

Colour quantisation using the adaptive distributing units algorithm

M. E. Celebi^{*1}, S. Hwang² and Q. Wen³

Colour quantisation (CQ) is an important operation with many applications in graphics and image processing. Most CQ methods are essentially based on data clustering algorithms one of which is the popular k-means algorithm. Unfortunately, like many batch clustering algorithms, k-means is highly sensitive to the selection of the initial cluster centres. In this paper, we adapt Uchiyama and Arbib's competitive learning algorithm to the CQ problem. In contrast to the batch k-means algorithm, this online clustering algorithm does not require cluster centre initialisation. Experiments on a diverse set of publicly available images demonstrate that the presented method outperforms some of the most popular quantisers in the literature.

Keywords: Colour quantization, Clustering, K-means, Competitive learning

Introduction

True-colour images typically contain thousands of colours, which makes their display, storage, transmission, and processing problematic. For this reason, CQ is commonly used as a pre-processing step for various graphics and image processing tasks. In the past, CQ was a necessity due to the limitations of the display hardware, which could not handle over 16 million possible colours in 24-bit images. Although 24-bit display hardware has become more common, CQ still maintains its practical value.¹ Modern applications of CQ in graphics and image processing include: compression,² segmentation,³ text localisation/detection,⁴ colour-texture analysis,⁵ watermarking,⁶ non-photorealistic rendering,⁷ and content-based retrieval.⁸

The process of CQ is mainly comprised of two phases: palette design (the selection of a small set of colours that represents the original image colours) and pixel mapping (the assignment of each input pixel to one of the palette colours). The primary objective is to reduce the number of unique colours, N' , in an image to K ($K \ll N'$) with minimal distortion. In most applications, 24-bit pixels in the original image are reduced to 8 bits or fewer. Since natural images often contain a large number of colours, faithful representation of these images with a limited size palette is a difficult problem.

CQ methods can be broadly classified into two categories:⁹ image-independent methods that determine a universal (fixed) palette without regard to any specific image,^{10,11} and image-dependent methods that

determine a custom (adaptive) palette based on the colour distribution of the images. Despite being very fast, image-independent methods usually give poor results since they do not take into account the image contents. Therefore, most of the studies in the literature consider only image-dependent methods, which strive to achieve a better balance between computational efficiency and visual quality of the quantisation output.

Numerous image-dependent CQ methods have been developed over the past three decades. These can be categorised into two families: pre-clustering methods and post-clustering methods.¹ Pre-clustering methods are mostly based on the statistical analysis of the colour distribution of the images. Divisive pre-clustering methods start with a single cluster that contains all N' image colours. This initial cluster is recursively subdivided until K clusters are obtained. Well-known divisive methods include median-cut,¹² octree,¹³ variance-based method,¹⁴ binary splitting,¹⁵ greedy orthogonal bipartitioning,¹⁶ centre-cut,¹⁷ and rwm-cut.¹⁸ More recent methods can be found elsewhere.^{19–23} On the other hand, agglomerative pre-clustering methods^{24–28} start with N' singleton clusters each of which contains one image colour. These clusters are repeatedly merged until K clusters remain. In contrast to pre-clustering methods that compute the palette only once, post-clustering methods first determine an initial palette and then improve it iteratively. Essentially, any data clustering method can be used for this purpose. Since these methods involve iterative or stochastic optimisation, they can obtain higher quality results when compared to pre-clustering methods at the expense of increased computational time. Clustering algorithms adapted to CQ include maxmin,^{29,30} k-means,^{31–35} k-harmonic means,³⁶ competitive learning,^{37,38} fuzzy c-means,^{39–45} rough c-means,⁴⁶ BIRCH,⁴⁷ and self-organising maps.^{48–53}

In this paper, we adapt Uchiyama and Arbib's adaptive distributing units (ADU) algorithm⁵⁴ to the CQ problem. The rest of the paper is organised as

¹Department of Computer Science, Louisiana State University, Shreveport, LA, USA

²Department of Computer Science, University of Illinois, Springfield, IL, USA

³School of Computer Science and Engineering, University of Electronic Science and Technology of China, Chengdu, China

*Corresponding author: M. E. Celebi, Department of Computer Science, Louisiana State University, Shreveport, LA, USA; email: ecelebi@lsus.edu

follows. The section on ‘BATCH VS ONLINE CLUSTERING ALGORITHMS’ describes the conventional batch k-means and ADU algorithms. The section on ‘EXPERIMENTAL RESULTS AND DISCUSSION’ presents the experimental setup and compares the proposed method to other CQ methods. Finally, the section on ‘CONCLUSION’ gives the conclusions.

Batch vs online clustering algorithms

Batch k-means algorithm

Given a data set $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \in \mathbb{R}^D$, hard partitional algorithms divide \mathcal{X} into K exhaustive and mutually exclusive clusters $\mathcal{S} = \{S_1, S_2, \dots, S_K\}$, $\bigcup_{k=1}^K S_k = \mathcal{X}$, $S_i \cap S_j = \emptyset$ for $1 \leq i \neq j \leq K$. These algorithms usually generate clusters by optimising a criterion function. The most intuitive and frequently used criterion function is the sum of squared error (SSE) given by

$$SSE = \sum_{k=1}^K \sum_{\mathbf{x}_i \in S_k} \|\mathbf{x}_i - \mathbf{c}_k\|_2^2 \quad (1)$$

where $\|\cdot\|_2$ denotes the Euclidean (\mathcal{L}_2) norm and $\mathbf{c}_i = 1/|S_i| \sum_{\mathbf{x}_j \in S_i} \mathbf{x}_j$ is the centroid of cluster S_i whose cardinality is $|S_i|$.

The number of ways in which a set of N objects can be partitioned into K non-empty groups is given by Stirling numbers of the second kind

$$\left\{ \begin{matrix} N \\ K \end{matrix} \right\} = \frac{1}{K!} \sum_{i=0}^K (-1)^{K-i} \binom{K}{i} i^N \quad (2)$$

which can be approximated by $K^N/K!$. It can be seen that a complete enumeration of all possible clusterings to determine the global minimum of (1) is clearly computationally prohibitive except for very small data sets. In fact, this non-convex optimisation problem is proven to be NP-hard even for $K=2$ (Ref. 55) or $D=2$ (Ref. 56). Consequently, various heuristics have been developed to provide approximate solutions to this problem.⁵⁷ Among these heuristics, Lloyd’s algorithm,⁵⁸ often referred to as the *batch k-means algorithm*, is the simplest and most commonly used one. This algorithm starts with K arbitrary centres, typically chosen uniformly at random from the data points. Each point is assigned to the nearest centre and then each centre is recalculated as the mean of all points assigned to it. These two steps are repeated until a predefined termination criterion is met. The pseudocode for this procedure is given in Algorithm (1) (**bold** symbols denote vectors). Here $m[i]$ denotes the membership of point \mathbf{x}_i , i.e. index of the cluster centre that is nearest to \mathbf{x}_i

