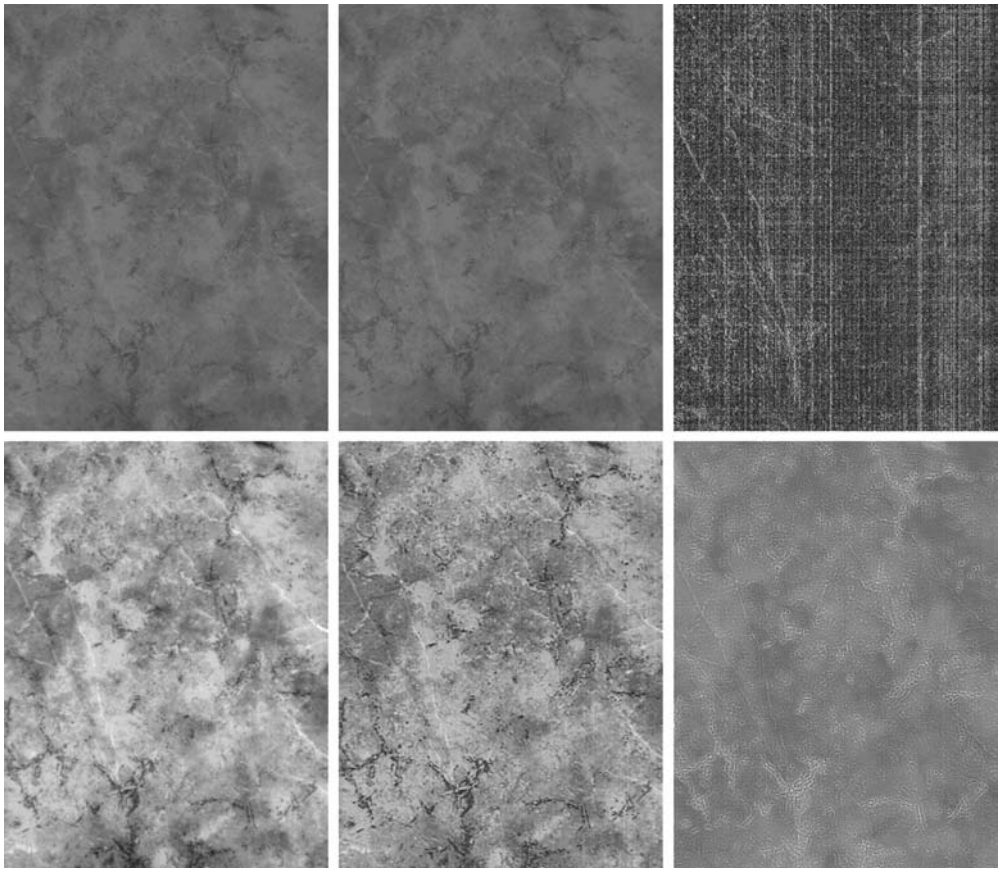

Fig. 6 Image reconstruction: (*top*) the original image, its reconstruction and their MSE difference; (*bottom*) the three eigenchannels of the reconstructed tile. The last channel shows texture structure, instead of being dominated by noise (cf. Fig. 2)

are the mean and standard deviation of the NCC values for the reference tile and $\gamma = 2, 3$ in common practise.

Figure 5 also illustrates the testing stage, in which at first, the LCV and feature vector of an unseen, novel tile image are extracted as in (5). These features are then projected onto the reference eigenspace derived from the "golden sample" using (7). Finally, the NCC measure can be applied as an indication of the similarity between the test image and the reference, which is then used in comparison to T_{NCC} to classify the tile.

For comparison and to demonstrate the function of Eqs. 6 and 7, we can reconstruct the tile image by mapping the colour features in the eigenspace back to the RGB space. Taking the leftmost image in Fig. 1 as the reference image providing X', the reconstructed colour features are $\hat{X} = \overleftarrow{PCA}(X', \Phi_{\bar{x},E_j})$. Then taking the first three feature elements $[\lambda_i, \beta_i, \gamma_i]^{\mathrm{T}}$, converting to RGB representation $[r_i, g_i, b_i]^{\mathrm{T}}$ and mapping back to the image domain gives the reconstructed tile image, as shown in Fig. 6(top-centre) along with the mean square error (MSE) between the original and the reconstructed images, in Fig. 6 (top-right). This clearly shows the noise associated with the less significant eigenvalues. Next, noise analysis in the reconstructed image is performed (as in Sect. 2.1 and in Fig. 2) showing that its third channel is *much less* noisy (bottom row of Fig. 6). The reconstruction shows that reliable, salient colour features are obtained by the proposed method.

