A hybrid Differential Evolution approach to colour map generation

Gerald Schaefer and Lars Nolle
School of Computing and Informatics, Nottingham Trent University
Nottingham, United Kingdom
{gerald.schaefer,lars.nolle}@ntu.ac.uk

Abstract— **Differential Evolution is an optimisation technique that has been successfully employed in various applications. In this paper we apply Differential Evolution to the problem of generating an optimal colour map for colour quantised images. The choice of entries in the colour map is crucial for the resulting image quality as it forms a look-up table for all pixels in the image. We show that Differential Evolution can be effectively employed as a method for deriving the entries in the map. In order to optimise the image quality our Differential Evolution approach is combined with a local search method that is guaranteed to find the local optimal colour map. This hybrid approach is shown to outperform various commonly used colour quantisation algorithms on a set of standard images.**

Keywords: Differential Evolution, colour map, colour quantisation, k-means clustering, hybrid optimisation

I. Introduction

The colours in digital images are typically represented by three 8-bit numbers representing the responses of the red, green., and blue sensors of the camera. Consequently, an image can have up to 2^{24} i.e. more than 16.8 million different colours. However, when displaying images on limited hardware such as mobile devices but also for tasks such as image compression or image retrieval [5] it is desired to representing images with one a limited number of different colours. Clearly, the choice of these colours is crucial as it determines the closeness of the resulting image to its original and hence the image quality. The process of finding such a palette or map of representative colours is known as colour map generation or colour quantisation. It is also know to constitute an np-hard problem [2]. In the image processing literature many different algorithms have been introduced that aim to find a palette that allows for good image quality of the quantised image [2], [1].

In this paper we apply Differential Evolution (DE) to the colour quantisation problem. As DE is a black-box optimisation algorithm it does not require any domain specific knowledge yet is usually able to provide a near optimal solution. We evaluate the effectiveness of our approach by comparing its performance to the results obtained by several purpose built colour quantisation algorithms [2], [1]. The results obtained show that even without any domain specific knowledge our DE based algorithm is able to provide similar image quality as standard quantisation algorithms. In a second step we combine DE with a standard clustering algorithm, k-means, which is guaranteed to find a local minimum. The resulting hybrid algorithm is shown to further improve the effectiveness of the search and to outperform other colour quantisation techniques in terms of image quality of the quantised images.

The rest of the paper is organised as follows: Section II provides the background for optimisation based on Differential Evolution. Section III explains our novel colour quantisation algorithm. Section IV provides experimental results based on a set of standard test images while V concludes the paper.

II. Differential Evolution

Differential Evolution (DE) [4] is a population-based optimisation method that works on real-number coded individuals. For each individual in the current generation $\vec{x}_{i,G}$, DE generates a new trial individual $\vec{x}'_{i,G}$ by adding the weighted difference between two randomly selected individuals $\vec{x}_{r1,G}$ and $\vec{x}_{r2,G}$ to a third randomly selected individual $\vec{x}_{r3,G}$. The resulting individual $\vec{x}'_{i,G}$ is crossed-over with the original individual $\vec{x}_{i,G}$. The fitness of the resulting individual, referred to as perturbated vector $\vec{u}_{i,G+1}$, is then compared with the fitness of $\vec{x}_{i,G}$. If the fitness of $\vec{u}_{i,G+1}$ is greater than the fitness of $\vec{x}_{i,G}$, $\vec{x}_{i,G}$ is replaced with $\vec{u}_{i,G+1}$, otherwise $\vec{x}_{i,G}$ remains in the population as $\vec{x}_{i,G+1}$.

Differential Evolution is robust, fast, and effective with global optimization ability. It does not require that the objective function is differentiable, and it works with noisy, epistatic and time-dependent objective functions.

