

Amplification of Textural Differences

D M Booth*, J E W Mayhew † and M K Pidcock ‡

The Research Initiative in Pattern Recognition,
DRA, RSRE, St. Andrews Road, Malvern, Worcs. WR14 3PS.

Abstract

The coefficients of a weighted summation of red, green and blue intensities are shown to effect textural characteristics. Thus multi-spectral images can be mutated into forms that facilitate their discrimination by a given set of statistical texture measures.

1 Introduction

The segmentation of colour imagery is often seen as an extrapolation of monochrome techniques, with the primary images being manipulated independently, and the results being combined in some way. However, success has been limited by the correlations between red (R), green (G) and blue (B) images [1, 2], since by implication, their respective segmentations are also highly correlated. The effectiveness of multi-spectral texture features [3] is similarly curtailed. As a result, the decorrelation of multi-spectral image data (using principal component analysis) has received much attention [2, 3]. Adopting a different approach, many road tracking systems employ labeled training samples to derive a single composite colour feature by classical discriminant analysis [4]. As a logical extension, multi-spectral texture discrimination is introduced by computing various texture statistics over each of the primary images, and then determining the Fisher linear discriminant [5]. However, there is a sense in which the colour and texture features are not co-operating.

2 Overview

That textural characteristics can depend on the weighting assigned to each primary image is demonstrated by examining the principal component images. The eigenvectors are computed by diagonalising the correlation matrix of a feature vector, each row of which contains the red, green and blue responses at a particular pixel location. The pixels in RGB space are projected onto each of the principal axes in turn, and the resulting images are histogram equalised to avoid representational problems. (Histogram equalisation can also remove the effects of unequal overall brightness and contrast [6].) Figures 1 and 2 show two multi-spectral (tweed-like)

*Seconded by DRA Electronics Division, RSRE, Malvern.

†AI Vision Research Unit, University of Sheffield, Sheffield, Yorks.

‡Dept. of Computing and Maths, Oxford Polytechnic, Headington, Oxford.




Figure 1. Red, green and blue images (histogram equalised to aid comparison)




Figure 2. Principal component images (PC1, PC2 and PC3) (256x256 pixels)

textures captured by a CCD camera in conjunction with a set of gelatin filters, and their principal components respectively.

The extent to which a pair of textures can be discriminated by a particular feature set is determined by the separation of their respective probability density functions. Here the feature set consists of a grey level mean statistic and five second order grey level difference statistics [6] (which are computed over 32 pixel square image blocks). A total of 53 training samples (originating from a preliminary segmentation) were made available, 27 from one class and 26 from the other. Thus, the number of features that can be employed at any one time is limited to five by the "curse of dimensionality" [7]. The performance of a parametric (multi-variate normal) Bayes classifier is predicted by the B-distance [8], which relates to an upper bound on error probability [8]. The B-distances for Figures 1 and 2 are : red = 11.7; green = 11.3; blue = 11.7; PC1 = 6.4; PC2 = 3.7; and PC3 = 26.0. So, in this case, the third principal component is much the easiest to segment.

More desirable than using principal component analysis would be a method of deriving a transformation that, given a particular set of texture measures, will find the best RGB transformation for classifying two multi-spectral textures. If mean grey level were the only available texture feature, then the problem could be solved using the Fisher linear discriminant. However, that axis is not optimal when other texture features are employed. The optimisation procedure is complicated further if it is also necessary to find the most discriminating texture features. This is because as the RGB transformation changes, the textures change, and consequently the features that best discriminate between the two textures may also change. The use

of a suboptimal feature selection algorithm should be avoided in order to prevent the introduction of local minima. An exhaustive search is employed here.

The cost function applies the colour transformation $C = u_0R + u_1G + u_2B$ where u_0 , u_1 and u_2 (the normalised red, green and blue coefficients) are the optimisation parameters. The resulting image is histogram equalised, the most discriminatory feature set determined, and the corresponding B-distance (i.e. the cost) is returned. Despite the nonconvexity of the cost function, optimisation is by a direct descent algorithm[9]. This was favoured over more sophisticated approaches, such as simulated annealing, for computational reasons. Thus the optimal solution is sacrificed.

