



Color Image Quantization by Minimizing the Maximum Intercluster Distance

ZHIGANG XIANG

The City University of New York, Queens College, NY

One of the numerical criteria for color image quantization is to minimize the maximum discrepancy between original pixel colors and the corresponding quantized colors. This is typically carried out by first grouping color points into tight clusters and then finding a representative for each cluster. In this article we show that getting the smallest clusters under a formal notion of minimizing the maximum intercluster distance does not guarantee an optimal solution for the quantization criterion. Nevertheless, our use of an efficient clustering algorithm by Teofilo F. Gonzalez, which is optimal with respect to the approximation bound of the clustering problem, has resulted in a fast and effective quantizer. This new quantizer is highly competitive and excels when quantization errors need to be well capped and when the performance of other quantizers may be hindered by such factors as low number of quantized colors or unfavorable pixel population distribution. Both computer-synthesized and photographic images are used in experimental comparison with several existing quantization methods.

Categories and Subject Descriptors: I.3.3 [Computer Graphics]: Picture/Image Generation—*display algorithms*; I.4.1 [Image Processing]: Digitization—*quantization*

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: Clustering, color quantization, image compression, indexed image

1. INTRODUCTION

The field of color image quantization can trace its origin to the transformation of a continuous-tone black and white picture into a discrete grayscale image. This digitization process maps intensity values from a continuous spectrum into a series of gray levels. Since we limit the number of gray levels between black and white, the two extremes of the intensity range, the question arises as to how to choose and use these gray levels to reproduce the original. As for the quantization of color pictures, the same question can be asked regarding each component of a continuous color

Author's address: Department of Computer Science, The City University of New York/Queens College, 65-30 Kissena Boulevard, Flushing, NY 11367.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1997 ACM 0730-0301/97/0700-0260 \$03.50

signal, with the additional complexity of the method of decomposition (e.g., trivariable color space) and the interplay of color components. This fundamental concept of color image quantization is further extended to include the reproduction of an already quantized color image, typically with several hundred to several hundred thousand colors in a 24-bit RGB format, using a much smaller color palette, typically with 256 or less 24-bit colors (not necessarily colors that appear in the original image) for an 8-bit frame buffer display [Heckbert 1982].

Quantization inevitably introduces distortion. Ideally, a quantization algorithm should distribute any visible distortion¹ “evenly” throughout the quantized image so that none stands out to be found particularly objectionable by an average human observer. This means that many factors have to be considered including color space, image context, viewing condition, and viewer’s experience and aesthetic judgment. Although a number of researchers have explored some of these factors,² a lot more needs to be done in order to understand all the intricacies involved.

On the other hand, a classical and less ambitious goal in quantizer design is to satisfy certain statistical measures regarding discrepancies between original pixel values and quantized pixel values (with the assumption that some suitable color space, e.g., one that is visually homogeneous, is used to define pixel values and discrepancies). One such numerical criterion is to minimize the variance [Lloyd 1957; Max 1960]. This is first used to quantize continuous-tone black and white pictures and later applied to the quantization of color images. However, an efficient optimal solution is probably nonexistent because the problem is proven to be NP-complete [Garey et al. 1982]. Various approximation algorithms have been investigated including Equitz [1989], Gray et al. [1980], Linde et al. [1980], Selim and Ismail [1984], Wan et al. [1990], and Wu [1991, 1992]. Although minimizing variance provides a much better balance between color discrepancy and pixel population than do such empirical algorithms as the popularity algorithm and the median-cut algorithm, there are cases where significant color shifts in a small image area result in eye-catching distortion that renders the quantized image unacceptable [Xiang and Joy 1994a].

Another numerical criterion is to minimize the maximum discrepancy between original and quantized pixel values. This means that no preference is given to image colors that appear more frequently than others. Palette colors are to be selected in such a way that the color shift of every pixel in the quantized image is restricted by a minimal cap. One nonadaptive [Kurz 1983] and several adaptive quantizers [Gervautz and Purgathofer 1988;

¹Note that visible distortion is a subjective/psychological notion whereas quantization error or pixel color discrepancy is a well-defined and measurable quantity in a given color space. The exact relationship between the two is not yet fully understood.

²Please see Balasubramanian et al. [1992, 1994], Crinon [1991], Gentile et al. [1990a, 1990b], Giusto et al. [1990], Joy and Xiang [1996], Kasson and Plouffe [1992], Mahy et al. [1991], Orchard and Bouman [1991], and Xiang and Joy [1994b].

Joy and Xiang 1993; Xiang and Joy 1994a] approximating this principle have been reported, along with experimental results indicating that gross mismapping of unpopular colors is largely under control and the overall quality of the quantized images can range from satisfactory to competitive.

The basic strategy employed by these recent adaptive minimum maximum-discrepancy quantizers is a two-step approach. The first step is to group original image colors in a hierarchical fashion into clusters that are as tight or as small as possible. The second step is to compute a cluster representative or quantized color for each cluster. The rationale here is that smaller clusters should lead to smaller maximum discrepancy between an original color and the corresponding quantized color.

