Colour Science on Computer Vision and Image Processing

Organised by the British Machine Vision Association

Friday, the 28th November 2003
Royal Statistical Society, 12 Errol Street, London

10:00	Registration and Coffee	

Morning Session
Chaired by Dr. Mitch Rosen (Rochester Institute of Technology, USA)

10:30	Efficient representation of spectral images	Dr. Stephen Westland (U. of Derby)
11:00	Spectral compression and decompression	Dr. Changjun Li (U. of Derby)
11:30	On the advantage of multi-scale filters for illumination compensation	Dr. Qiang Zhou (UMIST)
12:00	Region based stereo matching through bounded irregular pyramids	Ms. Rebeca Marfil, *et al* (U. of Malaga)
12:30	Lunch	

Afternoon Session
Chaired by Professor Ronnier Luo (University of Derby)

13:30	Simplifying gamut mapping colour constancy	Professor Graham Finlayson (U. of East Anglia)
14:00	A combined physical and statistical approach to computer colour constancy	Dr. Gerald Schafer, *et al* (Nottingham Trent U.)
14:30	Colour displays for categorical images	Dr. Chris Glasbey, *et al* (BIOSS)
15:00	Tea Break	
15:20	Evaluation of colour management systems to achieve colour communication fidelity	Dr. Xiaohong Gao (U. of Middlesex)
15:50	Computing image difference in colour and spatial domain	Dr. Guowei Hong (U. of Derby)
16:20	Close	

Efficient Representation of Spectral Images

Stephen Westland
Colour & Imaging Institute, University of Derby

Trichromatic images represent the colour information at each pixel in the image by three values such as Red, Green and Blue responses. A spectral image is one where the colour information at each pixel location is represented by a vector of reflectance or radiance values at each of many wavelength intervals in the visible spectrum. Typically, 31 values are required for each pixel to record the spectral information at wavelength bands centred at 400, 410, 420, …, 700 nm. It is known that certain properties of the spectral vectors of real scenes can be exploited to allow substantially less than 31 values to be used whilst preserving most of the useful information. This paper will present the results of psychophysical and computational studies that have been carried out in the Colour & Imaging Institute to explore the extent to which such efficient spectral representations may be used to store and represent spectral images. Two important and sometimes conflicting aims have been considered: the preservation of the trichromatic information of the original image and the preservation of the spectral information. A basis for an efficient spectral file format will be proposed that perfectly preserves the trichromatic information of the original spectral image (under a primary illuminant) whilst also preserving most (or all) of the spectral information.

Spectra Compression and Decompression

Changjun Li and M. Ronnier Luo
*Colour & Imaging Ins*titute, University of Derby

Spectral imaging is gaining attention during in recent years. A spectral image contains huge information since each pixel of it is represented by a reflectance vector in a series of narrow wavebands throughout the visible spectrum. However, it is not limited to the human visual range of the spectrum. It can include the data in IR or UV regions for a wide range of imaging applications. Spectral image has the advantage of accurately reproducing the image under any illuminant environment over the traditional RGB colour image. An effective image compression method would be extremely useful to reduce the image size for database storage and retrieval and data communication.

Several new compression methods have been proposed recently for multi-spectral images. They can be divided into spatial and spectral compressions. Only the latter will be considered in this talk. We first mainly review some existed methods such as Vector Subspace Method, Wavelet Method, CIE co-ordinates. Advantages and disadvantages of them are then given. Finally we propose some new approaches such as multi-illuminants, multi-coordinates plus fix points, spectral reflectance abridgement method.

On the Advantage of Multi-scale Filters for illumination compensation

Qiang Zhou and John P. Oakley
Dept. of Electrical Engineering & Electronics, UMIST

A multi-scale Gaussian filter has been proposed for use in illumination compensation is incorporated in [1]-[3]. This filter is one component of the Multi-scale Retinex algorithm, which now forms parts of a commercial image processing package [4]. The aim of illumination compensation is to mitigate the effect of uneven illumination on perceived image quality. For example when strong shadows are present it may be difficult to discern key image detail in the dark parts of an image. The image processing techniques referenced above can make a striking improvement in such cases. However the basis for the multi-scale filter is mainly empirical and the explanation given in [2] and [3] for the benefit of the approach is superficial.