Algorithm 1: batch k-means algorithm

```

input  $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \in \mathbb{R}^D$ 
      ( $N \times D$  input data set)
output  $\mathcal{C} = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_K\} \in \mathbb{R}^D$  ( $K$  cluster centres)
Select a random subset  $\mathcal{C}$  of  $\mathcal{X}$  as the initial set of
cluster centres
while termination criterion is not met do
  for ( $i = 1; i \leq N; i = i + 1$ ) do
    Assign  $\mathbf{x}_i$  to the nearest cluster;
     $m[i] = \underset{k \in \{1, 2, \dots, K\}}{\operatorname{argmin}} \|\mathbf{x}_i - \mathbf{c}_k\|_2^2$ 
  end
  Recalculate the cluster centres;
```

```

for ( $k = 1; k \leq K; k = k + 1$ ) do
  Cluster  $S_k$  contains the set of points  $\mathbf{x}_i$  that
  are nearest to the centre  $\mathbf{c}_k$ 
   $S_k = \{\mathbf{x}_i | m[i] = k\}$ 
  Calculate the new centre  $\mathbf{c}_k$  as the mean of
  points that belong to  $S_k$ 
   $\mathbf{c}_k = \frac{1}{|S_k|} \sum_{\mathbf{x}_i \in S_k} \mathbf{x}_i$ 
end
end
```

From a CQ perspective the batch k-means algorithm has two main drawbacks. First, due to its batch nature, i.e. cluster centres are updated after the presentation of all input vectors, the algorithm might get stuck in a local minimum. Second, the algorithm is highly sensitive to the selection of the initial cluster centres.⁵⁹ Adverse effects of improper initialisation include empty clusters (dead units), slower convergence, and a higher chance of getting stuck in bad local minima.³⁵

Adaptive distributing units (ADU) algorithm

The ADU algorithm is an online clustering algorithm based on the competitive learning paradigm, which is closely related to neural networks.⁶⁰ According to Rumelhart and Zipser,⁶¹ a competitive learning scheme consists of the following three basic components:

1. Start with a set of units that are all the same except for some randomly distributed parameter, which makes each of them respond slightly differently to a set of input patterns.
2. Limit the ‘strength’ of each unit.
3. Allow the units to compete in some way for the right to respond to a given subset of inputs.

The pseudocode for the ADU algorithm is given in Algorithm (2). Here, at any given time, n and $wc[i]$ denote the number of units (clusters) determined so far and the number of times that the i -th unit, represented by \mathbf{c}_i , won the competition in the past, respectively. The algorithm parameters θ , t_{\max} , and γ denote the maximum number of times a particular unit can win, maximum number of iterations, and the learning rate, respectively. The procedure starts with a single unit whose centre, \mathbf{c}_1 , is given by the centroid of the data set \mathcal{X} . In each iteration, an input vector \mathbf{x} is randomly selected from \mathcal{X} and the unit that is nearest to \mathbf{x} with respect to Euclidean distance is declared as the winner. This unit is then updated by moving its centre closer to \mathbf{x} and incrementing its win count. New units are added by splitting existing units that reach the threshold number of wins, θ , until the number of units reaches K . This splitting rule prevents certain units from monopolising the input vectors and thus effectively avoids the dead unit problem. When a new unit is generated by splitting an existing unit, the two units are temporarily colocated. Whenever an input vector that is nearest to these two units is presented, one of them becomes the winner and gets updated, while the other one remains unchanged. Hence, the problem of two units moving together cannot occur. Following,⁵⁴ the algorithm parameters were set to $\theta = 400K^{1/2}$, $t_{\max} = (2K - 3)\theta$, and $\gamma = 0.015$.

Algorithm 2: adaptive distributing units algorithm

```

input  $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \in \mathbb{R}^D$  ( $N \times D$  input data
set)
output  $\mathcal{C} = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_K\} \in \mathbb{R}^D$  ( $K$  cluster centres)
```

The first unit is given by the centroid of the input vectors

$n=1$

$$\mathbf{c}_n = 1/N \sum_{i=1}^N \mathbf{x}_i$$

Initialise the win counts to 0;

for ($i=1; i \leq K; i=i+1$) do

$wc[i]=0$;

end

for ($t=1; t \leq t_{\max}; t=t+1$) do

Select an input vector \mathbf{x} randomly from \mathcal{X}

Determine the winner (nearest) unit;

$$winner = \underset{k \in \{1, 2, \dots, n\}}{\operatorname{argmin}} \|\mathbf{x} - \mathbf{c}_k\|^2$$

Update the winner and its win count;

$$\mathbf{c}_{winner} = \mathbf{c}_{winner} + \gamma(\mathbf{x} - \mathbf{c}_{winner});$$

$$wc[winner] = wc[winner] + 1;$$

Split the unit if its win count exceeds the threshold;

if $wc[winner] = \theta$ and $n < K$ then

$n = n + 1$

$$\mathbf{c}_n = \mathbf{c}_{winner};$$

$$wc[n] = wc[winner] = 0;$$

end

end

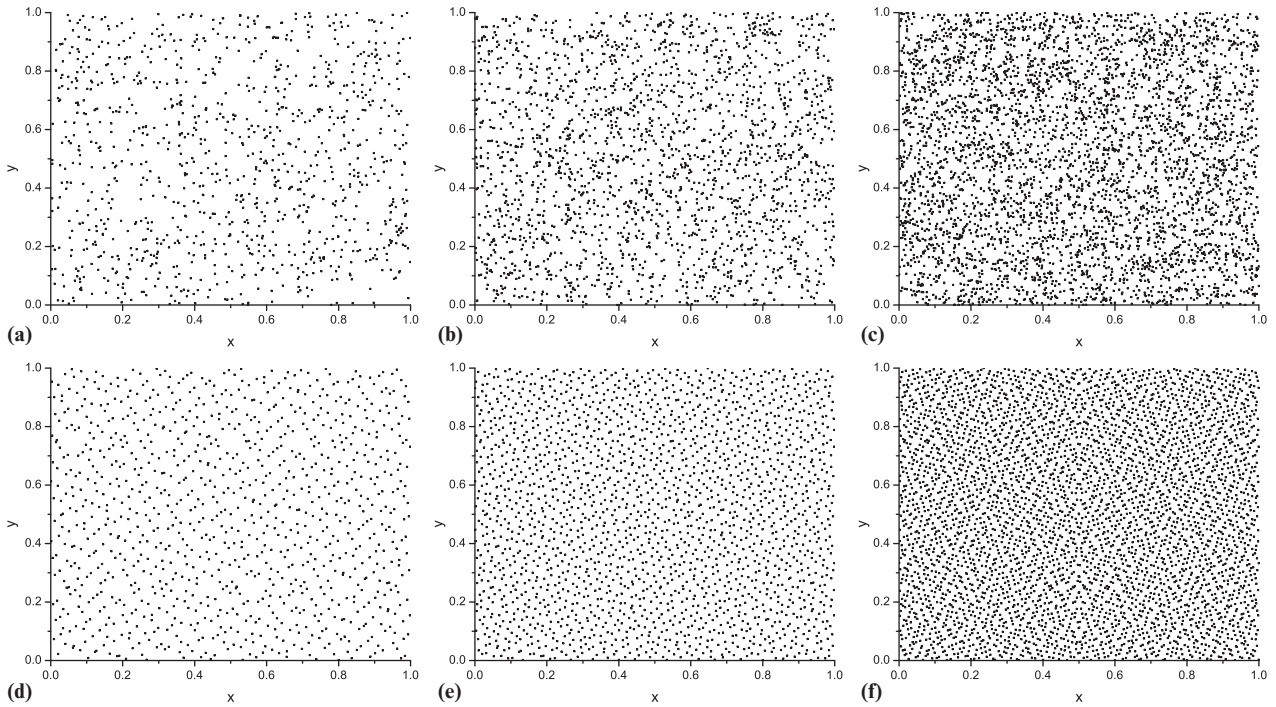
Due to their online formulation, i.e. cluster centres are updated after the presentation of each input vector, clustering algorithms based on the competitive learning paradigm are generally more adaptive and thus more likely to escape local minima when compared to batch algorithms. The online nature of such clustering algorithms, however, presents two drawbacks. First, these algorithms are order-dependent, that is, different presentation orders of the data vectors induce different partitions. Second, random selection of the input vectors renders these algorithms nondeterministic, i.e. each run