Figure 7 shows another example of a different type of tile pattern, which has a little less intensity variation and a bit more chromatic variation than the running example case used so far. The first row shows a good sample that is a reference tile X, the reconstructed image $\hat{X}$ which is projected back from its own (reference) eigenspace $\Phi_{\bar{x},E_j}$ and the reconstruction MSE error map, respectively. The second row shows the three eigenchannels of the original tile image. The textural information of the last eigenchannel is swamped by the image noise. However, the eigenchannels of the reconstructed tile image, shown in the third row, clearly exhibit the textural structure in less noisy detail. The tile image shown in Fig. 8, for example let it be an unseen tile image Y, is in fact an off-shade defect case of the normal shade example in Fig. 7. The reconstruction, shown in the top-centre, is based on the reference eigenspace $\Phi_{\bar{x},E_j}$ derived from the reference tile image in Fig. 7, i.e. $\hat{Y} = \overleftarrow{PCA}(Y', \Phi_{\bar{x},E_j})$. We can see that the eigenchannels of the reconstructed tile image are much less consumed by image noise. However, as we would like, the MSE shows that the reference eigenspace is not an ideal representation of this off-shade tile image as its distribution in the 9D feature space differs from that of the reference tile image. Thus, the NCC measure stands a good chance to discriminate the difference.

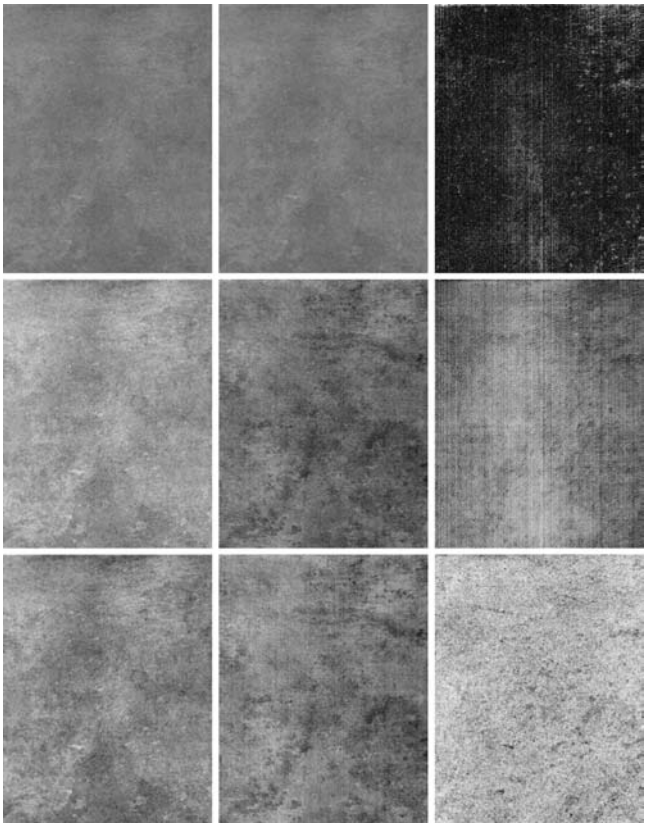

Fig. 7 Another example of tile image: (*top row*) the original tile image, the reconstructed tile image and the difference between them using MSE; (*middle row*) the three eigenchannels of the original tile image; (*bottom row*) the three eigenchannels of the reconstructed tile image

