Suboptimality of the Karhunen–Loève Transform for Transform Coding

Michelle Effros, *Senior Member, IEEE*, Hanying Feng, *Member, IEEE*, and Kenneth Zeger, *Fellow, IEEE*

***Abstract*—We examine the performance of the Karhunen–Loève transform (KLT) for transform coding applications. The KLT has long been viewed as the best available block transform for a system that orthogonally transforms a vector source, scalar quantizes the components of the transformed vector using optimal bit allocation, and then inverse transforms the vector. This paper treats fixed-rate and variable-rate transform codes of non-Gaussian sources. The fixed-rate approach uses an optimal fixed-rate scalar quantizer to describe the transform coefficients; the variable-rate approach uses a uniform scalar quantizer followed by an optimal entropy code, and each quantized component is encoded separately. Earlier work shows that for the variable-rate case there exist sources on which the KLT is not unique and the optimal quantization and coding stage matched to a "worst" KLT yields performance as much as 1.5 dB worse than the optimal quantization and coding stage matched to a "best" KLT. In this paper, we strengthen that result to show that in both the fixed-rate and the variable-rate coding frameworks there exist sources for which the performance penalty for using a "worst" KLT can be made arbitrarily large. Further, we demonstrate in both frameworks that there exist sources for which even a best KLT gives suboptimal performance. Finally, we show that even for vector sources where the KLT yields independent coefficients, the KLT can be suboptimal for fixed-rate coding.**

***Index Terms*—Bit allocation, lossy source coding, quantization, transform coding.**

I. Introduction

THE Karhunen–Loève transform (KLT) plays a fundamental role in a variety of disciplines, including statistical pattern matching, filtering, estimation theory, and source coding, and has been used, for example, in certain image coding applications (e.g., [1]). In many of these applications, the KLT is known to be "optimal" in various senses. This paper investigates the optimality of the KLT for source coding.

The main application of the KLT in source coding is in scalar quantized transform coding. In this type of transform code, an input vector is linearly transformed into another vector of the same dimension; the components of that vector are then described to the decoder using independent scalar quantizers on the coefficients. We consider both fixed-rate and variable-rate codes. The fixed-rate code uses an optimal fixed-rate scalar quantizer; the variable-rate code uses uniform scalar quantization followed by optimal entropy coding on each quantized component separately. The decoder reconstructs the quantized transform vector and then uses a linear transformation to get an estimate of the original input vector. In both cases, the goal is to find the pair of linear transforms and the allocation of an average bit budget among the scalar quantizers that together minimize the end-to-end distortion. In this work, we measure distortion as mean squared error (MSE). Unlike the best possible vector quantizer, a best possible transform code is suboptimal in general. Nonetheless, transform codes have served as important models for gaining an understanding of both optimal quantization and low-complexity code design.

In [2], Huang and Schultheiss show that if the vector source is Gaussian and the bit budget is asymptotically large, then the KLT and its inverse are an optimal pair of transforms for fixed-rate coding. In a more recent paper, Goyal, Zhuang, and Vetterli ([3, p. 1628, Proof 1 of Theorem 6]) and Telatar ([3, p. 1629, Proof 2 of Theorem 6]) improve that result by showing that the KLT is optimal for Gaussian inputs without making any high resolution assumptions. Their results apply to both the fixed-rate and the variable-rate coding models.

Intuitively, the optimality of the KLT in transform coding of Gaussian sources is typically explained by the assertion that scalar quantization is better suited to the coding of independent random variables than to the coding of dependent random variables. Thus, the optimality of the KLT for transform coding of Gaussian sources is believed to be a consequence of the fact that the KLT of a Gaussian vector yields independent transform coefficients. The application of the KLT in transform coding of non-Gaussian sources is then justified using the intuition that the KLT's coefficient decorrelation is, for general sources, the best possible approximation to the desired coefficient independence. In [4], Koschman shows that if we forgo optimal bit allocation and instead force a fixed number of the transform coefficients to be quantized at rate zero and the remaining components to be quantized with infinite accuracy, then for any stationary source the KLT minimizes the MSE over all possible choices of orthogonal transforms (this is known as "zonal coding"). While this result does not address the bit allocation problem, it seems to further support the above intuition. Over the years, this intuition has been taken for granted, and numerous references to

Manuscript received April 25, 2002; revised November 6, 2003. This material is based upon work supported in part by the National Science Foundation under Grants CCR-9909026, CCR-0073489, and CCR-0105734, an Intel equipment grant, the UCSD Center for Wireless Communications, and the Lee Center for Advanced Networking. The material in this paper was presented in part at the IEEE Data Compression Conference, Snowbird, UT, March 2003.

M. Effros is with the Department of Electrical Engineering, MC 136-93, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: effros@caltech.edu).

H. Feng is with the Department of Electrical Engineering, Mail Stop 4075, Stanford University, Stanford, CA 94305 USA (e-mail: fhy@stanford.edu).

K. Zeger is with the Department of Electrical and Computer Engineering, University of California, San Diego, La Jolla CA 92093-0407 USA (e-mail: zeger@ucsd.edu).

Communicated by R. Zamir, Associate Editor for Source Coding.

Digital Object Identifier 10.1109/TIT.2004.831787

0018-9448/04$20.00 © 2004 IEEE

the "optimality" of the KLT for transform coding appear in texts and scholarly journals.

In this paper, we investigate this intuition piece by piece, demonstrating both its failures and its successes. All results apply in the high rate limit (although some also hold at lower rates) and to non-Gaussian sources, including some "near-Gaussian" sources. First, we consider the question of sample decorrelation, showing that sample decorrelation is neither sufficient nor necessary for transform optimality in either fixed-rate or variable-rate transform coding. Then, we show for the fixed-rate case that even for examples where decorrelation yields coefficient independence, the KLT may fail to yield the optimal performance. Finally, we show for the variable-rate case that for examples where decorrelation yields coefficient independence, the KLT guarantees the optimal performance in the high resolution limit for suitably smooth distributions. Note that in a variable-rate coding framework, each quantized component is encoded separately.

We demonstrate that sample decorrelation is not sufficient for transform coding optimality by considering examples where the KLT is not unique and demonstrating that for these examples some choices of the KLT give suboptimal transform coding performance. The first such example appears in [5], which shows that an optimal variable-rate transform code (i.e., the quantization and coding stage) matched to a "worst" KLT can give a signal to quantization noise ratio (SQNR) as much as 1.5 dB worse than an optimal variable-rate transform code matched to a "best" KLT and also demonstrates that some previous intuition for more general sources is not valid.[1] In this paper, we extend that result to show that the performance penalty associated with using a worst KLT rather than a best KLT can be made arbitrarily large in both the fixed-rate and the variable-rate transform coding scenarios. This result occurs at distortions close to zero, where small distortion changes can cause arbitrarily large differences in SQNR. The basic idea for the proof is that when a vector is composed of independent and identically distributed (i.i.d.) random variables, then a variety of linear transformations tend to make the source appear more Gaussian, and thus "harder" to compress (e.g., see [6]).

We then show that even a best (or only) KLT can give suboptimal performance. We prove this by demonstrating that there exist non-Gaussian sources for which it is better to scalar quantize the correlated source symbols than to scalar quantize the uncorrelated coefficients of a best KLT. These results apply to both fixed-rate and variable-rate coding.

To test the intuition that scalar quantization gives better performance for independent random variables than for dependent random variables, we consider examples of dependent source vectors for which the KLT yields independent transform coefficients. For the fixed-rate coding scenario, we find an example that proves that this intuition is also incorrect. In particular, we demonstrate that there exist sources for which optimal fixed-rate scalar quantization of the dependent source symbols yields better performance than optimal fixed-rate scalar quantization of the independent KLT coefficients. In contrast, for the variable-rate coding scenario, we show that a transform that yields independent coefficients is optimal for variable-rate coding when the high-rate approximation applies.

The remainder of this paper is organized as follows. Section II introduces background material, notation, and definitions. Section III lists our main results. The proofs of these theorems appear in Sections IV–VI.

II. Preliminaries

Denote the entropy of a discrete random variable Z taking on outcomes in $\{z_1, z_2, \ldots\}$ by[2]

$$H(Z) = -\sum_i P(Z = z_i) \log P(Z = z_i)$$

and denote the differential entropy of a continuous random variable Z with probability density function (pdf) f by

$$h(Z) = -\int f(z) \log f(z) dz.$$

If Z is a continuous random vector with components $Z_1, \ldots, Z_n$, then define the quantity

$$\overline{h}(Z) = \frac{1}{n} \sum_{i=1}^{n} h(Z_i).$$

The usual notation $\mathcal{N}(m, \sigma^2)$ denotes the pdf of a scalar Gaussian random variable with mean m and variance σ^2, namely

$$f(z) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(z-m)^2/(2\sigma^2)}.$$

Let source X^n be an n-dimensional random vector (viewed as a column vector) with real components $X_1, \ldots, X_n$. Without loss of generality, we assume that each component has mean zero, giving covariance matrix

$$\Phi_X = E[X^n (X^n)^t].$$

Let transform T be an $n \times n$ orthogonal matrix with real elements, and let

$$Y^n = TX^n$$

denote the transformed random vector with coefficients $Y_1, \ldots, Y_n$. We restrict attention to orthogonal transforms since it suffices to demonstrate the suboptimality of the KLT over this class.

A scalar quantizer with resolution r bits is a mapping $Q : \mathbb{R} \to \mathbb{R}$ whose range (called a *codebook*) has cardinality 2^r. The rate of a *fixed-rate* scalar quantizer Q with resolution r is

$$R^{(\mathrm{fr})}(Q) = r.$$

The rate to describe source X_i with a *variable-rate* scalar quantizer Q is

$$R^{(\mathrm{vr})}(Q) = H(Q(X_i)).$$

[1]The optimal transform code "matched" to a transform is defined as the bit allocation, scalar quantizer, and fixed-rate or variable-rate encoder that yield the optimal performance given the fixed transform.

[2]All logarithms in this paper whose bases are omitted are base 2.