III. Hybrid Differential Evolution for colour map generation

In this paper we apply Differential Evolution algorithm as described in Section II as novel approach for colour map generation. For colour quantisation the objective is to minimise the total error introduced through the application of a colour map. The colour map C for an image I, a codebook of k colour vectors, should then be chosen so as to minimise the error function

$$\text{error}(C, I) = \frac{1}{\sum_{j=1}^{k} l_j} \sum_{i=1}^{k} \sum_{j=1}^{l_i} ||C_i - I_j|| + p(C, I) \quad (1)$$

with

$$p(C, I) = \sum_{i=1}^{k} \delta a_i, \quad a_i = \begin{cases} 1 & \text{if } l_i = 0 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

Proceedings 20th European Conference on Modelling and Simulation
Wolfgang Borutzky, Alessandra Orsoni, Richard Zobel © ECMS, 2006
ISBN 0-9553018-0-7 / ISBN 0-9553018-1-5 (CD)

where l_i is the number of pixels I_j represented by colour C_i of the colour map, $||.||$ is the Euclidean distance in RGB space, and δ is a constant ($\delta = 10$ in our experiments). The objective function error(C, I) used is hence a combination of the mean Euclidean distance and a penalty function. The penalty function $p(C, I)$ was integrated in order to avoid unused palette colours by adding a constant penalty value to the error for each entry in the map that is not used in the resulting picture.

There is always some variation in error values within a population. This indicates that although DE is able to find fairly good solutions, i.e. solutions from within the region around the global optimum, it rarely exploits that region completely. Therefore, in a second step, we combine the DE approach with a standard k-means clustering algorithm [3] to provide a stacked hybrid optimisation method. K-means clustering is guaranteed to converge towards the local clustering minimum by iteratively carrying out the following two steps:

- Each input vector should be mapped to its closest codeword by a nearest neighbour search.
- The input vectors assigned in each class (i.e. for each codeword) are best represented by the centroid of the vectors in that class.

In this hybridised algorithm the DE component is hence responsible for identifying the region in the search space that will contain the global optimum while the k-means component will then descend into the minimum present in that region.

IV. Experimental results

In our experiments we have taken a set of three standard images commonly used in the colour quantisation literature (*Sailboat, Airplane, and Pool* - see Figure) and applied both our pure DE-based and the hybrid DE colour map generation algorithm with a target palette size of 16 colours.

We have also implemented three popular colour quantisation algorithms so that we can generate comparative results with our new algorithm. The algorithms we have tested were:

- Popularity algorithm [2]: Following a uniform quantisation to 5 bits per channel the n colours that are represented most often form the colour palette.
- Median cut quantisation [2]: An iterative algorithm that repeatedly splits (by a plane through the median point) colour cells into sub-cells..
- Neuquant [1]: A one-dimensional self-organising Kohonen neural network is applied to generate the colour map.

For all algorithms, pixels in the quantised images were assigned to their nearest neighbours in the colour palette to provide the best possible image quality.

The results are listed in Tables I and II, expressed in terms of MSE (the mean-squared error, Table I) given as

$$\mathrm{MSE}(I_1, I_2) = \frac{1}{3nm} \sum_{i=1}^{n} \sum_{j=1}^{m} [(R_1(i,j) - R_2(i,j))^2 + \quad (3)$$

	Sailboat	Pool	Airplane	all
Popularity alg.	8707.41	669.87	1668.06	3681.78
Median cut	409.35	226.85	240.74	292.31
Neuquant	135.42	127.24	97.59	120.08
DE	223.29	89.56	116.63	143.16
hybrid DE	105.13	44.74	45.57	65.15

TABLE I: Quantisation results, given in terms of MSE.

	Sailboat	Pool	Airplane	all
Popularity alg.	8.73	19.87	15.91	14.84
Median cut	22.01	24.57	24.32	23.63
Neuquant	26.81	27.08	28.24	27.38
DE	24.64	28.61	27.46	26.90
hybrid DE	27.91	31.62	31.52	30.35

TABLE II: Quantisation results, given in terms of PSNR [dB].

$$(G_1(i,j) - G_2(i,j))^2 + (B_1(i,j) - B_2(i,j))^2]$$

(where $R(i,j)$, $G(i,j)$, and $B(i,j)$ are the red, green, and blue pixel values at location (i,j) and n and m are the dimensions of the images) and peak-signal-to-noise-ratio (PSNR) defined as

$$\mathrm{PSNR}(I_1, I_2) = 10 \log_{10} \frac{255^2}{MSE(I_1, I_2)} \quad (4)$$

From Table I and II we can see that Differential Evolution clearly outperforms the Popularity and Median cut algorithms while doing slightly worse than Neuquant. In general, it can therefore be judged as providing image quality similar to specialised colour quantisation results.

