Modeling Colour Refinement for Progressive Image Coding

Anthony J. Maeder and Daniel M. Bell
School of Electrical and Electronic Systems Engineering
Queensland University of Technology, Brisbane, Australia

Abstract

Progressive compression schemes allow successive refinement of image contents during reconstruction. Here an approach for modeling colour refinement of image region interiors is presented. The model is based on simple colour variations related to illumination, surface shape and texture.

Introduction

Image compression schemes aim to represent image data more compactly than specifying the individual pixel values, for efficient storage or transmission purposes. Many different philosophical approaches to this challenging problem have been explored, all of which attempt to model higher level structure in the image data. Some models exploit statistical redundancy and others psychovisual redundancy[1, 2]. It has become widely accepted that conventional statistical modeling for compression is far from optimal and so much attention has been directed towards 'second generation' techniques[3], which attempt to encode information for large spatial features or regions of the image separately.

Sometimes it is useful to represent an image using a constructive sequence of data values providing increasingly detailed information. This structure allows partial reconstruction of the image, using comparatively little data, to produce a simple version of the scene for early visual recognition and subsequent further refinement. This 'progressive compression' capability can be useful in many different image viewing situations, such as image database queries, telebrowsing, teleconferencing, low bandwidth channels, reduced resolution display devices.

Effective progressive modeling of image contents can be derived from knowledge of human psychovisual attributes. In such an approach, dominant scene contents over the whole image are extracted and encoded first, giving a coarse resolution representation of the image. These contents are associated with information identified by our pre-attentive or early vision, such as large relatively homogeneous regions, adjacent contrasting areas and distinctive boundary shape properties. More complex visual properties are encoded separately to allow successive refinements of the reconstructed image. These properties include subtleties of texture, shading, colour, small regions and small boundary features.

We have recently proposed a general progressive compression technique[4] which proceeds by first representing major regions approximately (e.g. as polygons of constant interior) followed by refinements of both region boundary and interior details (e.g. by triangular boundary extension and runlength encoded interior improvement). The work described in this paper is related to the choice of an early refinement strategy for the interior variations, in particular for pixel colour. Various refinements for the overall changes in colour that mimic different stages of visual cognition of increasing attentiveness are proposed.

Method

The influence of human perception of colour on our understanding of visual scenes has been widely studied, along with many other psychovisual phenomena such as brightness, textures, edges. Of these, colour has been hardest to model in general terms and consequently colour approximation techniques tend to be based on clustering or colour space subdivision approaches (e.g. Heckbert's method[5]). An alternative approach has been to model the physical world characteristics giving rise to the scene, with potentially vast computational expense (e.g. ray tracing).

The first stage in the image representation consists of choosing a set of prescribed colours for the coarsest progressive image version. Fast discrimination of colours occurs best when well-separated hues are used, particularly if the hues are primary or secondary in nature[6, 7]. Adapting this rule of thumb to suit colour representation, the primary or secondary hue closest to the average hues of colours being represented for a region can be used. Thus six prescribed hues were selected for use in the initial representation.

Lightness has as influential an effect on colour discrimination as hue, so a set of seven equidistant lightness levels was selected in an analogous way. Saturation matters least, so only three saturation levels (0.0, 0.5 and 1.0) were selected. In addition, two values were specified to represent pure white and pure black. All initial colours were therefore able to be represented using exactly 7 bits in this scheme. Fewer prescribed colours could have been adopted, at a cost of less flexibility in modeling those parts of scenes where colours were comparatively close.

For the second stage of representation, the regions were approximated by more accurate versions of the colours, obtained by computing the average hue, lightness and saturation for each region and specifying these values accurately. At most this would require as many new colours to be added to the colour table as the number of distinct regions to be represented in the coarsest version of the image. The regions appear more realistic in colour but uniform in appearance at this stage.

The third stage needs to represent the dominant changes in colours which model the non-uniformity in colour within each region. These changes are caused by illumination and by surface curvature, prompting shadows and highlights. Under many typical conditions, the changes are fairly smooth and so can be modeled well by low order polynomials. Linear interpolation is unsuitable when the light source is not diffuse and when object surfaces are not flat. A quadratic interpolant was therefore used to adapt the lightness and saturation values independently within each region. The hue values could have been interpolated similarly, but in the results presented here were left fixed at the average value, since the region segmentation had been performed with a preference for hue similarity.

The next stages of colour refinement would involve more detailed correction of the approximating colours through each region. This could be achieved by decomposing the regions into disjoint sub-regions within which colour variations could be modeled more accurately piecewise using the technique described above. It could also be achieved by incorporating corrections at neighbourhood and pixel resolution to the colours used.

Some structure or ordering may need to be imposed on which of the above types of corrections would be undertaken first, but this decision is not crucial as the major information of visual importance would already be present. Appreciation of further detail by the viewer at this stage would require strongly directed observation. Higher order models of region colour variations could also be considered to permit representation of texture, for which cognition is also strongly dependent on attentive vision.

The process described above necessitates the formation of a sequence of nested colour subsets of the representation colour space, to allow easy selection of the colours at successive stages of reconstruction. The subsets at each stage after the first can be derived from those of the previous stage and the appropriate modeling parameters for the types of changes represented by the next stage. This is a departure from the conventional approach to subsetting, where a bottom-up merging or similar process is applied to the full range of actual colours, a very expensive process computationally.