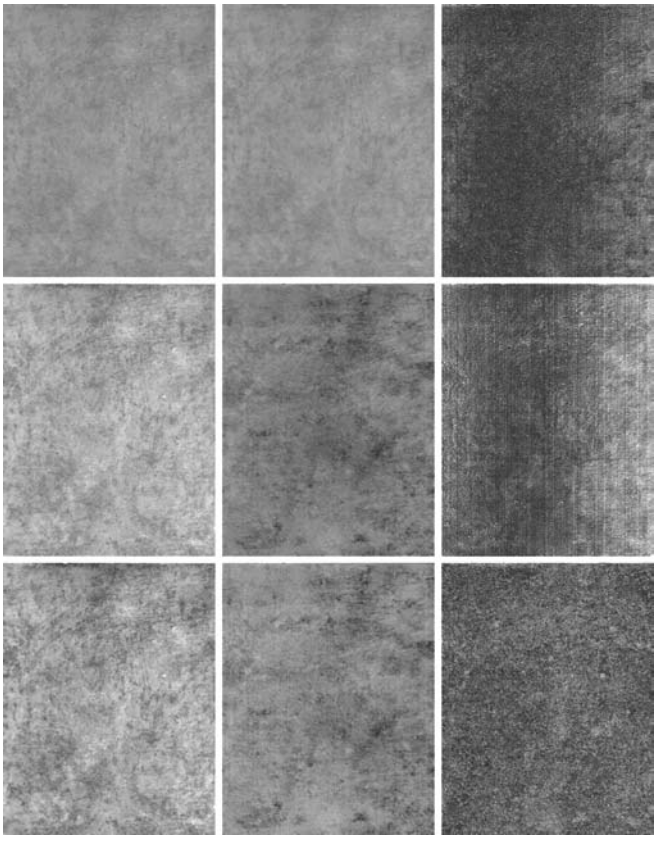

Fig. 8 Reconstruction of an off-shade tile image: (*top row*) the original off-shade tile image, the reconstructed tile image based on an eigenspace derived from a normal shade tile image and the difference between them using MSE; (*middle row*) the three eigenchannels of the original off-shade tile image; (*bottom row*) the three eigenchannels of the reconstructed tile image

4 Implementations and results

Our test data comprises eight tile types, totalling 456 tiles, with known groundtruth obtained from manual classification by experts. Within each set, one-third of tiles are standard colour shade and two-thirds off-shade. We compare the proposed method with a standard histogram-based method, with and without using vector directional preprocessing.

The inspection starts with the selection of a reference tile using a voting scheme amongst the best of the first few tiles in the production run (visually confirmed to be of high quality) as described in Sect. 3. The threshold T_{NCC} is chosen during this process as the criteria for colour shade variation.

Results are quantified using *specificity* to show how accurately good tiles are classified, *sensitivity* to show how accurately defective tiles are classified and *accuracy* as the correct classification rate of all tiles. They are defined as:

$$\begin{cases} spec.(\%) = \frac{P_t \cap P_g}{P_g} \times 100 \\ sens.(\%) = \frac{N_t \cap N_g}{N_g} \times 100 \\ accu.(\%) = \frac{P_t \cap P_g + N_t \cap N_g}{P_g + N_g} \times 100 \end{cases} \quad (9)$$

where P is the number of normal tiles, N is the number of tiles with colour tonality defects and the subscripts "t" and "g" denote the results by testing and groundtruth respectively.

Table 1 Testing results of colour histogram methods (values are percentages)

	RGB colour space		
Tile type	Specificity	Sensitivity	Accuracy
1	93.8	96.9	95.8
2	90.0	95.0	93.3
3	95.0	95.0	95.0
4	90.0	87.5	88.3
5	93.8	84.4	87.5
6	90.0	80.0	83.3
7	90.0	87.5	88.3
8	80.0	72.5	75.0
Overall	90.1	87.2	88.2

We first applied RGB colour histogram based method similar to that described in [10]. The colour images were projected onto the RGB space to produce the 3D histogram. For efficiency, each bin of the histogram was converted into a unique integer so that the 3D histogram can be economically stored in a binary tree. The selection of template samples, the training and testing stages were the same as described in Sect. 2.3. Table 1 shows the results of the RGB colour histogram based method, providing an average accuracy of 88.2%. We then applied the proposed algorithm including the vector directional smoothing, but only using the $(\lambda_i, \beta_i, \gamma_i)$ representation of the colour pixels instead of the full 9D feature vector in (5). Hence, the global histogram comparison is based on the distribution of magnitudes and vector directions only. The results are presented in Table 2. Different window sizes were tested, from 3×3 to 9×9 with the best results achieved using a window size of 5×5 at 91.2% accuracy. Smoothing the colour images proved to be beneficial as it decreased the negative effects introduced by noise in chromaticity. When we used the full 9D feature vector of our proposed method, significantly better tonality defect detection results were obtained as shown in Table 3. By incorporating the local colour information and comparing the dominant colour features using a high-dimensional histogram based method, an overall accuracy of 94.7% is achieved using a 7×7 window. Since the textures in the different tile types are of a various nature in terms of coarseness and density, clearly each may perform better with a particular window size. Here we conclude that on average and in the interest of computational expense at a marginally less accurate rate of 93.9%, a window size of 3×3 can be employed throughout. The results of tile set 8 were consistently lower than those of others for all the three methods. This is due to the nature of the texture printed on these tiles which have slightly larger intra-class variations. Using multiple 'golden samples' may improve the performance.