In this article we first show that getting the smallest clusters under a formal notion of minimizing the maximum intercluster distance³ does not in theory guarantee an optimal solution for the minimum maximum-discrepancy criterion. Nevertheless, we want to bring an efficient nonhierarchical clustering algorithm by Teofilo F. Gonzalez [1985] and some theoretical results to the attention of researchers and practitioners working on color image quantization. This clustering algorithm achieves the best possible approximation to the optimal solution (with regard to the clustering problem), if $P \neq NP$. Our implementation of this algorithm using 24-bit color vectors has resulted in a fast quantizer that is very effective in capping maximum quantization error and in limiting visible distortion in quantized images.

2. MINIMIZING THE MAXIMUM INTERCLUSTER DISTANCE

The problem of m -dimensional clustering to minimize the maximum intercluster distance can be formally stated as finding a partition of n points in m -dimensional Euclidean space into k disjoint clusters B_1, B_2, \dots, B_k such that $\max(M_1, M_2, \dots, M_k)$, where M_i is the maximum distance between two points in cluster B_i , is minimized. A corresponding decision problem that is computationally not harder can be stated as deciding if there is a partition into B_1, B_2, \dots, B_k such that $\max(M_1, M_2, \dots, M_k) \leq$ some given w .

Intuitively a partition with minimal maximum intercluster distance consists of tight or small clusters where datapoints in each cluster are close to each other, providing an effective cap on the distance (discrepancy) between any datapoint and its corresponding cluster representative (e.g., the cluster's geometric center). In fact, this optimization clustering problem is equivalent to the minimal maximum-discrepancy quantization criterion when $m = 1$. Since all datapoints in cluster B_i are now distributed along a line segment and M_i is the length of the line segment (i.e., the distance between the datapoint at one end of the line segment and the datapoint at

³Since we are trying to minimize the maximum distance between color points in each cluster it might feel more appropriate to use the word intracluster. However, if we view each color point as a singleton cluster we are indeed minimizing the maximum intercluster distance. We adopt this second view in order to be consistent with the existing literature on clustering.

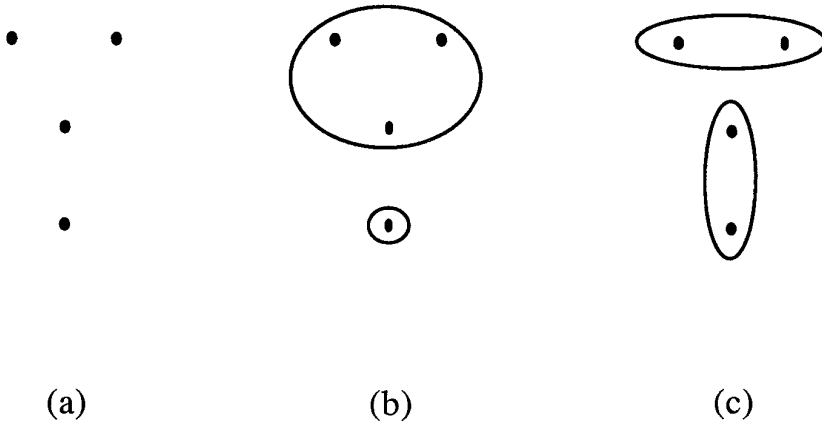


Fig. 1. (a) Example with four datapoints; (b) forming two clusters with minimum maximum intercluster distance; (c) forming two clusters with minimum maximum discrepancy between datapoints and their respective cluster representative.

the other end), we can simply choose the midpoint of the line segment to be the cluster representative. Clearly, minimizing the length of the line segment is equivalent to minimizing the distance between the midpoint and the endpoints.

However, these two minimization problems are not exactly the same when $m > 1$. Figure 1(a) shows a 2-dimensional example with four datapoints. The two points on the top and the one in the middle form an equilateral triangle of edge length L . The distance between the point in the middle and the one at the bottom is greater than L but less than $L/\cos(\pi/6)$. Now if we partition these four points into two clusters with minimal maximum intercluster distance we will have Figure 1(b), where the maximum intercluster distance is L and the maximum distance between a datapoint and the corresponding cluster representative is at least $L/2\cos(\pi/6)$, which is the minimal value we can get when we choose the geometric center of the equilateral triangle to be a cluster representative. On the other hand, if we partition these four points into two clusters using the minimum maximum-discrepancy criterion we will have Figure 1(c), where each cluster representative is the midpoint of the line segment connecting the two datapoints in the cluster, with the maximum distance between a datapoint and the corresponding cluster representative being less than $L/2\cos(\pi/6)$ and the maximum intercluster distance being greater than L .

