High Rate Vector Quantization for Detection

Riten Gupta and Alfred O. Hero, III *

November 23, 2002

Submitted to *IEEE Transactions on Information Theory*, Sep. 2001
Revised Nov. 2002

Abstract

We investigate high rate quantization for various detection and reconstruction loss criteria. A new distortion measure is introduced which accounts for global loss in best attainable binary hypothesis testing performance. The distortion criterion is related to the area under the receiver-operating-characteristic (ROC) curve. Specifically, motivated by Sanov's theorem, we define a performance curve as the trajectory of the pair of optimal asymptotic Type I and Type II error rates of the most powerful Neyman-Pearson test of the hypotheses. The distortion measure is then defined as the difference between the area-under-the-curve (AUC) of the optimal pre-encoded hypothesis test and the AUC of the optimal post-encoded hypothesis test. As compared to many previously introduced distortion measures for decision making, this distortion measure has the advantage of being independent of any detection thresholds or priors on the hypotheses, which are generally difficult to specify in the code design process. A high resolution Zador-Gersho type of analysis is applied to characterize the point density and the inertial profile associated with the optimal high rate vector quantizer. The analysis applies to a restricted class of high-rate quantizers that have bounded cells with vanishing volumes. The optimal point density is used to specify a Lloyd-type algorithm which allocates its finest resolution to regions where the gradient of the pre-encoded likelihood ratio has greatest magnitude.

Keywords: compression, binary hypothesis testing, discrimination, error exponents, Chernoff information, receiver operating characteristic (ROC), area under the curve (AUC)

Corresponding author:

Prof. Alfred O. Hero, III
Room 4229, Dept. of EECS University of Michigan
1301 Beal Avenue
Ann Arbor, MI 48109-2122
Tel. (734) 763-0564
Fax: (734) 763-8041
email: hero@eecs.umich.edu

*R. Gupta is with TRW Space and Electronics, Redondo Beach, CA. A. Hero is with the Dept. of EECS, University of Michigan, Ann Arbor, MI 48109-2122. This research was supported in part by AFOSR-AASERT grant F49620-96-0028 and ARO-MURI grant DAAH04-96-1-0337. Portions of this work were presented at the 2001 International Symposium on Information Theory, Washington, DC, June 2001

List of Figures

1 Introduction

In many applications, a source must be transmitted from a sensor to an end-user who will make decisions on the source from the received data. For example, an imaging radar or a video camera might transmit information to a user interested in the likelihood of presence of a particular target or object in the sensor's field of view. In such an application, it is often essential to reduce transmitted data rates by encoding the source prior to transmission at the cost of introducing a small amount of distortion at the decoder. The most common distortion measure is the mean square error (MSE) which forms the basis for the vast majority of lossy compression algorithms in use today [1, 2]. The MSE is a norm on the distortion incurred in reconstructing the source from its encoded values. However, it has long been recognized that MSE is not the most pertinent distortion measure when one is interested in the effect of compression on decision making performance. Indeed many different distortion measures have been previously proposed for assessing compression algorithms relative to detection, classification and other decision objectives [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 22, 13, 14, 15, 16].

This paper makes two contributions: 1) we apply a variant of the Zador-Gersho method of asymptotic high rate analysis to many-cell, multi-dimensional quantizers incorporating a Kullback-Liebler (KL) type of detection criterion; and 2) we introduce a new design criterion of this type which is closely related to the area under the receiver-operating-characteristic (ROC) curve. The new detection criterion is the area under the curve (AUC) specifying the optimal Type I and Type II error exponents specified by Sanov's theorem on asymptotic (large sample size) false alarm and miss probabilities. We compare the AUC criterion to other detection criteria including the information discrimination exponent of Stein's Lemma and the Chernoff information exponent of the Chernoff bound. For each of these criteria our high resolution analysis yields expressions for the optimal point density of the encoder which minimizes the information losses over all similarly constrained quantizers of fixed rate. By using the AUC optimal density in the LBG algorithm we show that one can obtain finite rate vector quantizers that have improved rate-distortion (in terms of AUC) characteristics relative to other quantizer designs.

Under a detectibility distortion measure such as probability of decision error, for minimum distortion the optimal quantizer must quantize the minimal sufficient statistic for detection, i.e. the likelihood ratio. In

this paper we call such a quantizer a log likelihood ratio (LLR) quantizer. As the minimal sufficient statistic is typically a non-linear and non-invertible function of the source, the induced source-domain quantization cells of such a detection-optimal quantizer may not be convex nor bounded, indeed some cells may have either zero or infinite volumes. A general high rate characterization of the source-domain cells of such a quantizer is difficult as it is strongly dependent on the underlying probability densities. For example, in some cases the distribution of the LLR is discrete valued with finite number of values. In these cases the LLR is itself an optimal finite rate encoder, incurring zero loss in detection performance, and higher rate quantization is pointless. While analysis of LLR quantizers is certainly worthwhile in this paper we focus on a restricted class of quantizers which we call small-cell quantizers. Small-cell quantizers have the property that the cells are bounded and have vanishing positive volume as the rate increases. With such quantizers the source can be reconstructed without distortion, under either detection hypothesis, as the number of cells goes to infinity, a condition that is only rarely satisfied by LLR quantizers. Furthermore, we conjecture that the high-rate small-cell assumption is not restrictive for mixed objectives consisting of a linear combination of MSE distortion and the detectibility distortion measures considered here.

We show that the optimal point densities of the high-rate small-cell quantizers are related to two important functions, called the *Fisher covariation profile* and the *discriminability*. Based on these optimal point densities a Lloyd-type compression algorithm is proposed under a congruent cell hypothesis and numerical comparisons are performed for several illustrative examples. A general characterization is that, as contrasted to MSE-optimal (minimum MSE) quantizers, detection-optimal quantizers should allocate finer resolution to regions where the gradient of the likelihood ratio has large magnitude.

Some background will be useful to place our contributions in the context of previous work. Quantization and source coding have been studied for many decades and the rich history is traced in [2]. Early research on asymptotic high rate quantization was reported by Zador [17] and Gersho [18]. In [19], Na and Neuhoff derived a formula for the asymptotic high-rate MSE of a vector quantizer in terms of two functions that characterize the quantizer, known as the *point density* and the *inertial profile*. These functions describe a quantizer's asymptotic distribution of points and cell sizes, respectively. In this paper we extend the results of [19] to distortion measures which incorporate information discrimination and other penalties for poor post-quantization detection performance.

The problem of optimal quantization for hypothesis testing has been analyzed for various quantization schemes and various distortion criteria. Kassam [3] considered quantization under an efficacy distortion measure for testing composite hypotheses $H_0 : \theta = 0$ versus $H_1 : \theta > 0$ under a parameterized density with scalar parameter θ. Poor and Thomas [4] investigated the quantization-induced loss in various Ali-Silvey distances between densities characterizing two simple hypotheses. Later Poor [5, 6] proposed the generalized f-divergence as a distortion measure and studied asymptotic high-rate quantization effects on this measure. From this work, it is seen that the loss in Kullback-Leibler distance due to quantization is a functional of a quantity called the *discriminability*, which plays a central role here. Picinbono and Duvaut [8] considered a deflection criterion similar to a signal-to-noise ratio (SNR) under one of two simple hypotheses. It was shown that maximization of this deflection criterion is achieved by a transform coder which quantizes the scalar likelihood ratio. Tsitsiklis [9] explores some properties of such likelihood ratio quantizers and he investigates optimality with respect to several divergence measures. Motivated by Chernoff's theorem, which bounds the exponential error rates of the Neyman-Pearson (NP) test, Benitz and Bucklew [7] proposed the loss in alpha entropy, also called Chernoff distance, as a distortion measure for scalar quantizers. Asymptotically optimal companding functions were then derived under a high resolution analysis. More recently Jain *et al* [15] proposed bounds on Chernoff distances as a distortion measure for quantization in the context of composite hypotheses. Flynn and Gray [10] consider a mixed distortion combining estimation and probability of detection for correlated observations in distributed sensing environments. Achievable rate-distortion regions are obtained for the case of two sensors which extend the lossless source coding analysis of Slepian and Wolf [25] to lossy source coding. These authors also presented non-asymptotic quantizer design for optimum detection performance via iterative maximization of the Chernoff distance. The distributed hypothesis testing problem with quantized observations is directly addressed in [11] by Longo, Lookabaugh, and Gray where an iterative algorithm for optimal scalar quantization is derived with loss in Bhattacharyya distance adopted as the distortion measure. Oehler and Gray [12] and Perlmutter *et al* [13] introduced a method of quantization and classification with a mixed distortion measure defined as a linear combination of MSE and Bayes risk. An iterative encoding algorithm was presented which minimizes this measure.

It is also worth mentioning the weighted MSE distortion metric which was introduced by Gardner and Rao [22], and later studied by Li *etal* [23] and Linder *etal* [24], to better account for perceptual distortion.

The main difference between Gardner and Rao's weighted MSE metric and the measures we consider is that our measures are specifically designed to capture the loss in detectibility due to quantization. The high rate results obtained for the weighted MSE distortion impose the condition of positive definiteness on the sensitivity matrix which weights the MSE norm [23, Condition 2, pp1083]. Our high rate expressions bear some similarity to the asymptotic expressions of [23]. However, in our expressions the weighting matrix is rank deficient and is a function of the gradient of the likelihood ratio. We avoid the positive definiteness condition assumed by [23] by imposing the aforementioned small-cell restriction.

A major difference between the approach of this paper to detection and previous approaches is that we tackle the problem of global optimization of the ROC curve and not just optimization at a point on the ROC curve. Both the AUC-difference and the Chernoff information considered in this paper are symmetric functionals of the hypothesized source densities under H_0 and H_1. Furthermore, the optimal encoder arising from minimizing the AUC-difference does not depend on the level of significance or the decision thresholds of the test. For example, the α-divergence criteria adopted in [7, 15] are indexed by α which specifies a threshold, a level of significance, and a particular point on the ROC curve. Likewise the mixed distortion criterion of [12, 13] depends on the Bayes risk which is parameterized by priors on the hypotheses which again specify a point on the ROC curve. Our high-resolution analysis follows along the lines of the approach taken in [19] for the MSE loss function. For the detection problem it turns out that the optimal high-rate quantizer depends on a matrix generalization of the inertial profile, called the *covariation profile*, that characterizes the cell shapes. The role of the covariation profile in our expressions for high-rate detectibility is similar to that of the normalized moment of inertia defined in the high-rate analysis of weighted MSE distortion of Li *etal* [23]. Our framework permits a lucid analysis of the merits of various quantizers with detection loss as they may be evaluated by their point densities and covariation profiles.

An outline of the paper is as follows. We briefly review elements of quantization for general tasks and discuss information discrimination criteria, including Chernoff information and AUC, in Section 2. In Section 3 we perform asymptotic high-rate analysis on these criteria. In Section 4 we obtain optimal point densities for various criteria. Finally, examples are presented in Section 5.

2 Vector Quantization

Let a k dimensional real valued source X take values $x \in \mathbb{R}^k$. A k-dimensional quantizer [1, 19] $Q = (\mathcal{S}, \mathcal{C})$ consists of a codebook $\mathcal{C} = \{x_1, \ldots, x_N\}$ and a set of cells $\mathcal{S} = \{S_1, \ldots, S_N\}$ that partition a bounded domain Ω which is a subset of $\mathbb{R}^k$. When $k = 1$ a quantizer is called a scalar quantizer while for $k > 1$ it is called a vector quantizer. Each codebook point x_i lies in cell S_i. The quantizer operator can be written as

$$Q(x) = x_i, \text{ for } x \in S_i$$

For a vector quantizer Q let $V_i = \int_{S_i} dx$ denote the volume of the ith cell. The *specific point density* [19] of Q is defined as

$$\zeta_s(x) = \frac{1}{NV_i}, \text{ for } x \in S_i.$$

This function is a normalized density as it is non-negative and its integral over Ω equals one. When integrated over a region $A \in \Omega$, it gives the approximate fraction of codebook points contained in A. Next, define the *diameter function* of the quantizer

$$d(x) = \sup\{\|u - v\| : u, v \in S_i\}, \text{ for } x \in S_i,$$

and the (scalar) *specific inertial profile* function $m(x)$ [18]

$$m(x) = \frac{\int_{S_i} \|y - x_i\|^2 dy}{V_i^{1+2/k}}, \text{ for } x \in S_i. \tag{1}$$

Note that $m(x)$ is invariant to scaling of S_i. This function contains partial information about the shapes of the cells of the quantizer. More information is provided by the following matrix valued function $M(x)$ which we call the *specific covariation profile*

$$M(x) = \frac{\int_{S_i} (y - x_i)(y - x_i)^T dy}{V_i^{1+2/k}}, \quad \text{for } x \in S_i.$$

It can be easily shown that if S_i is an ellipsoidal cell of the form $\{x : (x - x_i)^T R(x - x_i) \leq c\}$, where R is $k \times k$ symmetric positive definite and $c > 0$, then

$$M(x) = m_k^* |R|^{1/k} R^{-1}, \quad \text{for } x \in S_i,$$

where $|R|$ is the determinant of R and m_k^* is the moment of intertia of a unit sphere in $\mathbb{R}^k$. As the function $M(x)$ does not depend on the size parameter c it is also scale invariant. Furthermore, for spherical cells $M(x) = m_k^* I$ which is a scaled $k \times k$ identity matrix.

In this paper we restrict our treatment to product quantizers of n independent identically distributed (i.i.d.) samples of a k-dimensional source. This restriction allows us to use average MSE and error exponents to determine the MSE- and detection-optimal k-dimensional component quantizers applied to each sample. The restricted framework applies to either of the following scenarios: 1) n repeated temporal measurements (snapshots) of a k-dimensional source using a single sensor; or 2) a single snapshot of a network of n spatially distributed sensors measuring the same k-dimensional source. In either case, using the framework of this paper one can evaluate the average loss in detection/estimation performance arising from this restriction to product quantizers. For more details and examples the reader is referred to [26].

