S. J. Wan
P. Prusinkiewicz
S. K. M. Wong

Department of Computer Science
University of Regina
Regina, Saskatchewan
Canada S4S OA2

Variance-Based Color Image Quantization for Frame Buffer Display

Color image quantization is a process of representing an image with a small number of well selected colors. In this article an algorithm for multidimensional data clustering (termed the variance-based *algorithm), based on the criterion of minimization of the sum-of-squared error, is applied to the problem of reducing the number of colors used to represent a given color image. The suitability of the sum-of-squared error criterion for measuring the similarity between the original and quantized images is examined using a digitized image and a computer-generated image. The experimental results indicate that this error measure is basically consistent with the perceived quality of the quantized image. The performance of the variance-based algorithm is compared with that of other algorithms for color image quantization in terms of quantization error, image quality, and computational time. Dithered images generated using the colors selected by the variance-based and the median-cut algorithms are also presented.*

Introduction

Color image quantization is a process of representing an image with a small number of carefully selected colors. The need for quantization commonly arises when small frame buffers are used for reproduction of color images. Although one can use a wide gamut of colors, only a limited number K of them can be displayed simultaneously.[1] Thus, if the number of colors in the original image exceeds K, it must be reduced to K before image display. This process has been referred to as *color image quantization*,[2] although the term *color compression* might be more suitable. The authors proposed an algorithm for multidimensional data clustering (termed the *variance-based* algorithm), based on the criterion of minimization of the sum-of-squared error.[3] In this article, the variance-based algorithm is applied to the problem of reducing the number of colors used to represent a given color image. The quantization of a given color image consists of two stages: the selection of an appropriate set of representative colors, and the mapping of the original color set into the reduced color set. The optimization goal is to make the perceived difference between the original image and its quantized representation as small as possible.

Human color vision is an extremely complicated and not yet fully understood process.[4,5] It is very difficult to formulate a definitive solution to the image quantization problem in terms of *perceived* image quality. In fact, there is no good objective criterion available for measuring the perceived image similarity. In this paper, we use the sum-of-squared error to measure the difference between the original image and its quantized representation. We show that this measure can reasonably capture the perceived image quality although it is derived under a simplifying assumption that the spatial correlations among pixels are not taken into account.

Let $\Omega = \{(r, g, b) | 0 \leq r, g, b \leq 225\}$ be the RGB color space.[1] Let $(x, y) \in I \times I$ be the spatial coordinate of a pixel, where I is the integer set. A digital image is defined as a mapping which assigns a color to each pixel:

$$h : I \times I \rightarrow C \subseteq \Omega, \tag{1}$$

where $C = \{\vec{c}_1, \vec{c}_2, \ldots, \vec{c}_N\}$ is a set of colors used in the digital image. A quantized image is a mapping defined by

$$f : I \times I \rightarrow R \subset \Omega, \tag{2}$$

where $R = \{\vec{r}_1, \vec{r}_2, \ldots, \vec{r}_K\}$ is a set of representative

© 1990 by John Wiley & Sons, Inc.

colors used in the quantized image. Typically, N is much larger than K. Note that R is not necessarily a subset of C.

The replacement of a pixel color $\vec{c}_i = (r, g, b) \in C$ in the original image by a representative color $\vec{r}(\vec{c}_i) = (r', g', b') \in R$ introduces a quantization error measured by the square of the Euclidean distance:

$$\| \vec{c}_i - \vec{r}(\vec{c}_i) \|^2 = (r - r')^2 + (g - g')^2 + (b - b')^2. \quad (3)$$

The total quantization error of image f is defined as the average error for all M pixels in the image:

$$D(h, f) = \frac{1}{M} \sum_{(x,y)\in I \times I} \| h(x, y) - f(x, y) \|^2. \quad (4)$$

It should be noted that in this measure each pixel is treated independently. The spatial correlations among pixels are not taken into account.