Despite this theoretical difference between the two minimization problems we want to continue with the aforementioned two-step quantization strategy where original image colors are first grouped into tight clusters and a cluster representative is then calculated for each cluster. Instead of using bounding boxes to measure clusters as in Gervautz and Purgathofer [1988], Joy and Xiang [1993] and Xiang and Joy [1994a], we now use the precise notion of maximum intercluster distance to define cluster size. Thus the first quantization step is clustering to minimize the maximum intercluster distance.

This minimization problem has been proven to be polynomial solvable when $m = 1$ [Brucker 1978]. However, when $m = 2$ the corresponding decision problem⁴ is NP-complete, making the minimization problem NP-hard [Gonzalez 1985]⁵. Furthermore, when $m = 3$, which is typically the case in color image quantization, even finding a partition with maximum intercluster distance less than two times the optimal solution value, referred to as the $(2 - \epsilon)$ -approximation problem, is NP-hard for all $\epsilon > 0$ [Gonzalez 1985; Sahni and Gonzalez 1976].

Hence the best we can look for, if $P \neq NP$, is an efficient 2-approximation algorithm. To this end we have found a solution by Gonzalez that has worst case time complexity $O(nk)$ [Gonzalez 1985]. The following is a slightly modified description of the algorithm, with $\{c_1, c_2, \dots, c_n\}$ being the set of color points to be clustered, h_i being a designated point in B_i (h_i is called the head of B_i), and $d(a, b)$ being the distance between point a and point b :

```

 $B_1 = \{c_1, c_2, \dots, c_n\};$ 
 $h_1 = c_1;$ 
for ( $x = 1; x < k; x++$ ) {
   $d = \max\{d(c_i, h_j) | c_i \in B_j \text{ and } 1 \leq i \leq n \text{ and } 1 \leq j \leq x\};$ 
   $c = \text{one of the points whose distance to the head of the cluster it}$ 
   $\text{belongs to is } d;$ 
  move  $c$  to  $B_{x+1}$ ;
   $h_{x+1} = c;$ 
  for each  $c' \in (B_1 \cup \dots \cup B_x)$  {
    let  $j$  be such that  $c' \in B_j$ ;
    if ( $d(c', h_j) \geq d(c', c)$ ) move  $c'$  from  $B_j$  to  $B_{x+1}$ ;5
  }
}
```

Initially all color points are in the same cluster B_1 and the selection of h_1 is arbitrary. The main “for” loop continues as long as the current number of clusters is less than k . A new cluster is formed by first selecting a point that is the farthest away from its corresponding cluster head (an arbitrary selection is made to break any tie). This point is moved to the new cluster as its head. Every other color point that is closer to this new cluster head than to its own corresponding cluster head (or equally far away from the two heads) is then moved to this new cluster.

Our current implementation of this nonhierarchical clustering algorithm handles full 24-bit RGB color vectors since the usual bit-cutting color-reduction technique⁶ can itself be the cause of some rather significant visible distortion [Xiang and Joy 1994a]. A 2-dimensional pointer array indexed by the red and green components serves as the basis of a hash table-type data representation. Every original image color is represented by

⁴If we use nonnegative weights instead of Euclidean distances the decision problem is NP-complete even for $k = 3$ [Brucker 1978]. See problem MS9 in Garey and Johnson [1979].

⁵There is a typo in the original article [Gonzalez 1985], where B_l should have been B_{l+1} .

⁶Taking a few least-significant bits off each of the three 8-bit color components is equivalent to a uniform quantization step. This greatly consolidates original color points into a sparse color space, making possible a practical implementation of some quantization algorithms that have high complexity and memory requirement.

a node in a simply linked list that starts from the corresponding red and green-indexed array pointer. The list is maintained in ascending order based on the blue component. The blue component is stored in the node, along with an index to the representation of the cluster to which the image color belongs, and the distance between the cluster's head and the image color. This means that unless an image color is moved to another cluster, there is no need to recalculate the to-head-distance for that image color.

Each cluster is represented by a record containing the RGB value of its head, the RGB value of a member color point that is the farthest away from the head, and the cluster's RGB bounding box. This means that unless the cluster has been changed, there is no need to re-search the cluster to find a candidate point for possible use as the head of a new cluster. Even if the search is necessary, only color points within the cluster's bounding box need to be considered (some of them may not belong to this cluster). In addition, the use of an RGB bounding box that is determined by the current maximum intercluster distance and the point chosen to be the head of a new cluster restricts the search of color points that may satisfy the criterion for moving to the new cluster.

Following the clustering step comes the selection of cluster representative or quantized color. For this we have been using cluster centroid. Although such a choice may seem somewhat inconsistent with the minimal maximum-discrepancy criterion, our experience indicates that although the outcome is mostly determined by the clustering step, treating popular colors a bit more favorably at this stage is a good compromise. After all, as we have noted before in Xiang and Joy [1994a] and early in this article, optimal quantization results need more than carrying this numerical criterion to its extreme.