On the other hand, the hybrid optimisation technique based on a combination of Differential Evolution and k-means clustering introduced in this paper was able to clearly further improve the performance of DE alone and provides the best image quality for all images with a mean PSNR of 30.35 dB, an improvement of about 3 dB over Neuquant, the next best performing algorithm.

In Figure 1 we provide an example of the performance of the different algorithms[1]. Shown is the original *Pool* image together with the images colour quantised by all algorithms. It is clear that the popularity algorithm performs very poorly on this image as it assigns virtually all of the colours in the palette to green and achromatic colours. Median cut is better but still provides fairly poor colour reproduction as most of the colours in the quantised image are fairly different from the original. The same holds true for the images produced by Neuquant. Here the most obvious artefact is the absence of an appropriate red colour in the colour palette. A far better result is achieved by the (pure) DE algorithm although the red is not very accurate either

[1]A colour version of the paper is available at `http://vision.doc.ntu.ac.uk/`.

and the colour of the cue is greenish instead of brown. Clearly the best image quality is achieved by applying our hybrid DE technique. Although the colour palette has only 16 entries all colours of the original image are accurately presented including the red ball and the colour of the billiard cue.

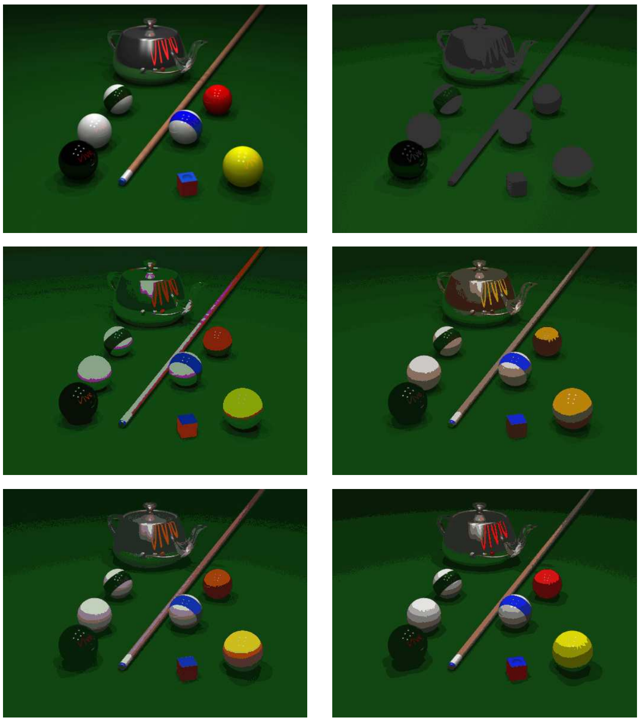

Fig. 1. Part of original *Pool* image (top-left) and corresponding images quantised with (from left to right, top to bottom): Popularity algorithm, Median cut, Neuquant, DE and hybrid DE algorithm.

V. Conclusions

In this work we have applied a hybrid Differential Evolution algorithm to the colour map generation problem. A standard DE approach was combined with a k-means clustering technique where the DE part identifies the region of a good minimum and k-means descends down to the local minimum. Experimental results obtained on a set of common test images have demonstrated that this approach can not only be effectively employed but clearly outperforms dedicated colour quantisation algorithms.

References

[1] A.H. Dekker. Kohonen neural networks for optimal colour quantization. *Network: Computation in Neural Systems*, 5:351–367, 1994.

[2] P. S. Heckbert. Color image quantization for frame buffer display. *ACM Computer Graphics (ACM SIGGRAPH '82 Proceedings)*, 16(3):297–307, 1982.

[3] Y. Linde, A. Buzo, and R.M. Gray. An algorithm for vector quanitzer design. *IEEE Trans. Communications*, 28:84–95, 1980.

[4] K. Price. An introduction to differential evolution. In D. Corne, M. Dorigo, and F. Glover, editors, *New Ideas in Optimization*. McGraw-Hill, 1999.

[5] G. Schaefer, G. Qiu, and G. Finlayson. Retrieval of palettised colour images. In *Storage and Retrieval for Image and Video Databases VIII*, volume 3972 of *Proceedings of SPIE*, pages 483–493, 2000.