2.1 Distortion Measures

Let $\mathbf{X} = [X^{(1)}, \ldots, X^{(n)}]$ be an i.i.d. sample from some probability density function (p.d.f.) $q(x)$. The class of N cell product quantizers $Q^{(n)}$ over $\mathbf{R}^{nk}$ is defined for a realization $\mathbf{x}$ of $\mathbf{X}$ by

$$Q^{(n)}(\mathbf{x}) = \left[Q\big(x^{(1)}\big), \ldots, Q\big(x^{(n)}\big)\right].$$

where Q is an N cell quantizer over $\mathbb{R}^k$.

The quality of a product quantizer is measured by an average loss function, also called an average distortion, $J(Q^{(n)})$ which is specified according to the particular task to be performed on the compressed data $Q^{(n)}(\mathbf{X})$. When the task is optimal reconstruction of the source $\mathbf{X}$ from $Q^{(n)}(\mathbf{X})$ it is appropriate to use the mean squared reconstruction error (MSE)

$$J(Q^{(n)}) = \mathrm{MSE}(Q) \stackrel{\mathrm{def}}{=} \sum_{i=1}^{n} E[\|Q(X^{(i)}) - X^{(i)})\|^2] = nE[\|Q(X^{(1)}) - X^{(1)})\|^2].$$

As the MSE is a measure of source estimation error we call the N cell quantizer Q that minimizes MSE the *MSE-optimal* N cell quantizer.

When the task is to decide between two hypothesized source distributions

$$\begin{aligned} H_0 &: X^{(i)} \sim q_0(x) \\ H_1 &: X^{(i)} \sim q_1(x) \end{aligned} \tag{2}$$

it is appropriate to use some combination of probability of false alarm (Type I error) $P_F(Q^{(n)})$ and probability of miss (Type II error) $P_M(Q^{(n)})$ of an optimal detector of H_0 vs. H_1 operating on quantized data $Q^{(n)}(\mathbf{X})$. Even though composite hypotheses can be easily handled by marginalization in a manner similar to [15], in this paper we focus on the case of simple hypotheses.

Several different distortion measures for optimal post-quantization detection are given below. For each of these measures the optimal post-compression detector is a likelihood ratio test (LRT) with a threshold chosen to satisfy a false alarm constraint, to reflect a particular pair of priors on H_0 or H_1, or to attain minimax detection performance. A few comments comparing MSE-optimal and detection-optimal quantizers are in order. The distortion function of an MSE-optimal quantizer typically depends on the cell partition as well as the domain and definition of the codewords, e.g., cell centers or centroids, and is a strictly decreasing function of the rate of the quantizer, i.e., increasing the number of cells N of the quantizer always decreases MSE. However, the probability of error of a detection-optimal quantizer only depends on the cell partition and cardinality of the codeword set and may not strictly decrease in N, e.g., the detection-optimal quantizer is a binary partition ($N = 2$) of the source space $\mathbb{R}^k$ for any fixed LRT threshold [26].

When one of the error probabilities, e.g., false alarm, must be constrained a natural criterion to consider is the probability of miss P_M of a LRT whose threshold T_1 is selected to meet the prespecified false alarm constraint $P_F(T_1) = \alpha$. Assuming a prior $p = P(H_1)$, another option is to consider the minimum average probability of error

$$P_e(p) = P_M(T_2)\, p + P_F(T_2)\, (1-p),$$

where $P_M(T_2)$ and $P_F(T_2)$ are the miss and false alarm probabilities of the LRT operating on the compressed data with LRT threshold $T_2 = p/(1-p)$. When p is unknown, the minimax post-compression probability of error can be adopted

$$P_e(p^*) = \min_{p \in [0,1]} \{P_M(T_2)\, p + P_F(T_2)\, (1-p)\},$$

which is achieved by the LRT with minimax threshold $T_3 = p^*/(1-p^*)$ where $p = p^*$ is the minimizing solution. The performance of any LRT is described by the receiver operating characteristic (ROC), given here in parametric form,

$$\{(P_F(T), P_D(T)) : T \in \mathbb{R}\}, \tag{3}$$

where $P_D = 1 - P_M$ is the probability of detection of H_1. Thus each of the above detection criteria arises from evaluating the ROC at a particular point $P_F(T)$, $T = T_1, T_2, T_3$ respectively, on the coordinate axis. The disadvantage of any of the aforementioned detection probability criteria is that they are local: they are only relevant to compression detection performance for a single LRT threshold, i.e., a given P_F. The AUC criterion discussed below is a global T-independent alternative which accounts for the entire range of attainable miss and false alarm probabilities of the MP-LRT.

When both reconstruction and detection performance of the quantizer are of interest Gray *et al.* [16, 12, 13] proposed using a mixed criterion equivalent to

$$J(Q^{(n)}) = (1-\rho) \cdot J_E(Q^{(n)}) + \rho \cdot J_D(Q^{(n)}) \tag{4}$$

where J_D and J_E are average distortion criteria which are minimized for detection-optimal and MSE-optimal quantizers, respectively. The weighting factor $\rho \in [0,1]$ is used to trade detection performance for estimation performance of the quantizer.

2.2 Distortion via Error Exponents of the LRT

Let $\alpha = P_F$ and $\beta = P_M$ denote the false alarm and miss probabilities of the LRT operating directly on the data $\mathbf{X}$. Then $1 - \beta(\alpha)$, $\alpha \in [0,1]$ is an equivalent but direct parameterization of the pre-quantization ROC curve (3), also known as the power of the test. For specified level α of false alarm, the most powerful (MP) pre-quantization test of level α of the hypotheses (2) is the LRT

$$\frac{1}{n}\sum_{i=1}^{n} \Lambda(x^{(i)}) \underset{H_1}{\overset{H_0}{\gtrless}} T$$

where

$$\Lambda(x) = \log q_0(x)/q_1(x), \tag{5}$$

is the single sample log likelihood ratio and the threshold T is set such that the probability of false alarm is equal to α, which may require randomization when the distribution of the LRT statistic is discrete [27].

Likewise, if $\hat{\alpha}$ and $\hat{\beta}$ denote the false alarm amd miss probabilities of the LRT operating on the output $Q^{(n)}(\mathbf{X})$ of an N cell product quantizer the MP post-quantization test of level α is the LRT

$$\frac{1}{n}\sum_{i=1}^{n}\hat{\Lambda}(x^{(i)}) \underset{H_1}{\overset{H_0}{\gtrless}} T$$

where

$$\hat{\Lambda}(x) = \log \overline{q}_{0,N}(x)/\overline{q}_{1,N}(x),$$

and $\overline{q}_{i,N} = \{\int_{S_j} q_i(x)dx\}_{j=1}^{N}$, $i = 0, 1$, are the probability mass functions (p.m.f.'s) of the output of the N cell component quantizer $Q(X^{(1)})$ with cells $\{S_j\}$.

For large sample size n the performance of the MP-LRT is completely characterized by a set of error exponents related to the Kullback-Leibler (KL) divergence, also called relative entropy or discrimination. The KL divergence between two discrete sources with p.m.f.'s $q_a(x)$ and $q_b(x)$ is [28, 29]

$$L(q_a\|q_b) = \sum_i q_a(x_i)\log\frac{q_a(x_i)}{q_b(x_i)}. \tag{6}$$

while for continuous sources with densities $q_a(x)$ and $q_b(x)$ the KL divergence is

$$L(q_a\|q_b) = \int q_a(x)\log\frac{q_a(x)}{q_b(x)}dx. \tag{7}$$

Stein's lemma gives a large n asymptotic expression for the probability of miss β_n of the LRT of (2) for arbitrary false alarm level $\alpha > 0$ [28]

$$\lim_{n\to+\infty}(\beta_n)^{1/n} = e^{-L(q_0\|q_1)}.$$

Hence, we have the large n approximation

$$\beta_n \approx e^{-nL(q_0\|q_1)}. \tag{8}$$

The intrinsic loss in probability of miss performance due to quantization can be expressed in terms of the loss incurred in the discrimination appearing in the Stein approximation (8)

$$\Delta L_N \overset{\text{def}}{=} L(q_0\|q_1) - L(\overline{q}_{0,N}\|\overline{q}_{1,N}). \tag{9}$$

This is monotonically related to the loss ratio $\hat{\beta}_n/\beta_n$ incurred in the miss probabilities due to quantization.

The Stein approximation (8) to the miss probability provides no information about the tradeoff between miss and false alarm probability. Sanov's theorem provides such information. Let β_n and α_n denote these respective probabilities. Then Sanov's theorem gives the following large n approximatons as a function of the LRT threshold T [28, 30, 31]:

$$\begin{aligned} \alpha_n &\approx e^{-nL(q_\lambda \| q_0)} \\ \beta_n &\approx e^{-nL(q_\lambda \| q_1)}. \end{aligned} \tag{10}$$

where the "tilted density" has been defined as

$$q_\lambda(x) = \frac{q_0(x)^{1-\lambda} q_1(x)^\lambda}{\int q_0(y)^{1-\lambda} q_1(y)^\lambda dy} \tag{11}$$

and the tilt parameter $\lambda \in [0,1]$ is defined implicitly in terms of T by

$$T = \int q_\lambda(x) \log \frac{q_0(x)}{q_1(x)} dx = L(q_\lambda \| q_1) - L(q_\lambda \| q_0). \tag{12}$$

Note that the Stein approximation (8) is a special case of the Sanov approximation (10) when $\lambda = 0$.

Similarly to the construction of the discrimination loss ΔL_N defined in (9), the Sanov approximation (10) allows us to quantify the effect of quantization on the ROC curve $\{(P_D(T), P_F(T)) : T \in \mathbb{R}\}$ by considering the difference between the pre-quantization error exponent curve

$$\{(L(q_\lambda \| q_0),\ L(q_\lambda \| q_1)) :\ \lambda \in [0,1]\}$$

to the post-quantization error exponent curve

$$\left\{ \left(L(\hat{q}_{\lambda,N} \| \overline{q}_{0,N}),\ L(\hat{q}_{\lambda,N} \| \overline{q}_{1,N}) \right) :\ \lambda \in [0,1] \right\}.$$

Here $\hat{q}_{\lambda,N}$ is the discrete tilted p.m.f. whose mass probabilities for $j = 1, \ldots, N$ are given by

$$\hat{q}_{\lambda,N,i} = \frac{\bar{q}_{0,N,i}^{1-\lambda} \cdot \bar{q}_{1,N,i}^{\lambda}}{\sum_{j=1}^N \bar{q}_{0,N,j}^{1-\lambda} \cdot \bar{q}_{1,N,j}^{\lambda}}. \tag{13}$$

Now let $P(H_1)$ and $P(H_0) = 1 - P(H_1)$ be priors on H_1 and H_0 and consider the large n approximation to average probability of error of a LRT with threshold T associated with tilt parameter λ

$$P_e = P(H_1) e^{-nL(q_\lambda \| q_1)} + P(H_0) e^{-nL(q_\lambda \| q_0)}.$$

The best achievable exponent in P_e is attained when λ equalizes the two error exponents, i.e., $L(q_\lambda \| q_0) = L(q_\lambda \| q_1)$ [29, Sec. 12.9]. By the equalizer property of the minimax Bayes test [32], this is also the value of λ which attains minimax probability of error performance over $P(H_1) \in [0, 1]$. For λ^* denoting this value of λ the common value of these two error exponents is called the *Chernoff information*.

2.3 Area-Under-Curve Detection Criterion

For a LRT based on n i.i.d. observations $\mathbf{X}$, the area under the ROC curve is defined as

$$\mathrm{AUC}_{\mathrm{ROC}}(\mathbf{X}) = \int_0^1 (1 - \beta(\alpha)) d\alpha \tag{14}$$

and has been widely used as a global measure for comparison of two different experiments. This criterion has a long history in signal detection theory, see Green and Swets [33]. Provost and Fawcett [34] call this a "whole-curve metric" to differentiate it from metrics which evaluate a single-point on the ROC curve like those discussed in Section 2.1. The area under the ROC curve has been applied to mathematical psychology [35, 36], diagnostic medical imaging [37, 38], and more recently to machine learning [39]. The area (14) is equivalent to the average power of the most powerful test under a uniform prior on the user's false alarm constraint. $\mathrm{AUC}_{\mathrm{ROC}}(\mathbf{X})$ is also equivalent to the probability of error of a Mann Whitney or Wilcoxon rank order test for randomly selected instances of H_0 vs. H_1 [40]. A large $\mathrm{AUC}_{\mathrm{ROC}}$ is better and $\mathrm{AUC}_{\mathrm{ROC}}$ is maximized by the MP-LRT. The whole-curve metric (14) is completely independent of the threshold and insensitive to the priors and/or Bayes costs which the end-user might associate with decision errors. The integral (14) can also be related to the equally likely probability of error $P_e(1/2)$ via the bounds of Barrett [41] and Shapiro [42].

For comparing quantizers for detection tasks a natural measure of quantizer distortion could be the loss in area under the ROC due to quantization

$$\Delta\mathrm{AUC}_{\mathrm{ROC}}(Q^{(n)}) = \mathrm{AUC}_{\mathrm{ROC}}(\mathbf{X}) - \mathrm{AUC}_{\mathrm{ROC}}(Q^{(n)}(\mathbf{X})), \tag{15}$$

where $\mathrm{AUC}_{\mathrm{ROC}}(\mathbf{X})$ and $\mathrm{AUC}_{\mathrm{ROC}}(Q^{(n)}(\mathbf{X}))$ are the areas under the ROC's of the LRT based on the unquantized sample $\mathbf{X}$ and the LRT based on the (product) quantized sample $Q^{(n)}(\mathbf{X}) = [Q(X^{(1)}), \ldots, Q(X^{(n)})]$, respectively. However, for purposes of asymptotic high-rate analysis of quantizer distortion it will be more

convenient to deal with the error exponent curves associated with the ROC's (see Fig. 1). As discussed in Section 2.2 these will be closely related to the ROC curves for large n. Define the following shorthand for the Sanov error exponents for pre-quantized and post-quantized data, respectively, using an N cell product quantizer:

$$\begin{aligned} L_0(\lambda) &= L(q_\lambda \| q_0), \quad L_1(\lambda) = L(q_\lambda \| q_1) \\ \hat{L}_0(\lambda) &= L(\hat{q}_{\lambda,N} \| \overline{q}_{0,N}), \quad \hat{L}_1(\lambda) = L(\hat{q}_{\lambda,N} \| \overline{q}_{1,N}). \end{aligned} \tag{16}$$

For large n the pre-quantization ROC curve is parameterized by the error exponent curve $\{(L_0(\lambda), L_1(\lambda)) : \lambda \in [0,1]\}$ which we also write in more direct form as the function $\{L_1(L_0) : \ L_0 > 0\}$. Similarly we can write the post-quantization error exponent curve as $\{\hat{L}_1(\hat{L}_0) : \ \hat{L}_0 > 0\}$. Analogously to (14) we define the area under the error exponent curve, more simply denoted as the area-under-the-curve (AUC) in this paper

$$\text{AUC} = \int_0^\infty L_1(L_0) dL_0 = \int_0^\infty L_1(\gamma) \frac{dL_0(\gamma)}{d\gamma} d\gamma. \tag{17}$$

Like AUC_{ROC}, AUC is maximized by implementing the MP-LRT. The AUC has the "threshold independent" attributes of a whole-curve metric that justify its use as a global distortion measure for quantizer detection performance.

