A Colour Importance Measure for Colour Image Analysis

Anthony Maeder
School of Electrical and Electronic Systems Engineering
Queensland University of Technology, Brisbane, Australia

Binh Pham
School of Computing and Information Technology
Griffith University, Nathan, Australia

Abstract

Techniques for digital image analysis tasks which take account of colour as well as intensity can provide more information about the image contents. The approach presented here for constructing such techniques is to define a measure expressing the "colour importance" at each pixel position in the image, to act as a weighting factor within the calculation. Results of image analysis tasks which are closer to human appraisal of the scene can be obtained using this approach.

1. Introduction

We use our human colour perception ability to assist us in many complex visual tasks, such as object and feature identification in scenes. It is thus natural to expect that digital image analysis techniques based on colour properties might provide additional information which cannot be obtained from conventional greyscale intensity based techniques. A major difficulty in extracting this additional information exists because data dependencies between the component colour bands are poorly defined. Image analysis techniques which take colour component correlation into account can outperform the equivalent greyscale techniques in some circumstances (e.g. image compression[1]). However, no general approach exists for incorporating colour correlation in arbitrary tasks. Furthermore, some behavioural aspects of the human visual system (e.g. non-linearity, simultaneous contrast) make it difficult to obtain results for a particular image analysis task which are in harmony with human expectations (e.g. colour correction[2]).

The above difficulties are aggravated by the commonly accepted practice in many image analysis techniques of treating each component colour band equally and not taking into account any special characteristics of the data. To address this problem, we introduce the concept of "colour importance," which provides a means to increase or reduce the effects of computational tasks in different places within the image. In section 2, we discuss some typical situations arising in colour image analysis tasks, indicating how our perception of the colours can influence our expectations of the results. We then develop the notion of colour importance formally in section 3 by proposing a measure composed of several factors, which can be used to modify colour image analysis tasks. Finally, in section 4 we present two examples of constructing such measures and the corresponding effects obtained on the image analysis results.

2. Colour Situations

The influence of colour in human understanding of a particular scene and thus on the results desired from various image analysis tasks, requires the identification of appropriate colour characteristics of the data. Once established, these characteristics can be used to derive computational techniques which incorporate elements of this influence. We describe here some situations which occur in several different colour image processing tasks, to indicate typical colour characteristics which might be considered.

2.1. Edge Detection

Edges are usually modelled as places where sudden local changes in pixel values occur. It is well-known that certain intensity characteristics of a scene, such as illumination levels, shadows, highlights and surface orientations, degrade the results of edge or contour detection processes.[3, 4, 5] However, colour is often considered to aid in edge detection performance.[6, 7]

In general, edges obtained using the luminance band are more pronounced than those produced from the chromatic bands. Hence the spectral differences become more important only in situations where the difference in luminance is negligible. In order to detect poorly-defined edges or boundaries between regions of equal luminance, we need to increase the influence of the chromatic bands.

On the other hand, shadows and highlights in an image cause sharp changes in luminance, producing undesirably strong edges at these places. Saturation can be used to detect shadows or highlights in situations where the hue is constant, as it gives an indication of the amount of white present in a colour. Hence we need to reduce the influence of luminance selectively in such situations.

2.2. Region Segmentation

Images are segmented by finding regions of adjacent pixels which are homogeneous with respect to some computed feature values. These feature values model certain perceptual characteristics of the region contents, often based on statistical properties of pixel intensity or colour. Repeated patterns or noise are modelled by considering several pixel values in combination as "texture" features. Individual pixel intensity or colour values alone are considered only when there is little variation in the texture.

If regions associated with different materials (or textures) are sought, we would again like to reduce the effects caused by shadows and highlights. In this case the influence of luminance in the texture feature calculation must be varied as appropriate. A related situation which often arises is a gradual variation in colour properties within a relatively large region, due to changes in illumination or surface orientation. In this case, when luminance and saturation are changing but do not have strong edges, more emphasis must be placed on the hue values in the feature calculations.

Another situation which requires special attention arises from the effects of simultaneous contrast in human vision. In cases where region boundaries are diffuse and not well defined by sharp edges, but region contents differ significantly in colour, boundary localization can be difficult. The range of intermediate colours in the diffuse area can lead to a perceived boundary position displaced towards one of the regions. For example, when a saturated colour blends with a pastel the boundary could be seen as being closer to the saturated region because small changes in saturated colours can be discerned more readily than in pastels. In some cases, such as adjacent regions of very different saturated colours, another region constituted by the colours in the diffuse area may even be perceived. These situations must be addressed by prescribing the influence of colours in the diffuse area on the homogeneity features.