3. EXPERIMENTAL COMPARISON

We now compare this new quantizer with three existing methods that use the same two-step quantization strategy to minimize the maximum discrepancy between original and quantized colors. All three take a hierarchical approach to clustering. The first is an agglomerative method (octree) that is based on a predetermined subdivision of the RGB color space into levels of octants [Gervautz and Purgathofer 1988]. The second is a divisive method (center-cut) that repeatedly splits the cluster whose bounding box has the longest side [Joy and Xiang 1993]. The third is an agglomerative method (agglomerative clustering) that merges neighboring clusters using their bounding boxes [Xiang and Joy 1994a]. In all cases we use 24-bit colors without bit-cutting (except where noted otherwise) and choose cluster centroids as cluster representatives.

Although conducting quantization in a standard reference color space such as the CIE $L^*u^*v^*$ or $L^*a^*b^*$ along with accurate display calibration should have positive impact (and make these experiments repeatable and comparable with others in a more rigorous sense), there are other factors such as image context that seem to affect quantization quality at least

equally, if not more profoundly. A case in point⁷ is that a color shift is likely to get the viewer's attention when it occurs to the primary or familiar object in the scene (e.g., the face of a portrait), whereas a shift of something unimportant (e.g., wallpaper in the background), with the same or even greater magnitude, may very well turn out to be much less objectionable. This means that the full benefit achievable by going through device-dependent transformation to make use of a reference color space awaits further study.

In the following experiments we simply apply these quantizers to each test image's RGB values. Although in most cases we treat the three color components equally (hereafter referred to as using the regular RGB color space), in some cases we use an approximate ratio (based on their relative contribution to luminance) to scale them in order to partially compensate for the nonuniform nature of the RGB color space (hereafter referred to as using the scaled RGB color space). Our experience indicates that using the scaled RGB space often leads to better quantization results.⁸ All images are displayed and photographed on a Sparc workstation color monitor. The major quantization artifacts are significant enough to be clearly visible on any ordinary color monitor in a typical office/lab setting.

We also want to highlight the difference between this minimum maximum-discrepancy criterion and the minimum variance criterion. The latter attempts to better preserve popular colors at the expense of mapping unpopular colors with relatively large quantization errors. Although this controlled tradeoff often works well, the bias against unpopular colors can lead to unpleasant artifacts. We can observe this phenomenon with all three of the following test images, using a representative variance-based quantizer [Wu 1991] to obtain experimental results.

Before proceeding to the test results we want to comment on the fact that our implementation of the variance-based method is in its original form, that is, with 3-3-3 bit-cutting and using the regular RGB space. Bit-cutting is necessary in order to have a practical memory requirement. It should actually help to put some adjacent colors in the same cluster as it does in the center-cut method. Using the scaled RGB space might alleviate some but not all of the following artifacts. The important point is that the reason for the problem we show lies with the minimum variance criterion, not a particular algorithmic and/or implementational approximation.

Figure 2(a) shows a projected view of the RGB color cube on a 512×512 black background (some contouring artifacts are introduced here and in Figure 2(h) by the production-printing process). With 49,027 distinct 24-bit colors this computer-synthesized image is very difficult, if not impossible, to quantize without introducing eye-catching distortion (false contour). However, it demonstrates clearly how these quantizers group original

⁷See another case concerning false contours in Xiang and Joy [1994b] and Joy and Xiang [1996].

⁸See the 3-2-4 bit-cutting color-reduction technique in Joy and Xiang [1993] and the 2:1:4 proportional cluster size limit in Xiang and Joy [1994a].

image colors into clusters. We first quantize Figure 2(a) in the regular RGB space to 32 colors. Such a small number makes the weakness of a hierarchical clustering method more noticeable. The octree method produces Figure 2(b), which actually has only 30 colors because in this approach each cluster merging operation may involve up to 8 suboctants, reducing the total number of clusters (i.e., quantized colors) by more than 1. The imbalance in cluster sizes is rather significant. The center-cut method produces Figure 2(c), which looks better but the size imbalance is still evident because a large cluster is halved to become two smaller clusters. The agglomerative clustering method produces Figure 2(d), which is only comparable to Figure 2(c) because the merging of two similar size clusters tends to create a cluster that is twice as large. In contrast, the new quantizer produces Figure 2(e), which is clearly more homogeneous in terms of both cluster size and cluster distribution.

Figure 2(f) shows the result from the variance-based quantizer. Since all colors except the background black have very low pixel population, this black color asserts great influence on how the quantizer partitions the cluster to which it belongs. At the level of 32 colors a small corner of the cube appears virtually black (same as the background) because the corresponding original colors are grouped into the same cluster as the original background black. When the image is quantized to 64 colors or more the results are more or less what one can reasonably expect, with the varying degree of unevenness in cluster size and distribution that is caused by the biased positioning of some cutting planes.