This motivates the new mixed detection-estimation metric for i.i.d. samples and product quantizers $Q^{(n)} = Q \times \ldots \times Q$

$$J(Q) = \rho \text{MSE}(Q) + (1-\rho) \Delta\text{AUC}(Q) \tag{18}$$

where $\text{MSE}(Q)$ is the mean square distortion of the constituent quantizer Q for a single sample, and, similarly to (15), $\Delta\text{AUC}(Q)$ is the single sample loss in AUC

$$\Delta\text{AUC}(Q) = \text{AUC}(X^{(1)}) - \text{AUC}(Q(X^{(1)})), \tag{19}$$

due to implementing product quantizer $Q^{(n)}$.

3 Asymptotic High-Rate Analysis

Asymptotic high-rate quantization analysis is commonly used to obtain interesting insights into the behavior of quantizers having many small cells, which we call *small-cell* quantizers. Bennet's integral [18, 19] is central

to this analysis. The most commonly used technique of asymptotic analysis is the sequence approach. The idea behind the sequence approach is to consider a sequence of quantizers $\{Q_N\}$. Each quantizer in the sequence has N cells and an associated specific point density, specific inertial profile, specific covariation profile, and diameter function. Assuming the first three of these sequences of functions converge to functions $\zeta(x)$, $m(x)$, $M(x)$, and that the sequence of diameter functions converges to zero, the limiting behavior of the quantizer sequence can be determined.

Many analyses, such as [19], are rigorous, measure-theoretic applications of the sequence approach. A very beneficial outcome of this strategy is the obviation of many assumptions regarding convergence of the sequences. This paper takes a more elementary, non-measure-theoretic approach and, as such, more assumptions are necessary. In Section A.1.1, some assumed conditions regarding sequence convergence are listed. Additionally, we assume that the densities q_0 and q_1 have compact support with nonempty interior, are twice continuously differentiable on an open set of probability 1, and are bounded away from zero on their support. Although these assumptions preclude the use of Gaussian densities, we include several Gaussian examples and assume that the densities are truncated to satisfy these assumptions.

3.1 Log-Likelihood Ratio Quantizers

The performance of the MP-LRT is unaffected by processing of the observations as long as the processing produces a sufficient statistic. For example, if for deciding between two hypothesized source densities q_0 and q_1 there exists a sufficient statistic, e.g. the minimal sufficient statistic $\log q_1(x)/q_0(x)$, which is discrete valued then this sufficient statistic is itself equivalent to a quantizer. In Gupta [26] this was called a sufficient quantizer and its distortion is equal to zero relative to any of the previously defined detection metrics. Sufficient quantizers rarely exist in practical problems and thus it is reasonable to quantize a sufficient statistic, such as the log-likelihood ratio [8, 9]. A *log-likelihood ratio quantizer* or *LLR quantizer* Q is a scalar quantizer applied to the log-likelihood ratio defined above (5). As the MP-LRT is a threshold test, the ROC curve of the MP-LRT implemented after N-level LLR quantization has an ROC curve which meets the unquantized ROC curve at exactly N false alarm points. Thus as the ROC is continuous, as N becomes large the loss in detection performance goes to zero over the entire range of false alarm. On the other hand, for vector valued data in $\mathbf{R}^k$ the k-dimensional cells induced by the N-level LLR quantizer are

the level sets of the log-likelihood ratio which may not be convex or bounded. For example, if $k = 2$ and the sources are Gaussian, $q_0 \sim \mathcal{N}([\mu_0, \mu_0], I)$ and $q_1 \sim \mathcal{N}([\mu_1, \mu_1], I)$, then the cells of the induced quantizer will be unbounded "strips" of slope -1 as shown in Figure 2 leading to very poor MSE performance. The mixed objective (4) can be used to attain a compromise between MSE and detection distortion of a quantizer and, for $\rho \in (0, 1)$ to enforce a small-cell quantizer as the number N of cells increases. Alternatively, we can use the sequence approach to enforce the small-cell constraint.

3.2 Stein Exponent Loss

We first consider the effect of quantization on the Type II error, for arbitrarily small Type I error, via the Stein exponent in (8) which is equal to the discrimination $L(q_0\|q_1)$ between p.d.f.'s q_0, q_1. The loss in discrimination incurred by quantization with the Nth product-quantizer in the sequence $\{Q^i\}_{i=1}^{\infty}$ is defined as $\Delta L_N = L(q_0\|q_1) - L(\bar{q}_{0,N}\|\bar{q}_{1,N})$ where, as above, $\bar{q}_{0,N}$ and $\bar{q}_{1,N}$ are the p.m.f.'s of the quantized source. In Appendix A.1 we use the sequence approach to show that for a small-cell quantizer with N cells

$$\begin{aligned} \lim_{N\to+\infty} N^{2/k}\Delta L_N &= \frac{1}{2}\int \frac{q_0(x)}{\zeta(x)^{2/k}}\mathrm{tr}(F(x)M(x))dx \\ &= \frac{1}{2}\int \frac{q_0(x)\mathcal{F}(x)}{\zeta(x)^{2/k}}dx \end{aligned} \tag{20}$$

where

$$\mathcal{F}(x) = \nabla\Lambda(x)^T M(x)\nabla\Lambda(x) \tag{21}$$

which we call the *Fisher covariation profile*. We adopt this nomenclature since $\mathcal{F}(x) = \mathrm{tr}\{\mathcal{I}(x)M(x)\}$ where $\mathcal{I} = \nabla\Lambda(x)\nabla\Lambda(x)^T$ and $E_0[\mathcal{I}] = \int \mathcal{I} q_0(x)dx$ is the Fisher information matrix associated with estimating a shift parameter in the density $q_0(x)/q_1(x)$, defined with respect to the measure q_0. The expression (20) will be used in Section 4.1 to derive discrimination-optimal quantizers which minimize the loss in the Stein error exponent.

3.3 Sanov Exponent Loss

We next consider the effect of quantization on the asymptotic high-rate Type I and Type II errors via the Sanov exponents (10). The losses incurred by quantization with the Nth product-quantizer in the sequence

$\{Q^i\}_{i=1}^{\infty}$ are defined as $\Delta L_{0,N} = L(q_\lambda \| q_0) - L(\hat{q}_{\lambda,N} \| \bar{q}_{0,N})$ and $\Delta L_{1,N} = L(q_\lambda \| q_1) - L(\hat{q}_{\lambda,N} \| \bar{q}_{1,N})$ where $\hat{q}_{\lambda,N}$ is the tilted quantized p.m.f. defined in (13). In Appendix A.2, we obtain the following

$$\lim_{N\to+\infty} N^{2/k} \Delta L_{0,N} = \frac{1}{2} \int \frac{q_\lambda(x)\mathcal{F}(x)}{\zeta(x)^{2/k}} \left[\lambda^2 + \lambda(1-\lambda)(L(q_\lambda \| q_0) - \Lambda_0(x))\right] dx \tag{22}$$

$$\lim_{N\to+\infty} N^{2/k} \Delta L_{1,N} = \frac{1}{2} \int \frac{q_\lambda(x)\mathcal{F}(x)}{\zeta(x)^{2/k}} \left[(1-\lambda)^2 + \lambda(1-\lambda)(L(q_\lambda \| q_1) - \Lambda_1(x))\right] dx \tag{23}$$

where

$$\Lambda_0(x) = \log \frac{q_\lambda(x)}{q_0(x)}, \text{ and } \Lambda_1(x) = \log \frac{q_\lambda(x)}{q_1(x)}. \tag{24}$$

4 Optimal Small-Cell Quantizers

Here we use the results of the previous section to obtain asymptotic expressions for the optimal point densities minimizing loss in error exponents. Even for the classical MSE high-rate quantization problem the determination of optimal cell shapes is a difficult open problem [18, 19]. The optimal cells of high-rate MSE quantizers are conjectured to be congruent, minimum-moment-of-inertia cells [18]. For the small-cell quantization-for-detection problem the determination of optimal cell shape appears no less difficult and is also an open problem. We will, however, obtain qualitative characterizations of the optimal cell shapes using attributes of the Fisher covariation profile. We define a Sanov-optimal quantizer as a quantizer that minimizes the loss in the Type II Sanov error exponent $L(q_\lambda \| q_1)$ for some fixed value of λ, e.g., λ determined to satisfy a Type I Sanov error exponent (false alarm) constraint.

4.1 Discrimination-Optimal Quantizers

Discrimination-optimal quantizers minimize the loss in the error exponent of Stein's lemma, equal to the discrimination between the sources q_0 and q_1 after quantization. The discrimination-optimal quantizer is a Sanov-optimal quantizer designed at the operating point $\lambda = 0$. To optimize a quantizer with respect to asymptotic discrimination loss, as given by (20), it is necessary to jointly optimize two functions, namely the point density $\zeta(x)$ and the covariation profile $M(x)$. First, the discrimination-optimal point density can

be obtained using calculus of variations or Holder's inequality in a manner analogous to [19]:

$$\zeta^d(x) = \frac{[q_0(x)\mathcal{F}(x)]^{\frac{k}{k+2}}}{\int [q_0(y)\mathcal{F}(y)]^{\frac{k}{k+2}} dy}. \tag{25}$$

The discrimination loss with the optimal point density is then

$$\Delta L_N \approx \frac{1}{2N^{2/k}} \left(\int [q_0(x)\mathcal{F}(x)]^{\frac{k}{k+2}} dx \right)^{\frac{k+2}{k}}. \tag{26}$$

This depends on the covariation profile $M(x)$ through $\mathcal{F}$ defined in (21).

If the quantizer's cells are congruent, the covariation profile $M(x)$ is constant independent of x. If in addition the cells have minimum moment of inertia, $M(x) = m_k^* I$ and the point density given by equation (25) becomes

$$\zeta^d(x) = \frac{[q_0(x)\|\nabla\Lambda(x)\|^2]^{\frac{k}{k+2}}}{\int [q_0(y)\|\nabla\Lambda(y)\|^2]^{\frac{k}{k+2}} dy}.$$

We call the function $\|\nabla\Lambda(x)\|^2$ the *discriminability function* which equals zero when the hypotheses have densities with identical zero-th and first order derivatives.

4.1.1 Ellipsoidal Cells

Ellipsoidal cells can not cover $\mathbb{R}^k$ without overlap and thus can not partition $\mathbb{R}^k$. However, if k is large, in analogy with the spherical cell approximation of [20] for lattice quantizers, it is plausible that a quantizer's cells can be close to ellipsoidal. Studying ellipsoidal quantizer cells yields important insights and can provide a lower bound on high rate distortion [21]. Accordingly, assume that in the neighborhood of some point x_i the cell is $S_i = \{x : (x - x_i)^T R(x - x_i) \leq c$. Then $M = m_k^* |R|^{1/k} R^{-1}$ has an eigendecomposition

$$M = \sum_{i=1}^{k} \frac{1}{\phi_i} v_i v_i^T,$$

where $\{\phi_1, \ldots, \phi_k\}$ are the positive eigenvalues of $M^{-1} = 1/m_k^* |R|^{-1/k} R$ corresponding to its orthonormal eigenvectors $\{v_1, \ldots, v_k\}$. Thus the Fisher covariation profile is

$$\mathcal{F} = \sum_{i=1}^{k} \frac{1}{\phi_i} \left(\nabla\Lambda^T v_i \right)^2.$$

Let $\phi_{\max}$ be a finite upper bound on the eigenvalues of M^{-1}. This upper bound restricts the minimum diameter of the cell to be positive, i.e., nondegenerate. The minimum of $\mathcal{F}$ over matrices R satisfying

$\max_i \phi_i \leq \phi_{\max}$ is achieved when: M has $1/\phi_{\max}$ as its minimum eigenvalue; and the corresponding minimizing eigenvector of M is $v_{\max} = \nabla\Lambda/\|\nabla\Lambda\|$, which is parallel to $\nabla\Lambda$. In this case the optimal Fisher covariation profile is

$$\mathcal{F} = \frac{1}{\phi_{\max}} \|\nabla\Lambda\|^2$$

Thus we conclude that if a cell centered at x_i is an eccentric ellipsoid which is nondegenerate, then its minor axis should be aligned along the direction of the normal vector to the log likelihood ratio surface. For large N, we see that this implies that any eccentric ellipsoidal cells should be aligned with the level sets of the log-likelihood ratio.

4.2 Chernoff-Optimal Quantization

The Chernoff-optimal quantizer is a Sanov-optimal quantizer designed at an operating point $\lambda = \lambda^*$ which minimizes the loss in Chernoff information due to quantization. Unfortunately, the asymptotic loss in Chernoff information can be very difficult to determine since the pre-quantization equalization condition $L(q_\lambda\|q_0) = L(q_\lambda\|q_1)$ and the post-quantization equalization condition $L(\hat{q}_\lambda\|\overline{q}_0) = L(\hat{q}_\lambda\|\overline{q}_1)$ are seldom satisfied for identical equalizer solution $\lambda = \lambda^*$. See Fig. 3 for illustration. Therefore, asymptotic Chernoff loss involves a complicated interaction between the pre-quantization and the post-quantization equalizer λ solutions. An exception which permits simple determination of the asymptotic Chernoff information loss occurs in the case where these equalizer solutions are identical.

