Research report 112

QUANTISATION OF DIGITISED COLOUR IMAGES

R W McColl & G R Martin

(RR112)

Abstract

Schemes for the quantisation of digitised colour images are investigated, in order to reduce their storage requirements. The schemes involve transforming the *RGB* vectors into CIE *XYZ* format, from which the luminance Y and chromaticity x,y values are drawn. The x, y values are further transformed to chromaticity vectors ξ, η in the MacAdam domain. Histograms are found for both the luminance and chromaticity data, and clustering algorithms are used to create the codebooks used to quantise the image. Careful attention is required to ensure that the quantised $Y \xi \eta$ vectors do not lie outside the *RGB* space. Compression ratios of around 6:1 are achieved, with the chromaticity component requiring less than 1 bit-per-pel.

Department of Computer Science
University of Warwick
Coventry CV4 7AL
United Kingdom

November 1987

QUANTISATION OF DIGITISED COLOUR IMAGES

R. W. McColl and G. R. Martin,

Department of Computer Science,
University of Warwick.

1. Introduction

Broadcast-quality images are usually digitised at 8 bits per colour plane so that the Red, Green and Blue components of the picture element have values in the range 0 to 255. For a typical reduced resolution image with 256 rows and 256 columns, the storage requirement is therefore

256 x 256 x 3 bytes = 196608 bytes,

approximately a quarter of a megabyte. For the more common 512x512 image, this requirement is almost 1 megabyte. Storage space for several of these images, or perhaps a sequence of frames for real-time applications, would obviously stretch the resources of many small systems. A suitable quantisation scheme could considerably reduce the storage requirement.

It is worth noting that what is important in a quantisation scheme is the perceived quality of the reconstructed image compared to the original. It is therefore desirable to identify a colour space that can be used to guarantee a minimum perceptual distortion at each pixel. Because of the non-homogeneous nature of the *RGB* colour space, and the high correlation between the Red, Green and Blue images, data processing takes place in a colour space in which the data are represented as luminance, or relative brightness, and chromaticity, or relative colour. The main interest lies in the compression of the chromaticity data, which occupies a much smaller bandwidth than the luminance in the broadcast signal.

Our method involves the use of the MacAdam chromaticity domain, a two-dimensional space possessing a nearly-homogeneous nature. The quantisation scheme includes the creation of histograms for the luminance and chromaticity at every pixel in the image. One and two-dimensional clustering algorithms are then used on the histograms in order to isolate the local centroids. These centroids are used as the codewords that give, for every pixel, the closest colour to the original *RGB* vector. Sixteen luminance levels and eight chromaticity levels form the compression limit of this scheme. With subsampling of the chromaticity and source coding of the luminance data, compression ratios of about 6 to 1 can be achieved with a resultant image quality perceptually close to the original.

An explanation of the homogeneous, or *psychometric,* colour space is given in the next section, including the reasons for choosing the MacAdam chromaticity domain and the ways of avoiding the high processing cost of transforming to this domain . The clustering algorithm is described in section three with a discussion on valid distortion measures. The fourth section describes the overall quantisation process and considers the problem of chosing codewords that lie outside the original *RGB* colour space and a fifth section provides a discussion of the results of the various schemes. A concluding section describes the pitfalls of the method chosen and considers future investigations.

2. A Metric Colour Space

In order to display a digitised image the data must be represented as vectors in *RGB* colour space. It makes sense then to assume *RGB* space to be the starting point in a quantisation scheme. Although some research indicates otherwise, it is widely accepted, after the theories of Young [1] and Helmholtz [2], that colour is a trivariant phenomenon, so the 3-dimensional *RGB* colour space cannot be reduced to a smaller dimension without irretrievable loss of information.

With a view to minimising subjective distortion, a colour space in which the Euclidean distance between any two points is perceived as equal is worth seeking. Such a space is a homogeneous and psychometric space because of its subjective properties. Unfortunately the *RGB* representation is ideal for driving television circuitry, but it is very non-homogeneous. The Red, Green and Blue planes of a digitised image also tend to be highly correlated, especially the Red and Green. This is due to the similar responses of the CIE $\bar{x}$ and $\bar{y}$ absorption functions [3], shown in figure 1, matched in the television camera. From a visual point of view, it can be explained partly from the theory that the second stage of the human visual system works by an opponent-colours process [4,5] involving the following tristimulus :

$$\begin{bmatrix} White \leftarrow Black \\ Green \leftarrow Red \\ Yellow \leftarrow Blue \end{bmatrix} \quad . \tag{1}$$

The origin in this space represents an achromatic grey. Transformation of the *RGB* vectors into luminance and chrominance components largely removes this correlation. The luminance, or brightness, component, Y, is a weighted sum of the Red, Green and Blue components, according to their compostion in a reference white standard. For most television standards [6] the relation is

$$Y = 0.299R + 0.587G + 0.114B \ . \tag{2}$$

A luminance value, then, represents a plane through the RGB space, such as shown in figure 2. The chrominance components are expressed as colour differences [6] that can be added to the luminance component to produce the *RGB* vector at the receiver. The chrominance coding scheme is designed to maximise the distance between the components to prevent interference [7].

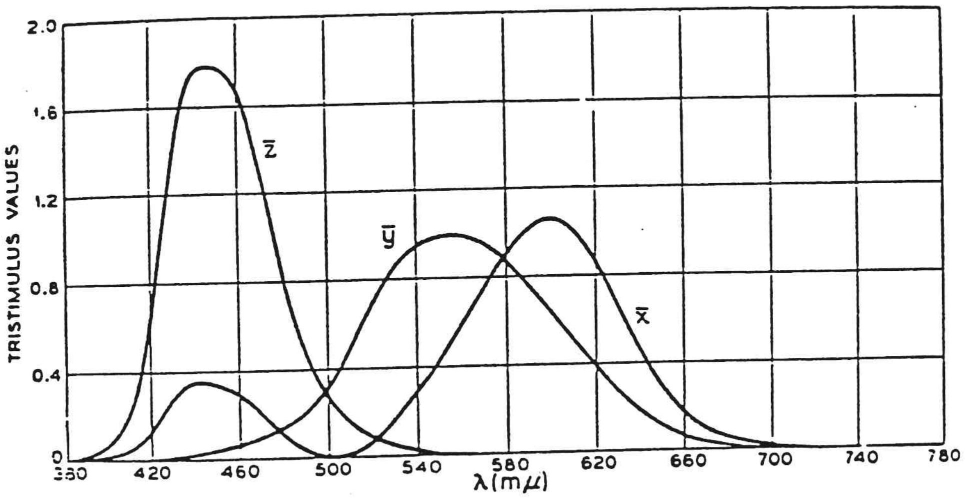

Figure 1. Colour matching functions for CIE Standard Observer in CIE coordinates. Source [3], Fig. 7.

