

Sequential Scalar Quantization of Vectors: An Analysis

Raja Balasubramanian, *Member, IEEE*, Charles A. Bouman, *Member, IEEE*, and Jan P. Allebach, *Fellow, IEEE*

Abstract— We propose an efficient vector quantization (VQ) technique that we call sequential scalar quantization (SSQ). The scalar components of the vector are individually quantized in a sequence, with the quantization of each component utilizing conditional information from the quantization of previous components. Unlike conventional independent scalar quantization (ISQ), SSQ has the ability to exploit intercomponent correlation. At the same time, since quantization is performed on scalar rather than vector variables, SSQ offers a significant computational advantage over conventional VQ techniques and is easily amenable to a hardware implementation. In order to analyze the performance of SSQ, we appeal to asymptotic quantization theory, where the codebook size is assumed to be large. Closed-form expressions are derived for the quantizer mean squared error (MSE). These expressions are used to compare the asymptotic performance of SSQ with other VQ techniques. We also demonstrate the use of asymptotic theory in designing SSQ for a practical application (color image quantization), where the codebook size is typically small. Theoretical and experimental results show that SSQ far outperforms ISQ with respect to MSE while offering a considerable reduction in computation over conventional VQ at the expense of a moderate increase in MSE.

I. INTRODUCTION

VECTOR quantization (VQ) has found increasing use in data compression applications such as image and speech coding. The technique is an extension of scalar quantization to the case of vectors and is motivated by the well-known result from Shannon's rate distortion theory that superior performance can always be achieved by coding vectors rather than scalars. As with scalar quantization, the objective in VQ design is to minimize a distortion criterion such as mean squared error (MSE). An iterative error minimization algorithm developed by Lloyd for scalar quantization was extended to the vector case by Linde *et al.* [1]. The iterative nature of this algorithm makes it computationally very intensive. A host of other suboptimal but computationally simpler VQ techniques have been reported in the literature either as an alternative to the Linde-Buzo-Gray (LBG) algorithm or as an initial step that may be refined by the LBG technique. Excellent reviews of these techniques may be found in [2]–[4]. Until recently, the general class of VQ techniques has not received much attention for practical implementation because of the high computational

cost and the lack of codebook structures that are amenable to hardware implementation.

In this paper, we investigate a VQ technique we call sequential scalar quantization (SSQ). As the name implies, the basic idea behind the technique is to sequentially quantize the scalar components of a vector, rather than to quantize the vector as a whole. The main computational savings arises from the fact that we are quantizing along scalar rather than vector dimensions. At the same time, due to its sequential nature, SSQ possesses the ability to exploit the correlation and statistical dependency between scalar components of a vector. SSQ falls under the framework of a generalized product code [5], where the features to be quantized are the scalar components of the input vector. More specifically, SSQ is a type of sequential search product code [6], where the quantization of a given feature depends on the input vector and the previously quantized features. The technique may also be viewed as a form of tree-structured VQ [3], [4].

As is the case with any other VQ technique, SSQ attempts to minimize a distortion measure. In order to analyze and optimize the performance of SSQ with respect to this measure, we appeal to asymptotic or high-rate quantization theory, where the number of output quantization levels is assumed to be very large. This theory allows us to derive closed form expressions for the distortion resulting from SSQ as a function of the quantizer design parameters and to find the optimum parameter values that minimize the distortion. It also proves to be a very useful tool in quantizer design even when the number of output levels is small. While SSQ may be used in any scenario that is amenable to vector quantization, we have investigated its use in the application of color image quantization, where a high-quality color image is to be displayed on a low-cost display device with a small palette of colors. Several VQ techniques have been applied to this problem. Braudaway [7] and Gentile *et al.* [8] used the LBG iterative algorithm for palette selection, Orchard and Bouman [9] utilized a tree-structured splitting VQ technique, and Balasubramanian and Allebach [10] reported a merging VQ approach. In [11], we describe in detail an algorithm that employs SSQ for color palette design. The algorithm uses the results of the asymptotic analysis developed in this paper to optimize the palette design with respect to a squared error criterion. We show that with the sequential technique, the palette design is performed very efficiently, whereas the resulting structure of the palette allows the mapping between image pixels and palette colors to be performed with no computation. In addition, the resulting image quality is comparable with or superior to that obtained

Manuscript received June 10, 1993; revised September 8, 1994. This work was supported by an NEC Faculty Fellowship (C.A.B.) and an Eastman Kodak Company Fellowship (R.B.). The associate editor coordinating the review of this paper and approving it for publication was Dr. Fredrick Mintzer.

R. Balasubramanian is with the Xerox Webster Digital Imaging Technology Center, Webster, NY 14580 USA.

C. A. Bouman and J. P. Allebach are with the School of Electrical Engineering, Purdue University, West Lafayette, IN USA.
IEEE Log Number 9413717.

from other color quantization algorithms. In this paper, we focus on a theoretical analysis of SSQ within a general VQ context and include a brief discussion of its application to color quantization. The reader is referred to [11] for details of this application.

II. VECTOR QUANTIZATION

In this section, we formally define the VQ problem and introduce some notation. We will use lower-case letters to denote real variables and vectors, whereas random variables and vectors will be written with upper-case letters. Vectors will be represented by boldface notation. We denote by $P(\cdot)$ the function that produces the probability of an event. We also use an informal notion where $p(\cdot)$ denotes the probability density function of the random variable or vector corresponding to the function's argument. \mathbb{R}^k refers to the k -dimensional space of reals.

Let \mathbf{X} be a random vector in \mathbb{R}^k with probability density function $p(\mathbf{x})$. An N -point k -dimensional vector quantizer $Q: \mathbb{R}^k \rightarrow \mathbb{R}^k$ is a function whose domain is the set of all possible values of \mathbf{X} and whose range is a set of N vectors $\mathcal{C} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ called a codebook. Such a quantizer defines a partition $S = \{S_1, \dots, S_N\}$ of N regions in \mathbb{R}^k , where $S_i = \{\mathbf{x} \in \mathbb{R}^k: Q(\mathbf{x}) = \mathbf{y}_i\}$. The quantization consists of two steps: the codebook design, which involves an appropriate selection of the output vectors $\mathbf{y}_1, \dots, \mathbf{y}_N$, and the mapping of each input vector to one of the output vectors according to the rule $Q(\mathbf{x}) = \mathbf{y}_i$ if $\mathbf{x} \in S_i$. In practice, the vector mapping consists of an encoder that assigns to each input \mathbf{x} a channel symbol and a decoder that maps each channel symbol to a unique output vector in the codebook. Define a distortion measure $d(\mathbf{x}, \mathbf{y})$, $d: \mathbb{R}^k \times \mathbb{R}^k \rightarrow [0, \infty)$. The quantizer is designed to minimize the expected distortion $D_k = E\{d[\mathbf{X}, Q(\mathbf{X})]\}$ between its input and output. Here, $E\{\cdot\}$ denotes expectation. The distortion measure that we will be using is the MSE

$$D_k = \frac{1}{k} E\{\|\mathbf{X} - Q(\mathbf{X})\|^2\} \quad (1)$$

where $\|\cdot\|$ denotes Euclidean distance. Later, we will also discuss weighted squared error measures. The two necessary conditions for a quantizer to be optimal with respect to the MSE distortion are that i) the \mathbf{y}_i are chosen to be the centroid of \mathbf{x} in S_i , i.e., $\mathbf{y}_i = E\mathbf{X}|\mathbf{x} \in S_i$, and ii) each input \mathbf{x} is quantized to one of the \mathbf{y}_i 's according to a nearest neighbor rule, i.e., $S_i = \{\mathbf{x} \in \mathbb{R}^k: \|\mathbf{x} - \mathbf{y}_i\|^2 \leq \|\mathbf{x} - \mathbf{y}_j\|^2, j = 1, \dots, N\}$ [4]. These two conditions are the basis for the iterative codebook design algorithm that was initially proposed by Lloyd for scalar quantization and then generalized to the vector case by Linde *et al.* [1]. The nearest neighbor mapping rule, which involves an exhaustive distance calculation between the input vector and each output vector in the codebook, introduces significant computation in the iterative scheme. Moreover, in general, there is no guarantee that the iterations converge to a global minimum in MSE. (In the 1-D case, the iterations will converge to a global minimum if the distribution $p(x)$ is known to satisfy certain properties such as log concavity [4]). Other suboptimal but more efficient

VQ techniques such as tree-structured VQ have been proposed [4]. Efficient strategies have also been devised to reduce the time taken for nearest neighbor searches [4]. We will use the term conventional VQ to refer to all methods that quantize a vector as a whole entity.

III. SCALAR QUANTIZATION OF VECTORS

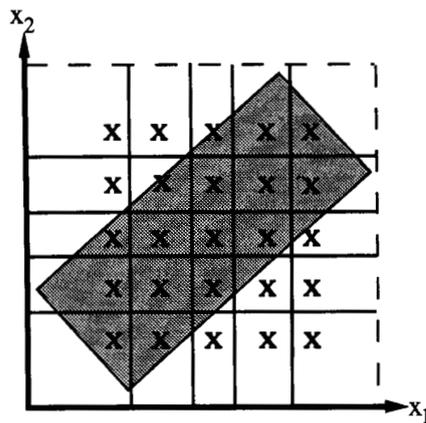
As alluded to above, the primary disadvantage of VQ is its associated complexity, which increases rapidly with the dimensionality of the vector and the codebook size. Another suboptimal but computationally simpler approach to quantize a vector $\mathbf{X} = [X_1, \dots, X_k]^t$ is to quantize each of its individual scalar components X_i , $1 \leq i \leq k$. This may be done either independently or in a sequential fashion.