If it so happens that the two equalizing λ are the same then the asymptotic expression (22) is valid, which we rewrite as follows

$$\Delta L_{0,N}(\lambda) \approx \frac{\lambda^2}{2N^{2/k}} \int \frac{q_\lambda(x)\mathcal{F}(x)}{\zeta(x)^{2/k}} dx + \frac{\lambda(1-\lambda)}{2N^{2/k}} D_{\lambda,0}$$

and

$$\Delta L_{1,N}(\lambda) \approx \frac{(1-\lambda)^2}{2N^{2/k}} \int \frac{q_\lambda(x)\mathcal{F}(x)}{\zeta(x)^{2/k}} dx + \frac{\lambda(1-\lambda)}{2N^{2/k}} D_{\lambda,1}$$

where for $i = 0, 1$ and Λ_i as defined in (24)

$$D_{\lambda,i} \stackrel{\text{def}}{=} \int \frac{q_\lambda(x)\mathcal{F}(x)}{\zeta(x)^{2/k}} \left(L(q_\lambda\|q_i) - \Lambda_i(x)\right) dx.$$

We denote the λ dependency explicitly by writing $\Delta L_{0,N}(\lambda)$ and $\Delta L_{1,N}(\lambda)$. The loss in Chernoff information is equal to $\Delta L_{1,N}(\lambda^*)$ where $\lambda = \lambda^*$ is the solution of $\Delta L_{0,N}(\lambda) = \Delta L_{1,N}(\lambda)$. Solving for λ^* can rarely be performed in closed form but may be accomplished using numerical root finding techniques on the difference $\Delta L_{0,N}(\lambda) - \Delta L_{1,N}(\lambda)$ which is equivalent to finding λ such that

$$-(1-2\lambda)\int \frac{q_\lambda(x)\mathcal{F}(x)}{\zeta(x)^{2/k}}dx + \lambda(1-\lambda)(D_{\lambda,0} - D_{\lambda,1}) = 0. \tag{27}$$

When $D_{\lambda,0} = D_{\lambda,1}$ then it is obvious that $\lambda = \lambda^* = 1/2$ is the equalization solution, and

$$\Delta L_{0,N}(\lambda^*) = \Delta L_{1,N}(\lambda^*) = \frac{1}{8N^{2/k}}\int \frac{q_{1/2}(x)\mathcal{F}(x)}{\zeta(x)^{2/k}}dx + \frac{D_{1/2,0}}{8N^{2/k}}.$$

A strategy for finding solutions to the asymptotic Chernoff information is to first find the pre-quantized equalizing solution λ^* which satisfies $L(q_\lambda\|q_0)(\lambda^*) = L(q_\lambda\|q_1)(\lambda^*)$ and then check if λ^* is also a solution to (27). If so then λ^* is a solution to $\Delta L_{1,N}(\lambda) = \Delta L_{0,N}(\lambda)$ which, as $L(q_{\lambda^*}\|q_0) = L(q_{\lambda^*}\|q_1)$, would imply that $L(\hat{q}_{\lambda^*}\|\overline{q}_0) = L(\hat{q}_{\lambda^*}\|\overline{q}_1)$, as required. We will follow this strategy in the Gaussian example considered below.

4.3 AUC Optimal Quantization

An alternative to the difficult Chernoff-optimal quantizer is the simpler AUC-optimal quantizer which minimizes the loss of area under the Sanov error-exponent curve.

Let $L_i(\lambda)$ and $\hat{L}_i(\lambda)$ be as defined in (16). Define $\hat{A}$ the area under the post-quantized error-exponent curve $\hat{L}_1(\hat{L}_0)$. Then

$$\hat{A} = \int_0^1 \hat{L}_1(\lambda)\frac{d}{d\lambda}\hat{L}_0(\lambda)d\lambda.$$

Define

$$\begin{aligned} f_0(x,\lambda) &= q_\lambda(x)\left[\lambda^2 + \lambda(1-\lambda)(L_0(\lambda) - \Lambda_0(x,\lambda))\right] \\ f_1(x,\lambda) &= q_\lambda(x)\left[(1-\lambda)^2 + \lambda(1-\lambda)(L_1(\lambda) - \Lambda_1(x,\lambda))\right]. \end{aligned} \tag{28}$$

Then

$$\begin{aligned} \hat{L}_0(\lambda) &= L_0(\lambda) - \frac{1}{2N^{2/k}}\int \frac{\mathcal{F}(x)}{\zeta(x)^{2/k}}f_0(x,\lambda)dx \\ \hat{L}_1(\lambda) &= L_1(\lambda) - \frac{1}{2N^{2/k}}\int \frac{\mathcal{F}(x)}{\zeta(x)^{2/k}}f_1(x,\lambda)dx \end{aligned}$$

and

$$\frac{d}{d\lambda}\hat{L}_0(\lambda) = \frac{d}{d\lambda}L_0(\lambda) - \frac{1}{2N^{2/k}}\int \frac{\mathcal{F}(x)}{\zeta(x)^{2/k}} \cdot \frac{\partial}{\partial\lambda} f_0(x,\lambda)dx.$$

Thus

$$\begin{aligned}\hat{L}_1(\lambda)\frac{d}{d\lambda}\hat{L}_0(\lambda) &= L_1(\lambda)\frac{d}{d\lambda}L_0(\lambda) - \\ &\frac{1}{2N^{2/k}}\int \frac{\mathcal{F}(x)}{\zeta(x)^{2/k}}\left[L_1(\lambda)\frac{\partial}{\partial\lambda}f_0(x,\lambda) + f_1(x,\lambda)\frac{d}{d\lambda}L_0(\lambda)\right]dx + \\ &o\left(\frac{1}{N^{2/k}}\right).\end{aligned}$$

The area $\hat{A}$ is thus

$$\hat{A} = A - \frac{1}{2N^{2/k}}\int \frac{\mathcal{F}(x)\eta(x)}{\zeta(x)^{2/k}}dx + o\left(\frac{1}{N^{2/k}}\right)$$

where

$$A = \int_0^1 L_1(\lambda)\frac{d}{d\lambda}L_0(\lambda)d\lambda$$

is the area under the pre-quantized error exponent curve $L_1(L_0)$ and

$$\eta(x) = \int_0^1 \left[L_1(\lambda)\frac{\partial}{\partial\lambda}f_0(x,\lambda) + f_1(x,\lambda)\frac{d}{d\lambda}L_0(\lambda)\right]d\lambda. \tag{29}$$

Finally, we obtain

$$\lim_{N\to+\infty} N^{2/k}(A - \hat{A}) = \frac{1}{2}\int \frac{\mathcal{F}(x)\eta(x)}{\zeta(x)^{2/k}}dx. \tag{30}$$

Note the resemblance of (30) to (20). Essentially, the source density $q_0(x)$ in (20) has simply been replaced by $\eta(x)$ in (30). Although $\eta(x)$ may not have a closed form expression the integral expression (29) can easily be evaluated numerically.

Analogous to the discrimination-optimal point density derived above, we can derive the AUC-optimal point density

$$\zeta^o(x) = \frac{[\mathcal{F}(x)\eta(x)]^{\frac{k}{k+2}}}{\int[\mathcal{F}(y)\eta(y)]^{\frac{k}{k+2}}dy} \tag{31}$$

and the resulting loss in area under the $L_1(L_0)$ curve, with the AUC-optimal point density is

$$\Delta A_N \approx \frac{1}{2N^{2/k}}\left(\int[\mathcal{F}(x)\eta(x)]^{\frac{k}{k+2}}dx\right)^{\frac{k+2}{k}}. \tag{32}$$

The congruent-cell quantizer is constructed analogously to Section 4.1 and is completely characterized by the optimal point density (31) which, in the case of minimum-moment-of-inertia cells, is given by

$$\zeta^o(x) = \frac{[\eta(x)\|\nabla\Lambda(x)\|^2]^{\frac{k}{k+2}}}{\int[\eta(y)\|\nabla\Lambda(y)\|^2]^{\frac{k}{k+2}}dy}.$$

For ellipsoidal cells the conclusions of the previous subsection equally apply to the AUC-optimal quantizer.

4.4 Optimal Quantizers for Mixed Objective Functions

It is simple to extend the high-rate analysis to mixed criteria such as (18). In particular, equation (30) indicates that the loss in AUC due to quantization by a sequence of N-point, small-cell quantizers converges to zero at the rate of $N^{-2/k}$. This is the same rate obtained by Na and Neuhoff [19] for the MSE under the sequential approach. Specifically, for an i.i.d. sample of k-dimensional vectors $\{X^{(i)}\}_{i=1}^n$ with marginal p.d.f. $q(x)$:

$$N^{2/k}\text{MSE} = \int \frac{q(x)}{\zeta(x)^{2/k}}dx \tag{33}$$

Let MSE_0 and MSE_1 denote the conditional MSE of the quantizer given $q = q_0$ and $q = q_1$, respectively, for a single sample ($n = 1$). Letting $(1-p), p$ be priors on hypotheses H_0, H_1 the average MSE is $\text{MSE} = \text{MSE}_0(1-p)+\text{MSE}_1 p$ and, using the results of the previous section the mixed measure (18), with appropriate normalization, satisfies

$$\lim_{N\to\infty}\{N^{2/k}J(Q)\} \quad = \quad \int \frac{\rho q(x) + (1-\rho)p(x)}{\zeta(x)^{2/k}}dx \tag{34}$$

where $q = q_0(1-p) + q_1 p$, ζ is the point density, and $p(x)$ is the density

$$p(x) = \frac{\eta(x)\mathcal{F}(x)}{\int \eta(y)\mathcal{F}(y)dy}.$$

The optimal point density for the mixed objective is simply

$$\zeta^J(x) = \frac{[\rho q(x) + (1-\rho)p(x)]^{\frac{k}{k+2}}}{\int[\rho q(y) + (1-\rho)p(y)]^{\frac{k}{k+2}}dy}, \tag{35}$$

which varies from the AUC-optimal point density for $\rho = 0$ to the MSE-optimal point density for $\rho = 1$.

5 Illustrative Examples

In this section, we demonstrate the concepts and procedures described in the previous section through some illustrative examples. The careful reader will note that the Gaussian densities considered for these examples have unbounded support and therefore do not satisfy the conditions of Appendix A.1.1. However, excepting overload distortion of order $O(\epsilon)$, our results apply to approximately Gaussian densities obtained by truncating the Gaussian's support to a large bounded region containing $1-\epsilon$ of the density's mass.

5.1 Scalar Gaussian Sources

As a first example, consider scalar, unit-variance Gaussian sources with different means $q_0 = \mathcal{N}(\mu_0, 1)$ and $q_1 = \mathcal{N}(\mu_1, 1)$. Assume the priors $P(H_1)$ and $1 - P(H_1)$ on H_1 and H_0 are equal to $1/2$. The point density minimizing the asymptotic MSE loss (33) is given by the formula (35) with the substitutions $q = (q_0 + q_1)/2$ and $\rho = 1$. The log-likelihood ratio is $\Lambda(x) = -\frac{1}{2}(\mu_0^2 - \mu_1^2) + (\mu_0 - \mu_1)x$ and the Fisher covariation profile is constant. The discrimination-optimal and AUC-optimal point densities are given by equations (25) and (31), respectively. From these equations, we see that the discrimination-optimal quantizer should concentrate its points according to density q_0 while the AUC-optimal quantizer concentrates its points according to the density $\eta(x)$.

Figure 4 shows the sources q_0 and q_1 with $\mu_0 = -2$ and $\mu_1 = 2$ along with the function $\eta(x)$ Note that $\eta(x)$ takes a maximum at $x = 0$ where the two source densities cross. In Figure 5, the AUC-optimal, discrimination-optimal, and MSE-optimal point densities are plotted. As the priors are equal, the MSE-optimal point density has peaks at the maxima of the source densities. With the constant discriminability function, the AUC-optimal and discrimination-optimal point densities are maximized at points where $\eta(x)$ and $q_0(x)$ are maximized, respectively.

In Figures 6, 7, and 8, the performances of scalar quantizers with the various optimal point densities are compared. The quantizers were obtained using the LBG algorithm, also known as the generalized Lloyd algorithm [1, 43, 44], applied to the relevant point densities. The essence of this approach is as follows: The MSE-optimal point density for source density $q(x)$ is $\zeta^{mse}(x) = T(q) = q(x)^{\frac{k}{k+2}} / \int q(y)^{\frac{k}{k+2}} dy$. Note that T is invertible and $T^{-1}(\zeta) = \zeta(x)^{\frac{k+2}{k}} / \int \zeta(y)^{\frac{k+2}{k}} dy$. Since the quantizer produced by the generalized

Lloyd algorithm is MSE-optimal, its point density is $T(q)$. Now, let $\tilde{\zeta}$ be an arbitrary point density. If the generalized Lloyd algorithm is run with source density $T^{-1}(\tilde{\zeta})$, the resulting quantizer will have point density $T(T^{-1}(\tilde{\zeta})) = \tilde{\zeta}$. For large N, this quantizer will have k-dimensional minimum-moment-of-inertia polytope cells. For further details, see [26]. Figure 6 shows the error exponent curves with and without quantization for the AUC-optimal, discrimination-optimal, and MSE-optimal quantizers with $N = 8$ cells. As expected, the AUC-optimal quantizer performs the best in terms of the area underneath the curve criterion. It is interesting to note that the error-exponent curve of the discrimination-optimal quantizer is quite poor. This quantizer minimizes the loss in the Type II error exponent $L(\bar{q}_0 \| \bar{q}_1)$, and is equivalent to a Sanov-optimal quantizer designed for the operating point $\lambda = 0$.

