

Inspecting Colour Tonality on Textured Surfaces

Xianghua Xie, Majid Mirmehdi, and Barry Thomas

Department of Computer Science, University of Bristol, Bristol BS8 1UB, England
{xie,majid,barry}@cs.bris.ac.uk

Abstract. We present a multidimensional histogram method to inspect tonality on colour textured surfaces, e.g. ceramic tiles. Comparison in the noise dominated chromatic channels is error prone. We perform vector-ordered colour smoothing and generate a PCA-based reconstruction of a query tile based on a reference tile eigenspace. Histograms of local feature vectors are then compared for tonality defect detection. The proposed method is compared and evaluated on a data set with groundtruth.

1 Introduction

The assessment of product surfaces for constant tonality is an important part of industrial quality inspection, particularly in uniform colour surfaces such as ceramic tiles and fabrics. For example with tiles, any changes in the colour shade, however subtle, will still become significant once the tiles are placed on a bathroom wall. This is a key problem in the manufacturing process, and quite 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.

Colour histograms have proved their worth as a simple, low level approach in various applications, e.g. [1,2,3]. 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. 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.

Numerous studies on tile defect detection are available, such as [4,5]. The only colour grading work known to us has been reported by Boukouvalas et al., for example in [6,2]. In the former work, 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 [2], 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.

In this paper, we present a multidimensional histogram approach to inspect colour tonality defects on randomly textured surfaces. The method combines local colour distribution with global colour distribution by computing the local common colours and local colour variations 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* 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 local common vector, and its local colour variance. To minimise the influence of noise, principal component analysis 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 into this eigenspace and used to form a multidimensional histogram. By projecting the colour features of an unseen tile into 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. We also demonstrate that the reconstructed image shows much less noise in the chromatic channels. Finally, we present our comparative results.

In Section 2, our proposed method is introduced, outlining noise analysis, local common vector smoothing, feature generation, eigenspace analysis, and histogram comparison using the linear correlation coefficient. Implementation details and results are shown in Section 3. Section 4 concludes the paper.

2 Proposed Approach

The difference between normal and abnormal randomly textured colour shades can be very subtle. Fig. 1 shows a particularly difficult example where the left and centre tiles belong to the same colour shade class and considered as normal samples, while the right one is an example of “off-shade” and should be detected as a defect.

2.1 Noise Analysis

While effort can be put into achieving uniform spatial lighting and temporal consistency during image capture, some attention must still be paid to the problem of image noise introduced in the imaging system chain. In this application, the tiles were imaged by a 2048 pixel resolution ‘Trillium TR-32’ RGB colour line-scan camera. The acquired image size varied from 600×800 pixels to 1000×1000 pixels corresponding to the physical size of the tiles.

To examine the noise, we performed Principal Component Analysis (PCA) directly on the RGB image. The pixel colours were then projected to the three




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.

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 is in most cases 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.




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.

Clearly, the noise can dominate in certain chromaticity channels, but 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 will make the distribution comparison unreliable when colour shade difference is small. 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, although the variation of

the image noise is small, 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 [7] to be an effective and robust approach for this purpose. We adopt its underlying principles to compute the Local Common Vector (LCV), which is later also used as an additional component of our colour feature set to characterise surface shade.

2.2 Vector Directional Median and LCV

Following the work in [7], a colour is represented as a vector in the 3D RGB space. The triangular plane connecting the three primaries in the RGB cube is known as the *Maxwell triangle*. The intersection point 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. As the position of the intersection point only depends on the direction of the colour vector, and not the magnitude, this direction then represents the chromaticity. The angle between any two colour vectors, e.g. f_1 and f_2 , represents the chromaticity difference between them. So, the directional median of the set of vectors f_1, \dots, 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 W are denoted as $\{g_i, i = 1, 2, \dots, n\}$. The element $f(g_i)$, denoted 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 $\{f_i, i = 1, 2, \dots, n\}$. Let α_i be the sum of the angles between the vector f_i and each of the vectors in the set. Then,