A *transform coder* is a system that quantizes X^n by transforming X^n by T and then applying an independent scalar quantizer Q_i to each transform coefficient Y_i, for $i \in \{1, \ldots, n\}$. Thus the per-symbol expected rate of a transform coder with transform T and quantizers $Q^n = (Q_1, \ldots, Q_n)$ is

$$R_T^{(\mathrm{fr})}(Q^n) = \frac{1}{n}\sum_{i=1}^{n} R^{(\mathrm{fr})}(Q_i)$$

for fixed-rate transform coding and

$$R_T^{(\mathrm{vr})}(Q^n) = \frac{1}{n}\sum_{i=1}^{n} R^{(\mathrm{vr})}(Q_i)$$

for variable-rate transform coding. The corresponding per symbol MSE is

$$D_T(Q^n) = \frac{1}{n}E[\|X^n - T^t Q^n(TX^n)\|^2]$$

where

$$Q^n(x^n) = (Q_1(x_1), \ldots, Q_n(x_n))^t$$

is the vector of scalar quantized components of vector $x^n = (x_1, \ldots, x_n)^t$ and the Euclidean distance between two arbitrary n-dimensional vectors $x^n = (x_1, \ldots, x_n)^t$ and $y^n = (y_1, \ldots, y_n)^t$ is denoted by

$$\|x^n - y^n\| = \sqrt{\sum_{i=1}^{n}(x_i - y_i)^2}.$$

Given an average rate budget of R bits per symbol and a fixed transform T, the operational distortion-rate function for a fixed-rate transform code based on T is defined as

$$D_T^{(\mathrm{fr})}(R) = \inf_{Q^n : R_T^{(\mathrm{fr})}(Q^n) \le R} D_T(Q^n).$$

For variable-rate coding, we restrict attention to uniform scalar quantizers. Thus, we define the operational distortion-rate function for a variable-rate transform code based on T as

$$D_T^{(\mathrm{vr})}(R) = \inf_{Q^n \in \mathcal{Q}_U^n : R_T^{(\mathrm{vr})}(Q^n) \le R} D_T(Q^n)$$

where we define $\mathcal{Q}_U^n$ such that $Q^n \in \mathcal{Q}_U^n$ if and only if Q_i is a *uniform* scalar quantizer whenever $1 \le i \le n$. When T is the identity matrix I, we drop the subscript and use $D^{(\mathrm{fr})}(R)$ and $D^{(\mathrm{vr})}(R)$ to denote the MSE corresponding to the classical fixed-rate scalar bit allocation problem and variable-rate uniform scalar bit allocation problem, respectively.

The optimality of a transform for transform coding is often considered in an asymptotic sense (e.g., [7]). Let

$$R_0^{(\mathrm{fr})}(T) = \sup\{R : D_T^{(\mathrm{fr})}(R) > 0\}$$
$$R_0^{(\mathrm{vr})}(T) = \sup\{R : D_T^{(\mathrm{vr})}(R) > 0\}.$$

Then $R_0^{(\mathrm{fr})}(T)$ and $R_0^{(\mathrm{vr})}(T)$ give the minimal rate required to perfectly describe the source, with $R_0^{(\mathrm{fr})}(T) = R_0^{(\mathrm{vr})}(T) = \infty$ for continuous sources. The *fixed-rate and variable-rate coding gains* obtained by using transform T_1 instead of transform T_2 are defined as

$$G_{T_1,T_2}^{(\mathrm{fr})} = \lim_{R \to R_0^{(\mathrm{fr})}(T_1)} D_{T_2}^{(\mathrm{fr})}(R)/D_{T_1}^{(\mathrm{fr})}(R)$$
$$G_{T_1,T_2}^{(\mathrm{vr})} = \lim_{R \to R_0^{(\mathrm{vr})}(T_1)} D_{T_2}^{(\mathrm{vr})}(R)/D_{T_1}^{(\mathrm{vr})}(R),$$

respectively. (Note that these two limits exist because the limits of the numerators and denominators exist since they are monotonically nonincreasing in R.) Each coding gain describes the asymptotic performance gap (in terms of SQNR) between the associated optimal transform coders. The gains measured in decibels (dB) are $10\log_{10} G_{T_1,T_2}^{(\mathrm{fr})}$ and $10\log_{10} G_{T_1,T_2}^{(\mathrm{vr})}$.

An orthogonal transform $T^\star$ is said to be *optimal* for fixed-rate transform coding on source X^n if $G_{T^\star,T}^{(\mathrm{fr})} \ge 1$ and for variable-rate transform coding on X^n if $G_{T^\star,T}^{(\mathrm{vr})} \ge 1$, for all orthogonal transforms T. All of the results that follow refer to optimality in this asymptotic sense.

A *KLT* of a source X^n is the linear map given by an $n \times n$ orthogonal matrix T^t such that $T^t \Phi_X T$ is a diagonal matrix. The matrix T decorrelates the random vector X^n since

$$\Phi_{T^t X^n} = E[(T^t X^n)(T^t X^n)^t] = T^t \Phi_X T.$$

A source may have more than one KLT (for example, any permutation matrix is a KLT of a source with uncorrelated components).

Lemma II.1: (Huang and Schultheiss, 1963 in [2]) Any KLT is optimal for fixed-rate transform coding of Gaussian sources.

It has been conjectured that Lemma II.1 holds for non-Gaussian sources as well.

We say that the pdf of a scalar random variable is *reasonably smooth* if at sufficiently high rate, the pdf of all symbols mapped to the same reproduction value approaches a uniform density. We use this somewhat loose definition, as was done in [7], to avoid burdensome and distracting calculations.

Given a vector source for which the marginal pdf on each dimension is reasonably smooth, the following high-rate bit allocation result of Huang and Schultheiss [2] (see also [7, pp. 228–232]) characterizes $D^{(\mathrm{fr})}(R)$. We use the notation

$$\|f\|_{1/3} = \left(\int_{-\infty}^{\infty} f^{1/3}(x)dx\right)^3.$$

Lemma II.2: Let $X^n = (X_1, \ldots, X_n)^t$ be a random vector with marginal pdf f_i and component variance $\sigma_i^2 > 0$ for $i = 1, \ldots, n$. If $f_1, \ldots, f_n$ are reasonably smooth, then

$$\lim_{R \to \infty} D^{(\mathrm{fr})}(R)2^{2R} = \frac{1}{12}\left(\prod_{i=1}^{n}\|f_i\|_{1/3}\right)^{1/n}.$$

It should be noted that Lemma II.2 allows rates for fixed-rate quantizers that cannot be achieved in practice (e.g., noninteger rates). Nevertheless, this result bounds the achievable performance and provides useful theoretical intuition.

The high-rate approximation for the performance of an optimal variable-rate scalar quantizer gives a similar result for bit allocation in variable-rate coding. This result also describes the asymptotic performance of a uniform scalar quantizer followed by entropy coding (e.g., [7, pp. 296], [8]). The proof of Lemma II.3, which is very similar to that of Lemma II.2, appears in the Appendix.

Lemma II.3: Let $X^n = (X_1, \ldots, X_n)^t$ be a random vector whose marginal densities are absolutely continuous with respect

to Lebesgue measure and have finite positive second moments. Then

$$\lim_{R\to\infty} D^{(\mathrm{vr})}(R)2^{2R} = \frac{1}{12}2^{2\bar{h}(X^n)}.$$

III. Summary of Results

We begin by showing, in Theorem III.1, that decorrelation is insufficient for optimality in transform coding. We prove this theorem in Section IV by examining sources for which the KLT is not unique. While every KLT is, by definition, a decorrelating transform, there exist sources for which the KLT is not unique and all KLTs are not equally good. In particular, we show the following result.

Theorem III.1: In both fixed-rate and variable-rate transform coding, there exist sources for which the coding gain of a best KLT for transform coding over a worst KLT for transform coding can be arbitrarily large.

Thus, decorrelation is insufficient for optimality in both fixed-rate and variable-rate transform coding.

Next, we demonstrate that even a best (or only) KLT can be suboptimal in both fixed-rate and variable-rate transform coding.

Theorem III.2: In both fixed-rate and variable-rate transform coding, there exist sources for which the coding gain of an optimal transform over a best KLT for transform coding is strictly greater than 1.

Thus, decorrelation is not only not sufficient for optimality, it is also not necessary for optimality. The examples used to prove this result in Section V demonstrate that choosing a decorrelating transform can *preclude* optimality for some sources.

Next, we consider sources for which the KLT is unique and decorrelation yields independent transform coefficients. The intuition described in Section I suggests that the KLT should be optimal in this coding framework. In Theorem III.3, we prove this intuition false for fixed-rate transform coding.

Theorem III.3: There exist sources for which the KLT is unique and produces independent transform coefficients and yet the KLT is not optimal for fixed-rate transform coding.

In contrast, we show that the KLT is optimal for variable-rate coding in Section IV-B.

Theorem III.4: If a KLT produces independent transform coefficients, then this KLT is optimal for variable-rate transform coding.

Note that all of the results refer to optimality in the high resolution sense.

IV. KLTs Are not Uniformly Good

In this section, we give the proof of Theorem III.1. That is, we show in both the fixed-rate and the variable-rate transform coding frameworks that there exist sources for which the KLT is not unique and a "best" KLT yields an infinite coding gain over a "worst" KLT. The results in both frameworks rely on the following family of examples.

Suppose that source $X^n = (X_1, \ldots, X_n)^t$ with marginal pdfs $(f_1, \ldots, f_n)$ is defined as

$$X^n = BU^n.$$

Here the components $U_1, \ldots, U_n$ of the random vector $U^n = (U_1, \ldots, U_n)^t$ are i.i.d. random variables with reasonably smooth marginal pdfs $(\hat{f}_1, \ldots, \hat{f}_n)$ and positive variances, and B is an $n \times n$ orthogonal matrix.

Let Φ_U and Φ_X denote the covariance matrices of U^n and X^n, respectively. Then for *any* orthogonal matrix B

$$\Phi_U = \Phi_X = \sigma^2 I$$

where I denotes the n-dimensional identity matrix. That is, any rotation of U^n creates an uncorrelated random vector X^n, and thus *any* transform matrix T is a legitimate KLT for X^n. While the KLT for X^n is not unique, practical implementations of the KLT (e.g., Householder reduction followed by the QL algorithm with implicit shifts based on the orthogonal–lower triangular decomposition of a matrix or Jacobi's algorithm [9]) give the identity matrix I as the KLT for X^n. Therefore, using the KLT in an optimal transform coder for X^n is in practice equivalent to optimal bit allocation followed by scalar quantization on the original source X^n. We therefore calculate the coding gain of the transform B^{-1} relative to the practically achieved KLT I.[3]

A. Fixed-Rate Coding

Using Lemma II.2, the fixed-rate coding gain of transform B^{-1} over transform I is

$$G^{(\mathrm{fr})}_{B^{-1},I} = \lim_{R\to\infty} \frac{D^{(\mathrm{fr})}(R)}{D^{(\mathrm{fr})}_{B^{-1}}(R)} = \frac{\lim_{R\to\infty} D^{(\mathrm{fr})}(R)2^{2R}}{\lim_{R\to\infty} D^{(\mathrm{fr})}_{B^{-1}}(R)2^{2R}} = \left(\prod_{i=1}^{n} \frac{\|f_i\|_{1/3}}{\|\hat{f}_i\|_{1/3}}\right)^{1/n}. \tag{1}$$

B. Variable-Rate Coding

Using Lemma II.3, if $U_1, \ldots, U_n$ are drawn i.i.d. according to a reasonably smooth pdf f_U with finite variance $\sigma^2 > 0$ and differential entropy $h(U)$, then the variable-rate coding gain of transform B^{-1} over an arbitrary transform T is

$$G^{(\mathrm{vr})}_{B^{-1},T} = \lim_{R\to\infty} \frac{D^{(\mathrm{vr})}_{T}(R)}{D^{(\mathrm{vr})}_{B^{-1}}(R)} = \frac{\lim_{R\to\infty} D^{(\mathrm{vr})}_{T}(R)2^{2R}}{\lim_{R\to\infty} D^{(\mathrm{vr})}_{B^{-1}}(R)2^{2R}} = \frac{2^{2\bar{h}(Y^n)}}{2^{2\bar{h}(U^n)}} = 2^{2(\bar{h}(Y^n)-\bar{h}(U^n))} \tag{2}$$

where $Y^n = TX^n$. Then

$$\overline{h}(U^n) = \frac{1}{n}h(U^n) = \frac{1}{n}h(Y^n) \le \frac{1}{n}\sum_{i=1}^{n} h(Y_i) = \overline{h}(Y^n)$$

since $U_1, \ldots, U_n$ are independent, TB is nonsingular, and

$$Y^n = TX^n = TBU^n.$$

[3]In practice, performance problems of the KLT are often exacerbated by the KLT's sensitivity to errors in estimating the off-diagonal terms of the covariance matrix when the covariance matrix is close to the identity matrix [5].