could potentially generate different clustering results. We avoid these problems by substituting the *pseudo-random* sampling scheme used in the ADU algorithm with *quasi-random* sampling. More specifically, we sample the image data by means of a quasi-random Sobol' sequence.⁶² A quasi-random sequence is a sequence of D -tuples that fill \mathbb{R}^D more uniformly than uncorrelated pseudo-random points.⁶³ This is illustrated in Fig. 1. Here the top row shows three pseudo-random sequences with increasing length from left to right generated by the MT19937 variant of the celebrated Mersenne Twister algorithm,⁶⁴ whereas the bottom row shows the corresponding quasi-random sequences generated by a Sobol' sequence. It can be seen that the pseudo-random sequences exhibit clumping, leading to rather uneven coverage of the sampled area,⁶⁵ whereas the quasi-random sequences result in a significantly more uniform point distribution. It should be noted that, despite their name, quasi-random sequences are in fact completely deterministic. This means that the proposed sampling scheme selects exactly the same set of pixels in each run. Since the theory and implementation of Sobol' sequences is mathematically involved, the interested reader is referred to the relevant literature^{62,63,65} for further information.

Experimental results and discussion

Image set and performance criteria

The proposed method was tested on a set of eight true-colour (24-bit) test images commonly used in the quantisation literature: Baboon (USC-SIPI Image Database, 512×512 , 230427 colours), Flowers & Sill (Kodak Lossless True Colour Image Suite, 768×512 , 37552 colours), Hats (Kodak Lossless True Colour Image Suite, 768×512 , 34871 colours), Lenna (USC-SIPI Image Database, 512×512 , 148279 colours),



1 Comparison of pseudo-random *a, b, c* and quasi-random sampling *d, e, f*. *a* random sequence (1024 pts); *b* random sequence (2048 pts); *c* random sequence (4096 pts); *d* Sobol' sequence (1024 pts); *e* Sobol' sequence (2048 pts); *f* Sobol' sequence (4096 pts)

Motocross (Kodak Lossless True Colour Image Suite, 768×512 , 63558 colours), Parrots (Kodak Lossless True Colour Image Suite, 768×512 , 72079 colours), Pills (Karel de Gendre, 800×519 , 206609 colours), and Sailboats (Kodak Lossless True Colour Image Suite, 512×768 , 24106 colours). These images are shown in Fig. 2.

The effectiveness of a quantisation method was quantified by the mean squared error (MSE) measure

$$MSE(\mathbf{X}, \hat{\mathbf{X}}) = \frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W \left\| \mathbf{X}(h,w) - \hat{\mathbf{X}}(h,w) \right\|_2^2 \quad (3)$$

where \mathbf{X} and $\hat{\mathbf{X}}$ denote respectively the $H \times W$ original and quantised images in the RGB colour space. MSE represents the average colour distortion with respect to the \mathcal{L}_2^2 (squared Euclidean) norm and is the most commonly used evaluation measure in the quantisation literature.¹

Comparison of ADU against other CQ methods

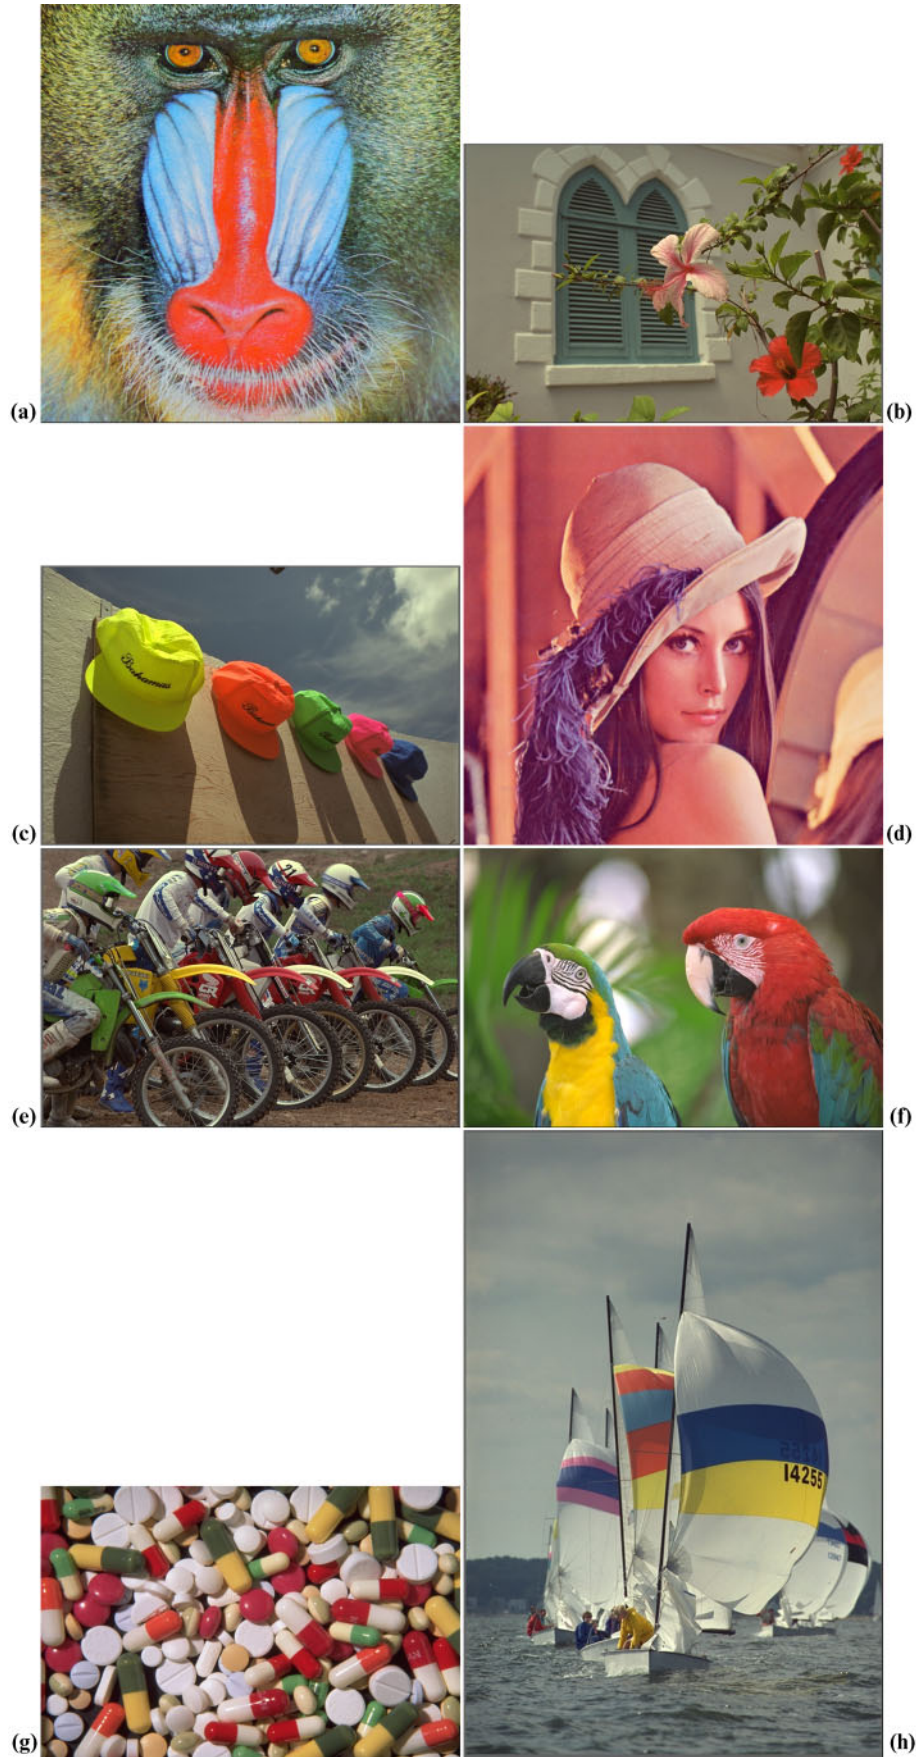
The proposed method was compared to 13 well-known CQ methods:

- **Popularity (POP):**¹² This method builds a $16 \times 16 \times 16$ colour histogram using 4 bits/channel uniform quantisation and then takes the K most frequent colours in the histogram as the colour palette.
- **Median-cut (MC):**¹² This method starts by building a $32 \times 32 \times 32$ colour histogram using uniform quantisation. This histogram volume is then recursively split into smaller boxes until K boxes are obtained. At each step, the box that contains the greatest number of colours is split along the longest axis at the median point, so that the resulting subboxes each contain approximately the same number of colours. The centroids of the final K boxes are taken as the colour palette.
- **Modified Popularity (MPOP):**⁶⁶ This method starts by building a $2^R \times 2^R \times 2^R$ colour histogram using R bits/channel uniform quantisation. It chooses the most frequent colour as the first palette colour \mathbf{c}_1 and then reduces the frequency of each colour \mathbf{c} by a factor of $(1 - e^{\alpha \|\mathbf{c} - \mathbf{c}_1\|_2^2})$, where α is a user-defined parameter. The remaining palette colours are chosen similarly. In the experiments, best results were obtained with the following parameter configuration: $\alpha = 0.25$ and $R = 4$ for $K \leq 64$ and $R = 5$ otherwise.
- **Octree (OCT):**¹³ This two-phase method first builds an octree (a tree data structure in which each internal node has up to eight children) that represents the colour distribution of the input image and then, starting from the bottom of the tree, prunes the tree by merging its nodes until K colours are obtained. In the experiments, the tree depth was limited to 6 to obtain the best results.
- **Variance-based method (WAN):**¹⁴ This method is similar to MC with the exception that at each step the box with the greatest SSE is split along the axis with the least weighted sum of projected variances at the point that minimises the marginal squared error.
- **Greedy orthogonal bipartitioning (WU):**¹⁶ This method is similar to WAN with the exception that at each step the box is split along the axis that minimises the sum of variances on both sides.

- **Centre-cut (CC):**¹⁷ This method is similar to MC with the exception that at each step the box with the greatest range on any coordinate axis is split along its longest axis at the mean point.
- **Self-organising map (SOM):**⁴⁸ This method utilises a one-dimensional self-organising map with K neurons. A random subset of N/f pixels is used in the training phase and the final weights of the neurons are taken as the colour palette. In the experiments, the sampling factor was set to $f = 1$ to obtain the best results.
- **Radius-weighted mean-cut (RWM):**¹⁸ This method is similar to WAN with the exception that the box is split along the vector from the origin to the radius-weighted mean (rwm) at the rwm point.
- **Modified maxmin (MMM):**³⁰ This method chooses the first palette colour \mathbf{c}_1 arbitrarily from the input image colours and the i -th colour \mathbf{c}_i ($i = 2, 3, \dots, K$) is chosen to be the colour that has the greatest minimum weighted \mathcal{L}_2^2 distance (the weights for the red, green, and blue channels are taken as 0.5, 1.0, and 0.25, respectively) to the previously selected colours, i.e. $\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_{i-1}$. Each of these initial palette colours is then recalculated as the mean of the colours assigned to it. In the experiments, the first colour was chosen as the centroid of the input image colours.
- **Pairwise clustering (PWC):**²⁷ This method is an adaptation of Ward's agglomerative hierarchical clustering method⁶⁷ to CQ. It builds a $2^R \times 2^R \times 2^R$ colour histogram and constructs a $Q \times Q$ joint quantisation error matrix, where Q is the number of colours in the reduced colour histogram. The clustering procedure starts with Q singleton clusters each of which contains one image colour. In each iteration, the pair of clusters with the least joint quantisation error are merged. This merging process is repeated until K clusters remain.
- **Split and Merge (SAM):**²⁸ This two-phase method first partitions the colour space uniformly into B partitions. This initial set of B clusters is represented as an adjacency graph. In the second phase, $(B - K)$ merge operations are performed to obtain the final K clusters. At each step of the second phase, the pair of clusters with the least joint quantisation error are merged. In the experiments, the initial number of clusters was set to $B = 20K$ to obtain the best results.
- **Cheng and Yang (CY):**²⁰ This method is similar to WAN with the exception that at each step the box is split along a specially chosen line defined by the mean colour and the colour that is farthest away from it at the mean point.

Two variants of ADU were implemented: one with pseudo-random sampling (ADU-PRS) and the other with quasi-random sampling (ADU-QRS). We also implemented the batch k-means (KM) algorithm with randomly chosen initial centres. Convergence of KM was determined by the following commonly used criterion: $(SSE_{i-1} - SSE_i) / SSE_i \leq \epsilon$, where SSE_i denotes the SSE (1) value at the end of the i -th iteration. The convergence threshold was set to $\epsilon = 0.001$.

Table 1 compares the effectiveness of the CQ methods on the test images. The best (lowest) error values are shown in **bold**. Note that since ADU-PRS and KM involve randomness, the quantisation errors for these methods are specified in the form of mean \pm standard deviation over 100 runs. The following observations are in order:

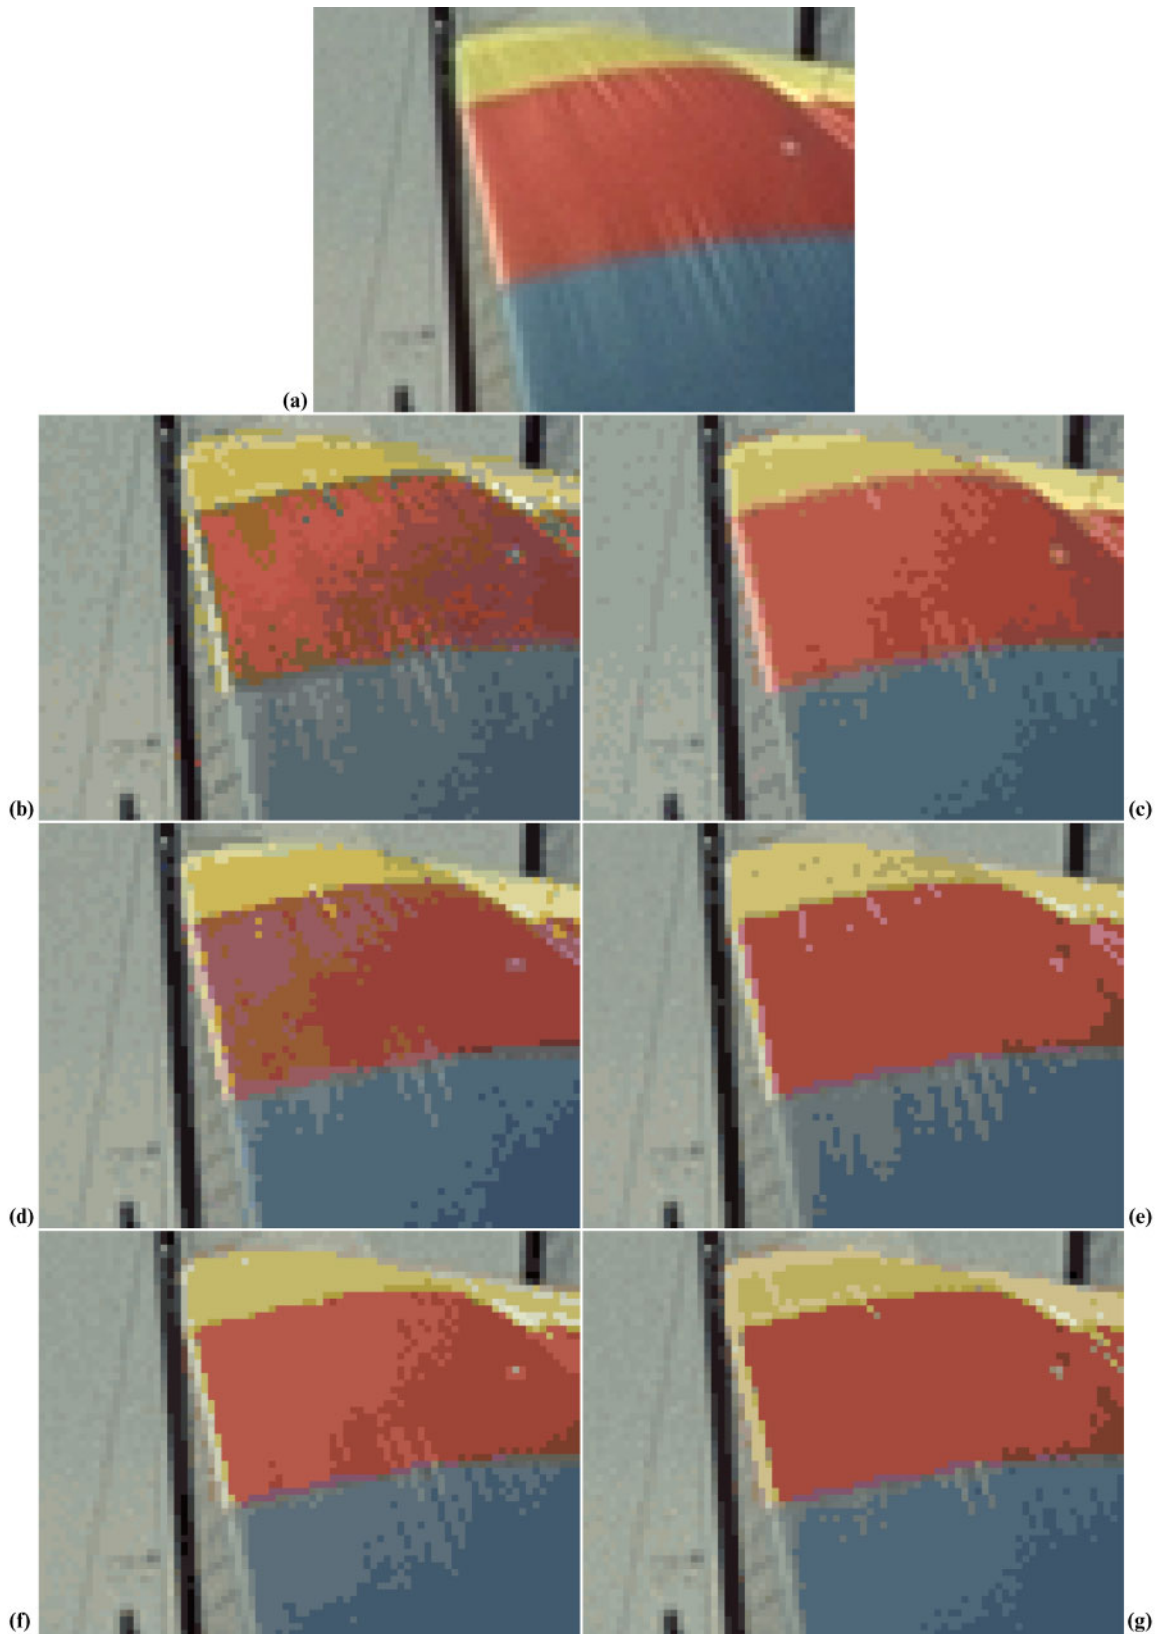


2 Test images: *a* Baboon; *b* Flowers & Sill; *c* Hats; *d* Lenna; *e* Motocross; *f* Parrots; *g* Pills; *h* Sailboats



3 Hats output images ($K=32$): *a* Original; *b* WAN output; *c* MC output; *d* SOM output; *e* KM output; *f* ADU-PRS output; *g* ADU-QRS output

- In general, post-clustering methods, i.e. SOM, MMM, KM, ADU-PRS, and ADU-QRS, are significantly more effective than the pre-clustering methods.
- KM, ADU-PRS, and ADU-QRS are the most effective post-clustering methods.
- Except in few cases, ADU-PRS is slightly more effective than KM. Furthermore, as expected, ADU-PRS is significantly less sensitive to randomness, which is evidenced by its low standard deviation.
- In about half of the cases, ADU-QRS is more effective than ADU-PRS, whereas in the remaining