Figure 7 shows the ROC curves of the MP LRT with $n = 2$ i.i.d. observations with and without quantization by various optimal quantizers with $N = 16$ cells. Note that the formulas (10) are accurate only as the number of observations n becomes large and therefore the AUC-optimal quantizer may or may not actually yield an optimum ROC curve. However, for this example we see that the AUC-optimal quantizer does indeed have the best performance. Finally, in Figure 8 the estimation performance of the three quantizers with $N = 16$ cells is compared. The reconstruction MSE of each quantizer is plotted versus the prior probability $P_0 \stackrel{\text{def}}{=} P(H_0)$. The MSE-optimal quantizer is assumed to have knowledge of the priors. As expected, the MSE-optimal quantizer yields the minimum reconstruction MSE of the three considered quantizers. Note the extremely poor performance of the discrimination-optimal quantizer for $P_0 < 1$. Recall that the discrimination-optimal quantizer concentrates its points mostly underneath density q_0. For $P_0 = 1$, the discrimination-optimal and MSE-optimal quantizers are the same. For $P_0 < 1$, however, the discrimination-optimal quantizer differs significantly from the MSE-optimal quantizer. See for example Figure 5, which shows the two point densities for the case $P_0 = 1/2$.

For equal-variance Gaussian sources the Chernoff-optimal quantizer can easily be obtained using the approach outlined in Section 4.2. We must show that the solution λ^* to the post-quantized equalization condition $L(\hat{q}_\lambda \| \overline{q}_1) = L(\hat{q}_\lambda \| \overline{q}_0)$, or equivalently the asymptotic version (27) of this condition, also satisfies the pre-quantized equalization condition $L(q_\lambda \| q_1) = L(q_\lambda \| q_0)$. First note that the pre-quantized tilted density is of Gaussian form: $q_\lambda \sim \mathcal{N}(\mu_\lambda, 1)$ where $\mu_\lambda = (1-\lambda)\mu_0 + \lambda\mu_1$. It is therefore easily verified [26] that the the value of λ which solves the pre-quantized equalization condition is $\lambda = 1/2$. Furthermore, the

log-likelihood ratios $\Lambda_0(x)$ and $\Lambda_1(x)$ given by (24) are linear in x and the Fisher covariation profile $\mathcal{F}$ is constant. Thus equation (27) is also solved for $\lambda = 1/2$.

For $\mu_0 = 0$ and $\mu_1 = 8$, Figure 9 shows the optimal point density for Chernoff information ζ^{Ch}, along with the AUC-optimal point density ζ^o. Both point densities are maximized at the point $x = 4$, where the two source densities cross. The point density of the Chernoff-optimal quantizer is more concentrated about this point, however. In Figure 10, the pre and post quantized error exponent curves $L_1(L_0)$ are plotted for both quantizers with $N = 8$ cells. Note that the intersection of each of these curves with the diagonal line gives the corresponding Chernoff information. The Chernoff-optimal curve lies above the AUC-optimal curve in a region close to the intersection with the unit-slope line, thus yielding greater Chernoff information. On the other hand, the area under the AUC-optimal curve is greater, as expected. Note that the Chernoff-optimal quantizer is optimized specifically for $\lambda = \lambda^* = 1/2$, and not for any other value of λ.

Finally, we remark that this analysis can be extended to obtain Chernoff-information-optimal vector quantizers for vector Gaussian sources with identity covariance matrices. For these cases, we must restrict attention to quantizers with point densities and covariation profiles that are symmetric about $\underline{\mu}_\lambda$, the mean of the tilted density. For example, restricted polar quantizers [45] and some shape-gain quantizers [1] satisfy this constraint.

5.2 Two-Dimensional Uncorrelated Gaussian Sources

Next, consider two-dimensional Gaussian sources with identity covariance matrices: $q_0 = \mathcal{N}\left(\underline{\mu}_0, I\right)$ and $q_1 = \mathcal{N}\left(\underline{\mu}_1, I\right)$ where $\underline{\mu}_0 = [\mu_0, \mu_0]$ and $\underline{\mu}_1 = [\mu_1, \mu_1]$. As in the scalar Gaussian example, the discriminability function is constant for two-dimensional Gaussian sources with identity covariance matrices. The discrimination-optimal and AUC-optimal point densities are given by equations (25) and (31), respectively. In addition to the vector quantizers considered in the previous 1D example, we investigated a 64 cell optimal scalar LLR quantizer under the AUC criterion, which we call the AUC-optimal LLR scalar quantizer, and an AUC-optimal mixed vector quantizer implemented by applying the LBG algorithm to the point density (35) with $q = (q_0 + q_1)/2$ and $\rho = 1/2$.

Figure 11 shows contours of the two source densities for $\mu_0 = -2$ and $\mu_1 = 2$. In Figures 12, 13, and 14, congruent-cell VQ's optimal for AUC, discrimination, and estimation, with $N = 64$ cells are shown. These

quantizers were again obtained using the LBG algorithm [26]. Similar to the one-dimensional case, the AUC-optimal quantizer's cells are concentrated between the source densities, the discrimination-optimal quantizer concentrates its cells underneath density q_0, and the MSE-optimal quantizer's cells are dense underneath the peaks of both densities.

The hypothesis testing performance of the 64-cell quantizers in Figures 12, 13, and 14 is compared in Figure 16. Similar to the scalar Gaussian example, the AUC-optimal quantizer performs the best, while the discrimination-optimal quantizer yields the largest discrimination between quantized sources $L(\bar{q}_0 \| \bar{q}_1)$, but performs poorly on average.

Figure 15 shows the optimal quantizer cells for the mixed estimation-detection objective function (18). This quantizer concentrates its points between the source density peaks as does the AUC-optimal quantizer in Figure 12, as well as underneath the peaks as does the MSE-optimal quantizer in Figure 14.

Figure 17 is a blowup of Figure 16 which shows the dominance in detection performance of: 1) the AUC-optimal LLR scalar quantizer, 2) the AUC-optimal vector quantizer, 3) the AUC-optimal mixed VQ, 4) the MSE-optimal vector quantizer, and 5) the discrimination-optimal vector quantizer, in that respective order. As expected the AUC-optimal LLR quantizer outperforms the rest in terms of detection performance, virtually attaining optimal unquantized performance in the blow-up region of L_o shown. The gap shown between the AUC-optimal LLR quantizer and the AUC-optimal vector quantizer is the small price paid by the AUC-optimal vector quantizer in order to attain improved MSE performance (not shown).

6 Conclusion

We have developed asymptotic theory for quantization for various measures of detection performance using the Sanov error exponents of binary hypothesis testing. This theory applies for a large number of observations, n, and a large number of quantization cells N. Under a small-cell assumption the asymptotic large N loss in the error exponent, called the discrimination, resembles Bennet's integral formula for the reconstruction MSE. Optimal small-cell quantizer point densities which minimize the loss in various functions of the Sanov exponents, including the discrimination, the Chernoff-information, and the area under the error exponent curve (AUC) were derived. Numerical examples of various optimal quantizers have been presented for

several types of scalar and two-dimensional sources. The Fisher covariation profile has significant influence on the placement of codebook points in quantizers optimal for binary hypothesis testing.

We think there are many worthwhile avenues for extending these results and we mention a few here. While our results are all limited to fixed rate quantizers, it is likely that variable rate high resolution analysis can also be performed on detectibility distortion criteria in a manner similar to Li *etal*'s analysis of the perceptual distortion measure of Gardner and Rao [22]. Our derivations of our high-rate expressions follow the classical but relatively informal Taylor series approach and it would be useful to refine and extend these results via more rigorous measure theoretic approaches, as described by Gray and Neuhoff in [2]. Finally, further investigation of the loss in detection performance due the restrictive small-cell condition is worthwhile. This can be accomplished by comparing the high rate performance of quantizers that operate directly on the log likelihood ratio (LLR) to the high rate small cell quantizers derived in this paper.

Appendix A: Derivation of Asymptotic Discrimination Losses

A.1 Asymptotic Loss in Discrimination Between Two Sources

To derive the asymptotic loss in discrimination (20) between q_0 and q_1, we follow the "sequence approach" used in [46, 47, 19]. Some assumptions not required in [46, 47, 19] are necessary for this analysis. Consider a sequence of quantizers $Q_N = (\mathcal{S}_N, \mathcal{C}_N)$ where the Nth quantizer contains the N cells $\mathcal{S}_N = \{S_{N,1}, \ldots, S_{N,N}\}$ and the N codebook points $\mathcal{C}_N = \{x_{N,1}, \ldots, x_{N,N}\}$.

The discrimination before quantization can be written in terms of the cells of the Nth quantizer:

$$L \triangleq L(q_0 \| q_1) = \sum_{i=1}^{N} \int_{S_{N,i}} q_0(y)\Lambda(y)dy.$$

The discrimination after quantization by the Nth quantizer can be written as

$$\hat{L}_N \triangleq L(\bar{q}_{0,N} \| \bar{q}_{1,N}) = \sum_{i=1}^{N} \bar{q}_{0,N,i}\bar{\Lambda}_{N,i}.$$

Since our goal is to maximize the discrimination after quantization, we will refer to the loss in discrimination as distortion. It is well known that discrimination can not increase with processing (i.e. quantization). Thus, the distortion is nonnegative. The distortion resulting from the Nth quantizer is thus

$$\Delta L_N \triangleq L - \hat{L}_N = \sum_{i=1}^{N} \int_{S_{N,i}} q_0(y)\Lambda(y)dy - \bar{q}_{0,N,i}\bar{\Lambda}_{N,i}. \tag{A.1}$$

Note that (A.1) is independent of the codebook $\mathcal{C}_N$. Therefore, we lose no generality by assuming that the codebook points are the centroids of their cells. That is, for each N

$$x_{N,i} = \frac{\int_{S_{N,i}} y dy}{V_{N,i}},\ i = 1, \ldots, N \tag{A.2}$$

where $V_{N,i}$ is the volume of the ith cell in the Nth quantizer. Note that (A.2) implies

$$\int_{S_{N,i}} (y - x_{N,i}) dy = 0,\ i = 1, \ldots, N.$$

A.1.1 Sequence Definitions

We define a few more sequences that will be necessary in analyzing the asymptotic behavior of the quantizer sequence.

1. The sequence of diameter functions is $d_N(x)$.

2. The sequence of specific inertial profile functions is $m_N(x)$.

3. The sequence of specific covariation profile functions is $M_N(x)$. We will write $M_{N,i} = M_N(x)$ for $x \in S_{N,i}$.

4. The sequence of specific point density functions is $\zeta_N(x) = \zeta_{N,i} = 1/(NV_{N,i})$ for $x \in S_{N,i}$.

The essence of the sequence approach are the following conditions: 1) $d_N(x)$ converges uniformly to zero; 2) $m_N(x)$ converges uniformly to a function $m(x)$, the specific inertial profile, that is uniformly bounded by m_B; 3) $M_N(x)$ converges uniformly to a full-rank matrix function $M(x)$, the covariation profile.

To facilitate the analysis, we define some simplifying notation. The density functions evaluated at codebook point $x_{N,i}$ will be denoted

$$\begin{aligned} q_{0,N,i} &= q_0(x_{N,i}) \\ q_{1,N,i} &= q_1(x_{N,i}). \end{aligned}$$

Similarly, the gradients and Hessians of q_0 and q_1 evaluated at $x_{N,i}$ will be denoted

$$\begin{aligned}
\nabla_{0,N,i} &= \nabla q_0(x_{N,i}) \\
\nabla_{1,N,i} &= \nabla q_1(x_{N,i}) \\
\nabla^2_{0,N,i} &= \nabla^2 q_0(x_{N,i}) \\
\nabla^2_{1,N,i} &= \nabla^2 q_1(x_{N,i})
\end{aligned}$$

and the log-likelihood ratio evaluated at $x_{N,i}$ is

$$\Lambda_{N,i} = \Lambda(x_{N,i}).$$

The following matrix functions will be useful in our analysis. The "Fisher" matrix function is defined to be the outer product of the log-likelihood ratio gradient:

$$F(x) = \nabla\Lambda(x)\nabla\Lambda(x)^T$$

and the matrix function $G(x)$ is

$$G(x) = \frac{\nabla^2 q_0(x)}{q_0(x)} - \frac{\nabla^2 q_1(x)}{q_1(x)}. \tag{A.3}$$

In keeping with the convention set forth above, we define

$$\begin{aligned}
F_{N,i} &= F(x_{N,i}) \\
G_{N,i} &= G(x_{N,i}).
\end{aligned} \tag{A.4}$$

A.1.2 Taylor Expansions

For all N, we can expand the function $q_0(x)$ in a Taylor series about the codebook points of quantizer Q_N. Therefore, for all N we can write

$$\begin{aligned}
q_0(x) &= q_{0,N,i} + \nabla^T_{0,N,i}(x - x_{N,i}) + \frac{1}{2}(x - x_{N,i})^T\nabla^2_{0,N,i}(x - x_{N,i}) \\
&\quad + o(\|x - x_{N,i}\|^2),\ \forall x \in S_{N,i}.
\end{aligned} \tag{A.5}$$

A similar expansion can be done for $q_1(x)$ and $\Lambda(x)$ as shown below:

$$\begin{aligned}
\Lambda(x) &= \Lambda_{N,i} + \nabla\Lambda^T_{N,i}(x - x_{N,i}) + \frac{1}{2}(x - x_{N,i})^T\nabla^2\Lambda_{N,i}(x - x_{N,i}) \\
&\quad + o(\|x - x_{N,i}\|^2),\ \forall x \in S_{N,i}.
\end{aligned} \tag{A.6}$$

The "o" terms in (A.5) and (A.6) are explained as follows. From the definition of the diameter function, we have $\|x - Q_N(x)\| \leq d_N(x)$ for all N and by assimption we have $\|x - Q_N(x)\| \to 0$ uniformly. Therefore, given $\epsilon > 0$ there is an integer N_0 such that for all $N \geq N_0$ and for all $x \in S_{N,i}$

$$\frac{o(\|x - x_{N,i}\|^2)}{\|x - x_{N,i}\|^2} < \epsilon.$$

A.1.3 Single-Cell Distortion

The distortion of the Nth quantizer given by (A.1) is a sum over the N quantizer cells of the quantity $\int_{S_{N,i}} q_0(y)\Lambda(y)dy - \bar{q}_{0,N,i}\bar{\Lambda}_{N,i}$. We call this term the single-cell distortion of cell $S_{N,i}$. The bulk of the analysis required to determine the distortion involves studying the single-cell distortion, which we do in this section.