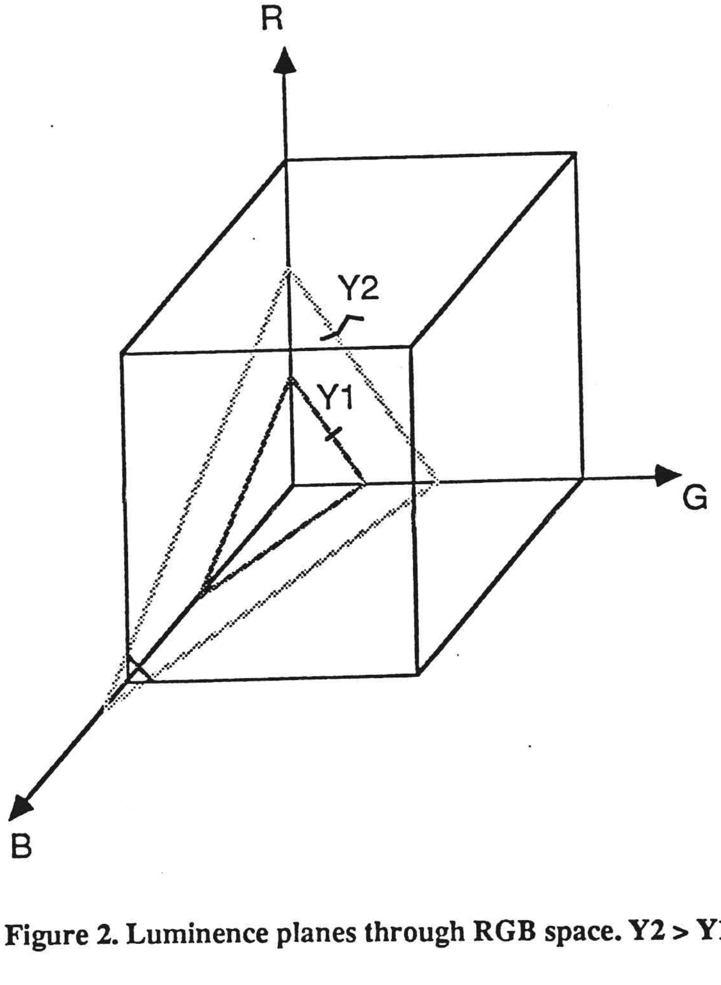

Figure 2. Luminence planes through RGB space. Y2 > Y1.

Since 1925 the Commission International d'Eclairage (CIE) has been involved in colour space standards. In 1931 they proposed the *XYZ* colour space [7], a linear transformation of the *RGB* space but where all light stimuli are represented by positive vectors. This is an artificial colour space but it circumvents the drawback of the *RGB* representation, where some light wavelengths must be matched by negative values of Red, Green or Blue. The relation between *RGB* and *XYZ* is [6]

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.607 & 0.174 & 0.000 \\ 0.299 & 0.587 & 0.114 \\ 0.201 & 0.006 & 1.117 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \quad (3)$$

This colour space is still non-homogeneous. The CIE have made a number of efforts to produce a more homogeneous colour space, such as the *UVW* colour space [8] proposed by MacAdam in 1969. Further work by MacAdam [9] demonstrated that a completely homogeneous colour space was in fact impossible to create by means of linear transformations. His research into metrics involved measuring the *just noticeable difference* between two colours. Tests of the abilities of a number of observers to determine the threshold at which two colours became noticeably different resulted in a series of ellipses whose circumferences represent the limits of observable similarity, shown on an x,y chromaticity histogram in figure 3. The chromaticity coordinates x and y are related to the *XYZ* space by

$$x=\frac{X}{X+Y+Z} \quad (4a)$$

and

$$y=\frac{Y}{X+Y+Z} \quad (4b)$$

These normalised values represent the chromatic content of the colour vector. The luminance axis is out of the page. It can be seen that the ellipses vary greatly in size across the chart. From these data, MacAdam derived a non-linear transformation of the x,y coordinates [9] to produce nearly-circular shapes of nearly-constant radii, as in figure 4. The transformations were achieved by the following :

$$\eta=404b-185b^2+52b^3+69a(1-b^2)-3a^2b+30ab^3, \quad (5a)$$

where

$$a=\frac{10x}{4.2y-x+1} \quad (5b)$$

and

$$b=\frac{10y}{4.2y-x+1}; \quad (5c)$$

$$\xi=3751a^2-10a^4-520b^2+13295b^3+32327ab-25491a^2b-41672ab^2+10a^3b-5227a^{1/2}+2952a^{1/4}, \quad (5d)$$

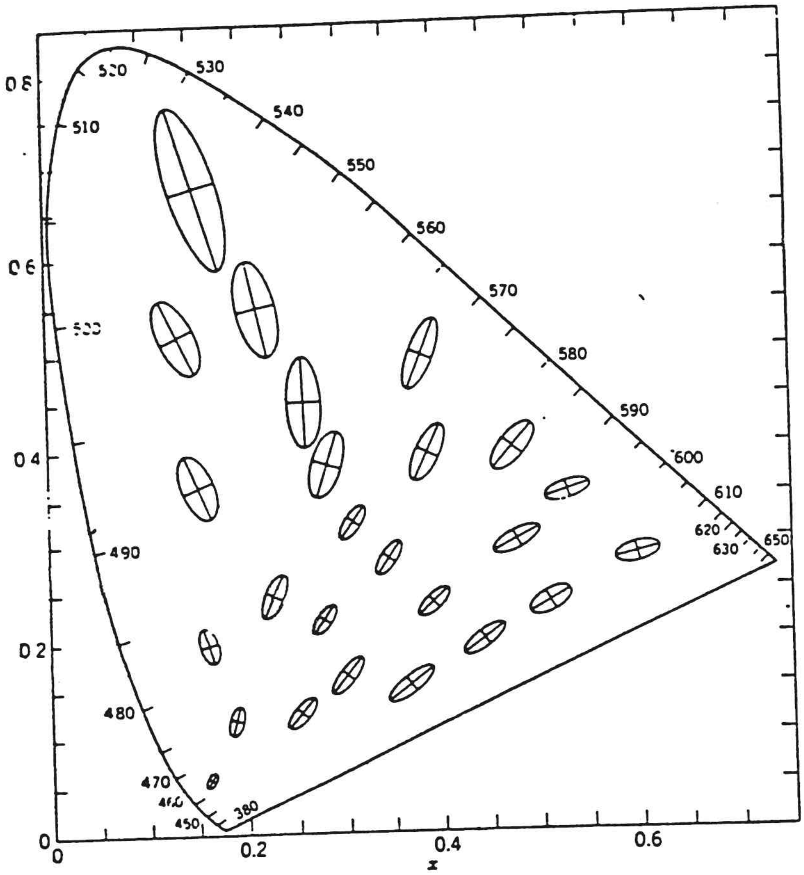

Figure 3. MacAdam's Ellipses on CIE Chromaticity Chart. Source [21], Fig. 6.36.

where

$$a=\frac{10x}{2.4x+34y+1} \tag{5e}$$

and

$$b=\frac{10y}{2.4x+34y+1}\,. \tag{5f}$$

The relationship between the CIE x-y and MacAdam chromaticity domains is shown in figure 5. On this chart the meaning of the distance between any two points is the least number of *just noticeable chromatic differences.* Note that this measure applies between colours of constant luminance. In fact, as the luminance level increases (the colour becomes brighter), the area in the MacAdam space shrinks (refer to figure 2) as the chromatic component variation is restricted. MacAdam recommended that the CIE use his formulae but they were not accepted and there is evidence of discrepancies in his data [10]. A more homogeneous space which is widely accepted in the literature has not been found however so the MacAdam space has been adopted here as the best compromise.

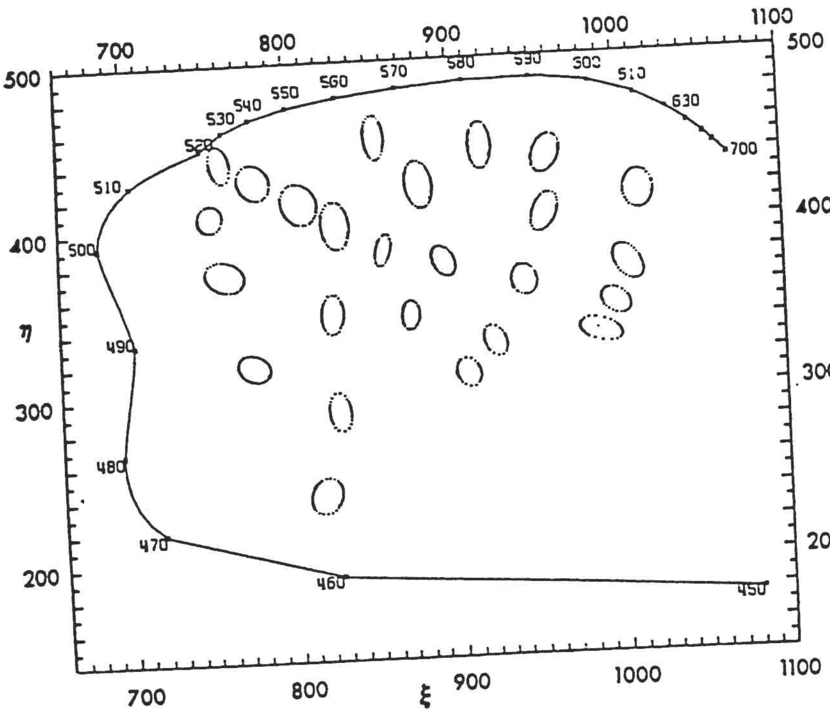

Figure 4. MacAdam's Ellipses on MacAdam Chromaticity Chart. Source [9], Fig. 3.