For practical implementation this technique needs to run at approximately 1 s/tile. Currently, the bottleneck in our system is in the LCV computation. Optimised 3D histogramming using ordered binary trees requires just less than 1 s per 1,000 pixel × 1,000 pixel tile. The proposed method requires a computational time in the order of 20 s/tile at

Table 2 Testing results of vector directional smoothing (values are percentages)

	Window size					
	3×3			5×5	7×7	9×9
Tile type	Specificity	Sensitivity	Accuracy	Accuracy	Accuracy	Accuracy
1	100	96.9	97.9	100	97.9	97.9
2	95.0	95.0	95.0	95.0	95.0	95.0
3	100	97.5	98.3	98.3	96.7	96.7
4	100	80.0	86.7	90.0	90.0	90.0
5	93.8	87.5	89.6	89.6	87.5	87.5
6	90.0	90.0	90.0	91.7	91.7	91.7
7	95.0	92.5	93.3	90.0	90.0	90.0
8	80.0	72.5	75.0	76.7	75.0	76.7
Overall	94.1	88.8	90.6	91.2	90.4	90.6

Table 3 Testing results of comparison in feature eigenspace (values are percentages)

	Window size					
	3 × 3			5 × 5	7 × 7	9 × 9
Tile type	Specificity	Sensitivity	Accuracy	Accuracy	Accuracy	Accuracy
1	100	100	100	100	100	100
2	95.0	97.5	96.7	96.7	96.7	96.7
3	100	95.0	96.7	98.3	100	100
4	95.0	92.5	93.3	93.3	91.7	91.7
5	100	96.9	97.9	97.9	97.9	97.9
6	95.0	92.5	93.3	91.7	93.3	93.3
7	95.0	90.0	91.7	93.3	91.7	91.7
8	85.0	82.5	83.3	83.3	88.3	86.7
Overall	95.4	93.1	93.9	94.1	94.7	94.5

present: 0.98 s for its histogramming, 18 s for LCV computation and smoothing and 0.94 s for NCC computation. This was computed on an AMD Athlon XP Processor (1.4 GHz) with 512 M memory. However, the computation time can be greatly reduced using dedicated hardware and optimised software, e.g. implementing the accelerating coding scheme as mentioned in [22] the computation costs for finding the LCV in a 7 × 7 window can result in a saving of 82% without compromising accuracy.

5 Conclusions

We presented an automatic colour shade defect detection algorithm for randomly textured surfaces. The shade problem is defined here as visual perception in colour, not in texture. We revealed the chromatic noise through eigenchannel analysis and proposed a method to overcome it using local and global colour information and PCA analysis on a new representative colour feature vector. The chromatic channels of the reconstructed image were found to be much less dominated by noise. A window size as small as 3 × 3 gives an overall accuracy of 93.9%. However, the increase in accuracy comes at a computational cost which is hoped will be overcome through more optimised code and faster hardware and memory.

While we present our work with respect to ceramic tiles, the proposed method should be applicable to other flat colour textured surfaces where the tonality issue is of importance, for example, wood and textiles.

Acknowledgements The authors thank Fernando Lopez for the tile database. This work is funded by EC project `G1RD-CT-2002-00783-MONOTONE`, and X. Xie is partly funded by the ORSAS, Universities UK.