The behavior of our new quantizer can be easily modified by scaling red:green:blue by, say, 0.5:1.0:0.25, in the computation of Euclidean color distances. Figure 2(g) shows the result of quantizing Figure 2(a) in this scaled RGB space to 256 colors. Figure 2(h) shows the result of Floyd-Steinberg error diffusion [Floyd and Steinberg 1975; Ulichney 1987] using those quantized colors (with some visible artifacts that vary according to the way the error terms are propagated).

Figure 3(a) shows a boy's portrait at 574×820 with 120,066 distinct 24-bit colors. The image has been used before to demonstrate the center-cut method with 3-2-4 bit-cutting and the agglomerative clustering method with 2:1:4 proportional cluster size limit. We now use each of the four minimum maximum-discrepancy methods to quantize Figure 3(a) to 64 colors in the regular RGB space. One can see that although the outcome of the octree method (Figure 3(b)) exhibits some significant distortion (e.g., the white of the boy's left eye being tinted blue and that of the right eye brown; see enlarged eye areas in Figure 4), the results from the center-cut method (Figure 3(c)), the agglomerative clustering method (Figure 3(d)), and the new quantizer (Figure 3(e)) seem quite reasonable, with perhaps the image from the agglomerative clustering method being the least objectionable and the one from the new quantizer being the second least. Slightly better results can be obtained for the center-cut, the agglomerative clustering (shown in Xiang and Joy [1994a]), and the new method when quantization is done in the scaled RGB space.

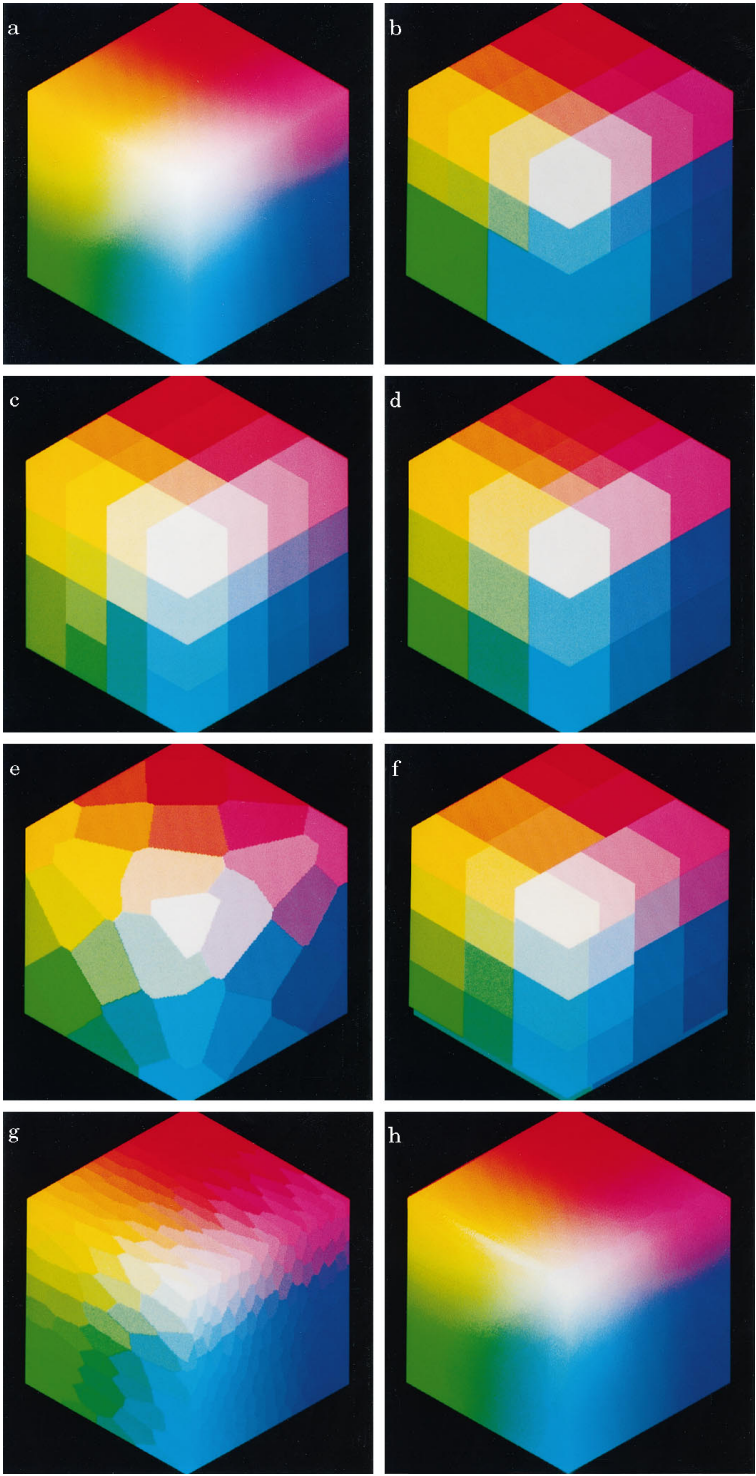


Fig. 2. (a) Projected view of RGB color cube. The image is quantized to 32 colors using (b) the octree method, (c) center-cut method, (d) agglomerative clustering method, (e) new quantizer, and (f) variance-based method. The result with 256 colors using the new method and the scaled RGB space is in (g), and that along with Floyd-Steinberg error diffusion is in (h).