A. Independent Scalar Quantization (ISQ)

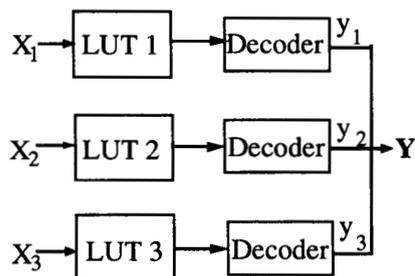
This is the conventional method of scalar quantization. A codebook \mathcal{C}_i of scalar outputs is designed independently for each scalar component X_i , $1 \leq i \leq k$, according to its marginal distribution $p(x_i)$. The final codebook is a k -fold Cartesian product of the k scalar codebooks and is therefore known as a product code. A 2-D example of this scheme is shown in Fig. 1(a) for a rotated uniform distribution. The symbols \mathbf{x} denote output vectors, which are taken to be the centroids within each region. In this example, the codebook size is 25. Vector mapping may be accomplished by independently encoding each X_i to a channel symbol through a set of k lookup tables (LUT's). This is depicted in Fig. 1(b) for $k = 3$. The outputs of the k LUT's are independently decoded to scalar outputs Y_1, \dots, Y_k , which constitute the output vector $\mathbf{Y} = [Y_1, \dots, Y_k]^t = Q(\mathbf{X})$. Since the codebook design only involves quantization of scalar variables, and the encoding operation only entails indexing into LUT's, ISQ involves far less computation than conventional VQ. However, with this scheme, many output vectors are wasted in regions where the input has zero probability of occurrence, as is seen in Fig. 1(a).

B. Sequential Scalar Quantization (SSQ)

With this approach, the first scalar X_1 is quantized to some predetermined number of levels N_1 based on its marginal distribution $p(x_1)$. Each subsequent X_i , $2 \leq i \leq k$ is then quantized based on a set of conditional distributions of X_i within regions formed from the quantization of the scalars X_1, \dots, X_{i-1} . A 2-D example of SSQ is shown in Fig. 2(a) for the same uniform distribution. In this example, first, X_1 is quantized to $N_1 = 5$ levels. This results in a partition of intervals B_{1j} , $1 \leq j \leq 5$, in \mathbb{R} or columns B_{2j} in \mathbb{R}^2 . Next, we quantize X_2 but confine the quantization to the columns formed from the quantization of X_1 . This results in a 2-D quantizer with $N_2 = 18$ output vectors in \mathbb{R}^2 . As shall be seen later, the relative allocations of the N_i , $1 \leq i \leq k$ are design variables whose choice is crucial to the success of the quantizer. The encoding of input vectors may be performed through a sequential or multistage LUT, as shown in Fig. 2(b) for 3-D vectors. The input to the first LUT is X_1 . The output symbol b_{i-1} of the $(i-1)$ th LUT $2 \leq i \leq k$ is then fed to the



(a)

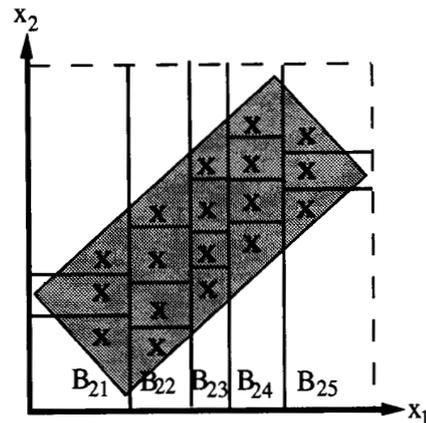


(b)

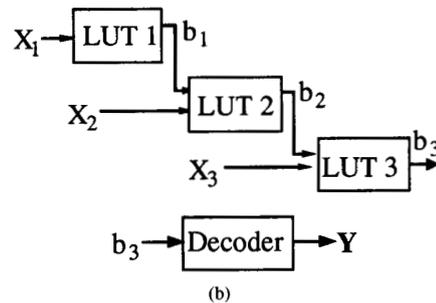
Fig. 1. (a) Two-dimensional example and (b) encoder-decoder operation in independent scalar quantization.

input of the i th LUT along with the i th scalar component X_i . Finally, the output symbol b_k of the last encoder is decoded to one of the output vectors in the codebook C . As is the case with ISQ, the codebook design only involves scalar quantization, whereas the encoding operation entails no computation and is easily amenable to a hardware implementation. Hence, there is a significant computational advantage to be gained by using SSQ rather than conventional VQ methods. In addition, SSQ places its output vectors only within the region of support of the input distribution, thus requiring fewer output codevectors than ISQ to achieve the same quality level.

The codebook storage requirement for any quantization technique will involve, as a minimum, the N output vectors; any additional storage is dictated largely by the encoding strategy. Typically, scalar quantization techniques utilize 1-D LUT's for encoding, whereas VQ techniques often employ tree structures for efficient nearest neighbor encoding. As is seen in Fig. 1(b), ISQ requires k 1-D LUT's to encode each of the components independently. (For the color quantization application, each LUT contains 256 elements.) The storage requirement for SSQ depends on the bit allocations N_i along each component. From Fig. 2(b), we see that the i th stage of the encoder $1 \leq i \leq k$ requires N_{i-1} 1-D LUT's, where $N_0 = 1$. Hence, there are a total of $1 + N_1 + \dots + N_{k-1}$ LUT's. For VQ encoding, binary tree structures such as K-D trees are often employed [4]. These require on the order



(a)



(b)

Fig. 2. (a) Two-dimensional example and (b) encoder-decoder operation in sequential scalar quantization.

of N tree nodes in addition to the N output vectors. Each node may contain several elements to store information about hyperplane decisions. In summary, SSQ demands considerably more storage than ISQ, whereas a comparison between SSQ and VQ will vary for different applications and encoding strategies. In the next section, we compare qualitatively the performance of ISQ and SSQ.

C. Qualitative Comparison of SSQ and ISQ

Makhoul *et al.* [3] provide an excellent qualitative discussion of the advantages of using conventional VQ over ISQ in terms of four properties of the input data: linear dependency, nonlinear dependency, shape of the input distribution, and vector dimensionality. They argue that conventional ISQ can only make use of two of these properties, namely, linear dependency and distribution shape, whereas VQ can exploit all four properties. Lookabaugh and Gray [12] quantified the VQ advantage in terms of the four properties. Here, we qualitatively compare ISQ and SSQ in terms of how well they exploit each of these four properties. Some of the examples of input distributions are taken from [3].

1) *Linear Dependency*: This refers to the statistical correlation between the vector components. The 2-D distribution in Fig. 1(a) represents data that are correlated. As was observed previously, ISQ places some vectors in regions of zero probability because it uses only marginal statistics and cannot take into account the intercomponent correlation. On

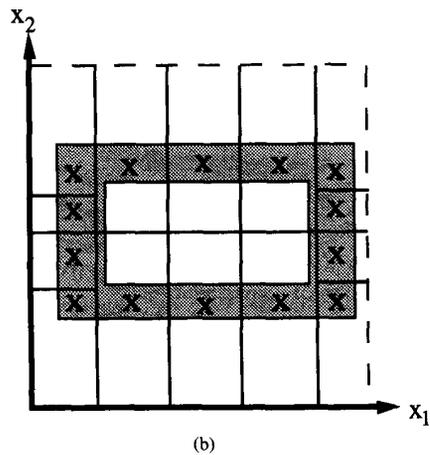
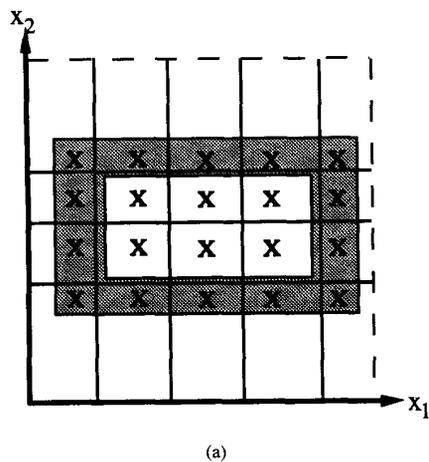


Fig. 3. (a) Independent and (b) sequential scalar quantization of uncorrelated data.

the other hand, since SSQ uses a combination of marginal and conditional statistics, or, effectively, the joint statistics, this technique takes intercomponent correlation into account and places all its output vectors within the support of the distribution. Note that if we rotate the uniform distribution to align with the two axes, then the data are both uncorrelated and independent. In this case, ISQ and SSQ offer equivalent performance.

2) *Nonlinear Dependency*: This is the residual dependency that remains after the correlation between components has been removed. Consider the distribution in Fig. 3, which has a constant value in the shaded area. It is easily shown that X_1 and X_2 are uncorrelated but not independent, i.e., there exists a nonlinear dependency between them. With the ISQ scheme of Fig. 3(a), each scalar is quantized according to its marginal distribution. Since the 2-D codebook is a Cartesian product of the two 1-D codebooks, some output vectors will fall inside the shaded rectangular annulus where the input distribution is zero. The SSQ scheme shown in Fig. 3(b) will, however, place all its vectors in the shaded area, thus again offering performance superior to that of ISQ.

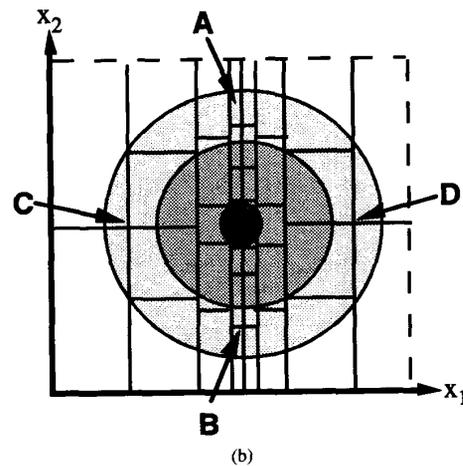
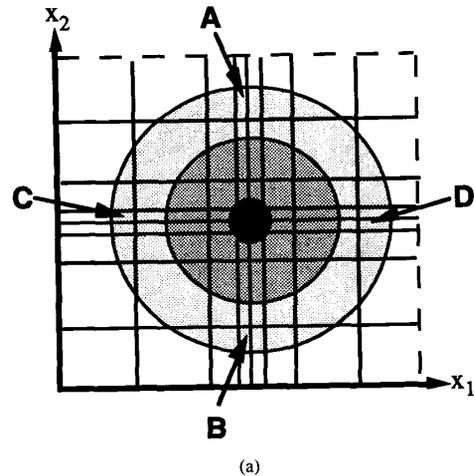


Fig. 4. (a) Independent and (b) sequential scalar quantization of 2-D jointly Gaussian independent data.