Thus,

$$G^{(\mathrm{vr})}_{B^{-1},T} \geq 1$$

with equality if and only if $Y_1, \ldots, Y_n$ are independent. In this case, a transform that makes the components independent is required for optimality in variable-rate coding.[4] This also proves Theorem III.4.

Now let us go back to the i.i.d. source, where

$$\begin{aligned} G^{(\mathrm{vr})}_{B^{-1},I} &= 2^{2(\bar{h}(Y^n)-\bar{h}(U^n))} = 2^{2(\bar{h}(X^n)-h(U))} \\ &\leq 2^{\log(2\pi e\sigma^2)-2h(U)} \end{aligned} \tag{3}$$

with equality if and only if the marginal pdf of each X_i is the Gaussian density $\mathcal{N}(0,\sigma^2)$. Given any choice of symmetric f_U such that the central limit theorem applies to

$$\frac{1}{\sqrt{n}}\sum_{i=1}^{n} U_i,$$

then for sufficiently large n and carefully chosen B, the marginal density of X_i can be made arbitrarily close to the above normal density.[5]

C. Examples

We next bound the fixed-rate and variable-rate coding gains for several example distributions.

Example IV.1: U_i is uniform on $[-a, a]$ with $a > 0$.

For this distribution, $\sigma^2 = a^2/3$, and $h(U) = \log(2a)$. For the fixed-rate coding calculations

$$\begin{aligned} \|f_U * f_U\|_{1/3} &= \frac{27a^2}{2} \\ \|f_U\|_{1/3} &= 4a^2 \\ \|\mathcal{N}(0,\sigma^2)\|_{1/3} &= 6\pi\sigma^2\sqrt{3} = 2\pi a^2\sqrt{3} \end{aligned}$$

where

$$(f_U * f_U)(x) = \int_{-\infty}^{\infty} f_U(t) f_U(x-t)\,dt$$

is the convolution of f_U with itself.

If $n = 2$ and

$$B = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

then

$$\begin{aligned} \|f_{X_1}\|_{1/3} &= \|f_{X_2}\|_{1/3} \\ &= \frac{\|f_U * f_U\|_{1/3}}{2} \end{aligned}$$

$$G^{(\mathrm{fr})}_{B^{-1},I} = \frac{\|f_U * f_U\|_{1/3}}{2\|f_U\|_{1/3}} = \frac{27}{16} \quad (\text{i.e., } 2.27 \text{ dB}).$$

If n is large and the marginal density of X_i is close to $\mathcal{N}(0,\sigma^2)$, then

$$G^{(\mathrm{fr})}_{B^{-1},I} \approx \frac{\|\mathcal{N}(0,\sigma^2)\|_{1/3}}{\|f_U\|_{1/3}} \approx 2.72 \quad (\text{i.e., } 4.35 \text{ dB}).$$

In this case, B^{-1} is not necessarily the optimal transform, so this is not necessarily the largest possible coding gain.

For variable-rate coding, B^{-1} is the optimal transform at high rate as shown earlier, and the coding gain between the best and worst KLTs for the worst case B is

$$\sup_B G^{(\mathrm{vr})}_{B^{-1},I} = \frac{2\pi e\sigma^2}{(2a)^2} \approx 1.41 \quad (\text{ i.e., } 1.5 \text{ dB})$$

which is consistent with [5].

Example IV.2: U_i is uniform on

$$[-a-\delta, -a+\delta] \cup [a-\delta, a+\delta]$$

with $0 < \delta < a$.

In this case, $\sigma^2 = a^2 + (\delta^2/3)$, and $h(U) = \log(4\delta)$. For the fixed-rate coding calculations

$$\begin{aligned} \|f_U\|_{1/3} &= 16\delta^2 \\ \|\mathcal{N}(0,\sigma^2)\|_{1/3} &= 6\pi\sigma^2\sqrt{3} = 6\pi a^2\sqrt{3} + 2\pi\delta^2\sqrt{3}. \end{aligned}$$

If the marginal distribution of X_i is close to $\mathcal{N}(0,\sigma^2)$, then

$$G^{(\mathrm{fr})}_{B^{-1},I} \approx \frac{\|\mathcal{N}(0,\sigma^2)\|_{1/3}}{\|f_U\|_{1/3}} = \frac{3\pi\sqrt{3}}{8}\left(\frac{a^2}{\delta^2} + \frac{1}{3}\right).$$

The coding gain can be made arbitrarily large by fixing a and letting $\delta \to 0$.

For variable-rate coding

$$\sup_B G^{(\mathrm{vr})}_{B^{-1},I} = \frac{2\pi e\sigma^2}{(4\delta)^2} = \frac{\pi e}{8}\left(\frac{a^2}{\delta^2} + \frac{1}{3}\right)$$

which can again be made arbitrarily large by fixing a and letting $\delta \to 0$.

The problem observed above for reasonably smooth, continuous random variables becomes even more pronounced for discrete random variables. For a discrete random variable, the previous high-rate approximation does not apply since the probability mass function (pmf) is not smooth. It is still relatively easy to calculate the coding gain for certain discrete random variables, as we show in the following example.

Example IV.3: U_i is discrete.

Consider $U_1, \ldots, U_n$ drawn i.i.d. according to

$$p(10) = p(-10) = 1/2.$$

In both fixed-rate and variable-rate coding, we can quantize each U_i with distortion $D = 0$ at rate 1 bit per symbol (bps). Achieving $D = 0$ for each X_i generally requires more rate. For example, choose B so that the marginal distribution of each X_i is the binomial distribution. For $n = 64$, achieving $D = 0$ requires approximately 4 bps in variable-rate coding and $\log(64+1) \approx 6$ bps in fixed-rate coding. At rate 1 bps, the MSE in (either fixed-rate or variable-rate) coding of U_i is approximately 36, giving infinite coding gain for both fixed-rate and variable-rate coding, i.e.,

$$\sup_B G^{(\mathrm{fr})}_{B^{-1},I} = \sup_B G^{(\mathrm{vr})}_{B^{-1},I} = \infty.$$

[4]While $U_1, \ldots, U_n$ are i.i.d. in this example, the result actually requires only that they be independent. In addition, this statement is true only under the high-rate assumption. It has been shown that for some sources, B^{-1} is suboptimal at certain low rates in variable-rate coding, i.e., a transform that yields independent coefficients is not always optimal for variable-rate coding when the high-rate approximation does not apply [10].

[5]This can be accomplished, for example, by letting B be $1/\sqrt{n}$ times a Hadamard matrix of order n. Hadamard matrices are known to exist at least for every n that is a power of 2. This would assure that B is orthogonal and has components all of equal magnitude, on a subsequence of $\{1, 2, \ldots, n\}$.
It should be noted that this tendency toward Gaussianity of the marginal pdf of a vector after linear transformation is not limited to this specific distribution or the Hadamard matrix. Related issues have been investigated in earlier literature (e.g., see [5] and [6]).

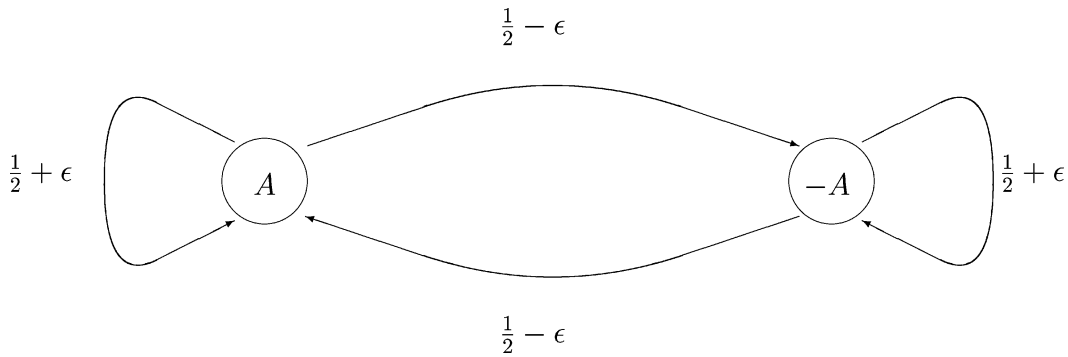

Fig. 1. Stationary Markov process S_n.

V. A Best KLT Can Be a Suboptimal Transform

In the previous section, we showed for both fixed-rate and variable-rate transform coding that when the KLT is not unique the coding gain of a best KLT over a worst KLT can be arbitrarily large. Thus, decorrelation is not sufficient for transform optimality. In this section, we prove Theorem III.2 by showing that even a best KLT can give suboptimal performance.

We first use the following source structure for both fixed-rate and variable-rate transform coding. Let $\{U_i\}_{i=1}^{\infty}$ be an i.i.d. real-valued random sequence with $E[U_i] = 0$ and $E[U_i^2] = \sigma^2$ for all i, and symmetric pdf f_U. Let $\{S_i\}_{i=1}^{\infty}$ be the stationary two-state Markov process shown in Fig. 1 with $A > 0$ and $0 < \epsilon < 1/2$. That is, $S_i = A$ with probability $1/2$, and $S_i = -A$ otherwise. The value of ϵ determines the tendency of the process to remain in its current state. Furthermore, assume that the processes U_i and S_i are independent of each other. Let

$$X_i = U_i + S_i$$

be a scalar source and define the two-dimensional random vector $X = (X_i, X_{i+1})^t$. Let

$$T = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \tag{4}$$

Then the vector $Y = (Y_1, Y_2) = TX$ (note that $T = T^t$) is a KLT for the two-dimensional stationary vector source X (see, e.g., [11, p. 335, Example 12.4.1]), and T simply rotates X through an angle of 45° clockwise.[6]

We next show for both fixed-rate and variable-rate transform coding that (in the high-rate limit) optimal scalar quantization of the components of the correlated random vector X produces a smaller MSE than optimal scalar quantization of the components of the decorrelated vector Y. (Note that Y_1 and Y_2 are uncorrelated, but they are not independent.)

Since the Markov process S_i is stationary, we assume without loss of generality that $X = (X_1, X_2)^t$. Notice that by symmetry the scalar components of X are identically distributed with pdf

$$f_{X_1}(x) = f_{X_2}(x) = \frac{1}{2}f_U(x+A) + \frac{1}{2}f_U(x-A). \tag{5}$$

[6]There are three other possible KLTs of X, namely,

$$\frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix}, \quad \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}, \quad \text{and} \quad \frac{1}{\sqrt{2}}\begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix}$$

which all give equal coding gains. Thus, we will refer to T as "the" best KLT in this example.