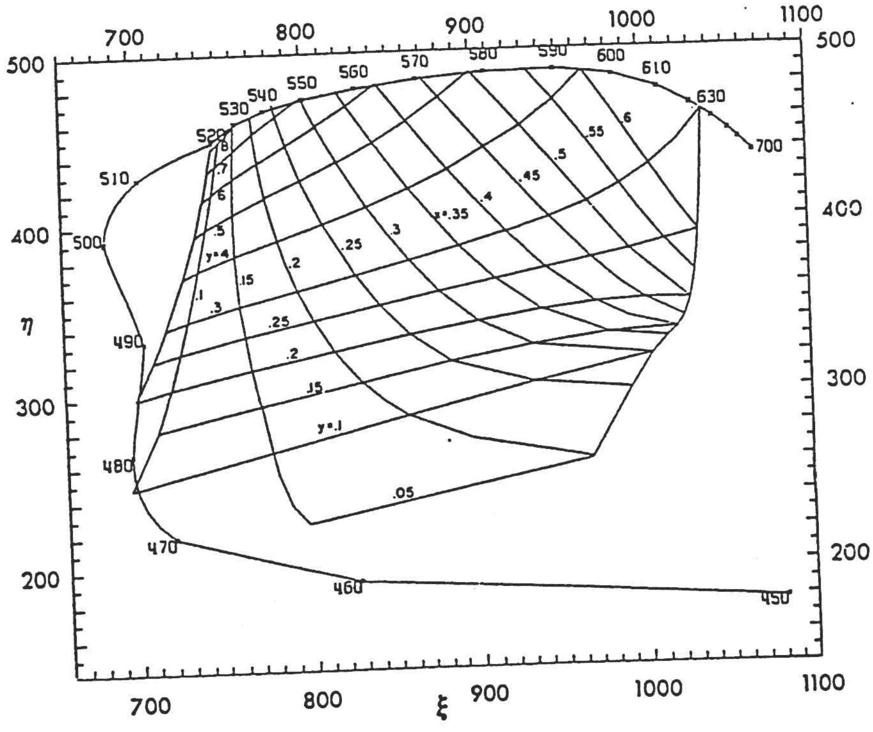

Figure 5. Relationship between CIE and MacAdam Chromaticity Domains. Source [9], Fig. 1.

The quantisation scheme deals separately with the luminance and chromaticity distributions in the image. It is assumed that the two datasets are uncorrelated and in fact the original, correlated *RGB* vectors cannot be predicted from the values of either dataset without the other, the chromaticity data being dimensionless (see equations (4)).

3. The Clustering Algorithm

3.1. Histograms - Distortion Measures

The purpose of creating histograms for the luminance and chromaticity content of an image is to provide statistics necessary to run some sort of quantisation engine. The criterion of importance in the engine is to minimise the distortion of the input image. The most common distortion measure is the mean-square error

$$MSE(x,\hat{x}) = \frac{1}{N}\sum_{i=0}^{N-1} |x_i - \hat{x}_i|^2 . \tag{6}$$

Referring to the previous section, the MSE in the chromaticity domain is exactly the perceptual distortion of the domain.

What about the distortion of the luminance ? Studies of the human visual system [11,12] suggest that retina responds to light according to the logarithm of its intensity. A number of models have made use of this theory [13,14] If they are correct, the distortion measure would be modified to

$$MSE(y,\hat{y}) = \frac{1}{N}\sum_{i=0}^{N-1} |\log(y_i) - \log(\hat{y}_i)|^2 . \tag{7}$$

Although there is considerable support for the logarithmic sensitivity of the eye, our experiments produced better results when the absolute values of luminance and chromaticity were used as the input signals to the histograms.

3.2. Clustering - Local Centroids

The clustering of the histogram data is analogous to *vector quantisation* . Following Lloyds ideas [15], Linde, Buzo and Gray [16] identified the *minimum distortion vector* in an arbitrary n-dimensional Euclidean space which they called the *centroid* or *centre of gravity* of the distributed signal. The centroid is defined as

$$c_i = \sum_{j=0}^{N-1} \frac{x_j p(x_j)}{p(x_j)} . \tag{8}$$

If the space is optimally split into a series of subspaces or sets, their local centroids form the reproduction alphabet providing the minimum distortion. The *k-means* test developed by

MacQueen [17] for cluster analysis, used in this report, is simple and easy to implement. The algorithm is an iterative one which relaxes to a solution in which k local centroids form the alphabet of codewords. The algorithm can be summarised as follows :

Assume the signal space is divided into k sets S_i with local centroids c_i, $i=0,1,..k-1$ i.e. k codewords are required.

Assume that there are N input vectors, or symbols, x_j $j=0,1,..,N-1$ in the signal.

Repeat until there is no change in the value of any of the centroids between iterations $n-1$ and n :

For each vector x_j, $j=0,1,..,N-1$

$$x_j \in S_i \text{ if } |x_j - c_i^{n-1}| < |x_j - c_l^{n-1}|,\ 0 \le l < k,\ l \ne i\ . \qquad (9a)$$

For each centroid c_i^{n-1}, $i=0,1,..,k-1$

$$c_i^n = \sum_{j=0}^{N-1} \frac{(x_j p(x_j) | x_j \in S_i)}{(p(x_j) | x_j \in S_i)}\ . \qquad (9b)$$

The final set boundaries would be retained as the partitions for each codeword in the alphabet.

4. The Quantisation Procedure

4.1. The Basic Outline

The main stages of the quantisation scheme are shown in figure 6. Although a number of variations of the scheme have been investigated, the figure applies to each one. The input to the process is assumed to be an *RGB* image, each plane containing 256 by 256 1 byte values. Using the matrix in (3) these vectors are transformed to the CIE coordinate system. the luminance values are recorded and the x,y chromaticities (from (4)) are transformed to their equivalents in the MacAdam domain. Given the orders of equations (5) this would appear to be an expensive procedure, but, as will be described later, a look-up table provides reasonable accuracy.

Histogram and cluster analysis is performed, the k-means test providing the appropriate codewords for the Y and ξ,η vectors. Because the luminance and chromaticity data are quantised independently, it is possible that their *RGB* equivalents may be outside the bounds of the input signal, so a check is performed and, where necessary, codewords are reassigned to a ***nearest neighbour*** in the chromaticity space whose $Y\xi\eta$ value is known to lie within the bounds of the RGB space.