3) *Shape of the Input Distribution*: This property refers to the ability of the quantizer to place its output vectors according to the shape of the input distribution in multidimensional space. Ideally, we would expect the output vectors to be more densely spaced where the distribution takes on larger values. Consider two jointly Gaussian random variables with a correlation coefficient of zero. These random variables are uncorrelated and independent. Their distribution is shown schematically in Fig. 4(a) with ISQ and in Fig. 4(b) with SSQ. Notice that ISQ results in densely spaced output vectors in the regions A, B, C, and D, even though the probability density takes on relatively small values in these areas. SSQ suffers from the same drawback only in regions A and B and not around C and D. This subtle difference in how the quantizers space their codevectors yields an interesting and surprising result, namely, that SSQ can outperform ISQ even when the input data is independent. In Section V-C, we will formally show this result for the Gaussian distribution in the asymptotic case where the number of quantization levels becomes large.

4) *Vector Dimensionality*: This property refers to the ability of VQ to pack arbitrarily shaped quantization regions into multidimensional space. The nearest neighbor condition for optimality can often only be achieved with nonrectangular polytopal quantization regions. However, ISQ and SSQ are confined to producing cells that are rectangular polytopes; therefore, they are both inferior to conventional VQ in this respect.

In summary, SSQ can be a powerful quantization technique because although it affords many of the performance advantages of conventional VQ, it also enjoys the computational simplicity of scalar quantization. In the sections below, we will quantify the performance of SSQ and compare it with ISQ and VQ.

IV. ASYMPTOTIC QUANTIZATION THEORY

As pointed out earlier, the minimization of MSE is, in general, a nonlinear iterative problem. Thus, the MSE cannot be written in closed form as a function of the number of quantization levels, except for the most trivial uniform distribution. However, if the number of output quantization levels N is allowed to become asymptotically large, then it is possible to arrive at approximate closed form error expressions [4]. The idea behind asymptotic theory is that the number of output quantization levels is assumed to be large enough (or, equivalently, the quantization cells small enough), and the probability distribution of the source is assumed to be locally smooth enough so that within any quantization cell, this distribution is approximately uniform. The MSE's within each quantization region, which are known in closed form for a uniform distribution, are appropriately summed to yield an approximation for the overall MSE. Conditions have been derived [15] that are less restrictive than the smoothness constraint on the input distribution. The asymptotic analysis not only provides intuition about the behavior of the quantizer but can also serve as a valuable guide in the design of the quantizer even for small N . In the section below, we briefly outline basic asymptotic results that have been developed for scalar and vector quantizers, and refer the reader to [4] for a comprehensive discussion of asymptotic theory. Detailed treatments of this subject may also be found in [12]–[19].

A. Asymptotic Scalar Quantization

Let X be a random variable with distribution $p(x)$. Consider a 1-D quantizer $Q: \mathbb{R} \rightarrow \mathbb{R}$ with N output points. Let $N(x) dx$ be the number of quantization levels that lie in the interval $[x, x + dx]$. We define the quantizer density function $\lambda(x)$ as

$$\lambda(x) = \lim_{N \rightarrow \infty} \frac{N(x)}{N}. \quad (2)$$

Thus, for sufficiently large N , the quantity $N\lambda(x) dx$ is approximately the number of quantization levels in the interval $[x, x + dx]$. It follows by definition that integrating $\lambda(x)$ over its entire domain will result in 1. It may be shown [4] that the MSE of the quantizer may be approximated by an integral

$$D = E\{[X - Y]^2\} \approx \frac{1}{12N^2} \int \frac{p(x)}{\lambda(x)^2} dx \quad (3)$$

where $Y = Q(X)$ is the quantizer output. Equation (3) is known as Bennett's distortion integral. Using Hölder's inequality or the calculus of variations, it may be shown that the function $\lambda(\cdot)$ that minimizes D is given by

$$\lambda(x) = \frac{p(x)^{1/3}}{\int p(x)^{1/3} dx} \quad (4)$$

and the resulting minimum distortion is given by

$$D \approx \frac{1}{12N^2} \|p(x)\|_{1/3}, \quad (5)$$

where we have used the notation

$$\|p(x)\|_m \equiv \left[\int p(x)^m dx \right]^{1/m}. \quad (6)$$

If we wish to find the conditional MSE $E\{[X - Y]^2 | A\}$ given the event A , we simply replace the marginal distribution $p(x)$ in (3)–(6) with the conditional distribution $p(x|A)$.

B. Asymptotic Vector Quantization

Extending the analysis to k -dimensional vectors, let \mathbf{X} be a random vector in \mathbb{R}^k with distribution $p(\mathbf{x})$; $p: \mathbb{R}^k \rightarrow \mathbb{R}$, and let Q be an N point quantizer $Q: \mathbb{R}^k \rightarrow \mathbb{R}^k$. We may define the quantizer density function $\lambda(\mathbf{x})$; $\lambda: \mathbb{R}^k \rightarrow \mathbb{R}$ in a manner analogous to (2) so that $N\lambda(\mathbf{x}) d\mathbf{x}$ is the number of quantization levels in an incremental volume $d\mathbf{x}$ around \mathbf{x} . Na and Neuhoff [14] extend Bennett's distortion integral (3) to the vector case yielding

$$D_k \approx \frac{1}{N^{2/k}} \int \frac{m_k(\mathbf{x})}{\lambda(\mathbf{x})^{2/k}} p(\mathbf{x}) d\mathbf{x} \quad (7)$$

where D_k is the MSE as defined in (1), and $m_k(\mathbf{x})$ is a function characterized by the shapes of the cells in the vicinity of \mathbf{x} known as the inertial profile function. Gersho [13] conjectures that for large N , the optimal quantizer is one whose quantization cells are tessellations of a congruent polytope. This result implies that $m_k(\mathbf{x}) = M_k$ is a constant with respect to \mathbf{x} and may be taken outside the integral in (7). As with the scalar case, (7) can be minimized with respect to $\lambda(\mathbf{x})$, yielding

$$\lambda(\mathbf{x}) = \frac{p(\mathbf{x})^{k/(k+2)}}{\int p(\mathbf{x})^{k/(k+2)} d\mathbf{x}} \quad (8a)$$

and

$$D_k \approx \frac{M_k}{N^{2/k}} \|p(\mathbf{x})\|_{k/(k+2)} \quad (8b)$$

where $\|p(\mathbf{x})\|_m$ is defined in a manner analogous to (6). Equation (8b) serves as a lower bound on MSE performance for an N -point VQ in k -dimensional space. It is easily shown that the inertial profile M_1 for an interval in \mathbb{R} is equal to 1/12 so that (8b) reduces to (5) for the case $k = 1$.

V. ASYMPTOTIC ANALYSIS OF SSQ

We now analyze the asymptotic behavior of SSQ in a manner similar to that used to derive the results in Section IV. Since SSQ is fundamentally different from either ISQ or conventional VQ, we cannot simply use the results of Section IV; rather, we will have to rederive the theory for the sequential structure. The basic results will provide us with the tools necessary for such a derivation. We begin in Section V-A by precisely defining the SSQ structure and making explicit the variables that we are free to choose in the design of the quantizer.

A. The SSQ Design Problem

Given an input random vector $\mathbf{X} = [X_1, \dots, X_k]^t$ in \mathbb{R}^k with distribution $p(\mathbf{x})$, we wish to design a codebook of N output vectors $\mathcal{C} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ by sequentially quantizing the scalar components X_1, \dots, X_k . To simplify the notation, we will assume that the components are quantized in the order X_1, X_2, \dots, X_k . The quantization of a scalar coordinate to a fixed number of levels requires some rule Φ to place the decision and reconstruction levels along that coordinate. This rule will be treated as a design variable.

We begin by designing an N_1 level quantizer for X_1 using the marginal distribution $p(x_1)$. N_1 is a design variable we shall momentarily assume to be fixed. This quantization creates a partition $P'_1 = \{B'_{11}, \dots, B'_{1N_1}\}$ of intervals in \mathbb{R} or a partition $P_2 = \{B_{21}, \dots, B_{2N_1}\}$ of cylinder sets in \mathbb{R}^2 . We denote by B'_{1j} the j th quantization region in \mathbb{R} and by B_{2j} the cylinder set in \mathbb{R}^2 given by $B_{2j} = B'_{1j} \times \mathbb{R}$. Next, we design a quantizer for X_2 , which creates a refinement $P'_2 = \{B'_{21}, \dots, B'_{2N_2}\}$ of P_2 containing N_2 quantization regions. Here, N_2 is again a design variable. The refinement is obtained by quantizing X_2 within each B_{2j} to some predetermined number of levels n_{2j} based on its conditional distribution $p(x_2|B_{2j})$. Referring to Fig. 2(a), we have $n_{21} = 3$, $n_{22} = 4$, etc. Note that the n_{2j} 's must sum to N_2 .

We may generalize this discussion to quantization of the i th scalar component X_i , $2 \leq i \leq k$. The result of quantization along the dimension x_{i-1} is a partition P'_{i-1} of N_{i-1} regions in \mathbb{R}^{i-1} or a partition P_i of N_{i-1} cylinder sets in \mathbb{R}^i . The i th quantizer produces a refinement P'_i of P_i with N_i regions, where $N_i \geq N_{i-1}$. This is accomplished by quantizing along x_i within each of the cylinder sets B_{ij} , $1 \leq j \leq N_{i-1}$ to n_{ij} levels according to the conditional distributions $p(x_i|B_{ij})$. The N_i 's are design variables that must satisfy the constraint $1 \leq N_1 \leq N_2 \leq \dots \leq N_k = N$, where N is the desired number of codewords in \mathbb{R}^k . Furthermore, for each i , $2 \leq i \leq k$, the n_{ij} 's are also design variables that must satisfy the constraint

$$\sum_{j=1}^{N_{i-1}} n_{ij} = N_i. \quad (9)$$

Our objective is to pick the design variables N_i , n_{ij} and the quantization rule Φ along each coordinate x_i to minimize the overall MSE D_k given in (1) for a fixed codebook size $N_k = N$. Since the MSE is a separable distortion measure,

we may write

$$D_k = \frac{1}{k} \sum_{i=1}^k d_i \quad (10)$$

where $d_i = E\{[X_i - Y_i]^2\}$ is the MSE along x_i resulting from the quantization $Y_i = Q(X_i)$. In Section V-B, we will obtain approximate analytical expressions for the d_i 's in the asymptotic case.