In this example, the KLT is not continuous in ϵ, since only $\epsilon = 0$ yields the identity matrix, while all $\epsilon > 0$ lead to 45° rotations.

We begin by setting up the coding gain calculations for both fixed-rate and variable-rate coding. We then calculate those coding gains for a variety of examples.

A. Fixed-Rate Coding

If $f_U(x) = 0$ for all $x \notin [-A/2, A/2]$, then either $f_U(x+A) = 0$ or $f_U(x-A) = 0$ for any x, and thus it can be shown that

$$\|f_{X_1}\|_{1/3} = \|f_{X_2}\|_{1/3} = 4\|f_U\|_{1/3}.$$

We can also show that

$$f_{Y_1}(y_1) = \sqrt{2}\Big[\beta(f_U * f_U)\left(y_1\sqrt{2} + 2A\right) \tag{6}$$
$$+ \beta(f_U * f_U)\left(y_1\sqrt{2} - 2A\right)$$
$$+2\alpha(f_U * f_U)\left(y_1\sqrt{2}\right)\Big] \tag{7}$$
$$f_{Y_2}(y_2) = \sqrt{2}\Big[\alpha(f_U * f_U)\left(y_2\sqrt{2} + 2A\right) \tag{8}$$
$$+ \alpha(f_U * f_U)\left(y_2\sqrt{2} - 2A\right)$$
$$+2\beta(f_U * f_U)\left(y_2\sqrt{2}\right)\Big] \tag{9}$$

where

$$\alpha = \frac{1}{4} - \frac{\epsilon}{2}$$
$$\beta = \frac{1}{4} + \frac{\epsilon}{2}.$$

From the marginal pdfs we obtain

$$\|f_{Y_1}\|_{1/3} = \frac{1}{2}(2^{1/3}\alpha^{1/3} + 2\beta^{1/3})^3\|f_U * f_U\|_{1/3}$$
$$\|f_{Y_2}\|_{1/3} = \frac{1}{2}(2^{1/3}\beta^{1/3} + 2\alpha^{1/3})^3\|f_U * f_U\|_{1/3}. \tag{10}$$

The details of these calculations appear in the Appendix.

By Lemma II.2, the coding gain obtained by quantizing the correlated scalar components instead of the uncorrelated components is

$$G_{I,T}^{(\mathrm{fr})} = \frac{\sqrt{\|f_{Y_1}\|_{1/3}\|f_{Y_2}\|_{1/3}}}{\sqrt{\|f_{X_1}\|_{1/3}\|f_{X_2}\|_{1/3}}}$$
$$= \frac{\|f_U * f_U\|_{1/3}}{8\|f_U\|_{1/3}}$$
$$\cdot \left[(2^{1/3}\beta^{1/3} + 2\alpha^{1/3})(2^{1/3}\alpha^{1/3} + 2\beta^{1/3})\right]^{3/2}. \tag{11}$$

The next lemma is useful in calculating the coding gain. Its proof also appears in the Appendix.

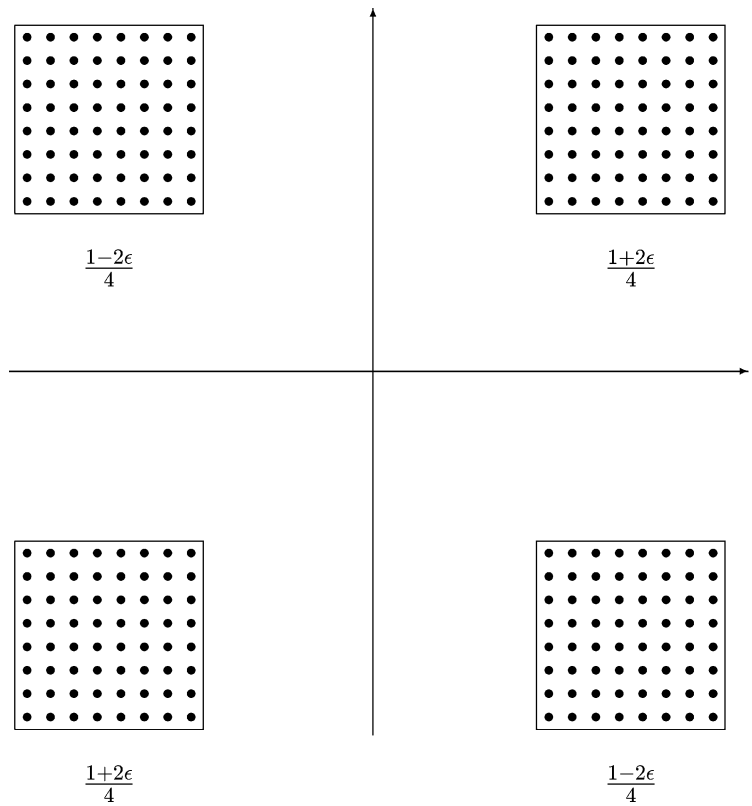

Fig. 2. The two-dimensional pdf of the correlated source X is uniform on each of the four squares with the heights $(1 \pm 2\epsilon)/4$ as indicated. The dots indicate reproduction locations associated with using the optimal 16-codeword scalar quantizer on each dimension.

Lemma V.1: Let

$$\gamma_\epsilon = \left[(2^{1/3}\beta^{1/3} + 2\alpha^{1/3})(2^{1/3}\alpha^{1/3} + 2\beta^{1/3})\right]^{3/2}$$

where $\alpha = (1/4) - (\epsilon/2)$ and $\beta = (1/4) + (\epsilon/2)$. Then on the interval $(0, 1/2)$, γ_ϵ is monotonically decreasing and

$$(2 + 2^{1/3})^3/4 = \gamma_0 \geq \gamma_\epsilon \geq \gamma_{1/2} = 2.$$

B. Variable-Rate Coding

In this case, if $f_U(x) = 0$ for all $x \notin [-A/2, A/2]$, then the differential entropy of X_1 is

$$\begin{aligned} h(X_1) &= -\int_{-\infty}^{\infty} f_{X_1}(x) \log f_{X_1}(x) dx \\ &= -\int_{-\infty}^{\infty} \frac{1}{2}[f_U(x+A) + f_U(x-A)] \\ &\quad \cdot \log\left(\frac{1}{2}[f_U(x+A) + f_U(x-A)]\right) dx \\ &= 1 - \frac{1}{2}\int_{-\infty}^{\infty} f_U(x+A) \log f_U(x+A) dx \\ &\quad - \frac{1}{2}\int_{-\infty}^{\infty} f_U(x-A) \log f_U(x-A) dx \\ &= 1 + h(U) \end{aligned}$$

where $h(U)$ is the differential entropy of U_i. Similarly,

$$h(X_2) = 1 + h(U).$$

The outcome of the KLT is $Y_1 = (X_1 + X_2)/\sqrt{2}$ and $Y_2 = (X_1 - X_2)/\sqrt{2}$. Thus,

$$\begin{aligned} h(Y_1) &= -\int_{-\infty}^{\infty} f_{Y_1}(y) \log f_{Y_1}(y) dy \\ &= -2\beta \log(\sqrt{2}\beta) + 2\beta h(f_U * f_U) \\ &\quad - 2\alpha \log(2\sqrt{2}\alpha) + 2\alpha h(f_U * f_U) \\ &= \beta - \alpha + \mathcal{H}(2\alpha) + h(f_U * f_U) \end{aligned}$$

where

$$\mathcal{H}(x) = -x\log(x) - (1-x)\log(1-x)$$

is the binary entropy function. Similarly,

$$h(Y_2) = \alpha - \beta + \mathcal{H}(2\alpha) + h(f_U * f_U).$$

By Lemma II.3, the coding gain obtained by quantizing the correlated scalar components instead of the uncorrelated components is

$$\begin{aligned} G_{I,T}^{(\mathrm{vr})} &= 2^{h(Y_1)+h(Y_2)-h(X_1)-h(X_2)} \\ &= 2^{2h(f_U * f_U) + 2\mathcal{H}(2\alpha) - 2h(U) - 2}. \end{aligned} \tag{12}$$

C. Examples

Example V.2: U_i is uniform on $[-a, a]$ with $0 < a < A/2$.

The solid lines in Figs. 2 and 3 show the regions of support for the two-dimensional pdfs of X and Y given a uniform distribution on U. Each pdf is uniform in its marked regions. The probability of each of those regions is marked in the figure. Fig. 4 shows the resulting marginal pdf f_{Y_2} of the decorrelated vector Y.

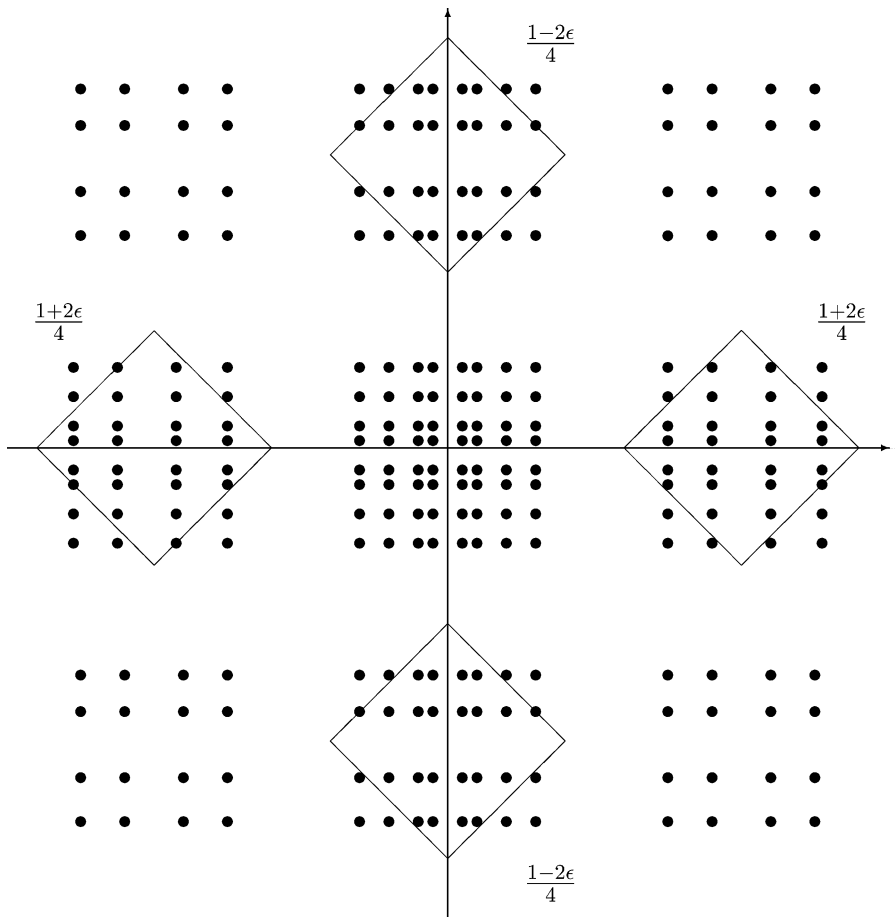

Fig. 3. The two-dimensional pdf f_Y of the uncorrelated source $Y = TX$ (i.e., the KLT output) equals the two-dimensional pdf of X rotated 45° about the origin. The dots indicate reproduction locations associated with using the optimal 16-codeword fixed-rate scalar quantizer on each dimension.