2.3. Colour Correction

When spatial sequences of images are viewed together, they often appear unnatural due to colour variations between the images. For example, when panoramic photographs are joined together, there may be large changes in the colours between individual photographs. These colour changes are the result of several separate effects. Time delays between capturing each image can allow the illumination of the scene to change, for example when the sun is obscured by clouds. The image colours may also vary due to adjustments performed by the capturing device.

To improve the visual appearance of a spatial image sequence as a whole, we can use a multiple regression technique to find a mapping from colours of spatially registered pixels from overlap region in one image to those in the adjacent image, such that the least square error is minimized.[8] Once the mapping is obtained, it is applied to one image in order to correct its colours to match the adjacent image.

This technique gives good results when the colours present in either image are also present in the overlap region. However, if some colours are present in one image but not in the overlap region, unusual results might occur because the multiple regression does not take those colours into account. Two important questions which need to be addressed are: (i) which colours should be incorporated into the regression and in what proportion? and (ii) how far can a colour be changed during correction without causing visual degradation? The answers to these questions depend on the contribution of a particular colour to the viewer's perception of the image, in terms of both abundancy and noticeability of the colour.

2.4. Colour Quantization

The process of colour quantization involves choosing a colour subset within the set of all colours occurring in an image, such that degradation of the visual appearance of the image displayed using this subset of colours is minimized. This process is useful when the display device is restricted to generating fewer colours than the raw data, or when synthetic images are generated for visualization of data values.[9] The process is undertaken in some colour image compression techniques where a saving in storage is achieved by reducing the number of colours used for representing the image. Quantization is often performed by identifying clusters of nearby colours in colour space and choosing a colour near the centre of a cluster to represent all colours in that cluster. This process ignores the significance of the colours in perception of different visual situations.

For quantization purposes, shadows and highlights can be treated as separate regions of the image from those of strong colour content, and can be represented by less accurate hue and saturation than the actual ones, as we are not as sensitive to these properties under extremes of illumination. This suggests that fewer, more neutral values could be used throughout the image to represent shadows and highlights. However, in the case of slowly changing illumination on strong colours, the variation in satur- ation of a colour must be quantized quite finely to avoid contouring.

Where saturated colours blend into one another (e.g. sunset against blue sky), the mixed colours lie close to each other but not necessarily on a smooth path in colour space. The accuracy of the representative colours (i.e. the shape of the path in colour space) is less significant here than the regular separation of intermediate colours and the accuracy of colours at the path ends. On the other hand, when pastel colours blend, the mixed colours lie on a far smoother path in colour space, which must be more accurately preserved because colour variations are more noticeable in such cases.

Even when many different saturated colours occur in an image, the representatives to be chosen can vary in accuracy depending on the situations in which the colours appear. When hues of nearby regions differ considerably, the representative colours chosen need not be as accurate as when the hues are more similar, as we perceive the gross difference in hues more strongly than the actual values. Larger regions can prompt more accurate choice of colours than smaller ones, as the area over which a colour occurs increases the accuracy of our perception. Furthermore, regions containing little variation in colour need more accurate colour representatives than regions containing many dispersed colours (i.e. patterns or noise), as we see uniform colours more accurately.

3. Colour Importance

The above situations all demonstrate that the influence of colour on the image analysis task being performed can vary widely within an image, depending on many different criteria. We should like to model this influence formally in order to reduce or increase certain effects while performing the image analysis task. We need to model the influence at each pixel position, since image analysis tasks use individual pixel values, but may wish to determine how we model it differently at different places in the image. We also need to be able to combine the effects of several influences during an image analysis task. We may wish to elect whether to make use of our knowledge of the influence for modifying colour component values prior to performing the tasks, or for adjusting intermediate or final results of tasks.

Our approach for modelling this influence has been to quantify it by constructing a measure which we term "colour importance" (denoted *IMP),* for a given image and image analysis task. The value of this measure is defined for every pixel position (i, j) in the image and provides a weight which can be used in conjunction with the image processing task to scale the calculations appropriately. The measure is constructed by combining a number of "factors" which express the individual influences we wish to consider. We derive these factors by considering what colour effects we wish to express in terms of human expectations. We can divide these factors into two categories, "global" and "local", depending on the domain over which the fac-tor exerts influence.