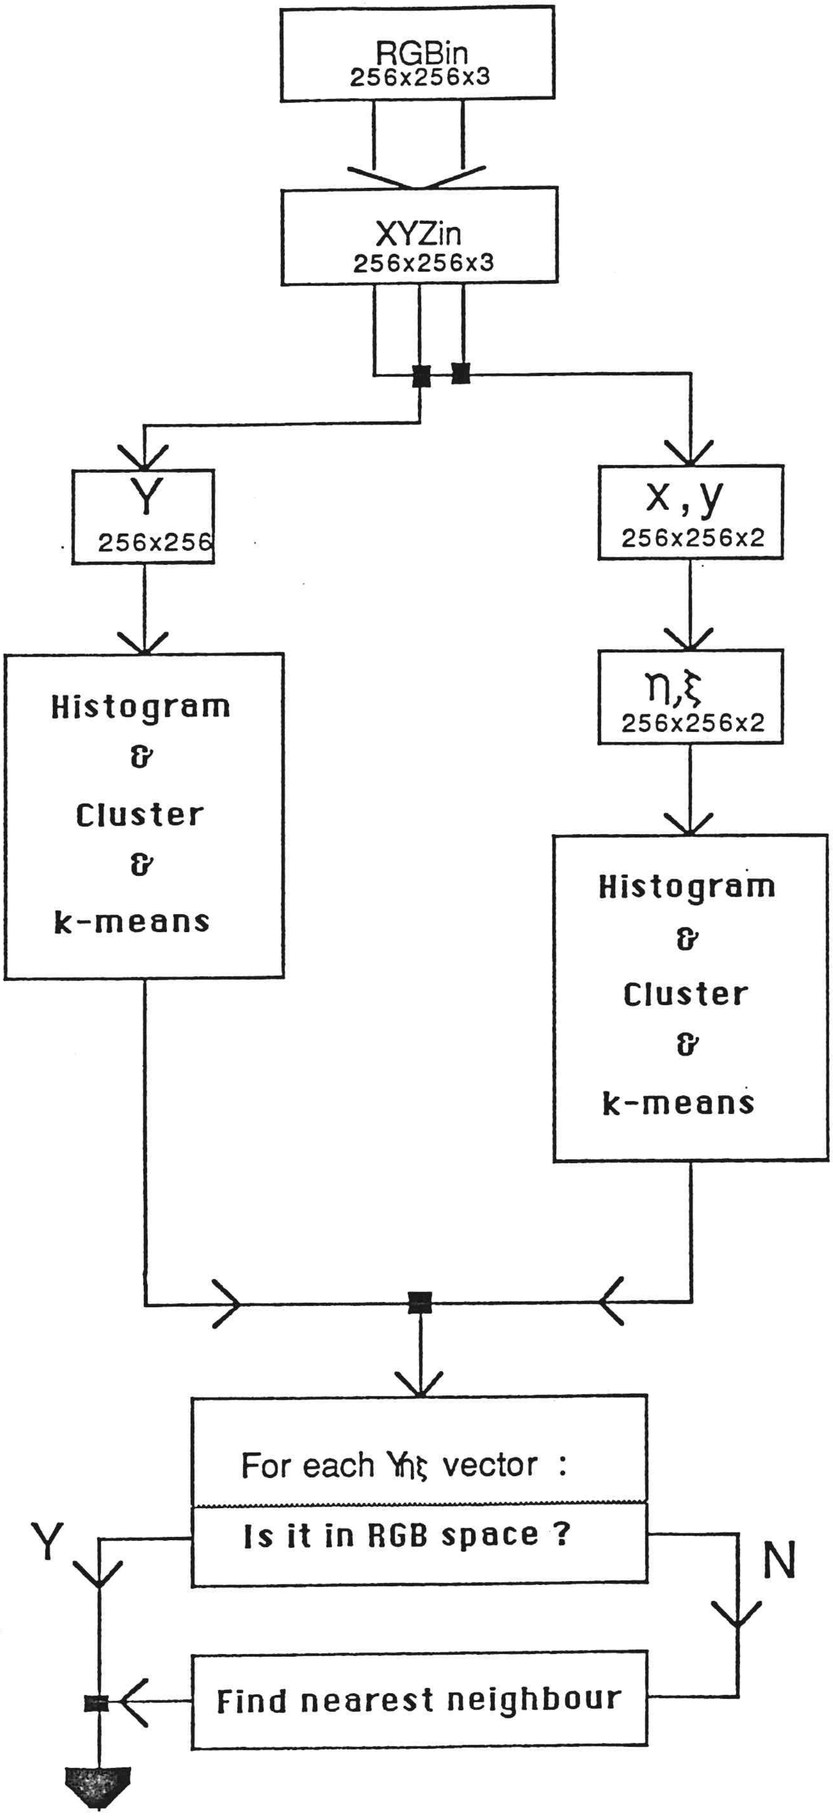
RGBin
256x256x3
XYZin
256x256x3
Y
256x256
x , y
256x256x2
η,ξ
256x256x2
Histogram
&
Cluster
&
k-means
Histogram
&
Cluster
&
k-means
For each Yηξ vector :
Is it in RGB space ?
Y
N
Find nearest neighbour

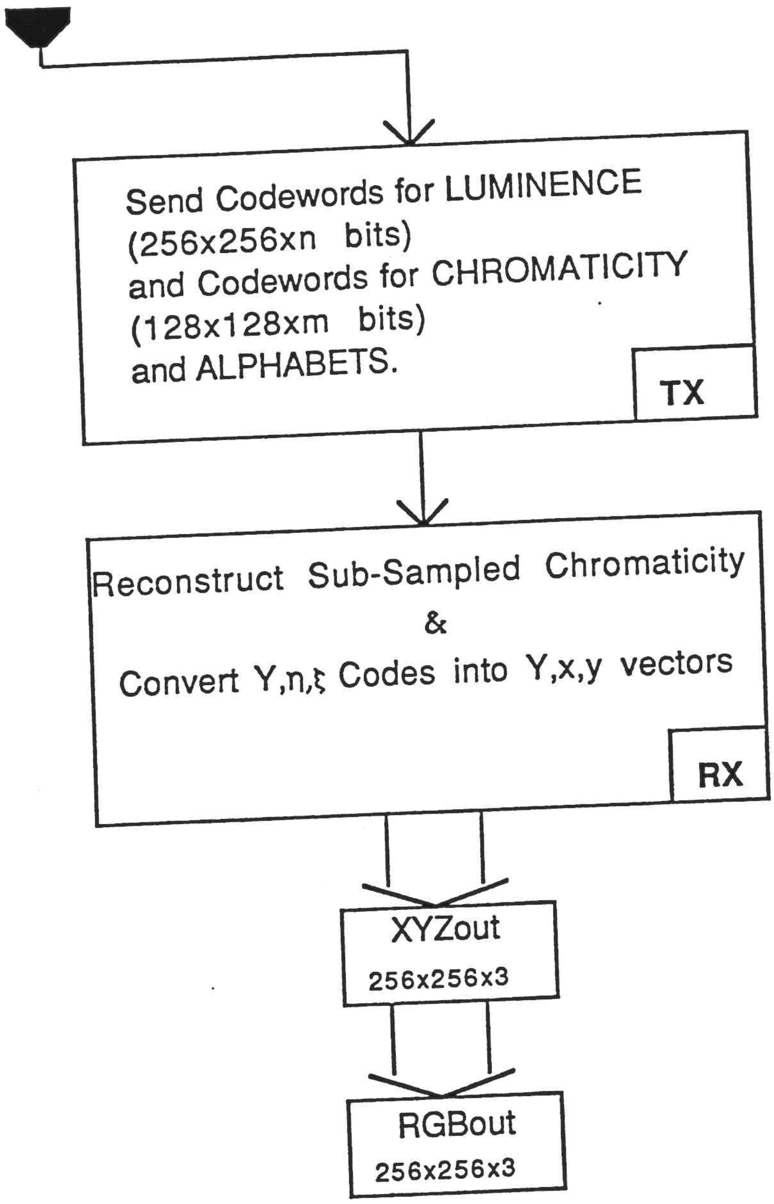

Figure 6. Flow Graph of the Quantisation process.