Let $p(\vec{c}_1), p(\vec{c}_2), \ldots, p(\vec{c}_N)$ denote the relative occurrence frequencies of the colors $\vec{c}_1, \vec{c}_2, \ldots, \vec{c}_N$ used in a digital image. Eq. (4) can then be rewritten as the sum-of-squared error [6]:

$$D(h, f) = \sum_{i=1}^{N} p(\vec{c}_i) \| \vec{c}_i - \vec{r}(\vec{c}_i) \|^2, \quad \vec{r}(\vec{c}_i) \in R. \quad (5)$$

Given a set of representative colors R, the best substitute for a color $\vec{c}_i$ is obviously the color in R nearest to $\vec{c}_i$, namely,

$$\vec{r}(\vec{c}_i) = \arg \min_{1 \le j \le K} \| \vec{c}_i - \vec{r}_j \|^2. \quad (6)$$

Our objective now is to find a set of K representative colors that minimizes the sum-of-squared error defined by Eq. (5). This problem appears to be at least as difficult as the construction of an optimal binary decision tree. The latter is known to be NP-complete[7] if the global minimum solution is sought.[8] A local minimum solution to this problem can be obtained by the k-means iterative procedure.[9,10] However, the computation required for the procedure to converge can easily become unacceptable for large problems. More efficient algorithms include the median-cut algorithm,[2,11,12] the mean-split algorithm[13] and the variance-based algorithm.[3] For detailed analysis of these algorithms see Wan et al.[3]

All of the above algorithms operate in the 3-dimensional *RGB* color space. In contrast, the Peano scan method[14,15] employs the Peano curve to transform a 3D *RGB* color histogram into a 1D histogram. Quantization is performed on the 1D histogram, using a monochromatic image quantization algorithm. The resulting values are then mapped back into the 3D color space using the one-to-one correspondence between the two spaces. An obvious disadvantage of this approach is that some of the neighboring relationships in the 3D color space are not properly preserved in the 1D space. Koo-Yam-Too[15] reported that using the algorithm described by Wong et al.[16] to quantize the 1D histogram, the Peano scan method is inferior to the variance-based method in the sense of quantization error and the perceived image quality.

In the subsequent sections, we first describe the implementation details of the variance-based algorithm for color image quantization. Next, we examine the suitability of the sum-of-squared error criterion for measuring the similarity between the original image and its quantized representation. The performance of the variance-based algorithm and other algorithms for color image quantization is evaluated using a digitized color image and a computer-generated color image. Finally, we address some open problems in color image quantization for further studies.

Implementation of the Variance-based Algorithm for Color Image Quantization

Let $\Omega_L \subseteq \Omega$ be a sub-box. The mean (centroid) and variance of Ω_L are defined by

$$\vec{\mu}_L = \sum_{\vec{c}\in\Omega_L} \vec{c}\, \frac{p(\vec{c})}{w_L}, \qquad \sigma_L^2 = \sum_{\vec{c}\in\Omega_L} \| \vec{c} - \vec{\mu}_L \|^2 \frac{p(\vec{c})}{w_L}, \quad (7)$$

where $w_L = \Sigma_{\vec{c}\in\Omega_L} p(\vec{c})$ is the weight of the sub-box Ω_L. The weighted variance $\hat{\sigma}_L^2$ of Ω_L is defined by

$$\hat{\sigma}_L^2 = \sum_{\vec{c}\in\Omega_L} \| \vec{c} - \vec{\mu}_L \|^2 p(\vec{c}) = w_L \sigma_L^2, \quad (8)$$

which gives the quantization error resulting from mapping all the colors within Ω_L into the mean $\vec{\mu}_L = (\mu_r, \mu_g, \mu_b)$ of the sub-box. Suppose that Ω_L is bounded by $r_1 \le r \le r_2$, $g_1 \le g \le g_2$ and $b_1 \le b \le b_2$. The weighted variance can be computed as follows:

$$\begin{aligned}\hat{\sigma}_L^2 &= w_L (\sigma_r^2 + \sigma_g^2 + \sigma_b^2) \\ &= \sum_{r=r_1}^{r_2} r^2 p_{r..} + \sum_{g=g_1}^{g_2} g^2 p_{.g.} \\ &\quad + \sum_{b=b_1}^{b_2} b^2 p_{..b} - w_L (\mu_r^2 + \mu_g^2 + \mu_b^2),\end{aligned} \quad (9)$$

where

$$\begin{aligned}p_{r..} &= \sum_{g=g_1}^{g_2} \sum_{b=b_1}^{b_2} p(r, g, b), \; p_{.g.} \\ &= \sum_{r=r_1}^{r_2} \sum_{b=b_1}^{b_2} p(r, g, b), \; p_{..b} \\ &= \sum_{r=r_1}^{r_2} \sum_{g=g_1}^{g_2} p(r, g, b)\end{aligned} \quad (10)$$

are the projected distributions, and $\sigma_r^2, \sigma_g^2, \sigma_b^2$ are the projected variances along the red, green, and blue axes, respectively.

Consider a sub-box Ω_L. Let $Q_L = (p_l, p_{l+1}, \ldots, p_m)$

TABLE I. Pseudo-code of the variance-based color image quantization algorithm.

```
ALGORITHM: Variance-based algorithm for color image quantization.
INPUT:     (1) A full-color image h,
           (2) K: the number of representative colors desired.
OUTPUT:    A quantized image f containing K colors.

Procedure VB(h, K);

  integer   colormap[K][3],i, j;
  real      frequency[32][32][32];
  structure    box {
      integer  r1, r2, g1, g2, b1, b2;
  };
  structure    record {
      real    weighted_variance;
      structure box boundaries;
  }  subbox[K];

begin
  /* step 1: initiation */
  read in the input image h pixel by pixel and calculate the frequency of each color used in this image (5 bits for each color component);
    the array frequency[32][32][32] forms the initial box;
  subbox[0].boundaries := boundaries of the initial box;
  subbox[0].weighted_variance := 100;

  /* step 2: iterative subdivision of the initial box */
  for i := 1 to K - 1 do begin
    select the subbox j with the largest subbox[j].weighted-variance;
    obtain the projected distributions by projecting all the points in the
      subbox j onto each of the color axes;
    for each projected distribution, calculate the optimial threshold from eq. (12) and compute the weighted sum of projected variances of
      the two intervals;
    the partition plane is chosen to be perpendicular to the axis with the minimum weighted sum of projected variances and passes
      through the optimal threshold;
    compute the weighted variance for each of the two new sub-boxes from eq. (9);
    subbox[i] := boundaries and weighted variance of the sub-box1;
    subbox[j] := boundaries and weighted variance of the sub-box2;
  end; /* for */

  /* step 3: finding the K representative colors */
  compute the centroids of the individual K sub-boxes, which form the colormap[K][3];

  /* step 4: displaying the quantized image f */
  replace the color of each pixel in the original image h by the closest color in the colormap [K][3];

end
```

be a projected distribution along an axis (for example, $P_{r..}$). Let t be a threshold which splits Q_L into two intervals: $(p_l, p_{l+1}, \ldots, p_{t-1})$ and $(p_t, p_{t+1}, \ldots, p_m)$. Let v_1, v_2 be the representatives assigned to the two intervals, respectively. The sum-of-squared error induced by this quantization is given by:

$$E(v_1, t, v_2) = \sum_{i=l}^{t-1} (i - v_1)^2 p_i + \sum_{i=t}^{m} (i - v_2)^2 p_i, \qquad \sum_{i=l}^{m} p_i = 1. \tag{11}$$

Wong et al.[16] has shown that to minimize $E(v_1, t, v_2)$, the representatives must be the means of the two intervals (i.e., $v_1 = \mu_1$, $v_2 = \mu_2$) and the optimal threshold is determined by

$$t_{opt} = arg \underset{\frac{\mu+l}{2} \leq t \leq \frac{\mu+m}{2}}{Max} \left[\frac{w_1}{w_2} (\mu - \mu_1)^2 \right], \tag{12}$$

where w_1, w_2 are the weights of the intervals and μ is the mean of the projected distribution Q_L.