B. Asymptotic Optimization of SSQ

We now derive an asymptotic theory for SSQ that provides us with closed-form expressions for the MSE of the quantizer in terms of the aforementioned design variables. This allows the optimization of the quantizer with respect to these variables. For the subsequent analysis, we make the following assumptions:

- 1) The number of quantization levels N_i , $1 \leq i \leq k$ and the relative allocations n_{ij} are large.
- 2) The probability distribution $p(\mathbf{x})$ of the source is relatively smooth.
- 3) The rule Φ for placing decision and reconstruction levels along a scalar dimension x is completely specified by the quantizer density function $\lambda(x)$.
- 4) Quantization along x_i will attempt to minimize only the MSE d_i along that dimension and, in particular, will not affect the MSE's d_1, \dots, d_{i-1} due to previous quantizations;
- 5) The probability of quantizer overload is assumed to be negligible [4].

We will derive in detail the error expressions for the case where \mathbf{X} is a 3-D vector. The results generalize to k dimensions in a straightforward manner. We begin by quantizing X_1 to N_1 levels using the marginal distribution $p(x_1)$. We obtain from (3) the asymptotic approximation to the MSE along x_1

$$d_1 \approx \frac{1}{12N_1^2} \int \frac{p(x_1)}{\lambda(x_1)^2} dx_1. \quad (11)$$

The optimal quantizer spacing $\lambda(x_1)$ that minimizes d_1 is given by (4) with $p(x)$ replaced by $p(x_1)$, and from (5), the resulting minimum d_1 is

$$d_1 \approx \frac{1}{12N_1^2} \|p(x_1)\|_{1/3}. \quad (12)$$

Next, we derive an expression for d_2 , which is the MSE along x_2 , by conditioning on the cylinder sets B_{2j} formed from quantizing X_1

$$\begin{aligned} d_2 &= E\{[X_2 - Y_2]^2\}, \\ &= \sum_{j=1}^{N_1} E\{[X_2 - Y_2]^2 | B_{2j}\} P(B_{2j}) \end{aligned} \quad (13)$$

where, for brevity, we have used $P(B_{2j})$ to denote $P([X_1, X_2]^t \in B_{2j})$. Note that within each B_{2j} , the random variable X_2 has conditional distribution $p(x_2|B_{2j})$

and is quantized to n_{2j} levels. We may use (3) to approximate the conditional distortion within each B_{2j} for large n_{2j}

$$E\{|X_2 - Y_2|^2 | B_{2j}\} \approx \frac{1}{12n_{2j}^2} \int \frac{p(x_2 | B_{2j})}{\lambda(x_2 | B_{2j})^2} dx_2 \quad (14)$$

where $\lambda(x_2 | B_{2j})$ denotes the relative density of quantization levels along x_2 within B_{2j} . We know from (4) that the optimal $\lambda(x_2 | B_{2j})$ that minimizes (14) is

$$\lambda(x_2 | B_{2j}) = \frac{p(x_2 | B_{2j})^{1/3}}{\int p(x_2 | B_{2j})^{1/3} dx_2} \quad (15)$$

Substituting (15) into (14), we get a minimum error expression similar to (5), which can then be substituted into (13) to yield

$$d_2 \approx \frac{1}{12} \sum_{j=1}^{N_1} \frac{1}{n_{2j}^2} \|p(x_2 | B_{2j})\|_{1/3} P(B_{2j}). \quad (16)$$

Now, we are left with the relative allocations n_{2j} , $1 \leq j \leq N_1$, which may be chosen to minimize (16) subject to the constraint (9) that they sum to N_2 . This is a straightforward constrained minimization problem that is easily solved by a Lagrangian method to yield the optimal allocations

$$n_{2j} = N_2 \frac{\|p(x_2 | B_{2j})\|_{1/3} P(B_{2j})^{1/3}}{\sum_{l=1}^{N_1} \|p(x_2 | B_{2l})\|_{1/3} P(B_{2l})^{1/3}} \quad (17)$$

The resulting minimum d_2 is given by

$$d_2 \approx \frac{1}{12N_2^2} \left\{ \sum_{j=1}^{N_1} \|p(x_2 | B_{2j})\|_{1/3} P(B_{2j})^{1/3} \right\}^3 \quad (18)$$

For large N_1 , we use the smoothness restriction on $p(\mathbf{x})$ to rewrite (18) as a Riemann sum and approximate it by an integral (see Lemma 1 of the Appendix)

$$d_2 \approx \frac{N_1^2}{12N_2^2} \cdot \left\{ \int \|p(x_2 | x_1)\|_{1/3} p(x_1) \lambda^2(x_1)^{1/3} dx_1 \right\}^3, \quad (19a)$$

$$= \frac{N_1^2}{12N_2^2} \left(\frac{\|p(x_2 | x_1)\|_{1/3} p(x_1)^{5/3} \|_{1/3}}{\left\{ \int p(x_1)^{1/3} dx_1 \right\}^2} \right) \quad (19b)$$

where (19b) is obtained by substituting $\lambda(x_1)$ from (4) into (19a). Looking at (12) and (19b), we see that the asymptotic assumptions allow us to express the MSE along each coordinate as a product of a term that depends on the number of quantization levels and a constant term that depends on the input statistics. This factorization is an important result, as will be seen shortly.

We can carry through a similar analysis to obtain the MSE d_3 along x_3 . This time, we condition on the cylinder sets

B_{3j} , $1 \leq j \leq N_2$ formed from quantization of X_2

$$\begin{aligned} d_3 &= E\{|X_3 - Y_3|^2\}, \\ &= \sum_{j=1}^{N_2} E\{|X_3 - Y_3|^2 | B_{3j}\} P(B_{3j}), \\ &\approx \frac{1}{12} \sum_{j=1}^{N_2} \frac{1}{n_{3j}^2} \|p(x_3 | B_{3j})\|_{1/3} P(B_{3j}), \end{aligned} \quad (20a)$$

$$\approx \frac{1}{12N_3^2} \left\{ \sum_{j=1}^{N_2} \|p(x_3 | B_{3j})\|_{1/3} P(B_{3j})^{1/3} \right\}^3 \quad (20b)$$

where (20b) is obtained by minimizing (20a) with respect to the n_{3j} subject to the constraint (9) that they sum to N_3 . Assuming that the joint distribution $p(x_1, x_2)$ is locally smooth, we may once again approximate (20b) by an integral

$$d_3 \approx \frac{N_2^2}{12N_3^2} \left\{ \int \int \|p(x_3 | x_1, x_2)\|_{1/3} \cdot p(x_1, x_2) \lambda^2(x_1, x_2)^{1/3} dx_2 dx_1 \right\}^3 \quad (21)$$

where $\lambda(x_1, x_2)$ is the 2-D quantizer density function that defines the relative spacing of output vectors in 2-D space. We may derive an approximate expression for $\lambda(x_1, x_2)$ in terms of our known 1-D marginal and conditional quantizer density functions $\lambda(x_1)$, $\lambda(x_2 | x_1)$ (see Lemma 2 of the Appendix), and substitute this into (21) to yield

$$\begin{aligned} d_3 &\approx \frac{N_2^2}{12N_3^2} \\ &\cdot \left(\|p(x_3 | x_1, x_2)\|_{1/3} p(x_1, x_2)^{5/3} p(x_1)^{4/9} \|_{1/3} \right. \\ &\left. \left\{ \int \int p(x_1, x_2)^{1/3} p(x_1)^{2/9} dx_1 dx_2 \right\}^2 \right). \end{aligned} \quad (22)$$

Once again, we have factored the error as the product of a constant and a term that depends only on the number of quantization levels. Substituting (12), (19b), and (22) into (10), we see that the overall 3-D MSE D_3 is of the form

$$D_3 \approx \frac{1}{36} \left(\frac{1}{N_1^2} \alpha_1 + \frac{N_1^2}{N_2^2} \alpha_2 + \frac{N_2^2}{N_3^2} \alpha_3 \right) \quad (23)$$

where α_i are the constant terms associated with d_i , $1 \leq i \leq 3$. Recall that $N_3 = N$ is the known codebook size. If $\alpha_i > 0$, $i = 1, 2, 3$ (i.e., if there is a probabilistic variation along each x_i), then (23) may be successively minimized with respect to N_1 and N_2 to yield the optimal allocations

$$\begin{aligned} N_1 &= N_2^{1/2} \left(\frac{\alpha_1}{\alpha_2} \right)^{1/4}; \\ N_2 &= N_3^{2/3} \left(\frac{\sqrt{\alpha_1 \alpha_2}}{\alpha_3} \right)^{1/3}. \end{aligned} \quad (24)$$

The resulting distortion is

$$D_3 \approx A_3 \frac{(\alpha_1 \alpha_2 \alpha_3)^{1/3}}{N_3^{2/3}} \quad (25)$$

where $A_3 = 1/12$.

Generalizing to k -dimensional vectors, it may be shown using arguments similar to the one presented above that the overall MSE D_k for an N -point SSQ is given by

$$D_k \approx \frac{1}{12k} \sum_{i=1}^k \frac{N_{i-1}^2}{N_i^2} \alpha_i \quad (26)$$

where $N_0 \equiv 1$ and $N_k = N$. The term α_i is a function of the conditional distribution $p(x_i|x_1, \dots, x_{i-1})$ and all lower order distributions. If $\alpha_i > 0$, $1 \leq i \leq k$, then we may optimize D_k with respect to the N_i 's and obtain expressions similar to (24), where each N_i is defined successively in terms of N_{i+1} . The resulting overall MSE is of the form

$$D_k \approx \frac{A_k}{N^{2/k}} \prod_{i=1}^k \alpha_i^{1/k} \quad (27)$$

where A_k is a dimension dependent number.