3.1. Global Factors

Global factors can be computed for the whole image independently of pixel position, so their values depend on some property of the pixel other than its position. Clearly, the property we are most interested in here is colour, so we discuss several global factors depending on colour. The most obvious influence on colour importance in an image is the probability with which the colour C occurs globally in the image, denoted $Pr_G(C)$. This factor gives an indication of the abundance of that particular colour within the image, enabling us to predict the prominence of the colour if all other influences are discounted, and so to compensate for it appropriately.

In some situations, a predefined grouping of colours is available (e.g. colour clusters) and we may wish to provide a factor related to these groups. We denote the particular colour group to which colour C belongs as K_C and hence define the probability with which a group occurs globally in the image as $Pr_G(K_C)$. The group provides a stronger indicator of influence than the raw colour in those images where a very wide range of colours is present and so each occurs with only small probability. We may also wish to define the probability with which a colour occurs within its colour group, defined by $Pr_G(C|K_C)$ Use of this factor would be appropriate when the influence of a colour depends on its abundance relative only to colours which are similar in some sense.

Another global influence on perceived effects in an image is the variability or dispersion of colours, denoted V_G. The colour variability may be characterized in various different ways, such as mean and variance of the colour probabilities, or the 3D Euclidean distance from one or more chosen positions in colour space. Since V_G provides a single value for an image (or perhaps for each band of an image) it can be used only as a scaling constant in the col-our importance formula.

Low variability indicates that the colours in the image are close to each other, so discrimination is limited and may need to be enhanced (e.g. to obtain a good segmentation). On the other hand, high variability implies that colours are more widely spread and may need to be used more selectively (e.g. for a good colour quantization scheme). As in the case of colour probabilities, considering variability relative to groups of colours may be more useful for modelling some effects. The variability amongst the groups (i.e. how well separated), denoted V_{GK}, and the variability within a particular group (i.e. how loosely clustered), denoted V_G/K_C, are two examples.

3.2. Local Factors

As image characteristics can vary widely within an image, global properties have a limited use in determining colour importance. Many colour influences depend on local contextual properties of image space, so weighting factors have to be determined adaptively, according to the context for the current pixel. Typical contexts are a regular neigh-bourhood of the pixel (e.g. mxn block of pixels), or an irreg-ular area to which the pixel belongs (e.g. a segmented region).

We can extend the definition of the global factors described above to include contextual restrictions. For a pixel with colour value C at position (i, j), let $Pr_{L(mxn)}(C)$ (i, j) denote the probability of C in the (mxn) local neighbourhood at position (i, j) in the image. $Pr_{L(mxn)}(C)$ (i, j) indicates the abundance of the colour C in the specified neighbourhood, and therefore the local importance of that colour. If a region was used in place of a neighbourhood, our equivalent notation would be $Pr_{LR}(C)(i, j)$, where R indicates the area which includes the pixel location (i, j).

We can also consider the restricted probability relative to only those colours within the neighbourhood which are in the same globally determined cluster as C, denoted $Pr_{L(mxn)}(C|K_C)(i, j)$ or similarly for the corresponding region, denoted $Pr_{LR}(C|K_C)(i, j)$. These factors indicate the influence of the current pixel colour amongst related colours within the current spatial context. We can define local variability factors in a similar fashion. Let $V_{L(mxn)}(i,j)$ denote the colour variability within an (mxn) neighbourhood at pixel location (i, j). This quantity gives an indication of the homogeneity in colour in the current spatial context.

3.3. Combining Factors

Once the relevant factors have been established for a given image and image analysis task, the formula for the colour importance measure must be constructed by combining the factors. By ensuring that each of the factors is within the range [0, 1], we can construct the final formula simply as the product of these factors. It is desirable for the factor to occupy as much of this range as possible, so that its contribution in the formula is most effective. Factors thus need to be normalized relative to the largest value they actually take for the given image. Where a negative correlation is required for the influence of a factor, the factor value can be subtracted from 1 to provide the appropriate transformation.

The distribution of values in [0, 1] for the different factors may differ, which can lead to a colour importance bias in favour of those factors skewed towards 1. Ideally, factors would be scaled so that they are mutually equivalent in overall influence, but retain individual dominant influences over certain parts of the range where necessary. In such cases, a non-linear transformation may be required to adjust the factor distributions. One such candidate transformation is the histogram matching process widely used for image enhancement [10, pp 171-184], which stretches distributions to produce equivalent cumulative distributions.