The *transmitter* represents the interface to the storage medium, which will store the codeword alphabets and the codewords for each pixel (only every fourth chromaticity codeword is transmitted). At the *receiver* the subsampled chromaticity data is reconstructed by an interpolative means, which will be described later.

At this stage in the process every pixel has a luminance and chromaticity value. The ξ,η vectors are transformed back to their x,y equivalents from the inverse look-up table and the quantised XYZ vectors may be derived by the following :

$$(X+Y+Z)_q = \frac{Y_q}{y_q} \; ; \; X_q = x_q(X+Y+Z)_q \; ; \; Z_q = (1-x_q-y_q)(X+Y+Z)_q \; . \tag{10}$$

The subscript q indicates a quantised value. Finally, the inverse of the matrix in (3) is used to transform XYZ_q into RGB_q vectors in the reconstructed image. The perceptual distortion introduced by the process may be measured by visually comparing the reconstructed image with the original.

4.2. Chromaticity Domain Conversion

To transform every pixel of the image to its chromaticity in the MacAdam domain, using equations (5), is a computationally expensive procedure. An inverse transform has not been identified (by the authors). It is therefore necessary that some sort of look-up table (LUT) approach be formulated to reduce the processing overhead and provide a valid inverse transform to the x,y chromaticity.

Consider the chart in figure 7 and compare this with figure 8 [10]. The latter displays the true colour limits of the television receiver (dotted lines) and the gamut of colours obtained from the commonly found dyes and pigments. Instead of varying between 0 and 1.0, the following restrictions apply to the realisable x,y chromaticity chart :

$$0.14 \le x \le 0.66 \; ; \; 0.07 \le y \le 0.72 \; ; \; y \ge 0.5x \; ; \; x+y \le 1.0 \; . \tag{11}$$

These inequalities describe the limits of the indices of a LUT for the forward transformation. Giving the x,y values an integer accuracy (to be the indices in the LUT) with a character range (0 - 255) gives the number of locations as

$$(0.0546 + 0.0735 + 0.0676) \times 256 \times 256 \;=\; 12825 \; \textit{locations} \; . \tag{12}$$

Some sort of offset will be subtracted from the actual chromaticities to make them valid indices. Note that the maximum error in the chromaticity values is ± 0.0039 in this case.

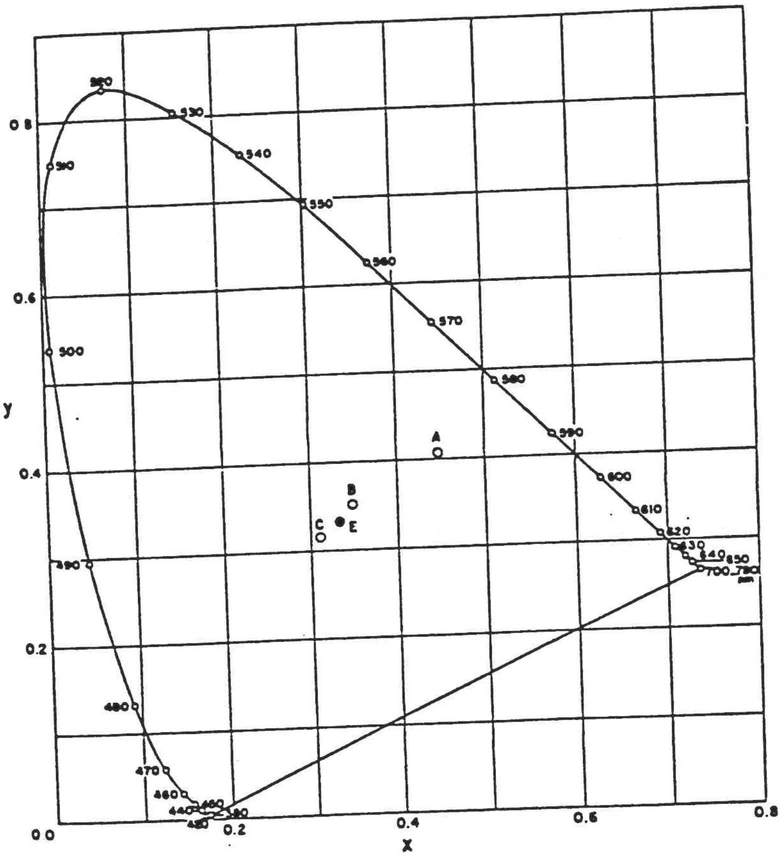

Figure 7. Visual Wavelengths (nm) and their Locations on the CIE Chromaticity Chart. Source [21], Fig. 3.10.

With the MacAdam domain, assuming again integer accuracy, the range of each ordinate is

$$720 \leq \xi \leq 1040 \,;\, 240 \leq \eta \leq 480 \,. \tag{13}$$

Unfortunately, the range of ξ is greater than 256. Given the approximation performed by rounding the x,y data, an associated rounding will be reflected in the MacAdam coordinates. Halving the resolution in the MacAdam domain introduces a further approximation, but it has been shown that a distance of one unit on the MacAdam chart is likely to be indistinguishable [10].

With a halved resolution, the number of locations in the inverse transform LUT is

$$160 \text{ x } 120 = 19200 \,. \tag{14}$$

The base of the areas of interest in each domain is not the origin, so offsets can be retained and applied to each index, so that the values contained within each location in both LUTs will be a pair of characters. The total space occupied by the LUTs will therefore be

$$(12825 + 19200) \text{ x } 2 = 64050 \;\; bytes \; \approx \; 64 \;\; kbytes \;, \tag{15}$$

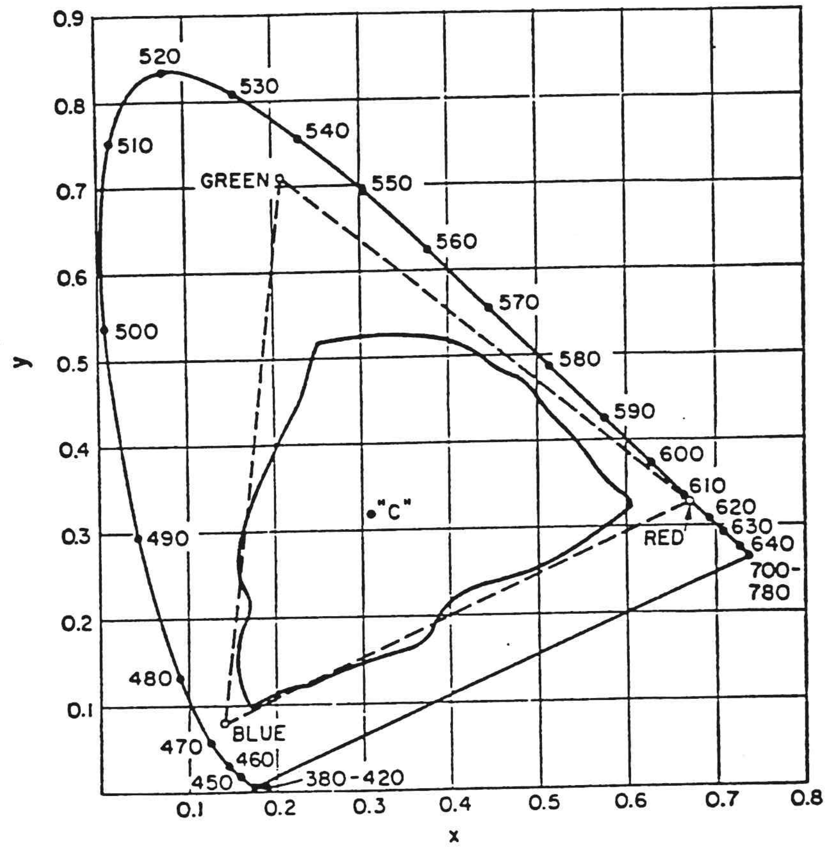

Figure 8. Chart showing RGB colour space (dotted triangle) and typical extent of pigments and dyes under Illuminant C. Source [22], Fig. 4.7.

which is small compared with the typical image size.