For the fixed-rate coding calculation

$$\begin{aligned}\|f_U * f_U\|_{1/3} &= \left(2\int_0^{2a}\left(\frac{1}{2a}-\frac{x}{4a^2}\right)^{1/3}dx\right)^3\\ &= \frac{27a^2}{2}\\ \|f_U\|_{1/3} &= 4a^2\end{aligned}$$

and therefore (11) implies that for all $\epsilon \in (0, 1/2)$

$$G_{I,T}^{(\mathrm{fr})} = \frac{27}{64}\left[(2^{1/3}\beta^{1/3}+2\alpha^{1/3})(2^{1/3}\alpha^{1/3}+2\beta^{1/3})\right]^{3/2}.$$

Lemma V.1 and numerical inversion of $G_{I,T}^{(\mathrm{fr})}$ establish that $G_{I,T}^{(\mathrm{fr})} > 1$ whenever $\epsilon < 0.4998$. Also, by Lemma V.1,

$$\lim_{\epsilon\to 0} G_{I,T}^{(\mathrm{fr})} = \frac{27}{256}(2+2^{1/3})^3 \approx 3.66 \quad (\text{ i.e., 5.63 dB}).$$

The dots in Figs. 2 and 3 show all possible two-dimensional reproductions when the individual components of the random vectors are quantized with the optimal fixed-rate scalar quantizers. The quantizer associated with the decorrelated random vector Y (shown in Fig. 3) is very inefficient and thus leads to higher MSE than the quantizer associated with the correlated random vector X (shown in Fig. 2).

The results for variable-rate transform coding are similar. In this case

$$h(U) = \log(2a)$$

$$\begin{aligned}h(f_U * f_U) &= -2\int_0^{2a}\frac{2a-x}{4a^2}\log\frac{2a-x}{4a^2}dx\\ &= \log(2a)+\frac{\log e}{2}.\end{aligned}$$

Therefore,

$$G_{I,T}^{(\mathrm{vr})} = 2^{\log e + 2\mathcal{H}(2\alpha)-2}$$

giving $G_{I,T}^{(\mathrm{vr})} > 1$ whenever[7]

$$\epsilon < \frac{1}{2} - \mathcal{H}^{-1}(1-\log\sqrt{e}) \approx .4518.$$

In the limit of small ϵ

$$\lim_{\epsilon\to 0} G_{I,T}^{(\mathrm{vr})} = e \approx 2.72 \quad (\text{i.e., 4.34 dB}).$$

In the previous example, for both fixed-rate and variable-rate coding, a best KLT is 4–5 dB worse than an optimal transform. In this case, it is better to scalar quantize correlated data than uncorrelated data. The following example gives a similar outcome for a Gaussian source.

Example V.3: U_i is Gaussian with mean 0 and variance $\sigma^2 \ll A$.

Given $f_U(x) = \mathcal{N}(0, \sigma^2)$, the marginal pdfs of X are again given by (5).

[7]We use $\mathcal{H}^{-1}$ to denote the inverse function of the restriction of $\mathcal{H}$ to $[0, 1/2]$.

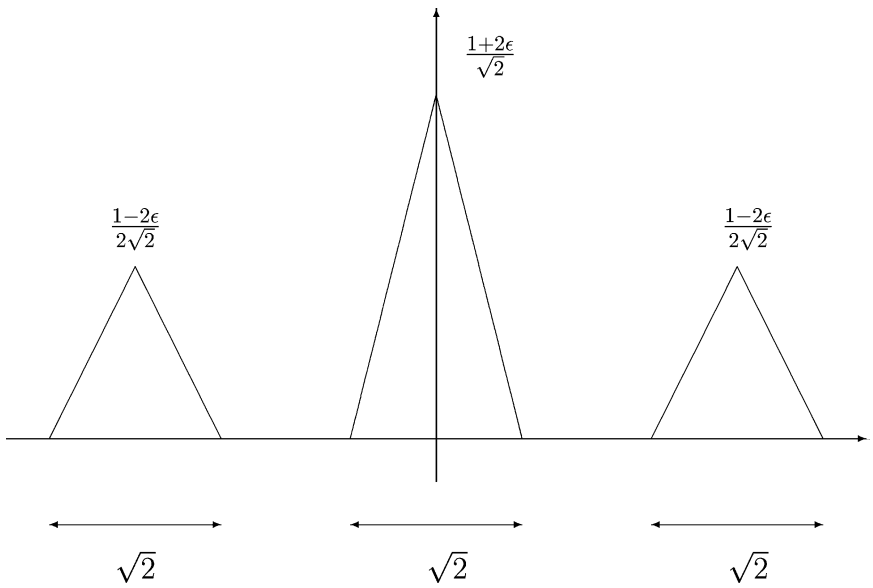

Fig. 4. The marginal pdf f_{Y_2} of the decorrelated random vector Y.

For fixed-rate coding, (13) does not apply directly because we violate the condition $f_U(x) = 0$ for all $x \notin [-A/2, A/2]$. However, (13) provides an approximation that becomes accurate as $\sigma \to 0$, giving

$$\begin{aligned}\lim_{\sigma\to 0}\frac{\|f_{X_1}\|_{1/3}}{\sigma^2} &= \lim_{\sigma\to 0}\frac{\|f_{X_2}\|_{1/3}}{\sigma^2}\\ &= 4\lim_{\sigma\to 0}\frac{\|f_U\|_{1/3}}{\sigma^2}\\ &= 24\pi\sqrt{3}.\end{aligned}$$

The marginal pdfs of Y are given by (7) and (9), where in this case

$$\begin{aligned}(f_U * f_U)(x) &= \frac{1}{2\sigma\sqrt{\pi}}e^{-x^2/(4\sigma^2)}\\ \|f_U * f_U\|_{1/3} &= 12\sqrt{3}\pi\sigma^2.\end{aligned}$$

Following the argument of (10), from the marginal pdfs we obtain

$$\begin{aligned}&\lim_{\sigma\to 0}\frac{\|f_{Y_1}\|_{1/3}}{\sigma^2}\\ &= \lim_{\sigma\to 0}\frac{\sqrt{2}}{\sigma^2}\left(\int_{-\infty}^{\infty}\Big[2^{1/3}\alpha^{1/3}(f_U * f_U)^{1/3}(y_1\sqrt{2})\right.\\ &\quad + \beta^{1/3}(f_U * f_U)^{1/3}\left(y_1\sqrt{2} - 2A\right)\\ &\quad \left. + \beta^{1/3}(f_U * f_U)^{1/3}\left(y_1\sqrt{2} + 2A\right)\Big]\,dy_1\right)^3\\ &= \lim_{\sigma\to 0}\frac{1}{2\sigma^2}(2^{1/3}\alpha^{1/3} + 2\beta^{1/3})^3\|f_U * f_U\|_{1/3}\\ &\lim_{\sigma\to 0}\frac{\|f_{Y_2}\|_{1/3}}{\sigma^2}\\ &= \lim_{\sigma\to 0}\frac{1}{2\sigma^2}(2^{1/3}\beta^{1/3} + 2\alpha^{1/3})^3\|f_U * f_U\|_{1/3}.\end{aligned}$$

Therefore, (11) holds in the limit as $\sigma \to 0$, and the limiting coding gain obtained by quantizing the correlated source symbols instead of the uncorrelated KLT coefficients is

$$\lim_{\sigma\to 0} G_{I,T}^{(\mathrm{fr})} = \frac{1}{4}\left[(2^{1/3}\beta^{1/3} + 2\alpha^{1/3})(2^{1/3}\alpha^{1/3} + 2\beta^{1/3})\right]^{3/2}.$$

Therefore, $G_{I,T}^{(\mathrm{fr})} > 1$ whenever $\epsilon < 0.485$, for some $\sigma > 0$. For $\epsilon \approx 0$, the coding gain obtained by not decorrelating the source X is

$$G_{I,T}^{(\mathrm{fr})} = \frac{1}{16}(2 + 2^{1/3})^3 \approx 2.17 \quad \text{(i.e., 3.36 dB)}.$$

Similarly, for variable-rate coding, (12) does not apply directly but provides an approximation that becomes accurate as $\sigma \to 0$. Moreover,

$$\begin{aligned}h(U) &= \frac{1}{2}\log(2\pi e\sigma^2)\\ h(f_U * f_U) &= \frac{1}{2}\log(4\pi e\sigma^2).\end{aligned}$$

Therefore,

$$\lim_{\sigma\to 0} G_{I,T}^{(\mathrm{vr})} = 2^{2\mathcal{H}(2\alpha)-1}.$$

In this case, $G_{I,T}^{(\mathrm{vr})} > 1$ whenever

$$\epsilon < \frac{1}{2} - \mathcal{H}^{-1}\left(\frac{1}{2}\right) \approx 0.3899.$$

We also get

$$\lim_{\epsilon\to 0}\lim_{\sigma\to 0} G_{I,T}^{(\mathrm{vr})} = 2 \quad \text{(i.e., 3.01 dB)}.$$

Example V.4: U_i is a constant.

Let $U_i = 0$ with probability 1. Then the pmfs of X_1 and X_2 are

$$\begin{aligned}p(X_1 = -A) &= p(X_1 = A)\\ &= p(X_2 = -A) = p(X_2 = A) = 1/2.\end{aligned}$$

At rate 1 bps, the MSE obtained by using the transform matrix I is 0 for both fixed-rate and variable-rate coding. In contrast, the pmfs of Y_1 and Y_2 are

$$\begin{aligned}&p(Y_1 = A\sqrt{2}) = p(Y_1 = -A\sqrt{2}) = \beta\\ &p(Y_1 = 0) = 2\alpha\end{aligned}$$

$$p(Y_2 = A\sqrt{2}) = p(Y_2 = -A\sqrt{2}) = \alpha$$
$$p(Y_2 = 0) = 2\beta.$$

For fixed-rate coding, a rate of $\log 3 > 1$ is required to achieve distortion 0 if $0 < \epsilon < 1/2$. For variable-rate coding, a rate of

$$\frac{\mathcal{H}(Y_1) + \mathcal{H}(Y_2)}{2} = 0.5 + \mathcal{H}(2\alpha)$$

is needed to achieve distortion 0. Here $0.5 + \mathcal{H}(2\alpha) > 1$ bps if $\epsilon < 0.3899$. Thus, in either of these cases, the coding gain associated with using transform I rather than the KLT T is infinite at rate 1 bps.

The previous examples leave us with the impression that the KLT may be suboptimal only for sources that are very different from Gaussians. The next example shows that even for a source whose distribution is very close to a Gaussian, the KLT may still be suboptimal.