References

1. Kittler, J., Marik, R., Mirmehdi, M., Petrou, M., Song, J.: Detection of defects in colour texture surfaces. In: IAPR Machine Vision Applications, pp. 558–567 (1994)
2. Wiltschi, K., Pinz, A., Lindeberg, T.: An automatic assessment scheme for steel quality inspection. Mach. Vision Appl. **12**, 113–128 (2000)
3. Chetverikov, D., Hanbury, A.: Finding defects in texture using regularity and local orientation. Pattern Recognit. **35**(10), 2165–2180 (2002)
4. Kumar, A., Pang, G.: Defect detection in textured materials using optimized filters. IEEE Trans. Syst. Man Cybernet. **32**(5), 553–570 (2002)
5. Baldrich, R., Vanrell, M., Villanueva, J.: Texture and color features for tile classification. In: SPIE Polarization and Color Techniques in Industrial Inspection, vol. 3826, pp. 124–135 (1999)
6. Lumbreras, F., Serrat, J., Baldrich, R., Vanrell, M., Villanueva, J.: Color texture recognition through multiresolution features. In: International Conference on Quality Control by Artificial Vision, vol. 1, pp. 114–121 (2001)
7. Ai, J., Liu, D., Zhu, X.: Combination of wavelet analysis and color applied to automatic color grading of ceramic tiles. In: International Conference on Pattern Recognition, vol. 3, pp. 23–26 (2004)
8. Penaranda, J., Briones, L., Florez, J.: Color machine vision system for process control in the ceramics industry. In: SPIE New Image Processing Techniques and Applications: Algorithms, Methods, and Components, vol. 3101, pp. 182–192 (1997)
9. Boukouvalas, C., Kittler, J., Marik, R., Petrou, M.: Automatic color grading of ceramic tiles using machine vision. IEEE Trans. Ind. Electron. **44**(1), 132–135 (1997)
10. Boukouvalas, C., Kittler, J., Marik, R., Petrou, M.: Color grading of randomly textured ceramic tiles using color histograms. IEEE Trans. Ind. Electron. **46**(1), 219–226 (1999)
11. Swain M and Ballard, D.: Indexing via color histograms. Int. J. Comput. Vision **7**(1), 11–32 (1990)
12. Pietikainen, M., Maenpaa, T., Viertola, J.: Color texture classification with color histograms and local binary patterns. In: International Workshop on Texture Analysis and Synthesis, pp. 109–112 (2002)
13. Jolliffe, I.: Principal Component Analysis. Springer-Verlag, New York (1986)
14. Nishino, K., Sato, Y., IKeuchi, K.: Eigen-texture method: appearance compression and synthesis based on a 3d model. IEEE Trans. Pattern Anal. Mach. Intell. **23**(11), 1257–1255 (2001)
15. Haenselmann, T., Effelsberg, W.: Texture resynthesis using principal component analysis. In: SPIE Human Vision and Electronic Imaging, vol. 4662, pp. 440–447 (2002)
16. Ade, F.: Characterization of texture by 'eigenfilter'. Signal Process. **5**(5), 451–457 (1983)
17. Randen T and, J.: Husϕy. Filtering for texture classification: a comparative study. IEEE Trans. Pattern Anal. Mach. Intell. **21**(4), 291–310 (1999)
18. Ade, F., Lins, N., Unser, M.: Comparison of various filter sets for defect detection in textiles. In: International Conference on Pattern Recognition, vol. 1, pp. 428–431 (1984)

19. Bharati, M., MacGregor, J.: Texture analysis of images using principal component analysis. In: SPIE Process Imaging for Automatic Control, vol. 4188, pp. 27–37 (2001)
20. Trahanias, P. and Venetsanopoulos, A.: Vector directional filters – a new class of multichannel image processing filters. IEEE Trans. Image Process. **2**(4), 528–534 (1993)
21. Pratt, W.: Digital Image Processing. Wiley, New York (1991)
22. Trahanias, P., Karakos, D., Venetsanopoulos, A.: Directional processing of color images: theory and experimental results. IEEE Trans. Image Process. **5**(6), 868–880 (1996)
23. Krzanowski, W.: Principles of Multivariate Analysis: A User's Perspective. Oxford University Press, Oxford (1988)

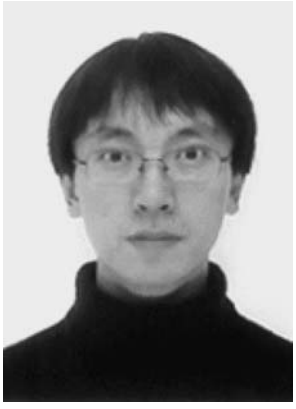

Xianghua Xie is currently a Ph.D. student and a research assistant in the Department of Computer Science, University of Bristol, U.K. Prior to this, he received an M.Sc. degree in advanced computing with commendation from the University of Bristol in 2002 and a B.Sc. degree in environmental engineering from the Tongji University, Shanghai, P.R. China, in 2000. His current research interests are texture analysis, image segmentation, surface inspection, deformable models and historical document analysis. He is a student member of the BMVA, the IEE and the IEEE.

Majid Mirmehdi received the B.Sc. (Hons.) and Ph.D. degrees in computer science in 1985 and 1991 respectively, from the City University, London. He has worked both in industry and in academia. He is currently a Reader in the Department of Computer Science at the University of Bristol, UK. His research interests include texture analysis, colour image analysis, medical imaging and document recognition. He has over 100 refereed conference and journal publications in these areas. He is an associate editor of the *Pattern Analysis and Applications Journal*. He is a member of the IEE, IEEE and a member and the Chairman of the British Machine Vision Association.