The result of quantizing the boy's picture to 64 colors using the variance-based method was shown before in Xiang and Joy [1994a], where the white of the boy's left eye is tinted blue. We include here the enlarged eye area of that image in Figure 4(f). Even at the commonly used quantization level of 256 colors, a careful observer can readily notice the discoloration (which leads to some perceived disfigurement) of the white of his eyes. On the other hand, with 256 colors and using the scaled RGB space, the center-cut method (with or without 3-2-4 bit-cutting),⁹ the agglomerative clustering method (shown in Xiang and Joy [1994a]), and the new quantizer all produce very good results. The result from the new quantizer is shown here in Figure 3(f), where one can see some slight false contours on the boy's face.

A third example here is an image from the Kodak Photo CD sampler disc: image #4 entitled *Portrait of Girl in Red*, with base resolution 512×768 and 32,195 distinct 24-bit colors (Figure 5(a)). It was used before in Xiang and Joy [1994b] and Joy and Xiang [1996] to demonstrate a feedback-based, context-sensitive approach to quantization. This high-quality photographic image features both texture (e.g., the hat and hair) and smooth shading (e.g., the satin backdrop and the skin tone). The result with 256 colors from the octree quantizer appears in Figure 5(b), where one can see significant quantization artifacts. Using the scaled RGB space, both the center-cut method (Figure 5(c)) (a better result can be obtained with 3-2-4 bit-cutting) and the agglomerative clustering method (Figure 5(d)) are able to transform the image into good 256-color reproduction. However, one does not have to be very picky to notice some cross-mapping between certain skin colors and some shadow shades of the background curtain. Only the new quantizer manages to achieve very good color separation (Figure 5(e)).

Although the variance-based method is capable of turning the girl's portrait into very good 256-color reproduction (with better texture on the underside of the hat near the rim on her right and slightly smoother shading in some skin areas), one can readily notice an unsightly aberration on her teeth (see enlarged teeth from the original, the octree method, the center-cut method, the agglomerative clustering method, the new quantizer, and the variance-based approach in Figure 6).

Overall, this new quantizer delivers competitive performance over a wide range of test images. It excels when the discrepancy between a pixel's quantized color and its original color needs to be well capped. This is particularly important in cases where some relatively large color deviation in a small image area catches a viewer's attention.

Implemented in C on a Sparc 10 workstation, this new quantizer achieves interactive or near-interactive speed. It quantizes the Kodak Photo CD image, the RGB color cube image, and the boy's image to 256

⁹Using 3-2-4 bit-cutting improves the performance of the center-cut method on this and some other images because such a preprocessing step prevents certain adjacent color points from being separated and grouped into different clusters by the cutting planes. See result with 3-2-4 bit-cutting in Joy and Xiang [1993].



Fig. 3. (a) Boy's portrait. The image is quantized to 64 colors using (b) the octree method, (c) center-cut method, (d) agglomerative clustering method, and (e) new quantizer. The result with 256 colors using the new method and the scaled RGB space is in (f).

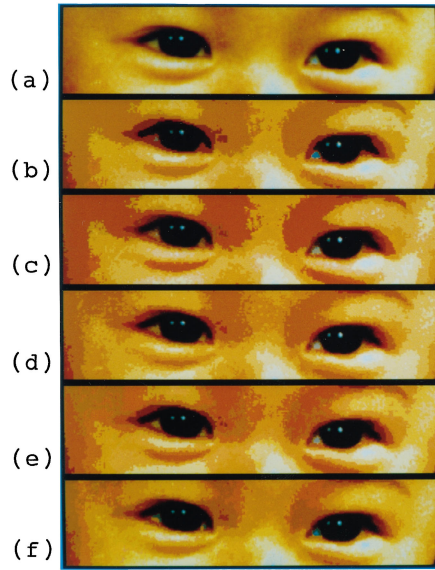


Fig. 4. Enlarged eye areas from the boy's image and the 64-color quantization results: (a) original, (b) octree method, (c) center-cut method, (d) agglomerative clustering method, (e) new quantizer, and (f) variance-based method.

colors using the scaled RGB space in 9.0, 17.0, and 32.0 seconds, respectively. Our implementation of the other quantizers can all perform at this speed or faster. Although the memory requirement of the octree method is $O(k)$, that of the center-cut method, the agglomerative clustering method, and the new method is $O(n)$. Dealing with truncated colors, the variance-based method has a memory requirement of $O(32 \times 32 \times 32)$.