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

where $0 \leq \mathcal{A}(f_i, f_j) \leq \pi/2$ denotes the angle between vectors f_i and f_j in a colour image. Then, the ascending order of all the α s gives

$$\alpha_{(1)} \leq \alpha_{(2)} \leq \dots \leq \alpha_{(k)} \leq \dots \alpha_{(n)} \quad (2)$$

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

$$f^{(1)} \leq f^{(2)} \leq \dots \leq f^{(k)} \leq \dots f^{(n)}. \quad (3)$$

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 of (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 $f^{(i)}, i = 1, 2, \dots, k$, we define a new simple metric so that the difference between any pair of vectors in the set is measured as

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

where λ denotes the magnitude of a vector. Thus, the vector that has the least sum of differences to other vectors is considered as the LCV. 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, as described in [8], can be used to select the value of k . Thus the LCV is computed in a running local window to smooth the image. The LCV will also then be used as a component of the colour feature vector applied for shade comparison.

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.

Setting up the Reference - A reference tile is selected using a simple voting scheme (details in Section 3). 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). Thus, we form a nine-element colour feature vector $F_i = [\lambda_i, \beta_i, \gamma_i, \lambda_i^0, \beta_i^0, \gamma_i^0, \sigma_i^\lambda, \sigma_i^\beta, \sigma_i^\gamma]$, comprising 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.

Let w and h denote the dimensions of the colour tile image, and X be a mean-centred $p \times q$ matrix containing the colour features, where $p = w \times h$ and $q = 9$. Then, PCA is performed to obtain the eigenvectors (principal axes) denoted by $e_i \in \mathcal{R}^q$. The matrix of eigenvectors are given as $E = [e_1, \dots, e_q] \in \mathcal{R}^{q \times q}$. The columns of E are arranged in descending order corresponding to the eigenvalues ω_i . Only $j, j < q$, eigenvectors with large eigenvalues are needed to represent X to a sufficient degree of accuracy determined by a simple threshold T :

$$T = \frac{\sum_{i=1}^j \omega_i}{\sum_{i=1}^q \omega_i}. \quad (5)$$

We refer to the subset thresholded with T as the *reference eigenspace* E' , where our colour features are well represented and surfaces with the desired shade should have a similar distribution. Characteristics not included in E' are small in variation and likely to be redundant noise. Colour feature comparison is then performed in this eigenspace for unseen tiles. The reference setup is completed by projecting the original feature matrix X into eigenspace E' , resulting in X' .

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 into the reference eigenspace E' , resulting in Y' . Note PCA is




Fig. 3. Image reconstruction - top: The original image, the reconstructed image, 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).

not performed on Y . This projection provides a mapping of the new tile in the reference eigenspace where defects will be made to stand out. Finally, 9D histogram comparison is performed to measure the similarity between X' and Y' . In [2], Boukouvalas et al. found that for comparing distributions of such kinds the Normalised Cross Correlation (NCC) 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 (x_i, y_i) , $i = 1, \dots, n$,

$$NCC = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2} \sqrt{\sum_i (y_i - \bar{y})^2}} \quad (6)$$

where \bar{x} and \bar{y} 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 [2], we found it unnecessary to quantise the histogram.

For comparison, 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

$\hat{X} = X'E'^T$. Then taking the first three elements, adding the deducted means and mapping back to the image domain gave the reconstructed tile image, as shown in Figure 3 along with the Mean Square Error (MSE) between the original and the reconstructed images. Next, noise analysis in the reconstructed image was performed (as in Section 2.1 and illustrated in Fig. 2) showing that its 3rd channel is *much less* noisy (bottom row of Figure 3).

3 Implementation and Results

Our test data comprises eight tile sets, totalling 345 tile images, with known groundtruth obtained from manual classification by experts. Some sets contain significant off-shade problems, while other sets have only very subtle differences. Within each set, one-third of tiles are standard colour shade and two-thirds off-shade. We use this data to evaluate the proposed method and compare it with a 3D colour histogram-based approach.