Using (A.5) and (A.6) along with the centroid condition we have

$$\begin{aligned}\int_{S_{N,i}} q_0(y)\Lambda(y)dy &= q_{0,N,i}\Lambda_{N,i}V_{N,i} + \int_{S_{N,i}} (y - x_{N,i})^T A_{N,i}(y - x_{N,i})dy \\ &\quad + \int_{S_{N,i}} o(\|y - x_{N,i}\|^2)dy \end{aligned} \tag{A.7}$$

where

$$A_{N,i} = \frac{1}{2}\left[\Lambda_{N,i}\nabla^2_{0,N,i} + q_{0,N,i}\nabla^2\Lambda_{N,i} + \nabla_{0,N,i}\nabla\Lambda^T_{N,i} + \nabla\Lambda_{N,i}\nabla^T_{0,N,i}\right]. \tag{A.8}$$

The last two terms in (A.8) arise due to the fact that the matrix in a quadratic form may be transposed without affecting the result [48]. After some algebra, (A.8) can be written

$$A_{N,i} = \frac{1}{2}\left[\Lambda_{N,i}\nabla^2_{0,N,i} + q_{0,N,i}(F_{N,i} + G_{N,i})\right] \tag{A.9}$$

where $F_{N,i}$ and $G_{N,i}$ are given in (A.4).

To simplify (A.7), we first focus on the last term. For $\epsilon > 0$ there is an integer N_0 such that for all $N \geq N_0$, the following two conditions hold:

$$\frac{o(\|y - x_{N,i}\|^2)}{\|y - x_{N,i}\|^2} \leq \frac{\epsilon}{2m_B},\ \forall y \in S_{N,i}$$

and

$$\begin{aligned}|m_N(y) - m(y)| &\leq m_B, \\ \Rightarrow m_N(y) &\leq m(y) + m_B \leq 2m_B,\ \forall y \in S_{N,i}.\end{aligned}$$

Therefore, for all $N \geq N_0$,

$$\begin{aligned}\left|\int_{S_{N,i}} o(\|y - x_{N,i}\|^2)dy\right| &\leq \int_{S_{N,i}} \left|o(\|y - x_{N,i}\|^2)\right| dy \\ &\leq \int_{S_{N,i}} \frac{\epsilon}{2m_B}\|y - x_{N,i}\|^2 dy \\ &= \frac{\epsilon}{2m_B} \cdot m_N(x) V_{N,i}^{1+2/k},\ \forall x \in S_{N,i} \\ &\leq \epsilon \cdot V_{N,i}^{1+2/k}.\end{aligned}$$

Therefore, the sequence

$$\frac{\left|\int_{S_{N,i}} o(\|y - x_{N,i}\|^2)dy\right|}{V_{N,i}^{1+2/k}}$$

converges to zero and we will thus write

$$\int_{S_{N,i}} o(\|y - x_{N,i}\|^2)dy = o\left(V_{N,i}^{1+2/k}\right).$$

Next, we rewrite the second term on the right-hand side of (A.7) as

$$\int_{S_{N,i}} (y - x_{N,i})^T A_{N,i}(y - x_{N,i})dy = \mathrm{tr}(A_{N,i}M_{N,i})V_{N,i}^{1+2/k}.$$

Therefore (A.7) becomes

$$\int_{S_{N,i}} q_0(y)\Lambda(y)dy = q_{0,N,i}\Lambda_{N,i}V_{N,i} + \mathrm{tr}(A_{N,i}M_{N,i})V_{N,i}^{1+2/k} + o\left(V_{N,i}^{1+2/k}\right). \tag{A.10}$$

We now turn our attention to the term $\bar{q}_{0,N,i}\bar{\Lambda}_{N,i}$ found in (A.1). From (A.5) and (A.6) we have

$$\bar{q}_{0,N,i}\bar{\Lambda}_{N,i} = q_{0,N,i}\bar{\Lambda}_{N,i}V_{N,i} + \mathrm{tr}(\hat{A}_{N,i}M_{N,i})V_{N,i}^{1+2/k} + o\left(V_{N,i}^{1+2/k}\right) \tag{A.11}$$

where

$$\hat{A}_{N,i} = \frac{1}{2}\bar{\Lambda}_{N,i}\nabla_{0,N,i}^2. \tag{A.12}$$

Combining (A.10) and (A.11) yields

$$\begin{aligned}\int_{S_{N,i}} q_0(y)\Lambda(y)dy - \bar{q}_{0,N,i}\bar{\Lambda}_{N,i} &= q_{0,N,i}\left(\Lambda_{N,i} - \bar{\Lambda}_{N,i}\right)V_{N,i} + \\ &\quad \frac{1}{2}\left(\Lambda_{N,i} - \bar{\Lambda}_{N,i}\right)\mathrm{tr}\left(\nabla_{0,N,i}^2 M_{N,i}\right)V_{N,i}^{1+2/k} + \\ &\quad \frac{1}{2}q_{0,N,i}\ \mathrm{tr}\left([F_{N,i} + G_{N,i}]M_{N,i}\right)V_{N,i}^{1+2/k} + \\ &\quad o\left(V_{N,i}^{1+2/k}\right).\end{aligned} \tag{A.13}$$

From the definitions of $\Lambda_{N,i}$ and $\bar{\Lambda}_{N,i}$ we have

$$\Lambda_{N,i} - \bar{\Lambda}_{N,i} = \log\left(\frac{q_{0,N,i} \cdot \bar{q}_{1,N,i}}{q_{1,N,i} \cdot \bar{q}_{0,N,i}}\right).$$

Using the Taylor expansion

$$\log a = (a-1) - \frac{1}{2}(a-1)^2 + o(|a-1|^2)$$

we have

$$\Lambda_{N,i} - \bar{\Lambda}_{N,i} = (l-1) - \frac{1}{2}(l-1)^2 + o(|l-1|^2)$$

where

$$l = \frac{q_{0,N,i} \cdot \bar{q}_{1,N,i}}{q_{1,N,i} \cdot \bar{q}_{0,N,i}}.$$

Next, using (A.5)

$$l = \frac{q_{0,N,i}q_{1,N,i}V_{N,i} + \frac{1}{2}q_{0,N,i}\mathrm{tr}(\nabla^2_{1,N,i}M_{N,i})V_{N,i}^{1+2/k} + o\left(V_{N,i}^{1+2/k}\right)}{q_{0,N,i}q_{1,N,i}V_{N,i} + \frac{1}{2}q_{1,N,i}\mathrm{tr}(\nabla^2_{0,N,i}M_{N,i})V_{N,i}^{1+2/k} + o\left(V_{N,i}^{1+2/k}\right)}$$

and

$$l - 1 = \frac{1}{2q_{1,N,i}}\mathrm{tr}(\nabla^2_{1,N,i}M_{N,i})V_{N,i}^{2/k} - \frac{1}{2q_{0,N,i}}\mathrm{tr}(\nabla^2_{0,N,i}M_{N,i})V_{N,i}^{2/k} + o\left(V_{N,i}^{2/k}\right). \tag{A.14}$$

Therefore, $(l-1)^2 = o\left(V_{N,i}^{2/k}\right)$ and using (A.14) and (A.3) we get

$$\Lambda_{N,i} - \bar{\Lambda}_{N,i} = -\frac{1}{2}\mathrm{tr}(G_{N,i}M_{N,i})V_{N,i}^{2/k} + o\left(V_{N,i}^{2/k}\right). \tag{A.15}$$

Finally, (A.13) and (A.15) give

$$\begin{aligned}\int_{S_{N,i}} q_0(y)\Lambda(y)dy - \bar{q}_{0,N,i}\bar{\Lambda}_{N,i} &= \frac{1}{2}q_{0,N,i}\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{1+2/k} + o\left(V_{N,i}^{1+2/k}\right) \\ &= \frac{1}{2}q_{0,N,i}\mathrm{tr}(F_{N,i}M_{N,i})\frac{V_{N,i}}{N^{2/k}\zeta_{N,i}^{2/k}} + \\ &\quad o\left(V_{N,i}^{1+2/k}\right).\end{aligned} \tag{A.16}$$

A.1.4 Total Distortion

Having calculated the single-cell distortion (A.16), the total distortion is obtained by summing over all quantizer cells. Using (A.1) and (A.16), the total distortion of quantizer Q_N is

$$\Delta L_N = \frac{1}{2N^{2/k}} \sum_{i=1}^{N} q_{0,N,i} \mathrm{tr}(F_{N,i} M_{N,i}) \frac{1}{\zeta_{N,i}^{2/k}} V_{N,i} + o\left(\frac{1}{N^{2/k}}\right) V_{N,i}.$$

Multiplying by $N^{2/k}$ and taking the limit, we obtain (20).

A.2 Asymptotic Loss in Sanov Exponents

We begin by writing the loss in discrimination between the tilted source q_λ and source q_0 due to quantization with an N-point vector quantizer as

$$\begin{aligned} \Delta L_{0,N} &\triangleq L(q_\lambda \| q_0) - L(\hat{q}_{\lambda,N} \| \bar{q}_{0,N}) \\ &= \sum_{i=1}^{N} \int_{S_{N,i}} q_\lambda(x) \Lambda_0(x) dx - \hat{q}_{\lambda,N,i} \hat{\Lambda}_{0,N,i} \end{aligned} \tag{A.17}$$

where

$$\Lambda_0(x) = \log \frac{q_\lambda(x)}{q_0(x)}, \qquad \hat{\Lambda}_{0,N,i} = \log \frac{\hat{q}_{\lambda,N,i}}{\bar{q}_{0,N,i}}.$$

In keeping with the notational convention we define

$$\begin{aligned} q_{\lambda,N,i} &= q_\lambda(x_{N,i}) \\ \nabla_{\lambda,N,i} &= \nabla q_\lambda(x_{N,i}) \\ \nabla^2_{\lambda,N,i} &= \nabla^2 q_\lambda(x_{N,i}) \end{aligned}$$

and

$$\Lambda_{0,N,i} = \Lambda_0(x_{N,i}).$$

Next we define

$$\begin{aligned} \mu &= \int q_0(x)^{1-\lambda} q_1(x)^\lambda dx = \sum_{i=1}^{N} \mu_{N,i} \\ \mu_{N,i} &= \int_{S_{N,i}} q_0(x)^{1-\lambda} q_1(x)^\lambda dx = \mu \int_{S_{N,i}} q_\lambda(x) dx \\ d_{N,i} &= \bar{q}_{0,N,i}^{1-\lambda} \cdot \bar{q}_{1,N,i}^{\lambda} - \mu_{N,i} \\ d_N &= \sum_{i-1}^{N} d_{N,i}. \end{aligned} \tag{A.18}$$

Thus we can write

$$\hat{q}_{\lambda,N,i} = \frac{\mu_{N,i} + d_{N,i}}{\mu + d_N}. \tag{A.19}$$

A.2.1 Expansions of $\mu_{N,i}$ and $d_{N,i}$

Expanding $q_\lambda(x)$ in a Taylor series about $x_{N,i}$ we get the following representation for $\mu_{N,i}$:

$$\mu_{N,i} = \mu q_{\lambda,N,i} V_{N,i} + \frac{\mu}{2} \int_{S_{N,i}} (x - x_{N,i})^T \nabla^2_{\lambda,N,i} (x - x_{N,i}) dx + o\left(V_{N,i}^{1+2/k}\right). \tag{A.20}$$

It can be straightforwardly shown that the Hessian of the tilted density is

$$\nabla^2 q_\lambda(x) = q_\lambda(x) \left[\lambda \frac{\nabla^2 q_1(x)}{q_1(x)} + (1-\lambda) \frac{\nabla^2 q_0(x)}{q_0(x)} - \lambda(1-\lambda) F(x) \right]. \tag{A.21}$$

Next, using the centroid assumption, we write

$$\begin{aligned} \bar{q}_{0,N,i} &= q_{0,N,i} V_{N,i} + \frac{1}{2} \int_{S_{N,i}} (x - x_{N,i})^T \nabla^2_{0,N,i} (x - x_{N,i}) dx + o\left(V_{N,i}^{1+2/k}\right) \\ \bar{q}_{1,N,i} &= q_{1,N,i} V_{N,i} + \frac{1}{2} \int_{S_{N,i}} (x - x_{N,i})^T \nabla^2_{1,N,i} (x - x_{N,i}) dx + o\left(V_{N,i}^{1+2/k}\right) \end{aligned} \tag{A.22}$$

and using the Taylor expansion

$$(x+y)^a = x^a + a x^{a-1} y + \frac{1}{2} a(a-1) x^{a-2} y^2 + o(y^2) \tag{A.23}$$

we obtain

$$\begin{aligned} \bar{q}_{0,N,i}^{1-\lambda} &= q_{0,N,i}^{1-\lambda} V_{N,i}^{1-\lambda} + \frac{1}{2}(1-\lambda) q_{0,N,i}^{-\lambda} V_{N,i}^{-\lambda} \int_{S_{N,i}} (x - x_{N,i})^T \nabla^2_{0,N,i} (x - x_{N,i}) dx + \\ &\quad o\left(V_{N,i}^{2/k+1-\lambda}\right) \end{aligned}$$

and

$$\bar{q}_{1,N,i}^{\lambda} = q_{1,N,i}^{\lambda} V_{N,i}^{\lambda} + \frac{1}{2} \lambda q_{1,N,i}^{\lambda-1} V_{N,i}^{\lambda-1} \int_{S_{N,i}} (x - x_{N,i})^T \nabla^2_{1,N,i} (x - x_{N,i}) dx + o\left(V_{N,i}^{2/k+\lambda}\right).$$