In our implementation, the color image quantization procedure consists of the following steps:

TABLE II. Total quantization errors.

Images	Colors	Median-cut	Mean-split	Variance-based	K-means
	8	6.2507	4.1359	3.0115	2.9286
Girl	64	0.4824*	0.3913*	0.2755*	0.2659*
	256	0.0832	0.0745	0.0379	0.0377
	32	5.1930*	4.4209*	3.0976*	2.8306*
Deck	64	2.3624	2.0840	1.0786	1.0259
	256	0.3809	0.3602	0.1823	0.1722

FIG. 1. 64-color quantized images of the *Girl*. (a) The original image; (b) the median-cut algorithm; (c) the mean-split algorithm; (d) the variance-based algorithm; and (e) the k-means algorithm.

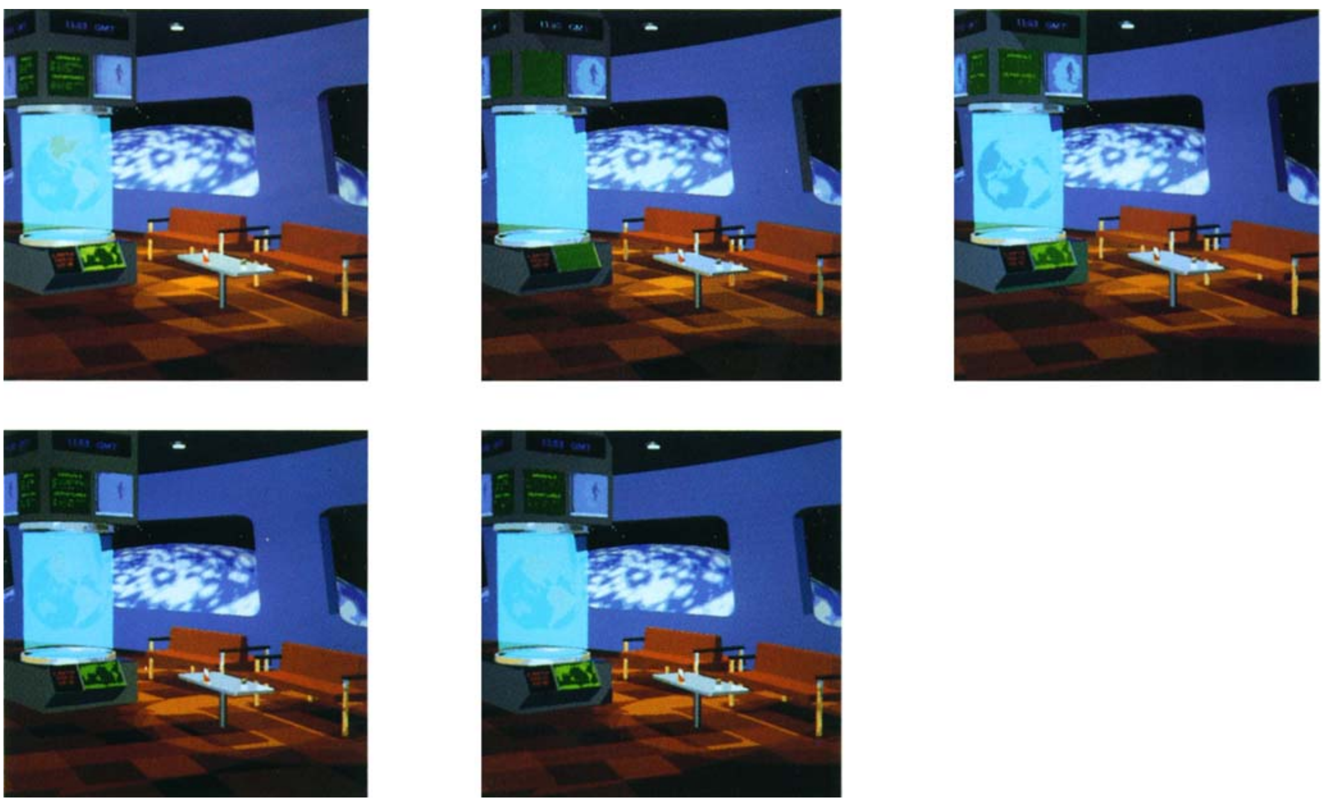

FIG. 2. 32-color quantized images of the *Observation Deck*. (a) The original image; (b) the median-cut algorithm; (c) the mean-split algorithm; (d) the variance-based algorithm; and (e) the k-means algorithm.

FIG. 3. 8-color dithered images of the *Girl.* (a) The median-cut algorithm; and (b) the variance-based algorithm.

1. represent the color of each pixel in the original image as a point in the RGB color space;
2. partition the color space into K disjoint sub-boxes by the variance-based algorithm;
3. choose the centroids of the K sub-boxes to be the representative colors;
4. replace the color of each pixel in the original image by the representative color closest to it;
5. and enhance the quality of the quantized image by the dithering technique (optional).

The occurrence frequencies of colors in an image define a 3D histogram in the *RGB* color space. To keep the histogram size within a reasonable limit (so that no paging is required to fit the histogram in main memory) and to reduce the computational time, we use a color resolution of 5 bits for each color component, that is, the frequency array has $32 \times 32 \times 32$ elements. As Heckbert[2] observed, this resolution is usually sufficient for color image quantization.