Under other circumstances, it may be desirable to transform some factors so that their influence is of different strength to others. For example, we may wish global colour factors to be more influential than local ones in a particular task. This effect can be achieved by linear scaling of a factor range into a subset of [0, 1], followed by translation of the scaled values towards the high or low end of the range as appropriate. Provided any such transformations map [0, 1] into [0, 1], the basic measure properties for colour importance will be preserved.

4. Examples

We provide two examples here of image analysis tasks performed on the commonly used image MANDRILL, in 256×256×24bit form. The factors involved in constructing colour importance measures for these two cases are discussed and the formulas used are presented. Results obtained using these formulas are compared visually with results from the basic image analysis techniques without taking account of colour importance.

4.1. Edge Detection

A common problem in edge detection tasks is how to reduce the occurrence of minor edges (e.g. isolated or spurious high edge strengths), yet retain major ones (e.g. large contiguous sets of high edge strengths). In MANDRILL, both types of edges occur where there are large colour changes, but minor edges occur almost exclusively where the colours are very mixed (e.g. around the periphery of the image) and within large regions of similar colours where surface effects interrupt the region homogeneity (e.g. creases on the cheeks).

The major edges we seek to emphasize in MANDRILL are those which yield the boundaries of the large regions of homogeneous but very different colours: the eyes, cheeks and nose. Noting that the colours blend at the boundaries over a few pixels, perhaps due to the digitization process, we expect that strong edges at places where unusual colours occur should be emphasized. This suggests that the edge strength should be weighted by the factor $(1 - Pr_G(C))$.

Next, we wish to emphasize edges where local variability of colours is high, which will be the case when regions with two very different homogeneous colours abut. For this effect, we used the factor $V_{L(9x9)}(i, j)$, taking the mean distance between the colour of the pixel at (i, j) and the colour of each other pixel in the neighbourhood, in RGB colour space, as the variability value. This factor has the added benefit of being low in value at positions inside regions, where colour is homogeneous, and so will reduce the edge strength of shadows caused by surface roughness within the regions.

Finally, we observe that variability is also high where there are many different colours in the neighbourhood (such as near the image periphery) so we must reduce edge strengths here. This suggests a factor $Pr_{L(9x9)}(C)(i, j)$ which will be low in such cases. The complete formula for the colour importance measure in this example can be written as follows:

$$IMP(i, j) = (1 - Pr_G(C)).\ V_{L(9x9)}(i, j).\ Pr_{L(9x9)}(C)(i, j)$$

Figure 1(a) shows the normalized result of using a Sobel 5×5 edge detector on the image, applying horizontal and vertical masks on each of the three separate RGB colour bands and taking the maximum edge strength obtained for all these instances, at each pixel position. The spurious edges in the cheeks and isolated edges in the hair are quite obvious. Figure 1(b) shows the normalised result of multiplying the edge strengths from Figure 1(a) with the values for $IMP(i, j)$ at each pixel position and normalizing. A considerable reduction in the density and strength of the isolated and spurious edges is evident.

4.1. Colour Quantization

The MANDRILL image is distinctive in that it contains very strong and fairly uniform colours over three large regions (eyes, cheeks and nose) and widely varying colours in the periphery (hair). There are also some regions of highlights and shadows which are small in area but are very noticeable to the observer. If we wish to quantize the colours in the image so as to retain a realistic visual impression of the scene using a minimal number of colours, we would expect to use at least six colours: yellow, blue and red for the face, white for the highlights, black for the shadows and grey for the hair.

Most sensible colour quantization techniques will select these six basic colour groups, but the actual representative colours used may differ quite widely. For example, the representatives for black, white and grey need not be neutral: the white which is chosen may be biased towards the bluish highlights on the cheeks or the reddish highlights on the nose. The degree of realism therefore depends on how these representative colours are chosen. Typically the representatives are chosen by weighting the colours associated with a particular group. By using colour importance to perform the weighting we can ensure that the influence we desire is placed on this choice.