Several points are worth noting. First of all, (27) does not, in general, give the globally minimum MSE over all sequential quantizers. This is due to our simplifying assumption that the quantization rule λ along each coordinate is chosen to minimize the MSE only along that coordinate, i.e., the optimization is performed in a greedy rather than joint fashion. Second, we have only considered the case where the scalar components are quantized in the order X_1, X_2, \dots, X_k . In general, since the terms α_i depend on conditional densities, they will be different for different quantization orderings. Hence, in theory, we must evaluate (27) for all $k!$ possible orderings, and pick the one that yields the minimum D_k . In practice, the specific application may impose a natural order of quantization (as we will see later). Failing this, some simple criterion such as variance along each coordinate may be used to determine an order that yields a possibly suboptimal but acceptable solution.

Third, comparing (27) with (8b), which is the lower bound on the MSE that is achieved by the optimal vector quantizer, we see that asymptotically, the MSE for both SSQ and optimal VQ in k dimensions decreases as $1/N^{2/k}$. This implies that a fixed percentage loss in performance that is independent of N will be incurred by using SSQ instead of the optimal VQ method. The loss will depend on the relative values of the constants in (8b) and (27), which in turn depend on the joint, conditional, and marginal statistics of the source X . The suboptimal performance of SSQ is explained by the rectangular cell shapes resulting from this technique and the fact that SSQ may not achieve the optimal density of quantization points in k -dimensional space given by (8a). Finally, we remark that the analysis developed above can be generalized to any r th power distortion measure for $r \geq 1$, where the only restriction is that the measure be separable along the scalar coordinates.

C. Analysis of SSQ for a Gaussian Source Distribution

We will use the analysis above to evaluate and compare the asymptotic performance of SSQ, ISQ, and VQ for the case where the input is a 2-D Gaussian random vector. Let $X = [X_1, X_2]^t$ be characterized by the jointly Gaussian density $N[\mu_1, \mu_2, \sigma_1, \sigma_2, \rho]$. We have used the usual convention that

μ_1 and μ_2 are the means along x_1 and x_2 ; σ_1 and σ_2 are the standard deviations, and ρ is the correlation coefficient. Since μ_1 and μ_2 do not affect the analysis in any way, we will assume them to be zero. We note first of all that if $p(x)$ is a 1-D Gaussian density with variance σ^2 , then $p(x)^m$ when properly normalized is also Gaussian with variance σ^2/m . Second, if $p(x_1, x_2)$ is jointly Gaussian, then the conditional density $p(x_2|x_1)$ is also Gaussian with mean $\rho\sigma_2x_1/\sigma_1$ and variance $\sigma_2^2(1-\rho^2)$. These two facts allow us to easily evaluate the integrals in (12) and (19b) and obtain the constants α_1 and α_2 . Noting that $N_2 = N$ is fixed, we may find the optimal N_1 from (24), substitute this into (12) and (19b), and average the two 1-D MSE's to yield a surprisingly simple expression for the 2-D MSE

$$D_2^{SSQ} \approx \frac{0.777\pi}{N} \sigma_1\sigma_2 \sqrt{1-\rho^2}. \quad (28)$$

Equation (28) tells us that the error is inversely related to the correlation coefficient ρ . This makes sense because the sequential scheme, which uses the conditional distribution $p(x_2|x_1)$ to quantize X_2 , will offer improved performance when X_1 provides more information about X_2 , i.e., when ρ is larger. Note also that for the particular case of a Gaussian source distribution, the performance of SSQ is independent of the order in which the scalar components are quantized.

We may compare SSQ with the lower bound given in (8b) for our example distribution. We will use $M_2 = 0.08$ as this corresponds to a hexagonal packing of cells, which is the optimal scheme in 2-D [13]. This yields the lower bound

$$D_2^{OPT} \approx \frac{0.642\pi}{N} \sigma_1\sigma_2 \sqrt{1-\rho^2}. \quad (29)$$

Since SSQ is restricted to rectangular cell shapes, we would obtain a tighter lower bound on its performance if we let $M_2 = 1/12$ in (8b), which corresponds to VQ with square cell shapes. This results in the fraction 0.642 in (29) being replaced by 0.667. Hence, there is a 16.5% loss in using SSQ over a VQ scheme with square cells. Finally, if we assume that X_1 and X_2 are quantized independently according to their marginal distributions with optimal bit allocations among the two coordinates, then we obtain the following asymptotic error [4]:

$$D_2^{IND} \approx \frac{0.886\pi}{N} \sigma_1\sigma_2. \quad (30)$$

As expected, the performance of ISQ does not depend on the interdata correlation coefficient ρ . As ρ increases, we see the increased advantage of using either SSQ or VQ rather than ISQ. If we let $\rho = 0$ (i.e., X_1 and X_2 are both uncorrelated and independent) and compare (28) and (30), we arrive at the interesting result that SSQ outperforms ISQ even when the components are independent. This result was alluded to earlier and is due to the ability of SSQ to exploit the shape of the input distribution.

D. Weighted Distortion Measures

Although the MSE is a mathematically tractable metric, it may not be an appropriate measure of distortion for a given application. This is certainly true in image processing

problems, where MSE often does not correlate well with visually perceived error. Variations of the squared error metric have been sought [3] to better reflect the application dependent distortion criterion. An example of this is the weighted squared error defined as

$$D_k^W = \frac{1}{k} E\{[X - Y]^t W [X - Y]\} \quad (31)$$

where W is a $k \times k$ positive definite weighting matrix. In general, (31) is not a separable distortion measure and, therefore, does not readily lend itself to the foregoing analysis. However, without loss of generality, we may assume that W is symmetric, in which case, we may factor it into the form $W = P^t W' P$, where W' is a diagonal matrix, and P is a $k \times k$ orthogonal matrix [3]. Letting $X' = P^t X$ and $Y' = P^t Y$, it is easily seen that D_k^W is given by

$$D_k^W = \frac{1}{k} \sum_{i=1}^k w'_{ii} d'_i \quad (32)$$

where w'_{ii} is the i th diagonal element of W' , and $d'_i = E\{[X'_i - Y'_i]^2\}$. Hence, the distortion measure is made separable by an orthogonal transformation of the data, and the foregoing analysis is valid with α_i being replaced by $\alpha'_i = w'_{ii} \alpha_i$ in (23)–(27).

E. Entropy Constrained SSQ

Thus far, we have looked at the problem of minimizing the distortion while keeping the codebook size N fixed. If VQ is the only compression scheme in a given application, then the rate of the coder is given by $R = \log_2(N/k)$ bits per source letter. In coding applications, however, the output of VQ is often additionally compressed with a lossless coder. From Shannon's noiseless source coding theorem, we know that the rate of a lossless coding scheme is bounded from below by the entropy of the input to the coder, which in our case is the output of the VQ. Hence, it would be in our interest to design the VQ to minimize the distortion while constraining the entropy of the output rather than the codebook size.

It has been shown [13] that the minimum entropy VQ scheme for asymptotically large N is one that achieves a uniform distribution of output points, i.e., $\lambda(\mathbf{x})$ is a constant in k -dimensional space. SSQ achieves this uniform distribution in the trivial case where all marginal and conditional quantizer density functions are identically uniform. This is equivalent to quantizing each component independently with a uniform quantizer. Hence, SSQ offers no performance advantage over ISQ if we wish to perform entropy-constrained quantization. On the other hand, VQ has an advantage over both ISQ and SSQ because for the same uniform distribution $\lambda(\mathbf{x})$ of output vectors, VQ with its arbitrary cell shapes can pack its cells in a more efficient way into a finite volume in k -dimensional space than is possible with either SSQ or ISQ.

VI. APPLICATION TO COLOR IMAGE QUANTIZATION

We will briefly describe how the analysis developed above may be used in a practical application, where the assumptions that the number of quantization levels is large and the input

distribution is smooth may no longer be valid. The reader is referred to [11] for details. The problem we address is that of quantizing a color image to a small palette of colors. Such quantization is often necessitated by hardware limitations with low-cost display and printing devices. We will focus on the display application, where the user is often allowed to choose a palette of (usually 256) colors from a much larger set of $2^{24} \approx 16$ million colors.

In the present context, the input to the quantizer is a 3-D RGB color vector X , and the output codebook is the desired palette. We wish to design the quantizer to minimize the MSE for a fixed palette size N . In order to perform the quantization in a visually meaningful space, we transform the image from RGB to $Y C_r C_b$ luminance-chrominance coordinates. A color histogram of the image is generated, which serves as the input probability distribution $p(\mathbf{x})$, and quantization is performed first on the chrominance components, followed by luminance. This ordering is motivated by the intuitively appealing idea of first assigning hues to the various objects in the image and then providing fine luminance shading to each hue.

In general, we have noticed that with color image data, the functional form in (23) is a fairly good model for the experimental MSE when the codebook size is large ($N \gg 256$). However, when the number of quantization vectors becomes small, the constant terms α_i are not well modeled by the integral expressions given in (12), (19b), and (22). To overcome this problem, a preliminary quantization is performed using predetermined initial values for N_1 and N_2 (these were obtained from an empirical study of the values of N_1 and N_2 that minimize the MSE for a large class of images). The experimental MSE's d_1, d_2, d_3 are then computed along the three coordinates, and empirical estimates of α_i are obtained by equating each $d_i, 1 \leq i \leq 3$ to the corresponding term in (23) [11]. Substituting these estimates into (24), we obtain our optimal N_1 and N_2 , which are then used for the final quantization. For the second coordinate X_2 , the number of quantization levels n_{2j} within each $B_{2j}, 1 \leq j \leq N_1$ (see Fig. 2(a)) is given by (17). An analogous expression may be derived for $n_{3j}, 1 \leq j \leq N_2$. We note that while expressions like (17) involve the integration of continuous probability functions, in the digital implementation, these are replaced by summations over discrete marginal and conditional histograms. Other practical considerations such as the integer representation of n_{ij} in (17) are discussed in detail in [11].