Finally we include some quantization statistics for the four minimum maximum-discrepancy quantizers. The data are collected from the results of quantizing the three test images in the regular RGB space without bit-cutting. The numbers provide a glimpse of the relative effectiveness of these techniques in their ability to bound quantization error. One can see from Table I that the new method is the overall winner in minimizing the maximum quantization error although other methods may occasionally do better. This should be no surprise since the underlying clustering algorithm in the new quantizer only guarantees two times the optimal solution value for the clustering step. Although other methods do not guarantee anything better, they may sometimes result in smaller clusters and as a consequence smaller maximum quantization error. Table II shows the average quantization error where one can see that the center-cut method, the agglomerative clustering method, and the new method are more or less in the same performance class.

4. CONCLUDING REMARKS

Although primarily of interest to researchers and practitioners in computer graphics and image processing, color image quantization has close ties to



Fig. 5. (a) Kodak Photo CD sampler image entitled *Portrait of Girl in Red*. The image is quantized to 256 colors using (b) the octree method, (c) center-cut method in the scaled RGB space, (d) agglomerative clustering method in scaled RGB space, and (e) new quantizer in scaled RGB space.



Fig. 6. Enlarged teeth from the girl's image and the 256-color quantization results: (a) original, (b) the octree method, (c) center-cut method, (d) agglomerative clustering method, (e) new quantizer, and (f) variance-based method.

Table I. Maximum Quantization Error in Regular RGB Space

Original Image	Quantization Level	Octree	Center-cut	Agglomerative Clustering	Min. Max. Intercluster Distance
rgbcube	32	118.40	71.55	68.61	71.63
boy	64	86.38	51.40	58.66	42.63
girl in red	256	39.94	25.18	23.45	19.65

color science, cluster analysis, and human vision. The recent proliferation of 256-color graphics in the mass computing market and on the Internet has also raised much curiosity on the subject.

What constitutes an optimal quantization method is still under investigation. Although there exist sophisticated quantizers that can produce some impressive results using “uniform” color space and/or image context, quantizers based on simple statistical criteria enjoy a great advantage in terms of speed and easy implementation. They also serve as important benchmarks for comparing quantization performance.

Table II. Average Quantization Error in Regular RGB Space

Original Image	Quantization Level	Octree	Center-cut	Agglomerative Clustering	Min. Max. Intercluster Distance
rgbcube	32	23.71	20.59	21.95	20.78
boy	64	14.26	12.08	11.78	11.13
girl in red	256	9.68	5.95	5.86	6.00

One of the numerical criteria is to minimize the maximum discrepancy between original pixel colors and their corresponding quantized colors. The basic philosophy here is that all image colors are equally important unless the context of the image says otherwise. This helps to prevent “surprises” from happening such as the corner of the RGB cube turning black or the white of the boy’s eye turning blue.

Although this minimization problem is shown in this article to be different from the problem of clustering for minimum maximum intercluster distance, our use of Gonzalez’s clustering algorithm has produced an efficient quantizer that effectively approximates the minimum maximum-discrepancy criterion. Since the clustering algorithm is optimal, with respect to the approximation bound of the clustering problem, this new quantizer should be a critical benchmark, especially for all other clustering first and finding representative second methods.

Our experimental results indicate that the new quantizer, striving for minimally sized clusters and using centroids as cluster representatives, is an excellent performer on a wide range of test images. On the other hand, the center-cut method with 3-2-4 bit-cutting and the agglomerative clustering method seem to be able to compete from time to time at more or less the same quality level and occasionally even do a little better. This means that despite the inherent weakness of these hierarchical methods we can expect them to deliver very good results in this practical application of color image quantization.

ACKNOWLEDGMENTS

The author would like to thank the Associate Editor and the reviewers for their valuable comments and suggestions.

REFERENCES

BALASUBRAMANIAN, R., BOUMAN, C. A., AND ALLEBACH, J. P. 1992. New results in color image quantization. In *Proceedings of SPIE 1657*, 289–303.

ACM Transactions on Graphics, Vol. 16, No. 3, July 1997.