A conversion routine may be created to produce the look-up tables as they are required. The inverse transform LUT is created simultaneously with the forward transform. Each transformed coordinate pair act as entries in the first LUT and indices in the second, the entries in the second being the indices in the first.

4.3. Sub-Sampling and Reconstruction of Chromaticity Data

As previously mentioned, the chrominance bandwidth in the television signal is much smaller than that of the luminance. Because the eye's response to spatial changes in the colour of light is less acute than to changes in the brightness of the light [18], all the television standards band-limit the chrominance as far as is possible to reduce transmission costs. Often, only one chromatic variable is transmitted per line (e.g SECAM [6]). It therefore seems reasonable to reduce the chromaticity resolution in our scheme. In fact, the deterioration is not discernible.

The quantisation scheme used in this report provides the full resolution of luminance values, but only every fourth pixel has its chromaticity component stored. At the receiver, or reconstruction stage, the pixels are arranged as shown in figure 9. Only every fourth pixel has both luminance and chromaticity (black circles), while every pixel has a luminance value, which may be used to predict the chromaticity component it should possess. It is not sufficient to simply provide the average of the surrounding chromaticities because if the pixel being assigned has a luminance value higher than the surrounding ones, it may be assigned outside the *RGB* solid. A modified interpolation is used as follows.

Calculate the average chromaticities ξ_{ave}, η_{ave} of the pixels on either side, vertically, horizontally or diagonally, of the pixel to be assigned. The luminance value of this pixel is used to determine which codebook to examine (there may be many, see section 5). Choose the codeword vector closest to the average chromaticities. In this way, a valid chromaticity vector is assigned to the pixel, reflecting the effect of the brightness of that pixel.

To reduce the time consuming search through the codebook, some sort of ordering of the codewords may be used, or a hash-type function may be used to automatically select the correct vector.

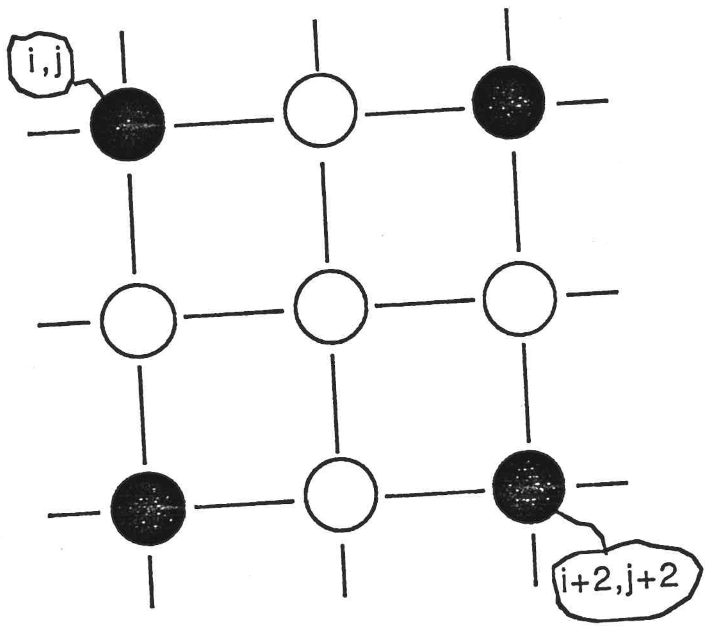

Figure 9. Pixel Arrangement at reconstruction. Black Circles - Luminence & Chromaticity Codewords Assigned. White Circles - Only Luminence Assigned.

4.4. Sources of Error

The implementation of the quantisation scheme described uses only integer arithmetic. This makes the program faster and more portable. Some error is obviously introduced, due to rounding. The main sources are those operations requiring real arithmetic, such as the following :

$$RGB \rightarrow XYZ\ ;\ XYZ \rightarrow Yxy\ ;\ xy \rightarrow \xi\eta \rightarrow \xi/2, \eta/2\ ;$$

$$\xi/2, \eta/2 \rightarrow xy\ ;\ xyY \rightarrow XYZ\ ;\ XYZ \rightarrow RGB\ . \tag{16}$$

The transformation of (3) is modified for integer arithmetic by multiplying each coefficient by 4096 and then right shifting the result 12 times. The maximum error in any row is 1.38×10^{-4}. For the maximum signal (255,255,255) the error is therefore less than 0.04, so the integer result may vary by 1. With the inverse transformation, the maximum error is 1.81×10^{-4}, thus the maximum error is less than 0.05. The true value may be different by 1.

Rounding the chromaticity values to the nearest character value (0 - 255) produces a maximum error of 0.0039. The maximum error reflected in the MacAdam domain (with reduced resolution) is 1 in each coordinate. This figure takes rounding in the MacAdam domain into account as well. The error induced in the $x_q, y_q \rightarrow \xi, \eta \rightarrow x_q, y_q$ transforms is zero because the indices and entries correspond exactly. The process of converting the integral chromaticity values back to fractions again introduces a maximum error of 0.0039.

Some errors will occur in the clustering process because, at each stage of the iteration process, the centroids will not always have integral values. This is not as important as the resulting error in the choice of the set boundaries, which are related to the error in the calculation of Euclidean distances. The distance calculation itself simply involves squaring and summing. The error is due to the fact that integers rather than real numbers are being squared.

4.5. Three Variations on Codebook Construction

Three variations on the diagram in figure 6 were tried. The basic scheme was modified in attempts to produce a better colour match, by considering local image variation. In all three procedures, the luminance codebook was derived from a histogram of the whole image.

Method I - Create one codebook for chromaticity, from a histogram of the entire image.

Method II - Create n chromaticity codebooks, where n is the number of luminance codewords, from n histograms of those pixels assigned the same luminance codeword. This method was intended to adapt to brightness-dependent chromatic variation.

Method III - Create n chromaticity codebooks, where n is also the number of blocks in the image, from n histograms of those pixels in each block. This method was intended to adapt to position-dependent chromatic variation.

Each modification gave better results than for method I, as will be seen later. The main interest in this report is in colour (or chromatic) compression, so no effort was made to take account of local effects in the luminance plane, as methods to do so appear widely in the literature.

5. Results

5.1. The Datasets

Three images were used as test data. They are called "Miss America", "Split Screen" and "Trevor" and are frames from digitised video sequences used for testing real-time image coding algorithms. The sequences were chosen by the CCITT committee on coding standards because of their content, such as textures, complex motion etc. The research covered in this report is connected with low bit-rate videoconferencing, hence these data sets. The original images named above are shown in figures 10, 11 and 12, after the references.

5.2. Reproduction of Results

Gauging the success of the quantisation schemes described earlier requires colour photographs, which are difficult to produce from digitised images. Some attempt has been made to indicate the relative success of each of the methods by providing an associated distortion measure. One of the disadvantages of working with images is that a calculable distortion measure which corresponds with subjective appraisal is difficult to establish [14]. A restricted distortion measure is shown, simply between original and reconstructed chromaticities in the MacAdam space. The luminance quantisation method for each of the three methods is the same so will not feature in the distortion measures.

In the following presentation of results, each image has been compressed using the following quantisations:

Luminance Codebook - 16 symbols (4 bits per pel);

Chromaticity Codebook - 8 symbols (3/4 bits per pel).

This was the limit of acceptable compression for the "best" method" with the "best" image, and in fact is too high for the "Split_Screen" image, as will be seen.