While the notion of a quantizer density function $\lambda(\cdot)$ is defined only in the asymptotic case, we use this function to design a 1-D quantizer in the low rate case as follows. In order to quantize X_1 to N_1 levels, we first compute $\lambda(x_1)$ from (4). We then partition the x_1 axis into N_1 intervals such that the area under $\lambda(x_1)$ within each interval is equal to $1/N_1$. It may be shown that with such a scheme, the fraction of quantization levels in an interval $[x, x + dx]$ will approach $\lambda(x) dx$ as N_1 becomes large [11]. The centroid of the data within each interval is then chosen to be the output level for that interval. The same procedure is used to quantize along the subsequent scalar dimensions $x_i, i = 2, 3$ using the conditional quantizer density functions $\lambda(x_i | B_{ij})$. This quantization rule works well when the marginal and

conditional distributions are relatively smooth. However, it can result in a nonoptimal positioning of decision and output levels when there are severe discontinuities in the 1-D histograms. We have developed techniques to detect and correct for such cases [11]. We remark that in the SSQ algorithm, any 1-D quantization strategy (e.g., the Lloyd–Max (LM) algorithm [4]) may be used in place of the aforementioned method, which we henceforth call the lambda technique. We chose the latter because it is consistent with the theory on which the algorithm is based. As we shall see in the results, the lambda technique can be quite competitive with the L-M algorithm for well-behaved source distributions. Furthermore, being noniterative, the former is far less computationally intensive than the latter, which is a consideration that is important for this application.

A common artifact that appears in color quantization is false contouring, where smooth color transitions in the original image are represented by a small number of palette colors. The contouring tends to be more visually objectionable along the luminance coordinate than along chrominance. To alleviate this problem, we use a simple weighted MSE measure, where the MSE along luminance is weighted K times as much as that along chrominance for some $K > 1$.

Finally, the mapping between input image pixels and output palette colors is performed through the sequential LUT shown in Fig. 2(b). In practice, the LUT structure is easily implemented in hardware, thus enabling the pixel mapping to be performed in real time.

VII. EXPERIMENTAL RESULTS

The results presented in this section are intended to verify the theoretical analysis developed in Section V and to assess the performance of SSQ on simulated and experimental data.

A. Verification of SSQ Analysis

We performed quantization using SSQ on 2-D simulated Gaussian data and compared the resulting experimental MSE with the theoretical expression given in (28). Fig. 5 shows a plot of the percentage difference between experimental and theoretical MSE as a function of the number of quantization levels N for two examples of jointly Gaussian distributions. We see that in both cases, as N becomes large, the experimental error converges to (28), thus confirming the validity of the analysis, at least in two dimensions.

B. Scalar Quantization According to $\lambda(x)$

In Section VI, we described an asymptotically optimal quantization rule to partition a scalar dimension into N intervals. Namely, the decision levels are positioned so that the area under the function $\lambda(x)$ given in (4) is equal within each interval. We compared this method with the scalar version of the binary splitting (BSP) algorithm [9] and the LM quantizer for a simulated Gaussian source distribution. Since the Gaussian density is log concave, the LM algorithm yields a globally minimum solution and serves as a lower bound on quantizer performance. The lambda technique was used as a starting point for the LM iterations. Fig. 6 compares the performance of the three quantizers for various N . In order

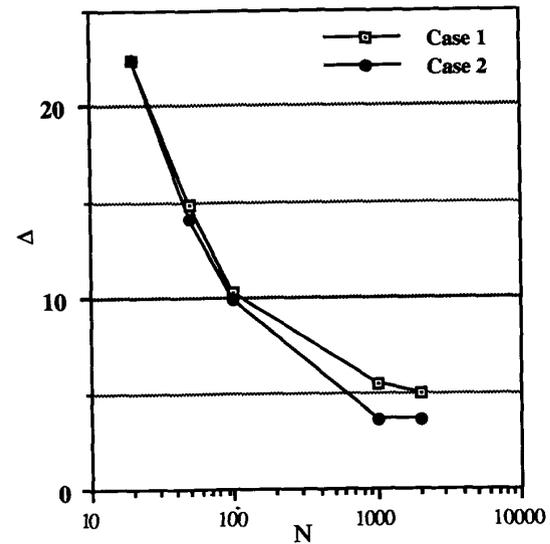


Fig. 5. Percentage discrepancy Δ between experimental and theoretical MSE's as a function of the codebook size N for 2-D Gaussian data (Case 1: $\sigma_1^2 = 1/3$, $\sigma_2^2 = 1/6$, $\rho = 0.8$; Case 2: $\sigma_1^2 = \sigma_2^2 = 1/4$, $\rho = 0$).

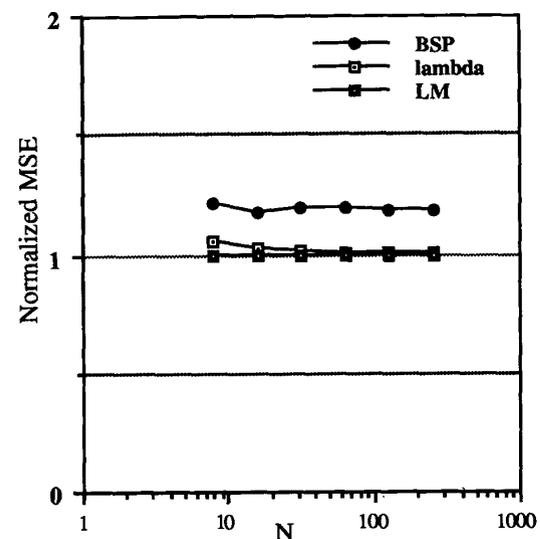


Fig. 6. Comparison of MSE performance of the lambda technique and binary splitting (BSP) relative to the Lloyd–Max (LM) quantizer as a function of the codebook size N for 1-D Gaussian data.

to obtain a meaningful comparison, we normalized the MSE of each algorithm by that of the LM algorithm. For small N , the MSE due to the lambda technique is only about 5% higher than the minimum MSE. As N increases, the lambda technique converges to the minimum MSE solution, verifying that this technique is indeed asymptotically optimal. The BSP algorithm results in a consistently higher relative MSE, suggesting that at least for 1-D Gaussian data, this is not an optimal strategy even in the asymptotic sense. We have also compared the lambda technique and BSP as 1-D quantization strategies in the SSQ algorithm when applied to 3-D color image data and have

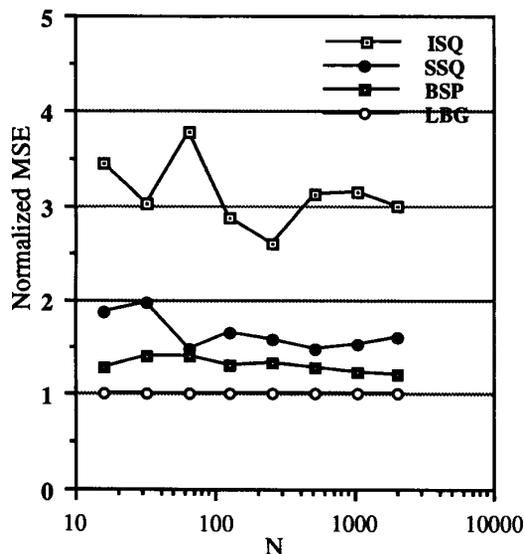


Fig. 7. Comparison of MSE performance of independent scalar quantization (ISQ), sequential scalar quantization (SSQ), and binary splitting (BSP) relative to the Linde-Buzo-Gray (LBG) quantizer as a function of the codebook size N for 3-D color image data.

observed that the lambda technique results in superior image quality.

C. Performance of SSQ on Experimental Data

We now present some results on 3-D color image data that lend support to the qualitative comparison of ISQ, SSQ, and VQ offered in Section III-C. The data was derived from a 512×512 color image, which is shown in Fig. 9. ISQ was designed to quantize each of the three color components independently according to its marginal distribution using the lambda technique. The relative number of quantization levels along each coordinate was chosen in proportion to the variance along that coordinate [3]. SSQ was implemented as described in Section VI, and a fast BSP algorithm [20] was chosen as the conventional VQ technique. Once again, we wished to evaluate these algorithms relative to some lower bound on MSE performance. Since an attempt to find the globally minimum MSE quantizer would be impractical, we used the LBG algorithm [1] to arrive at a local minimum in MSE. We ran the LBG algorithm with ISQ, SSQ, and BSP as starting points and chose the minimum of the three resulting MSE's as a lower bound. Table I compares the MSE performance of the aforementioned four algorithms as a function of the codebook size N . Fig. 7 shows a plot of the MSE's of ISQ, SSQ, and BSP normalized by the minimum MSE from LBG for various N . As anticipated, the performance of SSQ is far superior to that of ISQ, whereas BSP offers a moderate improvement over SSQ. We deduce that the superior performance of SSQ over ISQ is a result of the exploitation by SSQ of linear and nonlinear dependencies in the data, whereas the improved performance of BSP over SSQ is due mainly to the fact that the former is not constrained to cells that are rectangular parallelepipeds. We suspect that the

TABLE I
COMPARISON OF MSE PERFORMANCE OF INDEPENDENT SCALAR QUANTIZATION (ISQ), SEQUENTIAL SCALAR QUANTIZATION (SSQ), BINARY SPLITTING (BSP), AND LINDE-BUZO-GRAY (LBG) QUANTIZER AS A FUNCTION OF THE CODEBOOK SIZE N FOR 3-D COLOR IMAGE DATA

$\text{Log}_2 N$	MSE			
	ISQ	SSQ	BSP	LBG
4	116	63.0	42.9	33.5
5	53.6	35.0	24.8	17.8
6	46.3	18.2	17.1	12.3
7	18.0	10.3	8.13	6.27
8	9.11	5.55	4.63	3.51
9	7.53	3.57	3.05	2.41
10	5.03	2.44	1.96	1.60
11	3.12	1.66	1.24	1.04

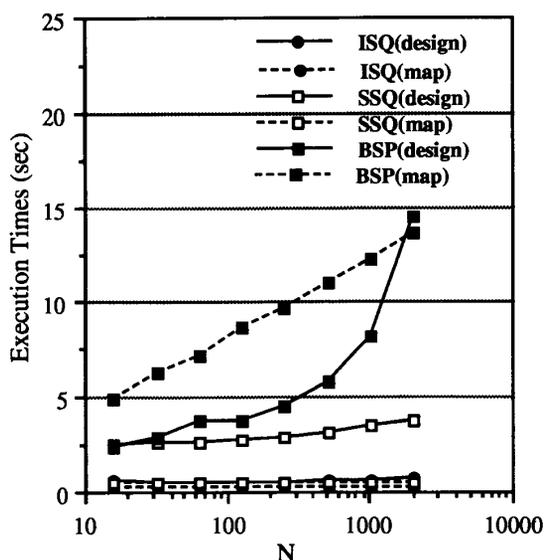


Fig. 8. Comparison of execution times for codebook design and pixel mapping on a SUN SPARC station 2 for independent scalar quantization (ISQ), sequential scalar quantization (SSQ), and binary splitting (BSP) as a function of the codebook size N for 3-D color image data.

improvement in performance of BSP over SSQ would increase with dimensionality, largely because there is more to be gained from the use of unconstrained cell shapes in higher dimensions. This trend can already be seen in going from the 1-D to the 3-D case in Figs. 6 and 7.