The variance-based quantization algorithm successively subdivides the color space Ω into smaller rectangular sub-boxes. A sub-box array is used to keep track of the partition of the color space. Each entry in the array contains the information about the boundaries and the weighted variance of a sub-box. At each step of the subdivision the sub-box with the largest weighted variance is selected to be partitioned into two smaller sub-boxes. To choose an appropriate partition plane, we first obtain the distribution of colors by projecting all points within the sub-box onto each axis. For each projected distribution, we determine the optimal threshold from eq. (12) and compute the weighted sum of projected variances of the two resulting intervals. The partition plane is chosen to be perpendicular to the axis with the smallest weighted sum of projected variances, and passes through the optimal threshold. The sub-box is then replaced by the two smaller sub-boxes. The above procedure is repeated until the total number of sub-boxes is equal to K, the desired number of representative colors used in the quantized image.

The centroids of the K resulting sub-boxes are chosen to be the representative colors. Truncation is performed for the centroids whose entries are not integers. The quantized image f is generated by assigning each color in the original image to its nearest representative color. (Note that the nearest representative color of a color is not necessarily the centroid of the subbox containing it.) This mapping operation can be performed efficiently by employing the locally sorted search technique.[2] Since contouring effects can become objectionable when a small number of representative colors is used, a dithering technique can be applied to enhance the perceived quality of the quantized image.

The detailed pseudo-code of the variance-based color image quantization algorithm is in Table I.

Experimental Results

Two full-color images were used in our experiments: *Girl:* a digitized image with 256×256 pixels and 15606 colors,* and *Observation Deck:* a computer-generated image with 512×512 pixels and 16522 colors.† Each pixel is represented by 24 bits, 8 bits per color component (*RGB*). The components are represented by integers ranging from 0 through 255.