In this paper a theoretical model is presented for the multi-scale filter. This model is used to predict the performance of the filter under step changes in illumination such as those caused by shadows. The advantage of the multi-scale filter is that it provides a relatively rapid response for transitions from shadow regions to well-lit regions. This rapid response is coupled with a better rejection of texture information than that provided by a single scale filter. These properties are verified using random Mondrian textures with simulated illumination changes. The estimated illumination profile is compared with that used to generate the synthetic image using the Pearson's correlation coefficient. The improved performance of the multi-scale filter is also confirmed by studying the mitigation of shadow artefact in real images. Several ways in which the performance of the multi-scale filter could be improved are identified.

[1] D. J. Jobson, Z. Rahman, and G. A. Woodell, "Properties and performance of a centre/surround Retinex", IEEE Transactions on Image Processing, vol. 6, no. 3, March 1997, 451-462.
[2] Z. Rahman, D. J. Jobson, and G. A. Woodell, " A Multiscale Retinex for Color Rendition and Dynamic Range Compression," Proceedings of the SPIE - The International Society for Optical Engineering, vol. 2847, 1996, 183-191.
[3] Jobson, .J.; Rahman, Z.; Woodell, G.A., "A multiscale retinex for bridging the gap between color images and the human observation of scenes", IEEE Transactions on Image Processing, Volume: 6 Issue: 7, July1997, 965 –976.
[4] TruView Imaging Company Homepage. http://www.truview.com

Region Based Stereo Matching through Bounded Irregular Pyramids

R. Marfil, A. Bandera, J.A. Rodriguez and F.Sandoval
University of Malaga, Spain

In this paper, we propose a region based stereo matching algorithm. Regions are chosen as matching features because: i) they present a large semantic content, which reduces mismatches, ii) their stability and tolerance to noise, and iii) pyramidal segmentation techniques yield to a reduced number of resulting regions, decreasing the computational load for the matching procedures. Though regions have been proposed in other works, they have been usually employed as an initial step which is independent to the matching process. However, the proposed method makes the segmentation and the matching cooperatively by considering the resulting regions from the left image segmentation as a reference to segment the right image. Besides, the right image segmentation is performed simultaneously with the matching process.

The algorithm is based on bounded irregular pyramids, which allow performing a fast multilevel hierarchical matching. The proposed method builds two pyramidal structures associated with the stereo pair images. When the left image is segmented, each main resulting region is searched in the higher right pyramid level presenting an adequate resolution. This level depends on the region shape and size.

The algorithm has proven to be fast by spending 500 ms to process a 256x256 colour stereo pair.

Simplifying gamut mapping colour constancy

Graham Finlayson
School of Computing Sciences, University of East Anglia

Gamut mapping color constancy algorithms (originally introduced by Forsyth [International Journal of Computer Vision, 1990]. attempt to map RGBs of surfaces viewed under an unknown light to corresponding RGBs under a known reference illuminant. With respect to a reference light source the set of all observable RGBs occupies a 3-dimensional convex region, or gamut, of RGB space. If a triple of 3 simple scalar factors defines the map from image colors to reference conditions then it has been shown that the set of all maps taking RGBs into the reference gamut is also a convex set. Gamut mapping algorithms work in 2 stages. First, the set of feasible maps is computed then in a second stage an optimal member of this map set is chosen.

Here we show that feasible map set does not need to be computed. Rather we combine the final map selection stage with the enforcement of the constraint that image colors are mapped inside the reference gamut. The main contribution of our research is to show that this combined computation can be set up as a simple convex programming problem. Two main advantages result from this reformulation of gamut mapping. First, several reasonable designations of `optimal' map can be formulated and tested within the convex programming framework. If we maximize the sum of the triplet of map parameters then color constancy is shown to be a linear programming problem. Maximizing the Euclidean magnitude of the mapping triplet gives a quadratic programming formulation and maximizing the volume of the mapped image gamut (Forsyth's original gamut mapping algorithm) is also possible The second advantage is that convex programming provides a fast solution to the color constancy problem. Indeed, we show that linear programming color constancy is a strictly more efficient implementation of gamut mapping compared to previous methods.

Experiments indicate that, for the Simon Fraser Data set (synthetic and real images), linear programming color constancy provides the best performance over all the convex programming methods tested.