Inspection starts with the selection of a reference tile using a voting scheme. First, a small number of good tiles are each treated as a reference and compared against each other. Each time the one with the largest NCC value is selected as the most similar one to the reference, and only its score is accumulated. Finally, the tile with the largest score becomes the reference. The NCC threshold is chosen during this process as the limit of acceptable colour shade variation.

Table 1. Comparative results for standard 3D histogramming and proposed method.

Test Num.	Colour histogram on 3D RGB Space			Proposed method				
	spe.	sen.	acc.	W: 3 × 3			5 × 5	7 × 7
				spe.	sen.	acc.	acc.	acc.
1	93.75	96.88	95.83	100	100	100	100	100
2	90.00	95.00	93.33	95.00	97.50	96.67	96.67	96.67
3	95.00	95.00	95.00	100	95.00	96.67	98.33	100
4	90.00	87.50	88.33	95.00	92.50	93.33	93.33	91.67
5	93.75	84.38	87.50	100	96.88	97.92	97.92	97.92
6	90.00	80.00	83.33	95.00	92.50	93.33	91.67	93.33
7	90.00	87.50	88.33	95.00	90.00	91.67	93.33	91.67
8	80.00	72.50	75.00	85.00	82.50	83.33	83.33	88.33
Overall	90.13	87.17	88.16	95.39	93.09	93.86	94.08	94.74

Table 1 shows the test results. For standard 3D colour histogramming in RGB the overall average accuracy was 88.16% (41 tiles misclassified). Specificity and sensitivity results are also shown. The processing (including NCC) requires about 1 second per 1000×1000 pixel image. The last five columns of Table 1 present the results of the proposed method for different W , including specificity and sensitivity results for the 3×3 version. The LCV computation proved to be beneficial as it decreased the negative effects introduced by noise in chromaticity. By incorporating the local colour information and comparing the dominant

colour features using a high dimensional histogram, an overall 93.86% accuracy was achieved (21 tiles misclassified). Different window sizes were tested, from 3×3 to 11×11 (not all shown), with the best results at 94.74% for 7×7 (18 tiles misclassified) at somewhat greater computational cost. Marked improvements in the specificity and sensitivity results were also observed.

For practical implementation this technique needs to run at approximately 1 second/tile. Currently, the bottleneck in our system is in the LCV computation. The proposed method requires a computational time in the order of 20 seconds per tile at present: 0.98 seconds for its 9D histogramming, 18 seconds for LCV computation and smoothing, and 0.94 seconds for *NCC* computation.

4 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 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 93.86%. 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.

Acknowledgments. 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 ORS, UK.

References

1. Swain, M., Ballard, D.: Indexing via color histograms. *IJCV* **7** (1990) 11–32
2. Boukouvalas, C., Kittler, J., Marik, R., Petrou, M.: Color grading of randomly textured ceramic tiles using color histograms. *IEEE T-IE* **46** (1999) 219–226
3. Pietikainen, M., Maenpaa, T., Viertola, J.: Color texture classification with color histograms and local binary patterns. In: *IWTAS*. (2002) 109–112
4. Kittler, J., Marik, R., Mirmehdi, M., Petrou, M., Song, J.: Detection of defects in colour texture surfaces. In: *IAPR Machine Vision Applications*. (1994) 558–567
5. Penaranda, J., Briones, L., Florez, J.: Color machine vision system for process control in the ceramics industry. *SPIE* **3101** (1997) 182–192
6. Boukouvalas, C., Kittler, J., Marik, R., Petrou, M.: Automatic color grading of ceramic tiles using machine vision. *IEEE T-IE* **44(1)** (1997) 132–135
7. Trahanias, P., Venetsanopoulos, A.: Vector directional filters - a new class of multichannel image processing filters. *IEEE T-IP* **2** (1993) 528–534
8. Trahanias, P., Karakos, D., Venetsanopoulos, A.: Directional processing of color images: theory and experimental results. *IEEE T-IP* **5** (1996) 868–880