Results

Computational experiments were performed using two standard test images, Lenna and Peppers, each 256×256 pixels of 24 bit colour. Lenna is a head-and-shoulders scene with simple illumination and a restricted range of hues, but with subtle variations in flesh tones. Peppers is a contrived still life scene containing several widely differing hues and many strong artifacts of lighting and surface curvature. Neither scene contains significant regions of strong texture content.

A simple colour segmentation was performed on each image using a region merging scheme developed for natural scenes[8]. The merging thresholds were adjusted to produce roughly the same number of regions for both test images. A second segmentation was then performed to produce roughly double that number of regions for both images. These segmentations provided basic regional descriptions of the image contents for use as the first level of spatial refinement in progressive compression of the image. Next, the multistage progressive colour selection process described above was applied. The initial (stage 1), average (stage 2) and interpolated (stage 3) image versions were computed and displayed for visual appraisal. No spatial progressive refinement was performed on these images, so as to confine visual assessment of the quality to colour aspects only.

Difference images with respect to the original image were computed for each of the three stages. Signal-to-noise ratios (SNR) in decibels for the difference images were calculated using the standard deviation values for the three HLS colour bands in each difference image, compared with those for the corresponding band in the original image. Approximate colour compression rates (CCR) in bits per pixel that were achieved by each stage were calculated from the number of values used to specify the region colouring and to define corresponding new colour table entries where necessary. The compression rates for complete representation of regions require region boundary specifications, which depend on the current level of spatial refinement and so should be considered separately from the region colours.

These quantitative assessments are not a good reflection of the visual quality of the reconstructed image versions, but provide a basis for limited objective comparison. As can be seen in Table 1, the increase in the number of regions led to a slight improvement in the SNR in both cases. The major improvement in SNR was obtained in converting the colours of stage 1 to those of stage 2, and a lesser improvement was obtained from stage 2 to stage 3. The relative rates of change of SNR and CCR indicate that the colour accuracy obtained by introducing more colours using interpolation is more expensive in terms of error per unit compression obtained as using the average colour values.

The effectiveness of this approach should really be judged by controlled human observers in a real-time reconstruction of the image data (e.g. over several frames of a video sequence), since that is the purpose for which it is intended. Both qualitative and quantitative indicators for this form of colour assessment are poorly developed however. In addition, the region boundaries would be undergoing corresponding refinements in accuracy at the same time as the colour improvement in a real application, which would further complicate human cognition of the changes.

Table 1. SNR and CCR results for test images

Image	Regions	Stage	SNR (H)	SNR (L)	SNR (S)	CCR (est)
Lenna	43	3	21.4	9.19	5.43	0.0072
		2	21.4	8.28	5.17	0.0033
		1	18.4	7.64	2.52	0.0006
Lenna	74	3	22.1	10.1	5.92	0.0124
		2	22.1	9.35	5.77	0.0056
		1	18.7	8.33	2.51	0.0011
Peppers	47	3	13.7	8.34	6.34	0.0079
		2	13.7	7.26	5.41	0.0036
		1	13.1	6.68	3.22	0.0007
Peppers	77	3	14.3	9.87	7.10	0.0129
		2	14.3	8.85	6.05	0.0059
		1	13.3	7.95	3.51	0.0012

Conclusion

The approach proposed here for progressive compression of colour in images is based on relating scene contents to simple aspects of increasingly attentive human vision. This contrasts with previous approaches to computational colour approximation, which tend to follow statistical or physical world modeling. The effectiveness of the approach must be established visually at this stage. Conventional quantitative measures indicate that the approach achieves the goal of progressive refinement in the test cases considered.

The opportunity exists to vary the models used at each stage in the process, and to subdivide stages into graded substages of partial refinement, based on properties of both colour and overall image reconstruction. The low order model of colour change adopted for ease of computation has already proven useful in other areas of colour approximation, suggesting that it is a reasonable choice. Higher order models or a general non-linear model based on colour importance[9] during cognition could be used to achieve more rapid convergence if the application warranted such computational expense.

This work was supported by the Australian Research Council and the Australian Space Office through the Space Industry Development Centre programme.

References

1. Netravali, A. N. and Haskell, B. G., Digital pictures: representation and compression, Plenum, 1988.
2. Rabbani, M. and Jones, P. W., Digital image compression techniques, *SPIE Optical Engineering* Press, 1991.
3. Kunt, M. Ikonomopoulos, A. and Kocher, M., Second-generation image-coding techniques, *Proceedings of the IEEE*, **73** (4) 1985, pp. 549-574.
4. Bell D. M. and Maeder, A. J., Applications of region-based progressive image compression, *DICTA-93 2nd APRS Conf. on Digital Image Computing, Techniques and Applications*, Sydney, Australia, 8-10 December 1993, pp. 155-162.
5. Heckbert, P., Color image quantization for frame buffer display, *ACM Computer Graphics (SIGGRAPH 82)* **16** (3) 1982, pp. 297-307.
6. Keller, P. R. and Keller, M. M., Visual cues: practical data visualization, *IEEE Press* 1993, pp. 26-33.
7. Maeder, A. J., Software tools for visual analysis of colour compression techniques, *IEEE Int'l. Workshop on Visual Signal Processing and Communications,* Melbourne, Australia, 20-22 September 1993, pp. 97-100.
8. Choo, A. P. Maeder, A. J. and Pham, B., Image segmentation for complex natural scenes, *Image and Vision Computing* **8** (2) 1990, pp. 155-163.
9. Maeder, A. and Pham, B., A colour importance measure for colour image analysis, I*S&T/SID Color Imaging Conference*, Scottsdale, AZ, USA, Nov. 1993, pp. 232-237.