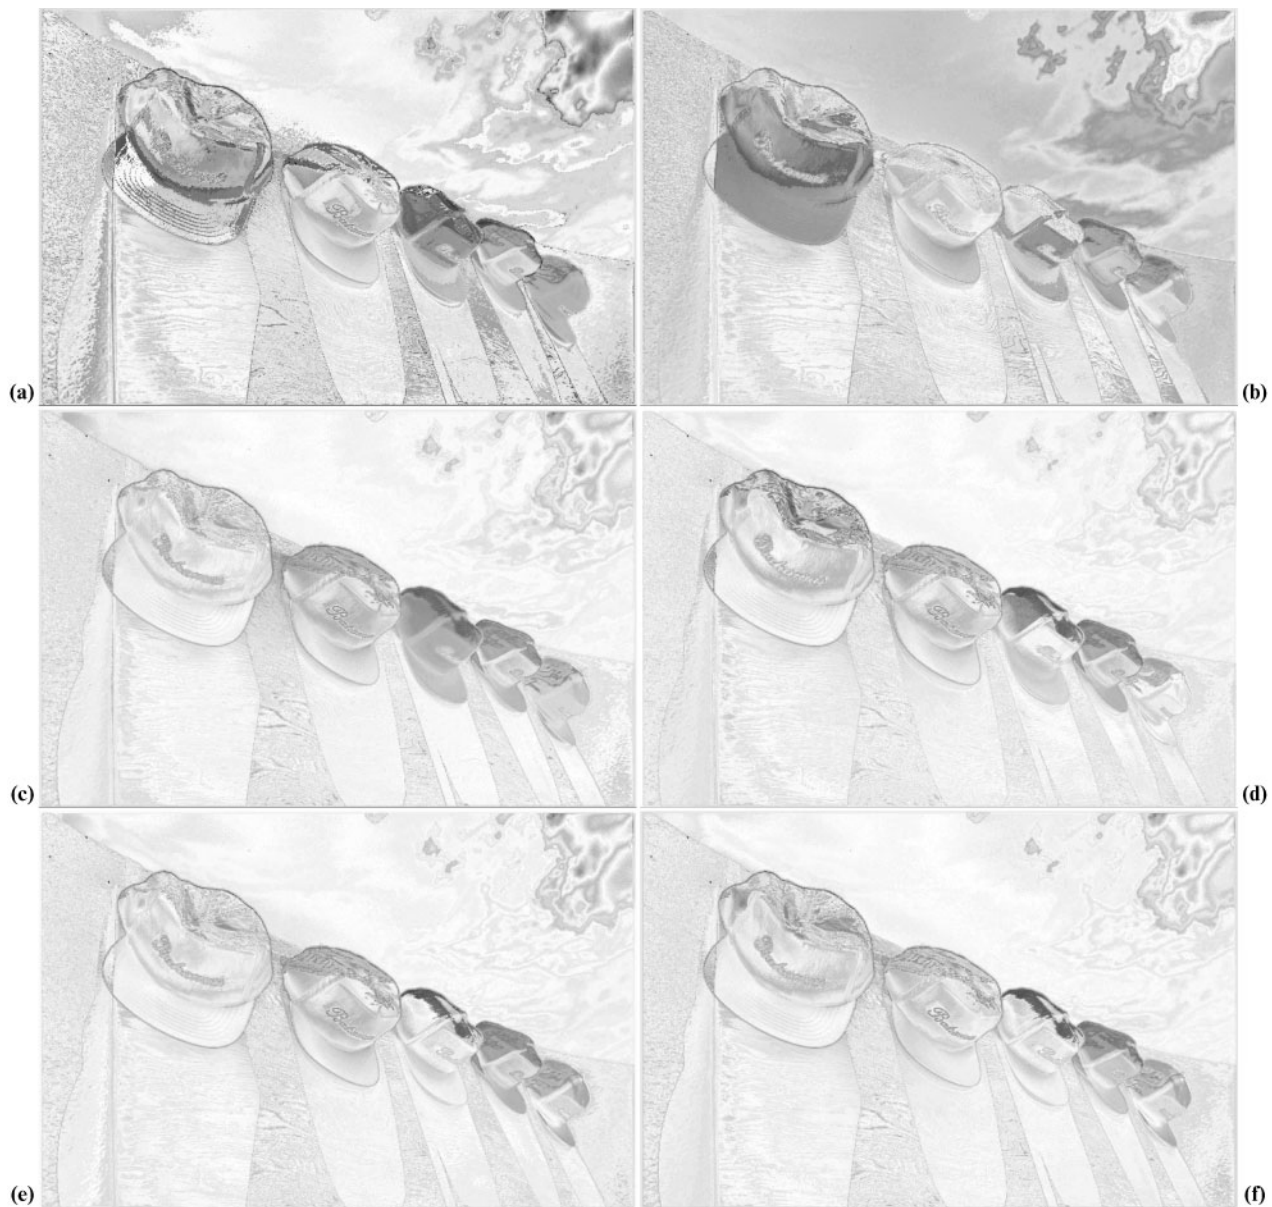
4 Sailboats output images ($K=64$): *a* Original; *b* WAN output; *c* MMM output; *d* OCT output; *e* KM output; *f* ADU-PRS output; *g* ADU-QRS output

cases the two methods have similar effectiveness. Note that the primary advantage of ADU-QRS over KM and ADU-PRS is its deterministic formulation, i.e. it gives exactly the same results in every run.

- POP is the least effective pre-clustering method. This is not surprising as this method disregards colours in

sparse regions of the colour space. Interestingly, despite being a simple modification of POP, MPOP performs surprisingly well, surpassing some of the better known methods such as MC, OCT, WAN, and CC.

- Despite its iterative nature, MMM performs poorly even when compared to pre-clustering methods. This



5 Hats error images ($K=32$): a WAN error; b MC error; c SOM error; d KM error; e ADU-PRS error; f ADU-QRS error

is because MMM tries to distribute the quantisation distortion more or less evenly throughout the image at the expense of increased mean distortion.

Figures 3 and 4 show sample quantisation results for close-up parts of the Hats and Sailboats images, respectively. Figures 5 and 6 show the full-scale error images for the respective images. The error image for a particular quantisation method was obtained by taking the pixelwise absolute difference between the original and quantised images. In order to obtain a better visualisation, pixel values of the error images were multiplied by 4 and then negated. It can be seen that the proposed method, ADU-QRS, performs exceptionally well in allocating representative colours to various image regions, resulting in cleaner error images.

Conclusions

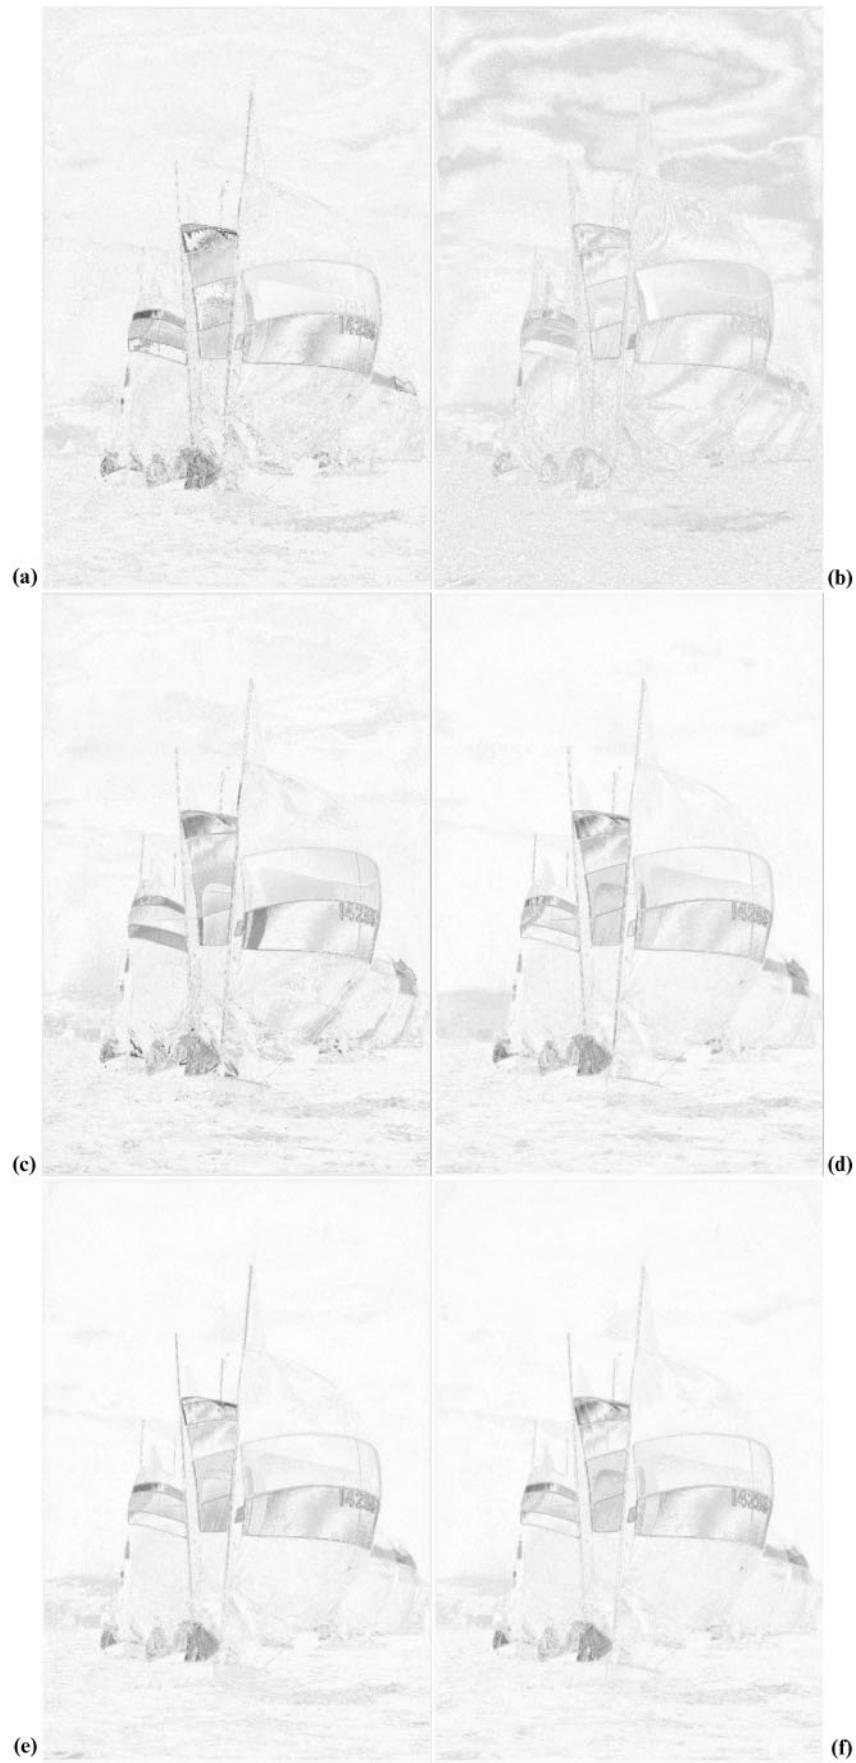
In this paper, an effective CQ method based on the ADU algorithm was introduced. This online clustering algorithm is advantageous over the conventional batch k -means clustering algorithm in that it does not require