Our first consideration should be that colours occurring often in the image should influence the choice of representatives proportionately. This suggests a factor of

Pr_G should be used. We also wish to ensure that colours which occur very frequently over a large neighbourhood (but perhaps not frequently globally) are favoured (e.g. the yellow eyes or the nose highlight). For this effect, a factor of $Pr_{L(9x9)}(C)(i, j)$ is used. Finally, we wish to place more emphasis on colours which occur in places where there is little colour variability, since the observer is more sensitive to colour shifts in such situations than where there is much variability (e.g. the hair). This suggests a factor of

$$(1-V_{L(9x9)}(i, j))$$

In this example, the median cut algorithm[11] was chosen for performing colour quantization. This algorithm already incorporates a weighting of $Pr_G(C)$ in its calculations, so this factor should be excluded from the colour importance measure. The final formula used here was:

$$IMP(i, j) - Pr_{L(9x9)}(C)(i, j).\ (1-V_{L(9x9)}(i, j))$$

and was applied to the set of colours used as input to the median cut process.

Figure 1(c) shows a normalised error image for a six colour median cut quantization of the colours. The error is obtained by computing the 3D Euclidean distance in RGB colour space between the raw pixel colour and the representative colour concerned. Error values are high for the cheeks and nose, and otherwise are distributed fairly uniformly through the image. Figure 1(d) shows a normalised error image obtained when the input colours to the median cut algorithm are weighted by the above colour importance measure. It is noticeable that the lower error values now occur in large regions of similar colour (e.g. cheeks and nose) and higher error values occur in some smaller regions (e.g. nose highlights) and in the periphery where there are mostly dispersed colours.