Our experiments were designed for the following purposes: (i) to examine the applicability of the sum-of-squared error to measure the perceived quality of the quantized image; (ii) to evaluate the performance of the variance-based algorithm and compare it with other color image quantization algorithms in terms of quantization error, perceived image quality and computational time; (iii) to compare the perceived quality of the dithered images generated using the colors selected by the variance-based algorithm and the well-known median-cut algorithm.

Quantization Error and Image Quality

The median-cut, mean-split, and variance-based algorithms are applied to quantize the two images using 8, 32, 64, and 256 colors, respectively. In all cases, color reso-

*The USC—Image Processing Institute Data Base, Image Processing Institute, University of Southern California, Los Angeles, 1981.

†Copyright by Darwyn Peachey, University of Saskatchewan, Canada, 1985.

lution of input images is reduced to 5 bits per component before the application of the quantization algorithms. The total quantization errors introduced by these algorithms are listed in Table II.

In contrast to the above algorithms, which terminate once the K representative colors are derived, the k-means procedure starts with K initial colors. They are updated step by step to reduce the quantization error. This process repeats until no further improvement occurs. The initial colors can be arbitrarily chosen, or obtained by another algorithm. The results listed in the last column of Table II were obtained using colors produced by the variance-based algorithm as the initial color set.

Figure 1 shows the 64-color quantized images of the *Girl* generated by different algorithms. The 32-color quantized images of the *Observation Deck* are given in Fig. 2. The quantization errors of all the displayed images are marked by * in Table II.

Contours are apparent on the face of the *Girl* in Fig. 1b (the median-cut algorithm) and on the sleeve in Fig. 1c (the mean-split algorithm). In contrast, Fig. 1d (the variance-based algorithm) exhibits very little contouring effects. The quantization error of Fig. 1e (the k-means algorithm) is slightly lower than that of Fig. 1d. However, the difference between the corresponding images can hardly be observed.

For the *Observation Deck,* the median-cut algorithm fails to reproduce the words on the screens and the map (see Fig. 2b). Furthermore, a contour surrounding the figure on one of the screens can be observed. Although a noticeable improvement is achieved by the mean-split algorithm as shown in Fig. 2c, the contour is still present around the figure, the words on the screens are blurred and part of the floor pattern disappears. On the other hand, the image in Fig. 2d, obtained from the variance-based algorithm, reproduces all the details in the original image. The k-means iterative procedure slightly reduces the quantization error over the initial color selection derived using the variance-based algorithm. However, the perceived image quality does not turn out this way. Fig 2e is apparently worse than Fig. 2d.

The above observations indicate that the quality of the quantized image is basically consistent with the corresponding quantization error. However, in some cases a slightly lower quantization error does not yield a perceivable improvement of image quality. Specifically, in our examples, the k-means iterative procedure reduced the quantization errors over the variance-based algorithm, but this did not lead to images of higher perceived quality. Two questions naturally arise: is it worthwhile to run the k-means algorithm over the color selection derived by the variance-based algorithm, and is there an error measure which would be more consistent with the perceived image quality than the sum-of-squared error? These questions are further discussed in the Conclusions section.

Computational Time

Our experiments were performed on a Silicon Graphics IRIS 3130 machine. Tables III through V present the total run time and the partition time required by different algorithms for different sizes of images and different numbers of representative colors. The same size of 3D histogram ($32 \times 32 \times 32$) is maintained throughout all the experiments. The total run time refers to the period from reading the image data file to completing the display of the quantized image, i.e., step 1 through step 4 in Procedure $VB(h, K)$. The partition time only includes the time required for subdividing the color space into K sub–boxes, i.e., step 2 in Procedure $VB(h, K)$. Most of the time is spent on reading an image from disk and redrawing the quantized image on the screen.

It can be seen from these tables that the total run time and partition time required by the median-cut and mean-split algorithms are approximately the same as that required by the variance-based algorithm. Thus, the improvement in image quality resulting from the use of the variance-based algorithm is achieved without a substantial increase in computational time.