- BALASUBRAMANIAN, R., BOUMAN, C. A., AND ALLEBACH, J. P. 1994. Sequential scalar quantization of color images. *J. Electron. Imag.* 3, 1, 45–59.
- BRUCKER, P. 1978. On the complexity of clustering problems. In *Optimization and Operations Research*, R. Henn, B. Korte, and W. Oettli, Eds., Springer-Verlag, Berlin, 45–54.
- CRINON, R. J. 1991. Picture compression based on two-dimensional adaptive sampling and adaptive quantization. *Optic. Eng.* 30, 10, 1490–1496.
- EQUITZ, W. H. 1989. A new vector quantization clustering algorithm. *IEEE Trans. Acoustics, Speech, Signal Process.* 37, 10, 1568–1575.
- FLOYD, R. AND STEINBERG, L. 1975. An adaptive algorithm for spatial gray scale. In *Society for Information Display 1975 International Symposium Digest of Technical Papers*, 36–37.
- GAREY, M. R. AND JOHNSON, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York.
- GAREY, M. R., JOHNSON, D. S., AND WITSENHAUSEN, H. S. 1982. The complexity of the generalized Lloyd-Max problem. *IEEE Trans. Inf. Theory IT-28*, 2, 255–256.
- GENTILE, R. S., ALLEBACH, J. P., AND WALOWIT, E. 1990a. Quantization of color images based on uniform color spaces. *J. Imag. Technol.* 16, 1, 11–21.
- GENTILE, R. S., WALOWIT, E., AND ALLEBACH, J. P. 1990b. Quantization and multilevel halftoning of color images for near original image quality. In *Proceedings of SPIE 1249: Human Vision and Electronic Imaging* (Santa Clara, CA, Feb. 12–14), 249–260.
- GERVAUTZ, M. AND PURGATHOFER, W. 1988. A simple method for color quantization: Octree quantization. In *New Trends in Computer Graphics*, N. Magnenat-Thalmann and D. Thalmann, Eds., Springer-Verlag, Berlin, 219–231.
- GIUSTO, D. D., REGAZZONI, C. S., SERPICO, S. B., AND VERNAZZA, G. 1990. A new adaptive approach to picture coding. *Annales Telecommun.* 45, 9–10, 503–518.
- GONZALEZ, T. F. 1985. Clustering to minimize the maximum intercluster distance. *Theor. Comput. Sci.* 38, 2–3, 293–306.
- GRAY, R. M., KIEFFER, J. C., AND LINDE, Y. 1980. Locally optimal block quantizer design. *Inf. Control* 45, 178–198.
- HECKBERT, P. 1982. Color image quantization for frame buffer display. *Comput. Graph.* 16, 3, 297–307.
- JOY, G. AND XIANG, Z. 1993. Center-cut for color image quantization. *Visual Comput.* 10, 1, 62–66.
- JOY, G. AND XIANG, Z. 1996. Reducing false contours in quantized color images. *Comput. and Graph.* 20, 2, 231–242.
- KASSON, J. M. AND PLOUFFE, W. 1992. An analysis of selected computer interchange color spaces. *ACM Trans. Graph.* 11, 4, 373–405.
- KURZ, B. J. 1983. Optimal color quantization for color displays. In *IEEE Proceedings of Computer Vision and Pattern Recognition*, 217–224.
- LINDE, Y., BUZO, A., AND GRAY, R. M. 1980. An algorithm for vector quantizer design. *IEEE Trans. Commun.* 28, 1, 84–95.
- LLOYD, S. P. 1957. Least squares quantization in PCM. Unpublished Bell Laboratories memorandum; also *IEEE Trans. Inf. Theory IT-28* (1982), 129–137.
- MAHY, M., VAN MELLAERT, B., VAN EYCKEN, L., AND OOSTERLINCK, A. 1991. The influence of uniform color spaces on color image processing: A comparative study of CIELAB, CIELUV, and ATD. *J. Imag. Technol.* 17, 5, 232–243.
- MAX, J. 1960. Quantizing for minimum distortion. *IRE Trans. Inf. Theory IT-6*, 7–12.
- ORCHARD, M. T. AND BOUMAN, C. A. 1991. Color quantization of images. *IEEE Trans. Signal Process.* 39, 12, 2677–2690.
- SAHNI, S. AND GONZALEZ, T. 1976. P-complete approximation problems. *J. ACM* 23, 3, 555–565.
- SELIM, S. Z. AND ISMAIL, M. A. 1984. K-means-type algorithms: A generalization convergence theorem and characterization of local optimality. *IEEE Trans. Pattern Anal. Mach. Intel. PAMI-6*, 1, 81–87.
- ULICHNEY, R. 1987. *Digital Halftoning*. MIT Press, Cambridge, MA.
- WAN, S. J., PRUSINKIEWICZ, P., AND WONG, S. K. M. 1990. Variance-based color image quantization for frame buffer display. *Color Res. Appl.* 15, 1, 52–58.

- WU, X. 1991. Efficient statistical computations for optimal color quantization. In *Graphics Gems II*, J. Arvo, Ed., Academic Press, San Diego, CA, 126–133.
- WU, X. 1992. Color quantization by dynamic programming and principal analysis. *ACM Trans. Graph.* 11, 4, 348–372.
- XIANG, Z. AND JOY, G. 1994a. Color image quantization by agglomerative clustering. *IEEE Comput. Graph. Appl.* 14, 3, 44–48.
- XIANG, Z. AND JOY, G. 1994b. Feedback-based quantization of color images. In *Proceedings of SPIE 2182: Image and Video Processing II* (San Jose, CA, Feb. 7–9, 1994), 34–42.

Received May 1996; accepted February 1997