Let us consider a two-dimensional stationary vector source $X = (X_1, X_2)^t$. Suppose the joint pdf of (X_1, X_2) is

$$f(x_1, x_2) = \alpha f_1(x_1, x_2) + (1-\alpha) f_2(x_1, x_2)$$

where $\alpha \in (0,1)$ is a constant, and f_1 and f_2 are each two-dimensional Gaussian pdfs. Let the means of f_1 and f_2 be

$$\mu_1 = (\mu_{11}, \mu_{12})^t$$
$$\mu_2 = (\mu_{21}, \mu_{22})^t$$

and the covariance matrices of f_1 and f_2 be

$$\Phi_1 = \begin{bmatrix} \sigma_{11}^2 & \rho_1\sigma_{11}\sigma_{12} \\ \rho_1\sigma_{11}\sigma_{12} & \sigma_{12}^2 \end{bmatrix}$$
$$\Phi_2 = \begin{bmatrix} \sigma_{21}^2 & \rho_2\sigma_{21}\sigma_{22} \\ \rho_2\sigma_{21}\sigma_{22} & \sigma_{22}^2 \end{bmatrix}$$

respectively. Without loss of generality, we assume that

$$\mu_{11} = \mu_{12} = \rho_1 = 0.$$

We also assume that $\sigma_{11}^2 \neq \sigma_{12}^2$. If α is very close to 1, f is close to the Gaussian pdf f_1 and we would expect the optimal transform matrix for X to be close to the identity matrix. However, the mean of X is $\mu = (1-\alpha)\mu_2$ and the covariance matrix is

$$\begin{aligned}\Phi &= \alpha(\Phi_1 + \mu_1\mu_1^t) + (1-\alpha)(\Phi_2 + \mu_2\mu_2^t) - \mu\mu^t \\ &= \begin{bmatrix} \alpha\sigma_{11}^2 + (1-\alpha)\sigma_{21}^2 + \alpha(1-\alpha)\mu_{21}^2 & (1-\alpha)\rho_2\sigma_{21}\sigma_{22} + \alpha(1-\alpha)\mu_{21}\mu_{22} \\ (1-\alpha)\rho_2\sigma_{21}\sigma_{22} + \alpha(1-\alpha)\mu_{21}\mu_{22} & \alpha\sigma_{12}^2 + (1-\alpha)\sigma_{22}^2 + \alpha(1-\alpha)\mu_{22}^2 \end{bmatrix}.\end{aligned}$$

It is possible to set

$$\alpha\sigma_{11}^2 + (1-\alpha)\sigma_{21}^2 + \alpha(1-\alpha)\mu_{21}^2 = \alpha\sigma_{12}^2 + (1-\alpha)\sigma_{22}^2 + \alpha(1-\alpha)\mu_{22}^2$$

by choosing μ_2, ρ_2, and the variances. Then, as long as

$$(1-\alpha)\rho_2\sigma_{21}\sigma_{22} + \alpha(1-\alpha)\mu_{21}\mu_{22} \neq 0$$

the transform matrix T given by (4) would be a KLT transform matrix for X, which is inefficient.

Example V.5: The source pdf is nonzero everywhere and is close to Gaussian.

Let $\alpha = 0.99$, $\mu_2 = (0.5, 3)^t$, and

$$\Phi_1 = \begin{bmatrix} 5 & 0 \\ 0 & 2 \end{bmatrix} \quad \text{and} \quad \Phi_2 = \begin{bmatrix} 1 & 0 \\ 0 & 289.3375 \end{bmatrix}.$$

As a result, $\mu = (0.005, 0.03)^t$ and

$$\Phi = \begin{bmatrix} 4.9625 & 0.0149 \\ 0.0149 & 4.9625 \end{bmatrix}$$

and $Y = (Y_1, Y_2) = TX$ is a KLT for X, where T is defined in (4).

We exploit the interesting fact that any 2×2 symmetric, nonsingular matrix with positive entries is decorrelated by a 45° rotation.

Using numerical methods, we find that the coding gain associated with using transform I rather than the KLT T is about 1.6426 dB for fixed-rate coding, and the coding gain for variable-rate coding is about 0.6124 dB. The joint pdf f is still very close to a Gaussian, e.g.,

$$D(f_1 \| f) \approx 0.0109$$

and

$$\|f - f_1\|_2 / \|f_1\|_2 \approx 9.456 \times 10^{-3}$$

where the relative entropy between f_1 and f is defined as

$$D(f_1 \| f) = \int_{-\infty}^{\infty} f_1(x)(\log(f_1(x)/f(x)))dx$$

and

$$\|f\|_2 = \sqrt{\int_{-\infty}^{\infty} f^2(x)dx}$$

for any function f. Some of the analysis can be found in the Appendix in Lemma IX.1.

If we leave α, μ_1, and μ_2 unchanged and set

$$\Phi_1 = \begin{bmatrix} 5 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \Phi_2 = \begin{bmatrix} 1 & 0 \\ 0 & 388.3375 \end{bmatrix}$$

then the coding gain is about 2.1609 dB for fixed-rate coding and 1.4572 dB for variable-rate coding with

$$D(f_1 \| f) \approx 0.0122$$

and

$$\|f - f_1\|_2 / \|f_1\|_2 \approx 9.666 \times 10^{-3}.$$

We have also ran a series of experiments with different parameters (allowing the KLT to deviate from a 45° rotation and α to grow to 0.5). The numerical results show that for most sources, the KLT is suboptimal with varying degrees of performance degradation.

This also suggests that the KLT may be suboptimal for most if not all non-Gaussian sources. Consider estimating an arbitrary source by a Gaussian mixture model (GMM). Even if a natural source is modeled by a single dominant Gaussian component mixed with a collection of far smaller Gaussian components, the above example suggests that the KLT may be suboptimal.

VI. KLTs With Independent Transform Coefficients Can Be Suboptimal for Fixed-Rate Transform Coding

The previous examples show that even a best KLT can be suboptimal for transform coding. In those examples, the transform vector Y has coefficients that are decorrelated but not independent. In this section, we prove Theorem III.3 by showing that even when the transform vector has independent coefficients, a KLT can be suboptimal for fixed-rate transform coding. Notice that this result applies only to fixed-rate coding since we showed in Section IV that when the coefficients of the KLT are independent, a KLT is optimal for variable-rate transform coding (for reasonably smooth pdfs).

We construct an example by using a two-dimensional Laplacian source $X = (X_1, X_2)^t$ with independent components. We show that the transform matrix T that rotates the source by 45° is a better transform for fixed-rate transform coding on X than the identity matrix I given by the KLT.

Let b and σ be positive constants and

$$f_{X_1}(x_1) = \frac{1}{\sigma\sqrt{2}} \exp\left(-\frac{|x_1|\sqrt{2}}{\sigma}\right)$$
$$f_{X_2}(x_2) = \frac{1}{b\sigma\sqrt{2}} \exp\left(-\frac{|x_2|\sqrt{2}}{b\sigma}\right).$$

Then

$$\|f_{X_1}\|_{1/3} = \left(2\int_0^\infty \left(\frac{1}{\sigma\sqrt{2}} e^{-x_1\sqrt{2}/\sigma}\right)^{1/3} dx_1\right)^3 = 54\sigma^2.$$

Similarly,

$$\|f_{X_2}\|_{1/3} = 54b^2\sigma^2.$$

Let the transform matrix T be the transform matrix defined in (4). Then the marginal pdfs of the transform vector $Y = (Y_1, Y_2)^t = TX$ are (details appear in the Appendix), for $b \neq 1$

$$f_{Y_1}(y) = f_{Y_2}(y) = \frac{1}{(b^2-1)\sigma}\left[b\exp\left(-\frac{2|y|}{b\sigma}\right) - \exp\left(-\frac{2|y|}{\sigma}\right)\right].$$

If we let $b = 4/3$, then

$$\|f_{Y_1}\|_{1/3} = \|f_{Y_2}\|_{1/3} = \frac{3\sigma^2}{7}\left(4^{4/3}-1\right)^3.$$

Therefore, the coding gain obtained by quantizing Y instead of the independent component source X is

$$G_{T,I}^{(\mathrm{fr})} = \frac{54b\sigma^2}{\frac{3\sigma^2}{7}\left(4^{4/3}-1\right)^3} = \frac{168}{\left(4^{4/3}-1\right)^3} \approx 1.097 \quad \text{(i.e., 0.40 dB} > \text{0 dB)}.$$

Note that the observed problem is not limited to $b = 4/3$. For example, if $b = 1$, the KLT is not unique, and it is still not optimal to quantize the independent components. The marginal pdfs of (Y_1, Y_2) become

$$f_{Y_1}(y) = f_{Y_2}(y) = \frac{1}{\sigma^2} e^{-2|y|/\sigma}\left(|y| + \frac{\sigma}{2}\right).$$

Therefore,

$$\|f_{Y_1}\|_{1/3} = \|f_{Y_2}\|_{1/3} = \frac{81e\sigma^2}{2}\Gamma^3(4/3, 1/3)$$

where the "incomplete Gamma function" is defined as

$$\Gamma(u,v) = \int_v^\infty t^{u-1}e^{-t}dt$$

(see [12, p. 317, eq. 3.381.3]). Therefore, the coding gain obtained by using transform T rather than the KLT I (that is quantizing Y instead of the independent component source X) is

$$G_{T,I}^{(\mathrm{fr})} = \frac{54\sigma^2}{\frac{81e\sigma^2}{2}\Gamma^3(4/3,1/3)} = \frac{4}{3e\Gamma^3(4/3,1/3)} \approx 1.167 \ \text{(i.e., 0.67 dB)}.$$

Since $G_{T,I}^{(\mathrm{fr})} > 1$ in each of these examples, the distortion in quantizing the independent Laplacian random variables X_1 and X_2 is greater than the distortion in quantizing the dependent random variables Y_1 and Y_2, giving the following corollary.

Corollary VI.1: For an i.i.d. Laplacian source $\{X_i\}$, it is asymptotically better to scalar quantize successive sums and differences than to scalar quantize the source directly if using fixed-rate coding. That is, for all i

$$\lim_{R\to\infty} 2^{2R} \cdot \left\{E\left(\frac{Q_+(X_i+X_{i+1})+Q_-(X_i-X_{i+1})}{2} - X_i\right)^2 + E\left(\frac{Q_+(X_i+X_{i+1})-Q_-(X_i-X_{i+1})}{2} - X_{i+1}\right)^2\right\} < \lim_{R\to\infty} 2^{2R}\left(2E[(\hat{Q}(X_i)-X_i)^2]\right)$$

where $\hat{Q}$, Q_+, and Q_- are optimal scalar quantizers for X_i, $X_i + X_{i+1}$, and $X_i - X_{i+1}$, respectively.

It should be noted that, in this scenario, the suboptimality of the KLT stems in part from the fixed-rate coding framework, which is inherently suboptimal to the variable-rate coding framework in terms of rate-distortion performance.

VII. Summary and Conclusion

In this paper, a family of sources has been demonstrated for which the KLT transform is suboptimal in a scalar quantized transform coding system (using either fixed-rate or variable-rate scalar quantizers).