Fig. 8 shows the execution times on a SUN SPARC station 2 of ISQ, SSQ, and BSP on the same image data set for various N . We have broken down the execution time of each algorithm into that required for codebook design and that needed to map input vectors (or in this case pixels) to codebook vectors. The latter is often the quantity of greater interest in practical applications. In terms of codebook design,

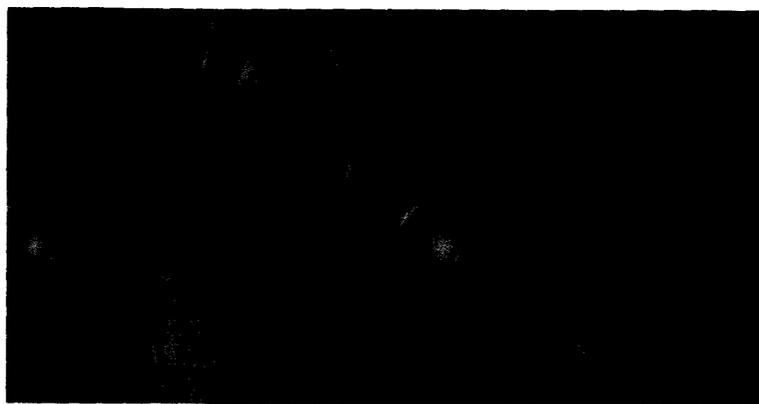


Fig. 9. Comparison of original color image (left) and image quantized to 256 colors using SSQ (right).

SSQ involves more computation than ISQ but is considerably faster than BSP. Moreover, as is characteristic of conventional VQ techniques, the required time for BSP increases with N , whereas this is not the case with ISQ; with SSQ, the rate of increase of computation with N is fairly small. In terms of the mapping step, both SSQ and ISQ use LUT's and thus require no computation. In contrast, the BSP technique uses a binary tree to perform mapping [9], hence requiring a considerable amount of computation that increases with the codebook size. These results are somewhat indicative of the asymptotic complexities of the three algorithms. If N_i is the number of input training vectors, then the complexity of the codebook design is $O(N_i \log_2 N)$ for BSP [9], [20], and $O(N_i)$ for both SSQ and ISQ [11]. If N_t is the number of test vectors to be quantized, then for the mapping, BSP requires $O(N_t \log_2 N)$ computation, whereas SSQ and ISQ require $O(N_t)$ operations (hence, the complexity graphs for the latter two techniques are almost coincident).

Fig. 9 compares the original image with an image that was quantized to 256 colors using the SSQ algorithm described in Section VI. A luminance weighting of $K = 4$ was used. This weighting has the effect of noticeably reducing the visibility of contouring artifacts, hence yielding a visual quality that is superior to that achieved by other techniques. A detailed comparison of the image quality resulting from the various quantization techniques may be found in [11].

VIII. CONCLUSION

We have proposed an efficient technique for quantizing vectors called sequential scalar quantization. A theoretical analysis of SSQ in the asymptotic case yields intuitive and useful results that allow us both to compare this method with other quantization strategies and to design a practical quantizer. We have theoretically and experimentally demonstrated the fact that SSQ yields considerably improved performance over conventional independent scalar quantization while offering a significant computational advantage over conventional VQ. Moreover, the resulting sequential structure of this technique

lends itself very easily to a hardware embodiment. Although we have successfully applied SSQ to the color image quantization problem, this technique is of potential use in any vector quantization application where computational cost is an important consideration, and a moderate loss in quantitative performance can be tolerated.

We remark finally that while the theory derived in this paper holds for arbitrary vector dimension k , a practical implementation of SSQ for the case of low rate and high dimensionality may be a nontrivial problem largely because of the bit allocation issue. It is likely that some of the components may be allocated less than 1 bit (i.e., $N_i = 1$). In such cases, the asymptotic assumption is severely violated, and one may need to resort to the use of heuristics, or the statistics of the data, in order to fine tune the algorithm for the application at hand.

APPENDIX

Lemma 1: Let \mathbf{X} be a random vector in \mathbb{R}^k with probability density function $p(\mathbf{x})$. For $p(\mathbf{x})$ sufficiently smooth and N_1 sufficiently large, the following approximation holds:

$$\sum_{j=1}^{N_1} [||p(x_2|B_{2j})||_{1/3} P(B_{2j})]^{1/3} \approx N_1^{2/3} \int [||p(x_2|x_1)||_{1/3} p(x_1)\lambda(x_1)^2]^{1/3} dx_1 \quad (\text{A.1})$$

where B_{2j} and $\lambda(x_1)$ are as defined in Section V.

Proof: Recall that quantization along x_1 divides this dimension into N_1 intervals B'_{1j} , $1 \leq j \leq N_1$. Let Δ_j be the width of interval B'_{1j} , and let y_j be any point within B'_{1j} . Given that N_1 is large and $p(\mathbf{x})$ is smooth, we may assume that the marginal density $p(x_1)$ is approximately constant within any interval Δ_j so that

$$\begin{aligned} P(B_{2j}) &= P([X_1, X_2]^t \in B_{2j}) \\ &= P(X_1 \in B'_{1j}) \approx p(y_j)\Delta_j. \end{aligned} \quad (\text{A.2})$$

Consider an interval $[y_j, y_j + dy_j]$ around the point y_j , where $dy_j \gg \Delta_j$. We have

$$\begin{aligned} \Delta_j &= \frac{(\text{length of interval } [y_j, y_j + dy_j])}{(\# \text{ of quantization levels in } [y_j, y_j + dy_j])}, \\ &= \frac{dy_j}{N_1 \lambda(y_j)} \\ &= \frac{1}{N_1 \lambda(y_j)}. \end{aligned} \quad (\text{A.3})$$

Now

$$P(B_{2j})^{1/3} \approx [p(y_j) \Delta_j]^{1/3} \quad [\text{from (A.2)}], \quad (\text{A.4a})$$

$$\begin{aligned} &= \left[\frac{p(y_j)}{\Delta_j^2} \right]^{1/3} \Delta_j \\ &= [N_1^2 p(y_j) \lambda^2(y_j)]^{1/3} \Delta_j \quad [\text{from (A.3)}]. \end{aligned} \quad (\text{A.4b})$$

In addition, for small Δ_j , we have

$$p(x_2 | B_{2j}) = p(x_2 | X_1 \in B'_{1j}) \approx p(x_2 | y_j). \quad (\text{A.5})$$

Substituting (A.4b) and (A.5) into the left hand side of (A.1), we have

$$\begin{aligned} &\sum_{j=1}^{N_1} [||p(x_2 | B_{2j})||_{1/3} P(B_{2j})]^{1/3} \\ &\approx N_1^{2/3} \sum_{j=1}^{N_1} [||p(x_2 | y_j)||_{1/3} p(y_j) \lambda(y_j)^2]^{1/3} \Delta_j. \end{aligned} \quad (\text{A.6})$$

The right-hand side of (A.6) is a Riemann sum that may be approximated by the integral on the right-hand side of (A.1).

Lemma 2: Let the quantizer density function $\lambda(x_1)$ along x_1 be given by (4) with $p(x)$ replaced by $p(x_1)$. For each fixed x_1 , let the quantizer density function $\lambda(x_2 | x_1)$ along x_2 be given by (4) with $p(x)$ replaced by $p(x_2 | x_1)$. Then, the 2-D quantizer density function $\lambda(x_1, x_2)$ is given by

$$\lambda(x_1, x_2) \approx \frac{p(x_1, x_2)^{1/3} p(x_1)^{2/9}}{\int \int p(x_1, x_2)^{1/3} p(x_1)^{2/9} dx_2 dx_1}. \quad (\text{A.7})$$

Proof: Consider an incremental rectangular area in 2-D space $\Delta(y_{1j}, y_{2k}) = \Delta(y_{1j}) \Delta(y_{2k})$ surrounding a quantization value $[y_{1j}, y_{2k}]^t$, $1 \leq j \leq N_1$, $1 \leq k \leq n_{2j}$. By definition of $\lambda(\cdot)$, we may make the following approximations:

- 1) The number of quantization levels along x_1 within this interval is $N_1 \lambda(y_{1j}) \Delta(y_{1j})$.
- 2) For each quantization level y_{1j} along x_1 within this interval, the number of quantization levels along x_2 is approximately $n_{2j} \lambda(y_{2k} | y_{1j}) \Delta(y_{2k})$, where n_{2j} is as defined in Section V-A.
- 3) The total number of quantization levels within $\Delta(y_{1j}, y_{2k})$ is $N_2 \lambda(y_{1j}, y_{2k}) \Delta(y_{1j}, y_{2k})$.