3 Results

The optimisation results (for Figure 1) appear in Table 1. The final column indicates which five of the six available features provided the best discrimination. Texture feature 0 corresponds to the grey level mean, which, as well as being a useful texture measure, also provides a means of discriminating nontextured regions on the basis of colour. Features 1 to 5 refer to grey level difference statistics : angular second moment (1), contrast (2), entropy (3), mean (4) and inverse difference moment (5). Significant increases in B-distance were achieved from all weight starts (including ten random starts not shown). Figures 3a and b show the natural contrast monochrome image (B/W) and the best of the transformed images respectively.

Table 1. Optimisation results

starting point					minimum				
type	W_R	W_G	W_B	B	W_R	W_G	W_B	B	Features
PC1	0.333	0.365	0.302	6.38	-0.386	0.563	-0.050	34.18	0,1,2,4,5
PC2	-0.390	-0.081	0.529	3.68	-0.393	0.529	-0.078	47.11	0,1,2,4,5
PC3	0.347	-0.469	0.184	25.98	0.427	-0.417	0.156	38.81	0,1,2,4,5
rnd	0.259	0.543	0.197	6.12	-0.391	0.559	-0.050	54.28	0,1,2,4,5
rnd	-0.372	0.384	0.244	13.58	-0.581	0.255	0.163	21.11	1,2,3,4,5
B/W	0.300	0.110	0.590	8.314	-0.387	0.553	0.060	33.64	0,1,2,3,5

4 Conclusions

The use of multispectral imagery has resulted in a considerable improvement in discrimination. It will be shown that in this example, it is sufficient for the size of the texture measurement windows to be halved (16x16), which pays-off in the form of improved boundary localisation. However, whether or not a particular transformation will remain effective at higher resolutions depends on several factors, most importantly, on the sizes of the texture primitives in comparison with the size of the measurement window. Rankings are therefore prone to change, e.g. B-distances at the 16 pixel resolution include : PC3 = 8.03; B/W = 4.1; and the previous best transformation = 7.9. Figure 4a shows the segmentation of a 64 by 64 pixel square constructed from the textures in Figure 3a. Similarly, Figure 4b is derived from the textures in 3b but segmented using half sized (16x16) measurement windows.




Fig 3(a) B/W $0.3R+0.11G+0.59B$
 (b) Best trans. $-0.39R+0.56G-0.05B$




Figure 4. Image segmentations
 a) B/W 32×32 windows
 b) Best 16×16 window

The colour-texture fusion mechanism described here suggests that once the RGB transformation has been determined, the texture statistics need be computed only over a single composite image, and therefore, the number of training samples, memory requirement, and processing time necessary to segment the image are no more than would be required of monochrome techniques. In future the intention is to compare the segmentation results achieved using this approach with those of other techniques, such as multi-spectral texture features. Also requiring some attention is the sensitivity of some transformed images to small variances in the original textures e.g. the left hand texture of the second principal component. This has similarities with observations made by Lowitz [10] regarding pronounced statistical noise on some principal images. He proposed a measure of metric content.

References

- [1] Nevatia, R. "A colour edge segmenter and its use in scene segmentation," *IEEE Trans, SMC-7*, No.11., 1977.
- [2] Ohta, Y. *Knowledge-based Interpretation of Outdoor Natural Color Scenes*, Pitman, London, 1985.
- [3] Rosenfeld, A. et al. "Multispectral texture," *IEEE Trans, SMC-12*, no.1, 1982.
- [4] Kuan, D., Phipps, G. and Hsueh, A.C. "Autonomous robot road following," *IEEE Trans. PAMI-10*, No.5, 1988.
- [5] Thorpe, C., Hebert, M., Kanade, T. and Shafer, S. "Vision and navigation for the Carnegie-Mellon Navlab," *Ann. Rev. Comput. Sci.*, 7, pp.521-556, 1987.
- [6] Weszka, J.S. et al. "A comparative study of texture measures for terrain classification," *IEEE Trans. SMC-6*, No.4, 1976.
- [7] Foley, D.H. "Considerations of sample and feature size," *IEEE Trans. vol. IT-18*, pp.618-626, 1972.
- [8] Fukunaga, K. *Introduction to statistical pattern recognition*, Academic Press, Inc. (London) Ltd., 1972.
- [9] Blake, A. and Zisserman, A. *Visual reconstruction*, The MIT Press, London, England, 1987.
- [10] Lowitz, G.E. "Stability and dimensionality of Karhunen-Loeve multispectral image expansions," *Proc. 3rd IJ CPR*, pp.673-677, 1976.