We considered three scenarios. First, there are sources for which the KLT is not unique, and the worst KLT may be arbitrarily worse than the optimal transform in terms of coding gain in both fixed-rate and variable-rate transform coding. Second, there are sources for which even a best (or only) KLT can give suboptimal performance in both fixed-rate and variable-rate transform coding. An independent work by Goyal also shows this result for variable-rate transform coding by using a different example [13] and Mallat also mentioned briefly in [14, Sec. 11.3.2] that for non-Gaussian circular stationary processes with piecewise smooth realizations, wavelet bases yield better distortion-rate performance than KLT bases. Finally, there exist

sources for which the KLT that yields independent components is suboptimal in fixed-rate coding. The transform that yields independent components is generally optimal in variable-rate coding under the high-rate assumption (note that this is not true at low rates in general).

In some cases, the suboptimality stems in part from the suboptimality of fixed-rate coding versus variable-rate coding, e.g., when the KLT yields independent components. In this respect, it may be better to consider entropy coding as a part of the standard model of transform coding as proposed by Goyal in [15]. However, the results on fixed-rate coding are still interesting since they give us a better picture of the KLT's weakness and fixed-rate coding has lower complexity and is less sensitive to transmission errors. In addition, the original work on the optimality of the KLT [2] uses this framework.

For real sources, some experimental results on the suboptimality of the KLT have been reported in [5] and [16]. Especially in [16], the transform matrices found by a new algorithm outperform the KLT by 0.3 to 1.2 dB on benchmark images like "Barbara" and "Goldhill." When the impact of the lack of stationarity in the source is reduced by using adaptive coding, the performance of the KLT is about 0.4–1 dB worse.

All these results indicate that decorrelation is not sufficient for optimality. There are two possible ways to pursue a truly optimal transform in variable-rate coding: we can either minimize the sum of the transformed component's entropy as (2) suggests, or minimize the distortion directly for a particular probabilistic model (e.g., the GMM model) as done in [16].

Another relevant issue is the suboptimality of the popular discrete cosine transform (DCT). The DCT has several advantages over the KLT since the DCT is low in complexity and is data independent. However, the DCT is generally inferior to the KLT in terms of rate-distortion performance, especially when the source contains more high-frequency content (e.g., text in an image) as shown in previous literature (e.g., [17]).

APPENDIX

Proof of Lemma II.3: If the bit allocation results in rate R_i and distortion $D_i(R_i)$ in coding Z_i, for $i = 1, \ldots, n$, then let

$$\overline{D}(R) = \frac{1}{n}\sum_{i=1}^{n} D_i(R_i).$$

The quantity $n\overline{D}(R)$ is the code's total distortion. If

$$n\overline{D}(R) \ll \sigma_{\min}^2 = \min\{\sigma_1^2, \ldots, \sigma_n^2\}$$

then

$$0 \le D_i(R_i) \le n\overline{D}(R) \ll \sigma_{\min}^2$$

i.e., $\overline{D}(R) \to 0$ is equivalent to each $R_i \to \infty$, or in order to minimize $\overline{D}(R)$, $R \to \infty$ is equivalent to each $R_i \to \infty$. From [18] (see also [7, pp. 295–302])

$$\lim_{R_i \to \infty} D_i(R_i) 2^{2R_i} = \frac{1}{12} 2^{2h(Z_i)}.$$

Thus,

$$\begin{aligned}\frac{1}{(12)^n} 2^{2n\bar{h}(Z^n)} &= \prod_{i=1}^{n} \frac{1}{12} 2^{2h(Z_i)} \\ &= \lim_{R\to\infty} 2^{2nR} \prod_{i=1}^{n} D_i(R_i) \\ &\le \lim_{R\to\infty} 2^{2nR} (\overline{D}(R))^n\end{aligned}$$

where the last inequality follows from the arithmetic/geometric mean inequality

$$\left(\prod_{i=1}^{n} D_i(R_i)\right)^{1/n} \le \frac{1}{n}\sum_{i=1}^{n} D_i(R_i) = \overline{D}(R)$$

with equality if and only if $D_i(R_i) = \overline{D}(R)$ for $i = 1, \ldots, n$. Thus the optimal bit allocation gives the desired result. □

Now we present detailed calculations of Section V-A.

Since $f_U(x) = 0$ for all $x \notin [-A/2, A/2]$, then either $f_U(x + A) = 0$ or $f_U(x - A) = 0$ for any x, and thus,

$$\begin{aligned}&\|f_{X_1}\|_{1/3} \\ &= \left(\int_{-\infty}^{\infty} f_{X_1}^{1/3}(x)dx\right)^3 \\ &= \left(\int_{-\infty}^{\infty} \left[\frac{1}{2} f_U(x+A) + \frac{1}{2} f_U(x-A)\right]^{1/3} dx\right)^3 \\ &= \frac{1}{2}\left(\int_{-\infty}^{\infty} \left[f_U^{1/3}(x+A) + f_U^{1/3}(x-A)\right] dx\right)^3 \quad (13) \\ &= 4\left(\int_{-\infty}^{\infty} f_U^{1/3}(x)dx\right)^3 \\ &= 4\|f_U\|_{1/3}.\end{aligned}$$

The joint cumulative distribution function and joint pdf of (X_1, X_2) are

$$\begin{aligned}F_{X_1,X_2}(x_1,x_2) &= P\{U_1 + S_1 \le x_1, U_2 + S_2 \le x_2\} \\ &= \beta[F_U(x_1+A)F_U(x_2+A) \\ &\quad + F_U(x_1-A)F_U(x_2-A)] \\ &\quad + \alpha[F_U(x_1+A)F_U(x_2-A) \\ &\quad + F_U(x_1-A)F_U(x_2+A)] \\ f_{X_1,X_2}(x_1,x_2) &= \beta[f_U(x_1+A)f_U(x_2+A) \\ &\quad + f_U(x_1-A)f_U(x_2-A)] \\ &\quad + \alpha[f_U(x_1+A)f_U(x_2-A) \\ &\quad + f_U(x_1-A)f_U(x_2+A)].\end{aligned}$$

Since

$$\begin{aligned}Y_1 &= (X_1 + X_2)/\sqrt{2} \\ Y_2 &= (X_1 - X_2)/\sqrt{2}\end{aligned}$$

and f_U is symmetric, the marginal pdfs of the decorrelated random vector Y are

$$\begin{aligned}f_{Y_1}(y_1) &= \sqrt{2}\int_{-\infty}^{\infty} f_{X_1,X_2}\left(y_2, y_1\sqrt{2} - y_2\right) dy_2 \\ &= \sqrt{2}\Big[\beta(f_U * f_U)\left(y_1\sqrt{2} + 2A\right) \\ &\quad + \beta(f_U * f_U)\left(y_1\sqrt{2} - 2A\right) \\ &\quad + 2\alpha(f_U * f_U)\left(y_1\sqrt{2}\right)\Big]\end{aligned}$$

$$\begin{aligned} f_{Y_2}(y_2) &= \sqrt{2}\int_{-\infty}^{\infty} f_{X_1,X_2}\left(y_1, y_1 - y_2\sqrt{2}\right) dy_1 \\ &= \sqrt{2}\Big[\alpha(f_U * f_U)\left(y_2\sqrt{2}+2A\right) \\ &\quad + \alpha(f_U * f_U)\left(y_2\sqrt{2}-2A\right) \\ &\quad +2\beta(f_U * f_U)\left(y_2\sqrt{2}\right)\Big]. \end{aligned}$$

From the marginal pdfs we obtain

$$\begin{aligned} \|f_{Y_1}\|_{1/3} &= \sqrt{2}\Bigg(\int_{-\infty}^{\infty}\Big[\beta^{1/3}(f_U * f_U)^{1/3}\left(y_1\sqrt{2}+2A\right) \\ &\quad + \beta^{1/3}(f_U * f_U)^{1/3}\left(y_1\sqrt{2}-2A\right) \\ &\quad +2^{1/3}\alpha^{1/3}(f_U * f_U)^{1/3}(y_1\sqrt{2})\Big] dy_1\Bigg)^3 \\ &= \frac{1}{2}(2^{1/3}\alpha^{1/3}+2\beta^{1/3})^3\|f_U * f_U\|_{1/3}. \end{aligned}$$

Similarly,

$$\|f_{Y_2}\|_{1/3} = \frac{1}{2}(2^{1/3}\beta^{1/3}+2\alpha^{1/3})^3\|f_U * f_U\|_{1/3}. \qquad \square$$

Proof of Lemma V.1: Since γ_ϵ is continuous on $[0, 1/2]$

$$\begin{aligned} \gamma_0 &= (2+2^{1/3})^3/4 \\ \gamma_{1/2} &= 2 \end{aligned}$$

follow by substitution. Let $J = \alpha^{1/3}$, $K = \beta^{1/3}$, and

$$H = \left(2^{1/3}K + 2J\right)\left(2^{1/3}J + 2K\right).$$

Then for all $\epsilon \in (0, 1/2)$

$$\begin{aligned} \frac{dH}{d\epsilon} &= \frac{\partial H}{\partial J}\frac{\partial J}{\partial \alpha}\frac{d\alpha}{d\epsilon} + \frac{\partial H}{\partial K}\frac{\partial K}{\partial \beta}\frac{d\beta}{d\epsilon} \\ &= \left(2^{7/3}J + (4+2^{2/3})K\right)\left(\frac{1}{3}J^{-2}\right)\left(-\frac{1}{2}\right) \\ &\quad + \left((4+2^{2/3})J + 2^{7/3}K\right)\left(\frac{1}{3}K^{-2}\right)\left(\frac{1}{2}\right) \\ &= \frac{1}{6J^2K^2}\Big[(4+2^{2/3})J^3 + 2^{7/3}J^2K \\ &\quad -2^{7/3}JK^2 - (4+2^{2/3})K^3\Big] \\ &= \frac{4+2^{2/3}}{6J^2K^2}(J-K) \\ &\quad \cdot\left[J^2 + \left(1+\frac{2^{7/3}}{4+2^{2/3}}\right)JK + K^2\right] \\ &< 0 \end{aligned}$$

since $J > 0$, $K > 0$, and $J - K < 0$. Thus, H is monotonically decreasing and therefore so is γ_ϵ. $\square$

The following lemma is useful in estimating the coding gain in Example V.5.

Lemma IX.1: If f_1 and f_2 are two pdfs, and

$$f(x) = \alpha f_1(x) + (1-\alpha)f_2(x)$$

then

$$\begin{aligned} \alpha h(f_1) + (1-\alpha)h(f_2) &\le h(f) \\ &\le \alpha h(f_1) + (1-\alpha)h(f_2) + \mathcal{H}(\alpha). \end{aligned}$$

Proof: The first inequality comes directly from the fact that $h(f)$ is a concave function of f (a discrete version appears in [19, Theorem 2.7.3]).