We assume that the quantity in 2) is constant for each y_{1j} within a small enough interval $\Delta(y_{1j})$. Therefore, we can approximate the quantity in 3) by a product of 1) and 2)

$$\begin{aligned} N_2 \lambda(y_{1j}, y_{2k}) \Delta(y_{1j}, y_{2k}) &\approx [N_1 \lambda(y_{1j}) \Delta(y_{1j})] \\ &\quad \cdot [n_{2j} \lambda(y_{2k} | y_{1j}) \Delta(y_{2k})] \\ \Rightarrow \lambda(y_{1j}, y_{2k}) &\approx \frac{N_1}{N_2} n_{2j} \lambda(y_{2k} | y_{1j}) \lambda(y_{1j}). \end{aligned} \quad (\text{A.8})$$

We may substitute n_{2j} with the optimum choice given in (17)

$$\begin{aligned} \lambda(y_{1j}, y_{2k}) &\approx N_1 \frac{[||p(x_2 | B_{2j})||_{1/3} P(B_{2j})]^{1/3}}{\sum_{l=1}^{N_1} [||p(x_2 | B_{2l})||_{1/3} P(B_{2l})]^{1/3}} \\ &\quad \cdot \lambda(y_{2k} | y_{1j}) \lambda(y_{1j}). \end{aligned} \quad (\text{A.9})$$

We now approximate the discrete probabilities $P(B_{2j})$ and $P(B_{2l})$ by continuous probability density functions. Substituting (A.3) into (A.2) to approximate $P(B_{2j})$ in the numerator above, and using (A.4b) to approximate $P(B_{2l})^{1/3}$ in the denominator, we have

$$\begin{aligned} \lambda(y_{1j}, y_{2k}) &\approx \frac{[||p(x_2 | B_{2j})||_{1/3} p(y_{1j})]^{1/3} \lambda(y_{1j})^{-1/3}}{\sum_{l=1}^{N_1} [||p(x_2 | B_{2l})||_{1/3} p(y_{1l}) \lambda(y_{1l})^2]^{1/3} \Delta_l} \\ &\quad \cdot \lambda(y_{2k} | y_{1j}) \lambda(y_{1j}). \end{aligned} \quad (\text{A.10})$$

As was done in Lemma 1, we argue that for small $\Delta(y_{1j})$ and D_l , we may replace B_{2j} and B_{2l} above with y_{1j} and y_{1l} , respectively. Finally, we replace all discrete quantization variables y_{1j}, y_{2k} with continuous variables x_1, x_2 and approximate the Riemann sum in the denominator of (A.10) by an integral to obtain

$$\begin{aligned} \lambda(x_1, x_2) &\approx \frac{[||p(x_2 | x_1)||_{1/3} p(x_1)]^{1/3}}{\int [||p(x_2 | x_1)||_{1/3} p(x_1) \lambda(x_1)^2]^{1/3} dx_1} \\ &\quad \cdot \lambda(x_2 | x_1) \lambda(x_1)^{2/3}. \end{aligned} \quad (\text{A.11})$$

If we now substitute the appropriate expressions for $\lambda(x_1)$ and $\lambda(x_2 | x_1)$ according to the hypothesis of the lemma, we obtain the desired result (A.7).

ACKNOWLEDGMENT

The authors are grateful to Eastman Kodak Company for the use of their image "balloon" in the experiments. They also thank the referees for their useful comments and suggestions.

REFERENCES

- [1] Y. Linde, A. Buzo, and R. M. Gray, "An algorithm for vector quantizer design," *IEEE Trans. Commun.*, vol. COM-28, pp. 84-95, Jan. 1980.
- [2] N. M. Nasrabadi and R. A. King, "Image coding using vector quantization: A review," *IEEE Trans. Commun.*, vol. 36, pp. 957-971, Aug. 1988.
- [3] J. Makhoul, S. Roucos, and H. Gish, "Vector quantization in speech coding," *Proc. IEEE*, vol. 73, pp. 1551-1588, 1985.
- [4] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*. Norwell, MA: Kluwer, 1991.
- [5] W.-Y. Chan and A. Gersho, "Generalized product codes: A framework for the design of structured vector quantizers," in *Proc. IEEE Int. Symp. Circuits Syst.*, San Diego, vol. 5, May 1992, pp. 2272-2275.

- [6] ———, "Enhanced multistage vector quantization with constrained storage," in *Proc. 24th Asilomar Conf. Circuits Syst. Comput.*, Nov. 1990, pp. 659–663.
- [7] G. Braudaway, "A procedure for optimum choice of a small number of colors from a large color palette for color imaging," in *Proc. Electron. Imaging '87*, San Francisco, CA, 1987.
- [8] R. S. Gentile, J. P. Allebach, and E. Walowit, "Quantization of color images based on uniform color spaces," *J. Imaging Technol.*, vol. 16, no. 1, pp. 12–21, Feb. 1990.
- [9] M. T. Orchard and C. A. Bouman, "Color quantization of images," *IEEE Trans. Signal Processing*, vol. 39, no. 12, pp. 2677–2690, Dec. 1991.
- [10] R. Balasubramanian and J. P. Allebach, "A new approach to palette selection for color images," *J. Imaging Technol.*, vol. 17, no. 6, pp. 284–290, Dec. 1991.
- [11] R. Balasubramanian, C. A. Bouman, and J. P. Allebach, "Sequential scalar quantization of color images," *J. Electron. Imaging*, vol. 3, no. 1, pp. 45–59, Jan. 1994.
- [12] T. D. Lookabaugh and R. M. Gray, "High-resolution quantization theory and the vector quantizer advantage," *IEEE Trans. Inform. Theory*, vol. 35, no. 5, pp. 1020–1033, Sept. 1989.
- [13] A. Gersho, "Asymptotically optimal block quantization," *IEEE Trans. Inform. Theory*, vol. IT-25, pp. 373–380, July 1979.
- [14] S. Na and D. L. Neuhoff, "Bennett's integral for vector quantizers, and applications," in *1990 IEEE Int. Symp. Inform. Theory*, Jan. 1990.
- [15] J. A. Bucklew and G. L. Wise, "Multidimensional asymptotic quantization theory with r th power distortion measures," *IEEE Trans. Inform. Theory*, vol. IT-28, no. 2, pp. 239–247, Mar. 1982.
- [16] J. A. Bucklew, "Companding and random quantization in several dimensions," *IEEE Trans. Inform. Theory*, vol. IT-27, no. 2, pp. 207–211, Mar. 1981.
- [17] H. Gish and J. N. Pierce, "Asymptotically efficient quantizing," *IEEE Trans. Inform. Theory*, vol. IT-14, pp. 676–683, Sept. 1968.
- [18] Y. Yamada, S. Tazaki, and R. M. Gray, "Asymptotic performance of block quantizers with difference distortion measures," *IEEE Trans. Inform. Theory*, vol. IT-26, no. 1, pp. 6–14, Jan. 1980.
- [19] P. L. Zador, "Asymptotic quantization error of continuous signals and the quantization dimension," *IEEE Trans. Inform. Theory*, vol. IT-28, no. 2, pp. 139–149, Mar. 1982.
- [20] R. Balasubramanian, C. A. Bouman, and J. P. Allebach, "New results in color image quantization," in *Proc. 1992 SPIE/SPSE Symp. Electron. Imaging-Sci. Technol.*, San Jose, CA, Feb. 10–13, 1992.



Raja Balasubramanian (M'95) received the B.S. degree from the University of Texas at Arlington, in 1987 and the M.S. and Ph.D. degrees from Purdue University, in 1988 and 1992, respectively, all in electrical engineering.

During his graduate studies, he conducted research in the areas of digital color image quantization and halftoning. During the summer of 1991, he worked at Eastman Kodak Company in the areas of color image quantization and rendering. He is currently employed at Xerox Webster Research

Center, where he is involved with research in color hardcopy reproduction. His other research interests include the use of human visual models in image processing and the use of color in computer vision and image understanding.

Dr. Balasubramanian is a member of IS&T, Eta Kappa Nu, and Tau Beta Pi.



Charles A. Bouman (M'89) was born in Philadelphia, PA. He received the B.S.E.E. degree from the University of Pennsylvania in 1981 and the M.S. degree in electrical engineering from the University of California at Berkeley in 1982. In 1987 and 1989, he received the M.A. and Ph.D. degrees in electrical engineering from Princeton University.

From 1982 to 1985, he was a staff member in the Analog Device Technology Group at the Massachusetts Institute of Technology, Lincoln Laboratory.

In 1989, he joined the faculty of the School of Electrical Engineering at Purdue University as an Assistant Professor. His research interests include problems in electronic imaging, image segmentation, analysis, and tomography.

Dr. Bouman is a member of SPIE and IS&T and is currently an associate editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING.

Jan P. Allebach (F'91) received the B.S. degree from the University of Delaware in 1972 and the M.S. and Ph.D. degrees from Princeton University in 1975 and 1976, respectively, all in electrical engineering.

He was with the Department of Electrical Engineering at the University of Delaware from 1976 to 1983. Since 1983, he has been with the School of Electrical Engineering at Purdue University, where he holds the rank of Professor. His research has primarily been in the areas of printing and display of images, scanning and sampling of multidimensional signals, and synthesis of digital diffractive elements. He has been a visiting summer faculty member at IBM Watson Research Center in Yorktown Heights, NY, Sandia Labs in Livermore, CA, and Hewlett-Packard Labs in Palo Alto, CA. In addition, he has consulted extensively for industry and government laboratories.

From 1972 to 1975, Prof. Allebach was a National Science Foundation Graduate Fellow. He is a member of the Optical Society of America and the IS&T. In 1987, he received the Senior Award from the IEEE Signal Processing Society for his paper on time-sequential sampling published in the February 1984 issue of the IEEE TRANSACTIONS ON SIGNAL PROCESSING. During 1991–1992, he served as the IS&T Visiting Lecturer. During 1994, he is a Distinguished Lecturer for the IEEE Signal Processing Society. He chaired the Image and Multidimensional Signal Processing Technical Committee of the IEEE Signal Processing Society from 1991 to 1992. He is a past Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING and currently serves on the Editorial Board for the *Journal of Electronic Imaging*. Presently, he is Chapters Vice-President of IS&T and Secretary for the IEEE Signal Processing Society.