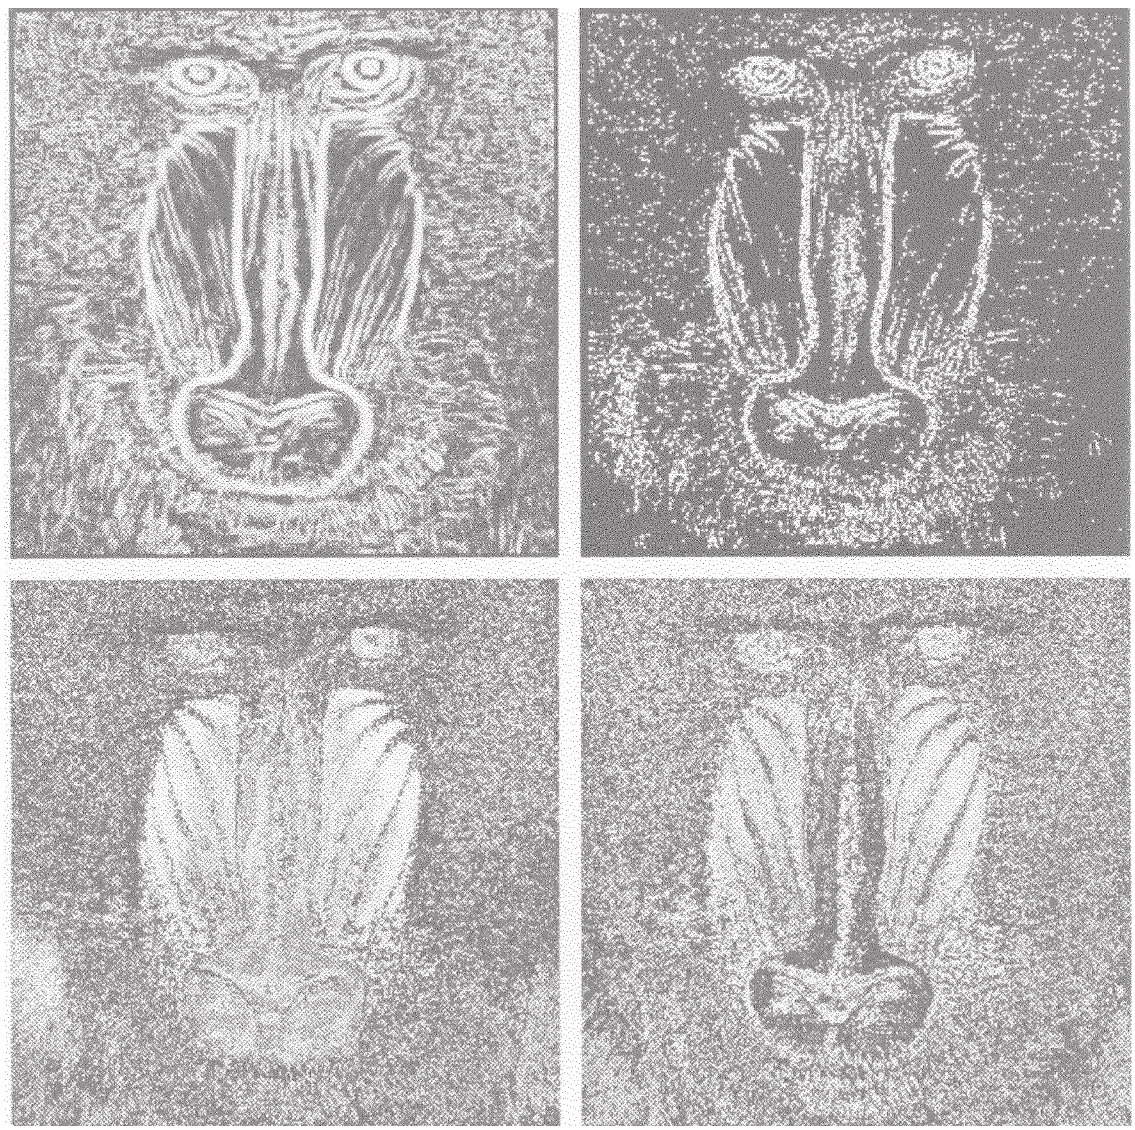

Figure 1. Results for examples: (a) Sobel 5x5 edge image (upper left); (b) Sobel 5x5 edge image weighted by colour importance (upper right); (c) Six colour median cut error image (lower left); (d) Six colour median cut weighted by colour importance error image (lower right).

5. Conclusion

The motivation for a general definition of a colour importance measure has been advanced and some examples of its use and effectiveness in image analysis tasks have been given. In both the examples above, other factors could be included to influence the task further, or alternative factors with fairly similar effects could be substituted. The definition of a colour importance measure for a given problem is neither unique nor demonstrably optimal as we have defined it. However, the approach is formal in that we can define factors unambiguously, and systematic in that it allows arbitrary colour influences to be modelled in conjunction with the image analysis task being performed, without the need to derive new techniques specifically for certain cases.

At present, the formulas used must be composed on a ease-by-ease basis, requiring some insight into the visual effects which it is desired to emphasize. The relative weights and linearity of the factors cannot be assured. The formulas are therefore useful only on an individual basis, within a given case. It would be helpful to formalize the derivation of the colour importance measure further so that general application of formulas and absolute comparisons of results could be undertaken for many different images and tasks.

The financial assistance of the Australian Research Council in supporting the work reported here is gratefully acknowledged.

References

1. A. J. Maeder, P. E. Tischer, and D. F. Chatterton, Colour space influence in context compression of images, *Proceedings of Workshop on 2 and 3 Dimensional Spatial Data: Representation and Standards,* Australian Pattern Recognition Society, Perth, Dec. 1992, pp. 50-55.
2. G. Pringle and B. Pham, Colour correction for an image sequence, *Proceedings* of *IS&T Color Imaging Conference: Transforms and Transportability of Color,* Phoenix, Nov. 1993, (to appear).
3. J. M. Wolfe, Hidden visual processes, *Scientific American,* **248**:94-103 (1983).
4. B. W. Tansley and R. M. Boynton, A line, not a space, represents the visual distinctions of borders formed by different colors, *Science,* **191**:954-957 (1976).
5. R. Gershon, Aspects of perception and computation in color vision, *Computer Vision Graphics and Image Processing* **32**:244-277 (1985).
6. R. Nevatia, A color edge detector and its use in scene segmentation, *IEEE Transactions on Systems Man and Cybernetics,* **SMC-7**:820-826 (1977).
7. M. Collins, G. Pringle, and B. Pham, Roles of colour spaces in image processing, *Proceedings of DICTA-91 Conference on Digital Image Computing: Techniques and Applications,* Australian Pattern Recognition Society, Melbourne, Dec. 1991, pp. 510-516.
8. G. Pringle and B. Pham, Colour equalization for image pairs, *Proceedings of Image and Vision Computing 93 Conference,* Auckland, Aug. 1993, pp. 229-236.
9. A. J. Maeder, Software tools for visual analysis of colour compression techniques, *IEEE Workshop on Visual Signal Processing and Communications,* Melbourne, Sept. 1993, (to appear).
10. R. M. Gonzalez and R. E. Woods, *Digital Image Processing,* Addison-Wesley, 1992.
11. P. Heckbert, Color Image quantization for frame buffer display, *SIGGRAPH '82 Proceedings,* 1982, pp. 297-307.