Multiplying the two formulas above yields

$$\begin{aligned} \bar{q}_{0,N,i}^{1-\lambda} \cdot \bar{q}_{1,N,i}^{\lambda} &= \mu q_{\lambda,N,i} \left(\frac{\lambda}{2 q_{1,N,i}} \int_{S_{N,i}} (x - x_{N,i})^T \nabla^2_{1,N,i} (x - x_{N,i}) dx \right. \\ &\quad \left. + \frac{1-\lambda}{2 q_{0,N,i}} \int_{S_{N,i}} (x - x_{N,i})^T \nabla^2_{0,N,i} (x - x_{N,i}) dx + V_{N,i} \right) \\ &\quad + o\left(V_{N,i}^{1+2/k}\right). \end{aligned} \tag{A.24}$$

Finally, using (A.20), (A.21), and (A.24) we get

$$d_{N,i} = \frac{\mu}{2}\lambda(1-\lambda)q_{\lambda,N,i}\int_{S_{N,i}}(x-x_{N,i})^T F_{N,i}(x-x_{N,i})dx + o\left(V_{N,i}^{1+2/k}\right). \tag{A.25}$$

We shall find the following formulas for $\mu_{N,i}$ and $d_{N,i}$ useful:

$$\mu_{N,i} = \mu q_{\lambda,N,i}V_{N,i} + \frac{\mu}{2}\mathrm{tr}\left(\nabla^2_{\lambda,N,i}M_{N,i}\right)V_{N,i}^{1+2/k} + o\left(V_{N,i}^{1+2/k}\right) \tag{A.26}$$

$$d_{N,i} = \frac{\mu}{2}\lambda(1-\lambda)q_{\lambda,N,i}\mathrm{tr}\left(F_{N,i}M_{N,i}\right)V_{N,i}^{1+2/k} + o\left(V_{N,i}^{1+2/k}\right). \tag{A.27}$$

A.2.2 Asymptotic Values of $\Delta L_{0,N}$ and $\Delta L_{1,N}$

From (A.7) and (A.8) we can write

$$\begin{aligned}\int_{S_{N,i}} q_\lambda(x)\Lambda_0(x)dx &= q_{\lambda,N,i}\Lambda_{0,N,i}V_{N,i} + \frac{1}{2}\Lambda_{0,N,i}\mathrm{tr}(\nabla^2_{\lambda,N,i}M_{N,i})V_{N,i}^{1+2/k} + \\ &\quad \frac{1}{2}q_{\lambda,N,i}\mathrm{tr}\left((F'_{N,i}+G'_{N,i})M_{N,i}\right)V_{N,i}^{1+2/k} + o\left(V_{N,i}^{1+2/k}\right)\end{aligned}$$

where

$$\begin{aligned}F'_{N,i} &= \nabla\Lambda_{0,N,i}\nabla\Lambda_{0,N,i}^T \\ G'_{N,i} &= \frac{\nabla^2_{\lambda,N,i}}{q_{\lambda,N,i}} - \frac{\nabla^2_{0,N,i}}{q_{0,N,i}}.\end{aligned} \tag{A.28}$$

Note that $F'_{N,i}$ can be written in terms of $F_{N,i}$:

$$F'_{N,i} = \lambda^2 F_{N,i}.$$

From (A.19), (A.26), and (A.27) we can write

$$\begin{aligned}\hat{q}_{\lambda,N,i} &= t_N\left(q_{\lambda,N,i}V_{N,i} + \frac{1}{2}\mathrm{tr}(\nabla^2_{\lambda,N,i}M_{N,i})V_{N,i}^{1+2/k} + \frac{1}{2}\lambda(1-\lambda)q_{\lambda,N,i}\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{1+2/k}\right) \\ &\quad + o\left(V_{N,i}^{1+2/k}\right)\end{aligned} \tag{A.29}$$

where

$$t_N = \frac{\mu}{\mu + d_N}.$$

Thus (A.17) becomes

$$\begin{aligned}\Delta L_{0,N} &= \sum_{i=1}^{N} q_{\lambda,N,i} V_{N,i} \left(\Lambda_{0,N,i} - t_N \hat{\Lambda}_{0,N,i}\right) + \\ &\quad \frac{1}{2}\mathrm{tr}(\nabla^2_{\lambda,N,i} M_{N,i}) V_{N,i}^{1+2/k} \left(\Lambda_{0,N,i} - t_N \hat{\Lambda}_{0,N,i}\right) + \\ &\quad \frac{1}{2} q_{\lambda,N,i} \mathrm{tr}\left((\lambda^2 F_{N,i} + G'_{N,i}) M_{N,i}\right) V_{N,i}^{1+2/k} - \\ &\quad \frac{\lambda(1-\lambda)}{2} t_N q_{\lambda,N,i} \hat{\Lambda}_{0,N,i} \mathrm{tr}(F_{N,i} M_{N,i}) V_{N,i}^{1+2/k} + o\left(V_{N,i}^{1+2/k}\right). \end{aligned} \tag{A.30}$$

Next we use the Taylor expansion

$$\log(x+y) = \log x + \frac{y}{x} - \frac{y^2}{2x^2} + o(y^2)$$

to write

$$\hat{\Lambda}_{0,N,i} = \Lambda_{0,N,i} + 2r_{0,N,i} - \frac{1}{2} r^2_{0,N,i} - \frac{3}{2} + o\left(\left(\frac{\hat{q}_{\lambda,N,i}}{\bar{q}_{0,N,i}} - \frac{q_{\lambda,N,i}}{q_{0,N,i}}\right)^2\right) \tag{A.31}$$

where

$$r_{0,N,i} = \frac{q_{0,N,i} \hat{q}_{\lambda,N,i}}{\bar{q}_{0,N,i} q_{\lambda,N,i}}.$$

To see that the last term in (A.31) is small, note that

$$\frac{\hat{q}_{\lambda,N,i}}{\bar{q}_{0,N,i}} - \frac{q_{\lambda,N,i}}{q_{0,N,i}} = \left(\frac{\bar{q}_{1,N,i}}{\bar{q}_{0,N,i}}\right)^{\lambda} \frac{1}{\mu + d_N} - \left(\frac{q_{1,N,i}}{q_{0,N,i}}\right)^{\lambda} \frac{1}{\mu}.$$

Using the Taylor expansions (A.22), after some algebra this becomes

$$\begin{aligned}\frac{\hat{q}_{\lambda,N,i}}{\bar{q}_{0,N,i}} - \frac{q_{\lambda,N,i}}{q_{0,N,i}} &= \left(\frac{q_{1,N,i}}{q_{0,N,i}} + o\left(V_{N,i}^{2/k}\right)\right)^{\lambda} \frac{1}{\mu + d_N} - \left(\frac{q_{1,N,i}}{q_{0,N,i}}\right)^{\lambda} \frac{1}{\mu} \\ &= \left[\left(\frac{q_{1,N,i}}{q_{0,N,i}}\right)^{\lambda} + o\left(V_{N,i}^{2/k}\right)\right] \frac{1}{\mu + d_N} - \left(\frac{q_{1,N,i}}{q_{0,N,i}}\right)^{\lambda} \frac{1}{\mu} \\ &= -\left(\frac{q_{1,N,i}}{q_{0,N,i}}\right)^{\lambda} \frac{d_N}{\mu(\mu + d_N)} + o\left(V_{N,i}^{2/k}\right) \end{aligned}$$

where the second equality follows from (A.23). From (A.27) it is easily seen that

$$o\left(\left(\frac{\hat{q}_{\lambda,N,i}}{\bar{q}_{0,N,i}} - \frac{q_{\lambda,N,i}}{q_{0,N,i}}\right)^2\right) = o\left(V_{N,i}^{2/k}\right).$$

Now, using (A.22) and (A.29), $r_{0,N,i}$ becomes

$$
\begin{aligned}
& r_{0,N,i} = \\
& \frac{t_N q_{0,N,i}\left(q_{\lambda,N,i} + \frac{1}{2}\mathrm{tr}(\nabla^2_{\lambda,N,i}M_{N,i})V_{N,i}^{2/k} + \frac{1}{2}\lambda(1-\lambda)q_{\lambda,N,i}\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{2/k}\right) + o\left(V_{N,i}^{2/k}\right)}{q_{\lambda,N,i}q_{0,N,i} + \frac{1}{2}q_{\lambda,N,i}\mathrm{tr}(\nabla^2_{0,N,i}M_{N,i})V_{N,i}^{2/k} + o\left(V_{N,i}^{2/k}\right)} \\
& = t_N\left(1 + \frac{\mathrm{tr}(\nabla^2_{\lambda,N,i}M_{N,i})}{2q_{\lambda,N,i}}V_{N,i}^{2/k} - \frac{\mathrm{tr}(\nabla^2_{0,N,i}M_{N,i})}{2q_{0,N,i}}V_{N,i}^{2/k} + \frac{1}{2}\lambda(1-\lambda)\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{2/k}\right) \\
& \quad + o\left(V_{N,i}^{2/k}\right) \\
& = t_N\left(1 + \frac{1}{2}\mathrm{tr}(G'_{N,i}M_{N,i})V_{N,i}^{2/k} + \frac{1}{2}\lambda(1-\lambda)\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{2/k}\right) + o\left(V_{N,i}^{2/k}\right) \qquad \text{(A.32)}
\end{aligned}
$$

and

$$
r^2_{0,N,i} = t_N^2\left(1 + \mathrm{tr}(G'_{N,i}M_{N,i})V_{N,i}^{2/k} + \lambda(1-\lambda)\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{2/k}\right) + o\left(V_{N,i}^{2/k}\right).
$$

Thus (A.31) becomes

$$
\begin{aligned}
\hat{\Lambda}_{0,N,i} &= \Lambda_{0,N,i} + 2t_N - \frac{1}{2}t_N^2 - \frac{3}{2} + \\
&\quad \left(\mathrm{tr}(G'_{N,i}M_{N,i})V_{N,i}^{2/k} + \lambda(1-\lambda)\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{2/k}\right)\left(t_N - \frac{1}{2}t_N^2\right) + o\left(V_{N,i}^{2/k}\right).
\end{aligned}
$$

Therefore

$$
\begin{aligned}
\Lambda_{0,N,i} - t_N\hat{\Lambda}_{0,N,i} &= \Lambda_{0,N,i}(1-t_N) + \frac{3}{2}t_N - 2t_N^2 + \frac{1}{2}t_N^3 - \\
&\quad \left(\mathrm{tr}(G'_{N,i}M_{N,i})V_{N,i}^{2/k} + \lambda(1-\lambda)\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{2/k}\right)\left(t_N^2 - \frac{1}{2}t_N^3\right) \\
&\quad + o\left(V_{N,i}^{2/k}\right). \qquad \text{(A.33)}
\end{aligned}
$$

Next, using (A.27), we note that

$$
\lim_{N\to+\infty} N^{2/k}\frac{d_N}{\mu} = \frac{1}{2}\lambda(1-\lambda)\int \frac{q_\lambda(x)\mathcal{F}(x)}{\zeta(x)^{2/k}}dx
$$

and thus

$$
\begin{aligned}
t_N &= 1 - \frac{d_N}{\mu + d_N} = 1 - \frac{d_N}{\mu} + o\left(\frac{1}{N^{2/k}}\right) \\
t_N^2 &= 1 - \frac{2d_N}{\mu} + o\left(\frac{1}{N^{2/k}}\right) \\
t_N^3 &= 1 - \frac{3d_N}{\mu} + o\left(\frac{1}{N^{2/k}}\right).
\end{aligned}
$$

Using this in (A.33) gives

$$\begin{aligned}\Lambda_{0,N,i} - t_N \hat{\Lambda}_{0,N,i} &= \Lambda_{0,N,i}\frac{d_N}{\mu} + \frac{d_N}{\mu} - \\ &\frac{1}{2}\left(\mathrm{tr}(G'_{N,i}M_{N,i})V_{N,i}^{2/k} + \lambda(1-\lambda)\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{2/k}\right) + \\ &o\left(V_{N,i}^{2/k}\right) + o\left(\frac{1}{N^{2/k}}\right). \end{aligned} \tag{A.34}$$

Next, (A.30) and (A.34) give

$$\begin{aligned}\Delta L_{0,N} &= \sum_{i=1}^{N} q_{\lambda,N,i}\Lambda_{0,N,i}V_{N,i}\frac{d_N}{\mu} - \frac{1}{2}\lambda(1-\lambda)q_{\lambda,N,i}\Lambda_{0,N,i}\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{1+2/k} + \\ &q_{\lambda,N,i}V_{N,i}\frac{d_N}{\mu} - \frac{1}{2}\lambda(1-\lambda)q_{\lambda,N,i}\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{1+2/k} + \\ &\frac{1}{2}\lambda^2 q_{\lambda,N,i}\mathrm{tr}(F_{N,i}M_{N,i})V_{N,i}^{1+2/k} + \\ &o\left(V_{N,i}^{1+2/k}\right) + o\left(\frac{V_{N,i}}{N^{2/k}}\right). \end{aligned} \tag{A.35}$$

Finally, by multiplying (A.35) by $N^{2/k}$ and passing to the limit, we obtain (22). By symmetry arguments, (23) can easily be obtained.

References

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

[2] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. on Inform. Theory*, vol. 44, no. 6, pp. 2325–2383, Oct. 1998.

[3] S. A. Kassam, "Optimum Quantization for Signal Detection," *IEEE Trans. on Communications*, vol. COM-25, pp. 479–484, May 1977.

[4] H. V. Poor and J. B. Thomas, "Applications of Ali-Silvey Distance Measures in the Design of Generalized Quantizers," *IEEE Trans. on Communications*, vol. COM-25, pp. 893–900, Sep. 1977.

[5] H. V. Poor, "A Companding Approximation for the Statistical Divergence of Quantized Data," in *Proc. of IEEE Conf. Decision and Control*, San Antonio, TX, Dec. 1983.

[6] H. V. Poor, "Fine Quantization in Signal Detection and Estimation – Part 1," *IEEE Trans. on Inform. Theory*, vol. 34, pp. 960–972, Sep. 1988.