5.3. Method I

Figures 13, 14 and 15 show reconstructed images using the first method. The worst of the three is by far the "Miss America" frame. The variety of subtle flesh tone on the face are lost in the clustering process. The false contours created in the "Split_Screen" image is due to the luminance quantisation. Using a codebook of 32 symbols rather than 16 removes the countouring problem. The lack of texture and variable light intensities incident on the curtain behind the figures causes a problem on this image alone.

5.4. Method II

The aim of this method was to improve the chromaticity quantisation of the faces of all the figures. Figures 16, 17 and 18 show the reconstructed images. "Miss America" has a more colourful face and so does "Trevor". Unfortunately, "Trevor's" handkerchief has not been reconstructed properly. This is due to the extreme chromaticity value of his red handkerchief being assigned to a centroid which is a more neutral colour. In fact, there are few pixels in the image with chromaticities in the red end of the spectrum, and thus it is impossible to force the clustering algorithm to pick out a small group such as that which contains the chromaticities of the tie. Overall, the results are more appealing compared with those from method I.

5.5. Method III

The aim of this method was to improve the chromaticity quantisation of local events such as "Trevor's" handkerchief, with the same overhead as method II. Figures 19, 20 and 21 show the results and we can see that the handkerchief has been successfully returned to its correct colour, but at the expense of facial tones, which are not as close to the original as those of method II.

5.6. A Distortion Measure

The Mean Square Errors (MSE) in the MacAdam chromaticity data for images reproduced in the preceding figures are shown in table 1. The square root of each of these numbers represents the average distance between original and reconstructed vectors in MacAdam coordinates (half resolution). Methods II and III in general perform better than Method I, but in the case of "Trevor" there is little improvement. Notice that, subjectively, figure 21 is closer to the original than figures 15 or 18, because the handkerchief is a striking feature in the image, but is not statistically very prominent. The trend in distortion figures does, however, agree with subjective appraisal of "Miss America". It is not appropriate, then, to use such a distortion figure as an absolute benchmark of reconstruction accuracy.

6. Summary and Conclusions

Three variations on a scheme for quantising digitised colour images have been presented. The modifications to the basic algorithm have been shown to provide perceptually better results, at the (slight) expense of processing and codebook overheads. The limit on compression ratio is about 6 to 1, i.e 4.75 bits per pel, but false contouring will be noticeable on some images due to the

Mean Square Error			
Method	Miss_America	Split_Screen	Trevor
I	53	15	5
II	15	7	6
III	12	4	4

Table 1. MSE Associated with each Quantisation Method.

luminance quantisation. This compression ratio compares favourably with those reported in the literature [19,20].

The strengths and weaknesses of the quantisation algorithm have been discussed in previous sections, but let us consider again the main source of distortion in the scheme which is the clustering algorithm. Although the algorithm is considered a robust one, the starting conditions do have an effect on the resulting codebook. A number of techniques were tried, to produce a set of starting centroids which were well separated and which guaranteed a convergence. The most reliable way turned out to be that suggested by Linde, Buzo and Gray [16], in which vectors drawn randomly from the input stream form the starting codewords. The clustering process is a statistical one and so extreme chromaticity vectors tend to be drawn further to the achromatic centre. In this way, the more unlikely a colour, such as "Trevor's" handkerchief, the more it will be distorted at reconstruction. Equal numbers in each partition would perhaps give a more pleasing result, but the perceptual and statistical requirements of the scheme, i.e.

A. Minimise the perceptual distortion

and

B. Minimise the average distortion

would seem to be conflicting. Some sort of knowledge-based clustering is called for, where a knowledge of perceptually important events is given some sort of priority. Alternatively, a non-Euclidean partition criterion may be used. This is an area of active research.

Further investigation into more efficient luminance quantisation would undoubtedly increase the compression ratio possible from this scheme. As stated previously, the main interest was in compressing the chromatic data as far as possible. In this respect, the success of the proposed scheme has been significant.

7. Acknowledgements

This research was carried out under the collaborative Alvey Project MMI085, "Image Analysis and Coding". The authors express their gratitude to the Alvey Directorate for making these funds available.

8. References

[1] Young, T., *Lectures on Natural Philosophy*, vol. 1, p440, 1807.

[2] von Helmholtz, *Treatise on Physiological Optics* (The Optical Society of America) vol. 2 pp 145, 235, 240, 412, 1924.

[3] Judd, D. B., "Fundamental Studies of Color Vision From 1860 to 1960," *Nat. Acad. of Sci.*, vol. 55, no. 6, pp 1311-1330, June 1966.

[4] Hering, E., *Zur Lehre vom Lichtsinn*, Berlin, 1878.

[5] Hurvich, L. M. and Jameson, D., "An Opponent-Process Theory of Color Vision," *Psychological Review*, vol. 64, no. 6, pp 384-404, 1957.

[6] Pratt, W., *Digital Image Processing*, Wiley, New York 1983.

[7] Wright, W. D., *The Measurement of Colour*, Hilger, London 1969.

[8] MacAdam, D. L., "Projective Transformations of ICI Color Specifications," *J. Opt. Soc. Am.*, vol. 27, no. 8, pp 294-299, August 1937.

[9] MacAdam, D. L., "Geodesic Chromaticity Diagram Based on Variances of Color Matching by 14 Normal Observers," *Applied Optics*, vol. 10, No. 1, pp 1-7, January 1971.

[10] Limb, J. O., Rubenstein, C. B. and Thompson, J. E., "Digital Coding of Color Video Signals - A Review," *IEEE Transactions on Communications*, vol. COMM-25, no. 11, pp 1349-1384, November 1977.

[11] Hecht, S., "The Visual Discrimination of Intensity and The Weber-Fechner Law," *J. Gen. Physiol.*, vol. 7, p 241, 1924.

[12] Cornsweet, T. N., *Visual Perception*, Academic Press, New York 1970.

[13] Stockham, T. G., "Image Processing in the Context of a Visual Model," *Proc. IEEE*, vol. 60, no. 7, pp 828-842, July 1972.

[14] Faugeras, O. D., "Digital Color Image Processing Within the Framework of a Human Visual Model," *IEEE Trans. Accoustics, Speech and Sig. Proc.*, vol. ASSP-27, no. 4, pp 380-393, August 1979.

[15] LLoyd, S. P., *Least squares estimation in PCM*, Bell Laboratories Technical Note, 1957.

[16] Linde, Y., Buzo, A. and Gray, R. M., "An Algorithm for Vector Quantizer Design," *IEEE Trans. Comms.*, vol. COMM-28, no. 1, pp 84-95, January 1980.

[17] MacQueen, J., "Some Methods for Classification and Analysis of Multivariate Observations," *Proc. of the Fifth Berkeley Symp. on Math., Stat and Prob.*, vol. 1, pp 281-296, 1967.

[18] McIlwain, K., "Requisite Color Bandwidth for Simultaneous Color-Television Systems," *Proc. IRE*, vol. 40, pp 909-912, Aug. 1952.

[19] Kurz, B. J., "Optimal Color Quantization For Color Displays," *Proc. IEEE Computer Soc. Conf. on Computer Vision and Pattern Recognition*, pp 217-224, Washington DC, USA, 19-23 June 1983.

[20] Stevens, R. J. and Preston, F. H., "Data Ordering and Compression of Multispectral Images Using the Peano Scan," *IEE Int. Conf. on Electronic Image Proc.*, pp 194-198, York, England, 26-28 July 1982.

[21] Wyszecki, G. and Stiles, W. G., *Color Science*, Wiley, New York 1967.

[22] Fink, D. G. ed, *Color Television Standards : Selected Papers and Records of the National Television Systems Committee*, McGraw-Hill, New York 1955.

Figures 10, 11 and 12. Original Images of "Miss America", "Split Screen" and "Trevor".