cluster centre initialisation and that it avoids the dead unit problem. Due to its online formulation, however, the ADU algorithm is both order-dependent and nondeterministic. These drawbacks were overcome using quasi-random sampling. The proposed ADU variant with quasi-random sampling, ADU-QRS, was shown to give comparable or better results than the original one based on pseudo-random sampling, ADU-PRS.

Experiments on a diverse set of classic test images demonstrated that the proposed method outperforms well-known CQ methods with respect to distortion minimisation. In addition, our method is significantly easier to implement when compared to dedicated CQ methods.

Acknowledgements

This publication was made possible by grants from the Louisiana Board of Regents (LEQSF2008-11-RD-A-12), US National Science Foundation (0959583, 1117457), and National Natural Science Foundation of China (61050110449, 61073120).



6 Sailboats error images ($K=64$): *a* WAN error; *b* MMM error; *c* OCT error; *d* KM error; *e* ADU-PRS error; *f* ADU-QRS error

References

1. L. Brun, A. Trémeau, Digital Color Imaging Handbook, 2002, Ch. Color Quantization, pp. 589–638 (CRC Press, Boca Raton, FL).
2. Yang, C.-K. and Tsai, W.-H. Color image compression using quantization, thresholding, and edge detection techniques all based on the moment-preserving principle. *Pattern Recognit. Lett.*, 1998, **19**, (2), 205–215.
3. Deng, Y. and Manjunath, B. Unsupervised segmentation of color-texture regions in images and video. *IEEE Trans. Pattern Anal. Machine Intell.*, 2001, **23**, (8), 800–810.
4. Sherkat, N., Allen, T. and Wong, S. Use of colour for hand-filled form analysis and recognition. *Pattern Anal. Appl.*, 2005, **8**, (1), 163–180.
5. Sertel, O., Kong, J., Catalyurek, U. V., Lozanski, G., Saltz, J. H. and Gurcan, M. N. Histopathological image analysis using model-based intermediate representations and color texture: follicular lymphoma grading. *J. Signal Process. Syst.*, 2009, **55**, (1–3), 169–183.
6. Kuo, C.-T. and Cheng, S.-C. Fusion of color edge detection and color quantization for color image watermarking using principal axes analysis. *Pattern Recognit.*, 2007, **40**, (12), 3691–3704.
7. Wang, S., Cai, K., Lu, J., Liu, X. and Wu, E. Real-time coherent stylization for augmented reality. *Visual Comput.*, 2010, **26**, (6–8), 445–455.
8. Deng, Y., Manjunath, B., Kenney, C., Moore, M. and Shin, H. An efficient color representation for image retrieval. *IEEE Trans. Image Process.*, 2001, **10**, (1), 140–147.
9. Xiang, Z. Handbook of Approximation Algorithms and Metaheuristics, 2007, Ch. Color Quantization, pp. 86–186-17 (Chapman & Hall/CRC, London).
10. Gentile, R. S., Allebach, J. P. and Walowit, E. Quantization of color images based on uniform color spaces. *J. Imag. Technol.*, 1990, **16**, (1), 11–21.
11. Mojsilovic, A. and Soljanin, E. Color quantization and processing by Fibonacci lattices. *IEEE Trans. Image Process.*, 2001, **10**, (11), 1712–1725.
12. Heckbert, P. Color image quantization for frame buffer display. *ACM SIGGRAPH Comput. Graph.*, 1982, **16**, (3), 297–307.
13. Gervautz, M. and Purgathofer, W. New Trends in Computer Graphics, 1988, Ch. A Simple Method for Color Quantization: Octree Quantization, pp. 219–231 (Springer, Berlin/New York).
14. Wan, S. J. and Wong, S. K. M. Prusinkiewicz, variance-based color image quantization for frame buffer display. *Color Res. Appl.*, 1990, **15**, 52–58.
15. Orchard, M. and Bouman, C. Color quantization of images. *IEEE Trans. Signal Process.*, 1991, **39**, 2677–2690.
16. Wu, X. Graphics Gems Volume II, 1991, Ch. Efficient Statistical Computations for Optimal Color Quantization, pp. 126–133 (Academic Press, San Diego, CA).
17. Joy, G. and Xiang, Z. Center-cut for color image quantization. *Visual Comput.*, 1993, **10**, 62–66.
18. Yang, C.-Y. and Lin, J.-C. RWM-cut for color image quantization. *Comput. Graph.*, 1996, **20**, (4), 577–588.
19. Hsieh, I.-S. and Fan, K.-C. An adaptive clustering algorithm for color quantization. *Pattern Recognit. Lett.*, 2000, **21**, 337–346.
20. Cheng, S. and Yang, C. Fast and novel technique for color quantization using reduction of color space dimensionality. *Pattern Recognit. Lett.*, 2001, **22**, (8), 845–856.
21. Lo, K. Chan, Y. and Yu, M. Colour quantization by three-dimensional frequency diffusion. *Pattern Recognit. Lett.*, 2003, **24**, (14), 2325–2334.
22. Sirisathitkul, Y., Auwatanamongkol, S. and Uyyanonvara, B. Color image quantization using distances between adjacent colors along the color axis with highest color variance. *Pattern Recognit. Lett.*, 2004, **25**, (9), 1025–1043.
23. Kanjanawanishkul, K. and Uyyanonvara, B. Novel fast color reduction algorithm for time-constrained applications. *J. Visual Commun. Image Represent.*, 2005, **16**, (3), 311–332.
24. Equitz, W. H. A new vector quantization clustering algorithm. *IEEE Trans. Acoust., Speech Signal Process.*, 1989, **37**, (10), 1568–1575.
25. Balasubramanian, R. and Allebach, J. A new approach to palette selection for color images. *J. Imag. Technol.*, 1991, **17**, (6), 284–290.
26. Xiang, Z. and Joy, G. Color image quantization by agglomerative clustering. *IEEE Comput. Graph. Appl.*, 1994, **14**, (3), 44–48.
27. Velho, L., Gomez, J. and Sobreiro, M. V. R. Color image quantization by pairwise clustering, Proc. 10th Brazilian Symp. on Computer Graphics and Image Processing, Campos do Jordao, Brazil, October 1997, IEEE, pp. 203–210.
28. Brun, L. and Mokhtari, M. Two high speed color quantization algorithms, Proc. 1st Int. Conf. on Color in Graphics and Image Processing, Saint-Etienne, France, October 2000, pp. 116–121.
29. Goldberg, N. Colour image quantization for high resolution graphics display. *Image Vis. Comput.*, 1991, **9**, (5), 303–312.