Dithering Techniques

The quality of the quantized images can be improved by dithering techniques. The main effect of dithering is to propagate the quantization error introduced at each pixel to its neighbors. As perceived by the human eye, the average intensity of several pixels in a small region produces an additional color. In this way, the contouring effects can be reduced considerably. The existing dithering techniques include error diffusion,[17,18] ordered dither,[19,20] dot diffusion,[21] Peano curves,[22] and blue-noise dither.[23] The dithering technique used in our experiments is from Floyd and Steinberg.[17] The 8-color dithered images of the *Girl* are shown in Fig. 3. It can be seen that the dithered Fig. 3b based on our selection of color representatives is much better than Fig. 3a produced by using the median-cut algorithm.

TABLE III. Computational time (in seconds) of the median-cut algorithm.

	Girl (256 × 256 pixels)		*Deck* (512 × 512 pixels)	
Colors	Total time	Partition time	Total time	Partition time
8	16	2	66	3
64	18	3	71	5
256	29	5	88	8

TABLE IV. Computational time (in seconds) of the mean-split algorithm.

	Girl (256 × 256 pixels)		*Deck* (512 × 512 pixels)	
Colors	Total time	Partition time	Total time	Partition time
8	17	2	69	4
64	19	4	75	7
256	30	6	91	11

TABLE V. Computational time (in seconds) of the variance-based algorithm.

	Girl (256 × 256 pixels)		*Deck* (512 × 512 pixels)	
Colors	Total time	Partition time	Total time	Partition time
8	17	2	67	3
64	19	4	74	6
256	32	7	93	10

Conclusions

We have shown that under the assumption of pixel independence, the color image quantization problem can be reduced to that of minimizing the sum-of-squared error. Our experimental results indicate that this criterion is basically consistent with the perceived image quality.

It is not surprising that the variance-based algorithm is able to produce quantized images of higher quality than those generated by other methods, because it is specifically designed for minimizing the sum-of-squared error.

Some problems open for further research are listed below:

- Although the *RGB* color space is widely used for color image quantization, it is not necessarily the best choice. An obvious problem with *RGB* is that it is not a perceptually uniform space. In other words, the distance between two colors in the *RGB* space is not necessarily consistent with the perceived difference between these two colors. On the other hand, it has been shown[24] that a 3D luminance/chromaticity space, formed by MacAdam's 2D chromaticity domain complemented with a logarithmic luminance coordinate,[25] exhibits better visual homogeneity and linearity properties. It would be interesting to investigate how much improvement can be gained by partitioning a perceptually uniform color space rather than the *RGB* space.
- In the median-cut, mean-split, and variance-based algorithms, the partition plane is always perpendicular to one of the color axes. This restriction greatly simplifies the algorithm. Obviously, using more general partition planes would lead to smaller quantization errors. It is not clear, however, if the improved image quality obtained by such a technique would warrant the higher computational cost.
- The total quantization error defined by Eq. (4) seems to be a reasonable measure of the difference between the original image and its quantized representation. However, it does not explain the dramatic improvement of the perceived image quality resulting from dithering. It would be useful to introduce a more suitable criterion to measure the similarity between images as perceived by the human eye. Intuitively, such a measure should take into account the spatial correlations among pixels. On the other hand, it seems unlikely that a perfect image similarity measure with only one quantity involved would exist. This is simply because an image usually carries a tremendous amount of information that can hardly be measured by a single parameter.