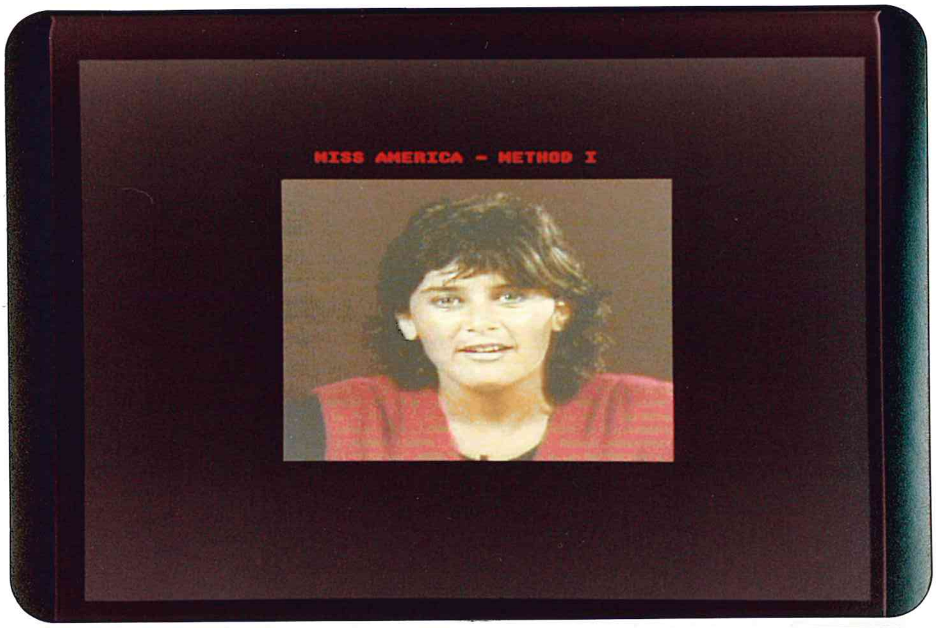

Figures 13, 14 and 15. Reconstructed Images using Method I.

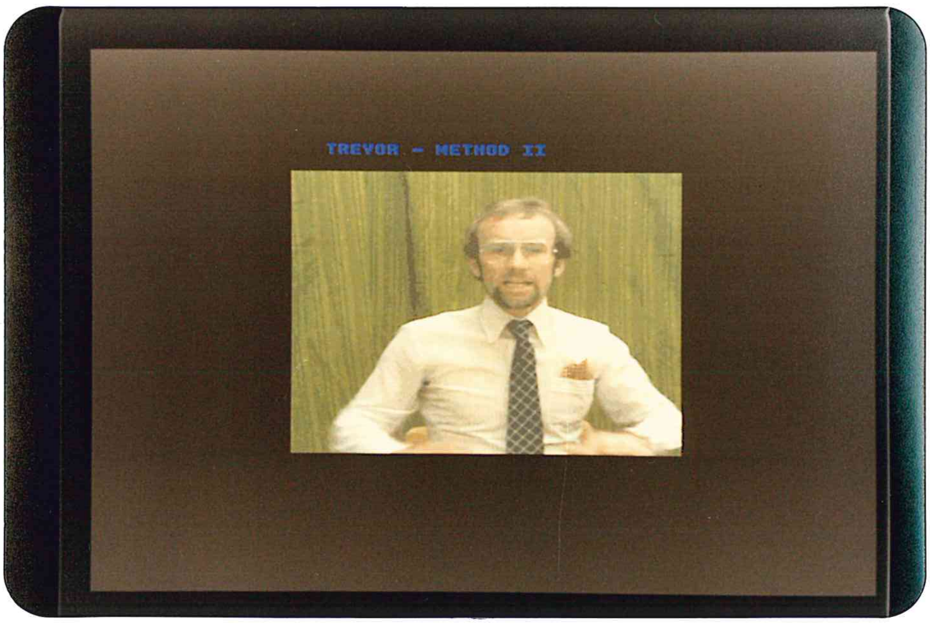

Figures 16, 17 and 18. Reconstructed Images using Method II.

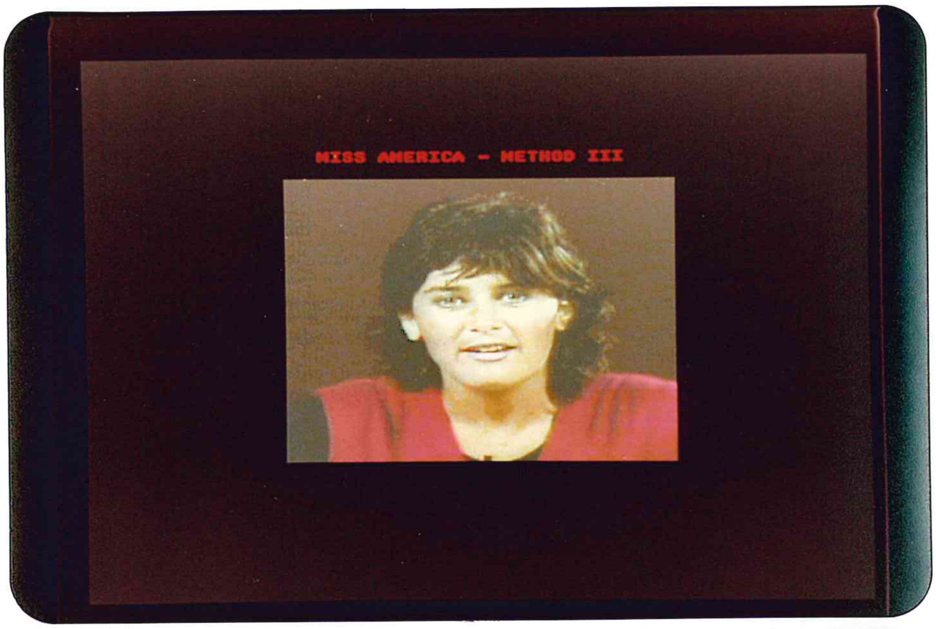

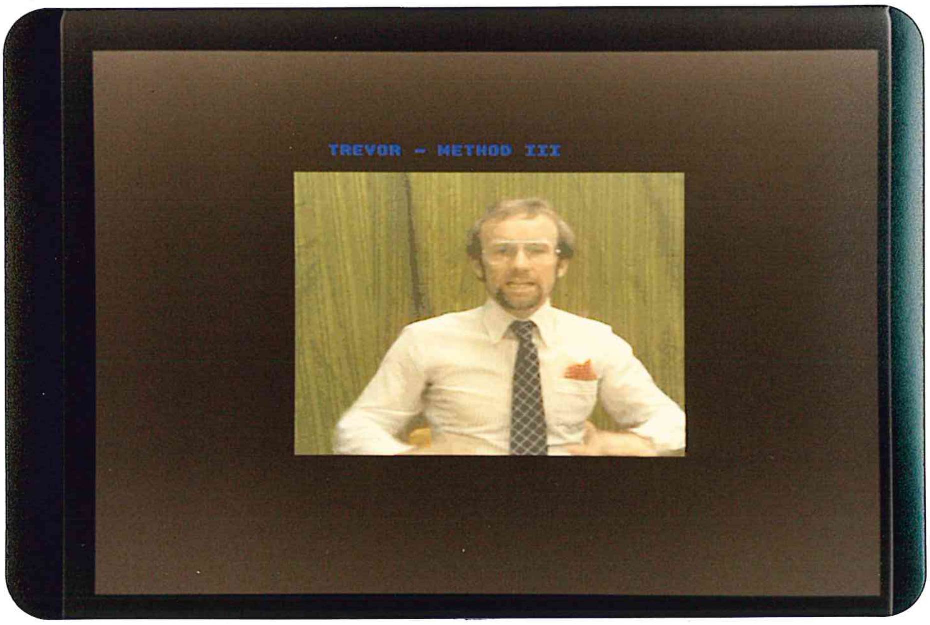

Figures 19, 20 and 21. Reconstructed Images using Method III.