A combined physical and statistical approach to computer colour constancy

Gerald Schaefer*, Steve Hordley+, and Graham Finlayson+
* School of Computing & Technology, Nottingham Trent University
+ School of Information Systems, University of East Anglia

It is a well known fact that the RGBs recorded by a device such as a digital camera do not solely depend on the characteristics of the objects being captured but also on the incident illumination. The same object taken under two different lights will produce two different images which in turn makes tasks such as colour based object recognition intrinsically difficult. Colour constancy algorithms try to address this problem by estimating the colour of the scene illumination - once, this is known images can be mapped to appear as taken under the same canonical light. In general, colour constancy algorithms can be divided into two classes: statistical and physics-based methods. While the former attempt to correlate the statistics of the colours in an image with statistical knowledge about light and surfaces, the latter are founded on an understanding of how physical processes such as specularities and inter-reflections manifest themselves in images.

In this work we developed a combined physical and statistical approach to colour constancy. On the physics-based side we make use of the dichromatic reflection model which predicts the colour signals of most objects to fall on a plane with the illumination vector embedded in the same plane. Crucially, we add a constraint on possible solutions based on statistical knowledge about common light sources. Estimates obtained from this method are combined with those achieved by the Colour by Correlation method, a purely statistical colour constancy approach.

Experimental results obtained from a large number of images taken under 11 lights confirm this new combined method clearly outperforms both purely physical and purely statistical algorithms.

Colour displays for categorical images

C.A.Glasbey+ , V. Toh++, A. J. Gray++ and J.Bigot+++
+ Biomathematics and Statistics Scotland
++ Department of Statistics and Modelling Science, University of Strathclyde
+++ Laboratoire LMC -IMAG, Grenoble

Suppose we have a categorical image, such as that produced by a segmentation algorithm, with *n* unordered labels. To produce a clear display of the image we need the colours for the *n* categories to be as distinct as possible. The problem has some similarities with that of colour quantisation. We can formalise it as: find a set of colours, $x_1, \ldots x_n$ such that $\min_{i\,j} \{D(x_i, x_j)\}$ is maximised, where $x_{\cdot}$ *(r,g,b)* denotes a point in the RGB colour cube and *D* is a distance metric. So, we find a set of colours such that the minimum distance between any pair of colours is made as large as possible. However, it is inconvenient to have distinct sets of colours for every different value of *n*. So, instead we define colours recursively. In increasing the set from *n* to *n+1* colours, we find the new colour that maximises the minimum distance between it and all existing colours. Therefore

$$x_{n+1} = \arg\max_x [\min_{i=1\ldots n} D(x_i,x)] ,$$

which can be solved by exhaustive search.

For *D*, we have considered Euclidean and perceptual colour spaces, such as that based on the CIELAB transformation. However, we were surprised to find a lack of agreement between CIELAB and our perceptions of both on-screen and printed colours. Therefore, we also derived an empirical metric, which appears to produce the best set of colours for displaying categorical images.

Evaluation of Colour Management Systems to Achieve Colour Communication Fidelity

Xiaohong W. Gao
School of Computing Science, Middlesex University

This paper reviews the necessary procedures to achieve colour fidelity over the internet and presents the evaluation results on Spyder with OptiCAL colour management system from PANTONE ColorVision Inc. Two LCD monitors are applied. One is with 15" monitor running Windows 98 operating system, the other with 17" monitor runing Windows XP system. Initial measurement and visual matching results show that both measurement and visual matching results are different for the primary colours including red, green, blue, white and black across two monitors. Some recommendations are given in order to improve the visual matching of colours across the monitors.

Computing Image Difference in Colour and Spatial Domain

Dr. Guowei Hong
*Colour & Imaging Ins*titute, University of Derby

This presentation describes a newly developed image difference metric based on the use of gamut mapped images. In our study, reproduced images with different levels of distortion caused by various gamut mapping algorithms were simulated on a CRT display and psychophysical experiments were carried out to obtain subjective visual assessment results. During the psychophysical experiment, it was found that observers often needed to make decisions to balance overall colour difference versus local spatial details. The proposed algorithm attempted to simulate this decision making by dividing overall image difference into two separate aspects: colour and spatial details. Colour difference is computed as an average of pixel-by-pixel difference between the low frequency components of the images while the difference of spatial details is calculated from the high frequency components of the images. The results obtained so far has shown that we can better understand how the human visual system compare images and judge the difference by calculating differences of colour and spatial details separately.