30. Xiang, Z. Color image quantization by minimizing the maximum intercluster distance. *ACM Trans. Graph.*, 1997, **16**, (3), 260–276.
31. Kasuga, H., Yamamoto, H. and Okamoto, M. Color quantization using the fast k-means algorithm. *Syst. Comput. Jpn.*, 2000, **31**, (8), 33–40.
32. Huang, Y.-L. and Chang, R.-F. A fast finite-state algorithm for generating RGB palettes of color quantized images. *J. Inf. Sci. Eng.*, 2004, **20**, (4), 771–782.
33. Hu, Y.-C. and Lee, M.-G. K-means based color palette design scheme with the use of stable flags. *J. Electron. Imag.*, 2007, **16**, (3), 033003.
34. Hu, Y.-C. and Su, B.-H. Accelerated k-means clustering algorithm for colour image quantization. *Imag. Sci. J.*, 2008, **56**, (1), 29–40.
35. Celebi, M. E. Improving the performance of k-means for color quantization. *Image Vis. Comput.*, 2011, **29**, (4), 260–271.
36. Frackiewicz, M. and Palus, H. Computational Vision and Medical Image Processing: Recent Trends, 2011, Ch. KM and KHM Clustering Techniques for Colour Image Quantisation, pp. 161–174 (Springer, Berlin).
37. Verevka, O. and Buchanan, J. Local k-means algorithm for colour image quantization, Proc. Graphics/Vision Interface Conf., Quebec Canada, May 1995, pp. 128–135.
38. Scheunders, P. Comparison of clustering algorithms applied to color image quantization. *Pattern Recognit. Lett.*, 1997, **18**, (11–13), 1379–1384.
39. Cak, S., Dizdar, E. N. and Ersak, A. A fuzzy colour quantizer for renderers. *Displays*, 1998, **19**, (2), 61–65.
40. Ozdemir, D. and Akarun, L. Fuzzy algorithm for color quantization of images. *Pattern Recognit.*, 2002, **35**, (8), 1785–1791.
41. Kim, D.-W., Lee, K. and Lee, D. A novel initialization scheme for the fuzzy c-means algorithm for color clustering. *Pattern Recognit. Lett.*, 2004, **25**, (2), 227–237.
42. Schaefer, G. and Zhou, H. Fuzzy clustering for colour reduction in images. *Telecommun. Syst.*, 2009, **40**, (1–2), 17–25.
43. Zhou, H., Schaefer, G., Sadka, A. and Celebi, M. E. Anisotropic mean shift based fuzzy c-means segmentation of dermoscopy images. *IEEE J. Sel. Top. Signal Process.*, 2009, **3**, (1), 26–34.
44. Schaefer, G. and Zhou, H. Foundations of Computational Intelligence Volume 2, 2009, Ch. An Overview of Fuzzy C-Means Based Image Clustering Algorithms, pp. 295–310 (Springer, Berlin/Heidelberg).
45. Wen, Q. and Celebi, M. E. Hard vs. fuzzy c-means clustering for color quantization, *EURASIP J. Adv. Signal Process.*, 2011, (1), 118–129.
46. Schaefer, G. Innovations in Intelligent Image Analysis, 2011, Ch. Intelligent Approaches to Colour Palette Design, pp. 275–289 (Springer, Berlin).
47. Bing, Z., Junyi, S. and Qinke, P. An adjustable algorithm for color quantization. *Pattern Recognit. Lett.*, 2004, **25**, (16), 1787–1797.
48. Dekker, A. Kohonen neural networks for optimal colour quantization. *Netw.: Comput. Neural Syst.*, 1994, **5**, (3), 351–367.
49. Papamarkos, N., Atsalakis, A. and Strouthopoulos, C. Adaptive color reduction. *IEEE Trans. Syst., Man, Cybern. Part B*, 2002, **32**, (1), 44–56.
50. Chang, C.-H., Xu, P., Xiao, R., and Srikanthan, T. New adaptive color quantization method based on self-organizing maps. *IEEE Trans. Neural Netw.*, 2005, **16**, (1), 237–249.
51. Rasti, J., Monadjemi, A. and Vafaei, A. Color reduction using a multi-stage kohonen self-organizing map with redundant features. *Expert Syst. Appl.*, 2011, **38**, (10), 13188–13197.
52. Chung, K.-L., Huang, Y.-H., Wang, J.-P. and Cheng, M.-S. Speedup of color palette indexing in self-organization of kohonen feature map. *Expert Syst. Appl.*, 2012, **39**, (3), 2427–2432.
53. Xiao, Y., Leung, C.-S., Lam, P.-M. and Ho, T.-Y. Self-organizing map-based color palette for high-dynamic range texture compression. *Neural Comput. Appl.*, 2012, **21**, (4), 639–647.
54. Uchiyama, T. and Arbib, M. An algorithm for competitive learning in clustering problems. *Pattern Recognit.*, 1994, **27**, (10), 1415–1421.
55. Aloise, D., Deshpande, A., Hansen, P. and Popat, P. NP-hardness of euclidean sum-of-squares clustering. *Mach. Learn.*, 2009, **75**, (2), 245–248.
56. Mahajan, M., Nimbhorkar, P. and Varadarajan, K. The planar k-means problem is NP-hard. *Theor. Comput. Sci.*, 2012, **442**, 13–21.

57. Tarsitano, A. A computational study of several relocation methods for k-means algorithms. *Pattern Recognit.* 2003, **36**, (12), 2955–2966.
58. Lloyd, S. Least squares quantization in PCM. *IEEE Trans. Inform. Theory*, 1982, **28**, (2), 129–136.
59. Celebi, M. E., Kingravi, H. and Vela, P. A. A comparative study of efficient initialization methods for the k-means clustering algorithm. *Expert Syst. Appl.*, 2013, **40**, 200–210.
60. Xu, R. and Wunsch, D. Clustering, 2008 (Wiley-IEEE Press, New York).
61. Rumelhart, D. and Zipser, D. Feature discovery by competitive learning. *Cogn. Sci.*, 1985, **9**, (1), 75–112.
62. Bratley, P. and Fox, B. L. Algorithm 659: implementing Sobol's quasirandom sequence generator. *ACM Trans. Math. Softw.*, 1988, **14**, (1), 88–100.
63. Press, W. H., Teukolsky, S. A., Vetterling, W. T. and Flannery, B. P. Numerical Recipes, 2007, 3rd edition, (Cambridge University Press, New York).
64. Matsumoto, M. and Nishimura, T. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.*, 1998, **8**, (1), 3–30.
65. Heath, M. T. Scientific Computing: An Introductory Survey, 2002, 2nd edition (McGraw-Hill, New York).
66. Braudaway, G. W. Procedure for optimum choice of a small number of colors from a large color palette for color imaging, Proc. Electronic Imaging Conf., San Francisco, CA February 1987, SPIE, pp. 71–75.
67. Ward, J. Hierarchical grouping to optimize an objective function. *J. Am. Stat. Assoc.*, 1963, **58**, (301), 236–244.