For the second inequality, since $f(x) \ge \alpha f_1(x)$ and $f(x) \ge (1-\alpha)f_2(x)$ for any x, one gets

$$\begin{aligned} h(f) &= -\int_{-\infty}^{\infty}(\alpha f_1(x) + (1-\alpha)f_2(x))\log f(x)dx \\ &= -\int_{-\infty}^{\infty}\alpha f_1(x)\log f(x)dx \\ &\quad -\int_{-\infty}^{\infty}(1-\alpha)f_2(x)\log f(x)dx \\ &\le -\int_{-\infty}^{\infty}\alpha f_1(x)\log(\alpha f_1(x))dx \\ &\quad -\int_{-\infty}^{\infty}(1-\alpha)f_2(x)\log((1-\alpha)f_2(x))dx \\ &= -\alpha\log\alpha + \alpha h(f_1) \\ &\quad -(1-\alpha)\log(1-\alpha) + (1-\alpha)h(f_2) \end{aligned}$$

that gives us the desired result. $\square$

This lemma gives a good estimate of $h(f)$ when α is close to 0 or 1.

We give here some analysis for Example V.5.

Since

$$\begin{aligned} T\mu_1 &= (0,0)^t \\ T\mu_2 &= (2.4749, -1.7678)^t \\ T'\Phi_1 T &= \begin{bmatrix} 3.5 & 1.5 \\ 1.5 & 3.5 \end{bmatrix} \\ T'\Phi_2 T &= \begin{bmatrix} 145.1687 & -144.1687 \\ -144.1687 & 145.1687 \end{bmatrix} \end{aligned}$$

the marginal pdfs of Y_1 and Y_2 are

$$\begin{aligned} f_{Y_1}(y_1) &= \alpha\mathcal{N}(0, 3.5) + (1-\alpha)\mathcal{N}(2.4749, 145.1687) \\ f_{Y_2}(y_2) &= \alpha\mathcal{N}(0, 3.5) + (1-\alpha)\mathcal{N}(-1.7678, 145.1687) \end{aligned}$$

where $\mathcal{N}(\mu, \sigma^2)$ represents a normal pdf with mean μ and variance σ^2. Thus,

$$\begin{aligned} h(Y_1) &\ge \alpha h(\mathcal{N}(0, 3.5)) + (1-\alpha)h(\mathcal{N}(2.4749, 145.1687)) \\ &> 2.9776 \\ h(Y_1) &\le \alpha h(\mathcal{N}(0, 3.5)) + (1-\alpha)h(\mathcal{N}(2.4749, 145.1687)) \\ &\quad + \mathcal{H}(\alpha) \\ &< 3.0585 \end{aligned}$$

by Lemma IX.1. Similarly, we can bound $h(Y_2)$ as

$$2.9776 < h(Y_2) < 3.0585$$

and numerical methods show that

$$h(Y_1), h(Y_2) \approx 3.0246.$$

On the other hand,

$$f_{X_1}(x_1) = \alpha\mathcal{N}(0,5) + (1-\alpha)\mathcal{N}(0.5,1)$$

and thus, by Lemma IX.1

$$\begin{aligned} 3.1964 < h(X_1) &< 3.2773 \\ 2.5829 < h(X_2) &< 2.6638. \end{aligned}$$

Numerical methods show that

$$\begin{aligned} h(X_1) &\approx 3.2026 \\ h(X_2) &\approx 2.6430. \end{aligned}$$

Therefore, for variable-rate coding, the coding gain associated with using transform I rather than the KLT T is

$$\begin{aligned} G_{I,T}^{(\mathrm{vr})} &= 2^{h(Y_1)+h(Y_2)-h(X_1)-h(X_2)} \\ &> 1.0098 \quad \text{(i.e., 0.0424 dB)}. \end{aligned}$$

The relative entropy $D(f_1(x_1,x_2)\|f(x_1,x_2))$ can be bounded using its convexity (see, e.g., the discrete version from Theorem 2.7.2 in [19]), i.e.,

$$\begin{aligned} D(f_1\|f) &= D(\alpha f_1 + (1-\alpha)f_1 \| \alpha f_1 + (1-\alpha)f_2) \\ &\le \alpha D(f_1\|f_1) + (1-\alpha)D(f_1\|f_2) \\ &= (1-\alpha)D(f_1\|f_2) \\ &\le 0.048. \end{aligned}$$

Numerical results show that $D(f_1\|f) \approx 0.0122$. □

Here we present detailed calculations of Section VI.

Let the transform matrix T be the transform matrix defined in (4). Then the joint pdf of the transform vector $Y = (Y_1, Y_2)^t = TX$ is

$$\begin{aligned} &f_{Y_1,Y_2}(y_1,y_2) \\ &= \frac{\partial^2}{\partial y_1 \partial y_2} \int_{-\infty}^{(y_1+y_2)/\sqrt{2}} \int_{x_1-y_2\sqrt{2}}^{y_1\sqrt{2}-x_1} f_{X_1}(x_1) f_{X_2}(x_2) dx_2 dx_1 \\ &= f_{X_1}\left(\frac{y_1+y_2}{\sqrt{2}}\right) f_{X_2}\left(\frac{y_1-y_2}{\sqrt{2}}\right) \\ &= \frac{1}{2b\sigma^2} \exp\left(-\frac{|y_1+y_2|}{\sigma} - \frac{|y_1-y_2|}{b\sigma}\right) \end{aligned}$$

and the marginal pdfs are (for $b \neq 1$)

$$\begin{aligned} f_{Y_1}(y) &= f_{Y_2}(y) \\ &= \frac{1}{(b^2-1)\sigma}\left[b\exp\left(-\frac{2|y|}{b\sigma}\right) - \exp\left(-\frac{2|y|}{\sigma}\right)\right]. \end{aligned}$$

Therefore, by setting

$$w = \exp(-2|y|/(3b\sigma))$$

one gets

$$\begin{aligned} \|f_{Y_1}\|_{1/3} &= \|f_{Y_2}\|_{1/3} \\ &= \left(2\int_0^\infty \left(\frac{1}{(b^2-1)\sigma}\left(b\exp\left(-\frac{2y}{b\sigma}\right)\right.\right.\right. \\ &\qquad \left.\left.\left. - \exp\left(-\frac{2y}{\sigma}\right)\right)\right)^{1/3} dy\right)^3 \\ &= \frac{8}{(b^2-1)\sigma}\left(\int_0^1 \frac{3b\sigma}{2}\left(bw^3 - w^{3b}\right)^{1/3}\frac{dw}{w}\right)^3 \\ &= \frac{27b^3\sigma^2}{b^2-1}\left(\int_0^1 \left(b - w^{3b-3}\right)^{1/3} dw\right)^3. \end{aligned}$$

If we let $b = 4/3$, then

$$\begin{aligned} \|f_{Y_1}\|_{1/3} &= \|f_{Y_2}\|_{1/3} \\ &= \frac{576\sigma^2}{7}\left(\int_0^1 \left(\frac{4}{3} - u\right)^{1/3} du\right)^3 \\ &= \frac{576\sigma^2}{7}\left(\frac{3}{4}\left(\left(\frac{4}{3}\right)^{4/3} - \left(\frac{1}{3}\right)^{4/3}\right)\right)^3 \\ &= \frac{3\sigma^2}{7}\left(4^{4/3} - 1\right)^3. \end{aligned}$$

If $b = 1$, the joint density of (Y_1, Y_2) becomes

$$f_{Y_1,Y_1}(y_1,y_2) = \frac{1}{2\sigma^2}\exp\left(-\frac{|y_1+y_2|}{\sigma} - \frac{|y_1-y_2|}{\sigma}\right)$$

and the marginal pdfs are

$$f_{Y_1}(y) = f_{Y_2}(y) = \frac{1}{\sigma^2}e^{-2|y|/\sigma}\left(|y| + \frac{\sigma}{2}\right). \qquad \square$$

ACKNOWLEDGMENT

The authors are grateful to the anonymous reviewers and the Associate Editor Ram Zamir for their detailed suggestions and comments and to Vivek Goyal and Tamás Linder for pointing out citations [10] and [8], respectively.

REFERENCES

[1] P. L. Dragotti, G. Poggi, and A. R. P. Ragozini, "Compression of multispectral images by three-dimensional SPIHT algorithm," *IEEE Trans. Geosci. Remote Sensing*, vol. 38, pp. 416–428, Jan. 2000.

[2] J.-Y. Huang and P. M. Schultheiss, "Block quantization of correlated Gaussian random variables," *IEEE Trans. Commun.*, vol. COM-11, pp. 289–296, Sept. 1963.

[3] V. K. Goyal, J. Zhuang, and M. Vetterli, "Transform coding with backward adaptive updates," *IEEE Trans. Inform. Theory*, vol. 46, pp. 1623–1633, July 2000.

[4] A. Koschman, "On the filtering of nonstationary time series," in *Proc. 1954 Electronics Conf.*, 1954, p. 126.

[5] H. Feng and M. Effros, "On the rate-distortion performance and computational efficiency of the Karhunen–Loeve transform for lossy data compression," *IEEE Trans. Image Processing*, vol. 11, pp. 113–122, Feb. 2002.

[6] K. Popat and K. Zeger, "Robust quantization of memoryless sources using dispersive FIR filters," *IEEE Trans. Commun.*, vol. 40, pp. 1670–1674, Nov. 1992.

[7] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*. Boston, MA: Kluwer Academic, 1992.

[8] R. M. Gray, T. Linder, and J. Li, "A Lagrangian formulation of Zador's entropy-constrained quantization theorem," *IEEE Trans. Inform. Theory*, vol. 48, pp. 695–707, Mar. 2002.

[9] W. H. Press, *Numerical Recipes in C : The Art of Scientific Computing*. Cambridge, U.K.: Cambridge Univ. Press, 1992, ch. 11, pp. 460–481.

[10] V. K. Goyal, "Theorems and counterexamples in transform coding," in *Proc. DIMACS Workshop on Source Coding and Harmonic Analysis*, New Brunswick, NJ, May 2002.

[11] K. Sayood, *Introduction to Data Compression*. San Francisco, CA: Morgan Kaufmann, 1996.

[12] I. S. Gradshteyn and I. M. Ryzhik, *Table of Integrals, Series, and Products*. New York: Academic, 1980.

[13] V. K. Goyal, private communication.

[14] S. Mallat, *A Wavelet Tour of Signal Processing*. San Diego, CA: Academic, 1998.

[15] V. K. Goyal, "Theoretical foundations of transform coding," *IEEE Signal Processing Mag.*, vol. 18, no. 5, pp. 9–21, Sept. 2001.

[16] C. Archer and T. K. Leen, "A generalized Lloyd-type algorithm for adaptive transform coder design," *IEEE Trans. Signal Processing*, vol. 52, pp. 255–264, Jan. 2004.

[17] H. Feng and M. Effros, "Separable Karhunen–Loeve transforms for the weighted universal transform coding algorithm," in *Proc. IEEE Int. Conf, Acoustics, Speech, and Signal Processing*, Phoenix, AZ, Mar. 1999.

[18] H. Gish and J. N. Pierce, "Asymptotically efficient quantizing," *IEEE Trans. Inform. Theory*, vol. IT-14, pp. 676–683, Sept. 1968.

[19] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.