1. J. D. Foley and A. Van Dam, *Fundamentals of Interactive Computer Graphics*, Addison-Wesley, Reading, 1982.
2. P. Heckbert, Color image quantization for frame buffer display, *Computer Graphics*, **16,** No. 3, 297–307 (1982).
3. S. J. Wan, S. K. M. Wong, and P. Prusinkiewicz, An algorithm for multidimensional data clustering, *ACM Trans. on Mathematical Software*, **14,** No. 2, 135–162 (1988).
4. E. H. Land, The retinex theory of color vision, *Scientific American*, **237,** No. 6, 108–128 (1977).
5. G. Montgomery, Color perception: Seeing with the brain, *Discover*, Dec. 1988, pp. 52–59.
6. R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*, Jone Wiley & Sons, Inc., 1973, pp. 217–237.
7. J. Garey, *Computers and Intractability: A Guide to the Theory of NP Completeness*, W. H. Freemand and Company, 1979.
8. L. Hyafil and R. L. Rivest, Construction optimal binary decision trees is NP-complete, *Information Processing Letters*, **5,** May 15–17 (1976).
9. J. B. MacQueen, Some methods for classification and analysis of multivariate observations, *Proc. Fifth Berkeley Symposium on Mathematical Statistics and Probability*, **1,** 281–297 (1967).
10. S. Z. Selim and M. A. Ismail, K-means-type algorithm: A generalized convergence theorem and characterization of local optimality, *IEEE Trans. on Pattern Analysis and Machine Intelligence*, **PAMI-6,** No. 1, 81–87 (1984).
11. J. L. Bentley and J. H. Frieman, Data structure for range searching, *Computing Surveys*, **11,** No. 4, 397–409 (1979).
12. J. H. Frieman, J. L. Bentley, and R. A. Finkel, An algorithm for finding best matches in logarithmic expected time, *ACM Trans. on Mathematical Software*, **3,** No. 3, 209–226 (1977).
13. X. L. Wu and I. H. Witten, *A Fast K-means Type Clustering Algorithm, Technical Report*, Dept. of Computer Science, Univ. of Calgary, Calgary, Canada, May 1985.
14. R. J. Stevens, A. F. Lehar, and F. H. Perston, Manipulation and presentation of multidimensional image data using the peano scan, *IEEE Trans. on Pattern Analysis and Machine Intelligence*, **PAMI-5,** No. 5, 520–526 (1983).
15. H. C. Koo–Yam–Too, *A Peano-scan Approach to Multivariate Data Clustering, with an Application*, Master Thesis, Dept. of Computer Science, Univ. of Regina, August 1988.
16. S. K. M. Wong, S. J. Wan, and P. Prusinkiewicz, Monochrome image quantization, *Proc. of Canadian Conference on Electrical and Computer Engineering*, Montreal, Canada, Sept. 1989, pp. 28–31.
17. R. W. Floyd and L. Steinberg, An adaptive algorithm for spatial gray scale, *SID. 75, Int. Symp. Dig. Tech. Papers*, 36 (1975).
18. J. F. Jarvis, N. Judice, and W. H. Ninke, A survey of techniques for the display of continuous tone pictures on bilevel displays, *Computer Graphics and Image Processing*, **5,** No. 1, 13–40 (1976).
19. J. O. Limb, Design of dither waveforms for quantized visual signals, *Bell Syst. Tech. J.*, **48,** 2555–2582 (1969).
20. B. Lippel and M. Kurland, The effect of dither on luminance quantization of pictures, *IEEE Trans. Commun. Tech.* **COM-19,** 879–888 (1971).
21. D. E. Knuth, Digital halftones by dot diffusion, *ACM Trans. on Graphics*, **6,** 245–273 (1987).
22. I. H. Witten and R. M. Neal, Using peano curves for bilevel display of continuous-tone images, *IEEE Computer Graphics and Applications*, **2,** No. 3, May 47–52 (1982).
23. R. A. Ulichney, Dithering with blue noise, *Proc. of the IEEE*, Jan, 56–79 (1988).
24. B. J. Kurz, Optimal color quantization for color displays, *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, Washington D.C., USA, Jan. 1983, pp. 217–224.
25. D. L. MacAdam, Geodesic chromaticity diagram based on the variances of color matching, *Applied Optics*, **10,** 1–7 (1971).

Received May 25, 1989; accepted Sept. 12, 1989