[7] G. R. Benitz and J. A. Bucklew, "Asymptotically Optimal Quantizers for Detection of I.I.D. Data," *IEEE Trans. on Inform. Theory*, vol. 35, no. 2, pp. 316–325, Mar. 1989.

[8] B. Picinbono and P. Duvaut, "Optimum Quantization for Detection," *IEEE Trans. on Communications*, vol. 36, no. 11, pp. 1254–1258, Nov. 1988.

[9] J. N. Tsitsiklis, "Extremal Properties of Likelihood Ratio Quantizers," *IEEE Trans. on Communications*, vol. 41, no. 4, pp. 550–558, Apr. 1993.

[10] T. J. Flynn and R. M. Gray, "Encoding of Correlated Observations," *IEEE Trans. on Inform. Theory*, vol. 33, no. 6, pp. 773–787, Nov. 1987.

[11] M. Longo, T. D. Lookabaugh, and R. M. Gray, "Quantization for Decentralized Hypothesis Testing under Communication Constraints," *IEEE Trans. on Inform. Theory*, vol. 36, pp. 241–255, March 1990.

[12] K. Oehler and R. M. Gray, "Combining Image Compression and Classification using Vector Quantization," *IEEE Trans. on Pattern Anal. and Machine Intell.*, vol. 17, no. 5, pp. 461–473, May 1995.

[13] K. O. Perlmutter, S. M. Perlmutter, et al., "Bayes Risk Weighted Vector Quantization with Posterior Estimation for Image Compression and Classification," *IEEE Trans. on Image Processing*, vol. 5, no. 2, pp. 347–360, Feb. 1996.

[14] S. Jana and P. Moulin, "Optimal design of transform coders for image classification," in *Proc. Conf. on Information Sciences and Systems*, Baltimore, MD, March 1999.

[15] A. Jain, P. Moulin, M. I. Miller, and K. Ramchandran, "Information-theoretic bounds on target recognition performance based on degraded image data," to be published, Dec. 1999.

[16] R. M. Gray, "Quantization, Classification, and Density Estimation," in *Proc. of IEEE Information Theory Workshop on Detection, Estimation, Classification, and Imaging*, Santa Fe, NM, Feb. 1999.

[17] P. Zador, *Development and Evaluation of Procedures for Quantizing Multivariate Distributions*, PhD thesis, Stanford University, Stanford, CA, 1963.

[18] A. Gersho, "Asymptotically Optimal Block Quantization," *IEEE Trans. on Inform. Theory*, vol. IT-25, pp. 373–380, July 1979.

[19] S. Na and D. L. Neuhoff, "Bennet's Integral for Vector Quantizers," *IEEE Trans. on Inform. Theory*, vol. 41, no. 4, pp. 886–900, July 1995.

[20] R. Zamir and M. Feder, "On Lattice Quantization Noise," *IEEE Trans. on Inform. Theory*, vol. 42, no. 4, pp. 1152–1159, July 1996.

[21] Y. Yamada, S. Tazaki, and R. M. Gray, "Asymptotic Performance of Block Quantizers with Difference Distortion Measures," *IEEE Trans. on Inform. Theory*, vol. 26, pp. 6–14, January 1980.

[22] W. Gardner and B. Rao, "Theoretical Analysis of the High-Rate Vector Quantization of LPC Parameters," *IEEE Trans. Speech and Audio Processing*, vol. 3, no. 5, pp. 367–381, September 1995.

[23] J. Li, N. Chadda, and R. M. Gray, "Asymptotic Performance of Vector Quantizers with a Perceptual Distortion Measure," *IEEE Trans. on Inform. Theory*, vol. 45, no. 4, pp. 1082–1091, May 1999.

[24] T. Linder, R. Zamir, and K. Zeger, "On Source Coding with Side-Information-Dependend Distortion Measures," *IEEE Trans. on Inform. Theory*, vol. 46, no. 7, pp. 2697–2704, November 2000.

[25] D. Slepian and J. K. Wolf, "Noiseless Coding of Correlated Information Sources," *IEEE Trans. on Inform. Theory*, vol. 19, pp. 471–480, July 1973.

[26] R. Gupta, *Quantization Strategies for Low-Power Communications*, PhD thesis, University of Michigan, Ann Arbor, MI, 2001.

[27] E. L. Lehmann, *Testing Statistical Hypotheses*, Wiley, New York, 1959.

[28] R. Blahut, *Principles and Practice of Information Theory*, Addison-Wesley, 1987.

[29] T. Cover and J. Thomas, *Elements of Information Theory*, Wiley, New York, 1987.

[30] J. A. Bucklew, *Large Deviation Techniques in Decision, Simulation, and Estimation*, Wiley, New York, 1990.

[31] A. Dembo and O. Zeitouni, *Large Deviation Techniques and Applications*, Springer-Verlag, 1998.

[32] H. L. Van-Trees, *Detection, Estimation, and Modulation Theory: Part I*, Wiley, New York, 1968.

[33] D. M. Green and J. A. Swets, *Signal Detection Theory and Psychophysics*, Wiley, 1966.

[34] F. Provost and T. Fawcett, "Robust classification for imprecise environments," *Machine Learning Journal*, vol. 42, no. 3, pp. 203–231, March 2001.

[35] P. M. Altham, "A non-parametric measure of signal discriminability," *Brit. J. Math. Statist. Psychol.*, vol. 26, pp. 1–12, 1973.

[36] D. Bamber, "The area above the ordinal dominance graph and the area below the receiver operating characteristic graph," *Journal of Mathematical Psychology*, vol. 12, pp. 387–415, 1975.

[37] C. E. Metz, "Basic principles of ROC analysis," *Seminars in Nuclear Medicine*, vol. 8, no. 4, pp. 283–298, 1978.

[38] J. A. Swets and R. M. Pickett, *Evaluation of Diagnostic Systems: Methods from Signal Detection Theory*, Kluwer Academic Press, 1982.

[39] A. Bradley, "The use of the area under the roc curve in the evaluation of machine learning algorithms," *Pattern Recognition*, vol. 30, no. 7, pp. 1145–1159, 1997.

[40] J. A. Hanley and B. J. McNeil, "The meaning and use of the area under a receiver operator characteristic (ROC) curve," *Radiology*, vol. 143, pp. 29–36, 1982.

[41] H. H. Barrett, C. K. Abbey, and E. Clarkson, "Objective Assessment of Image Quality. III. ROC Metrics, Ideal Observers, and Likelihood-Generating Functions," *J. Opt. Soc. Am.*, vol. 15, no. 6, pp. 1520–1535, June 1998.

[42] J. Shapiro, "Bounds on the area under the ROC curve," *J. Opt. Soc. Am.*, vol. 16, pp. 53–57, Jan. 1999.

[43] Y. Linde, A. Buzo, and R. M. Gray, "An Algorithm for Vector Quantizer Design," *IEEE Trans. on Communications*, vol. COM-28, pp. 84–95, Jan. 1980.

[44] K. Sayood, *Introduction to Data Compression*, Morgan-Kaufman, 1996.

[45] P. W. Moo, *Asymptotic Analysis of Lattice-Based Quantization*, PhD thesis, University of Michigan, Ann Arbor, MI, 1998.

[46] J. A. Bucklew, "Two Results on the Asymptotic Performance of Quantizers," *IEEE Trans. on Inform. Theory*, vol. IT-30, pp. 341–348, March 1984.

[47] J. A. Bucklew and G. L. Wise, "Multidimensional Asymptotic Quantization Theory with r-th Power Distortion Measures," *IEEE Trans. on Inform. Theory*, vol. IT-28, pp. 239–247, March 1982.

[48] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge, 1985.

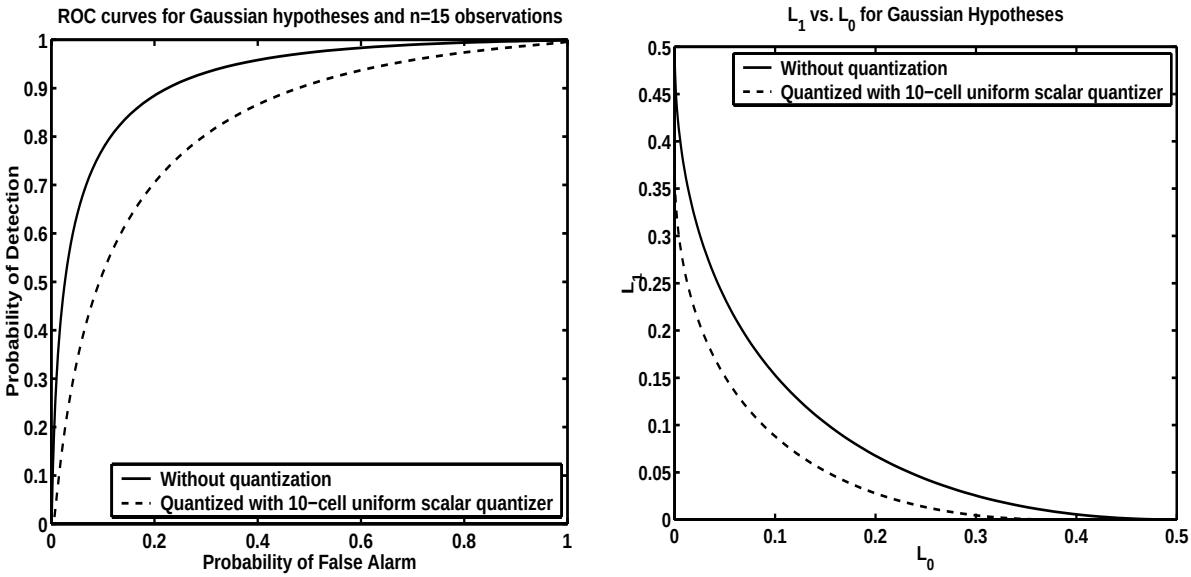

Figure 1: ROC curves and associated error exponent curves for Gaussian hypotheses.

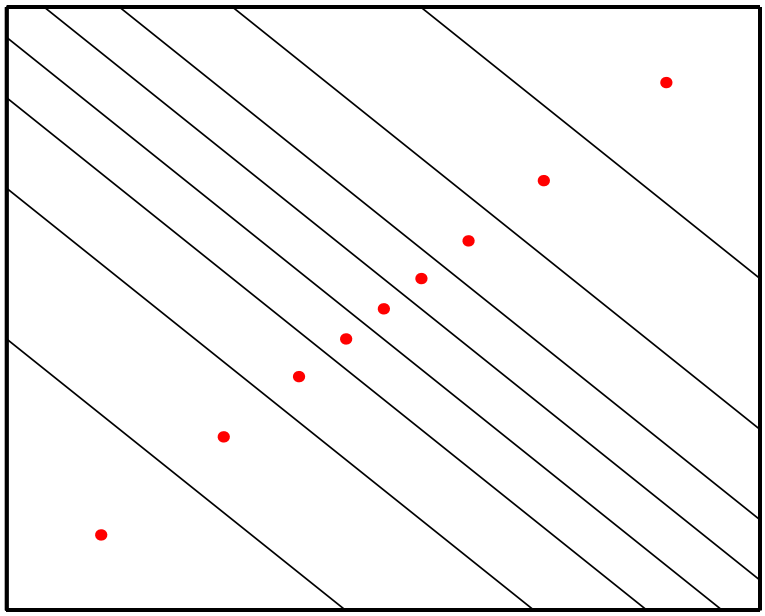

Figure 2: Log-likelihood ratio quantizer for two-dimensional Gaussian sources with identity covariance matrices.

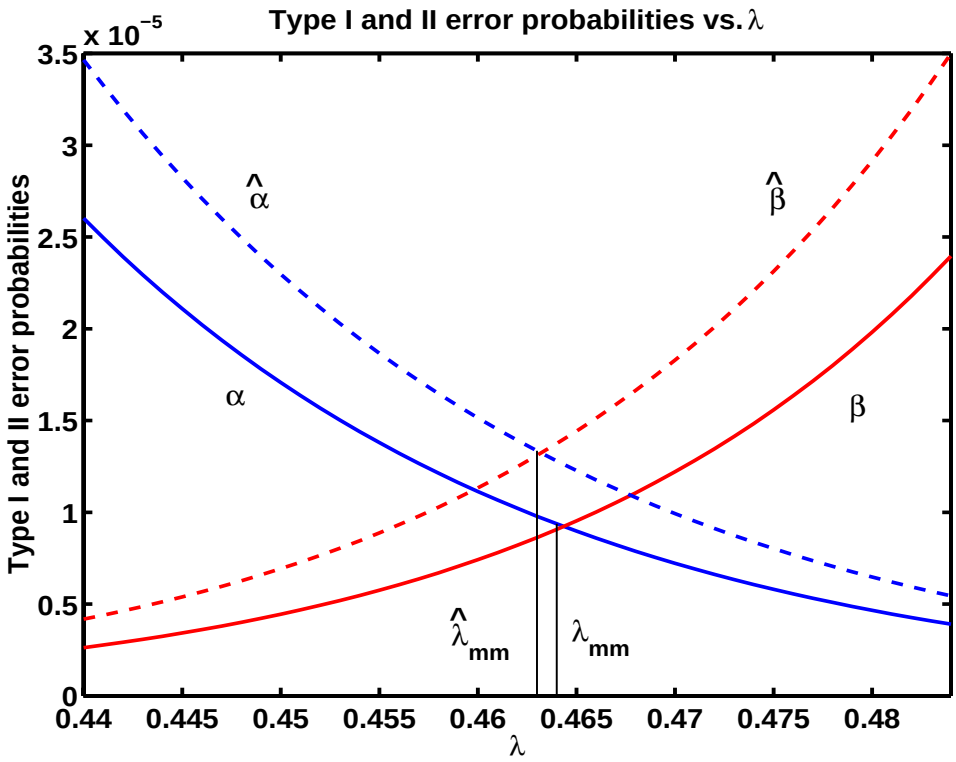

Figure 3: Sanov approximations to Type I and Type II errors indexed by λ before and after quantization for a one dimensional Gaussian example. The point of intersection of Type I and Type II error probabilities define the Chernoff information and the minimax operating point over λ.

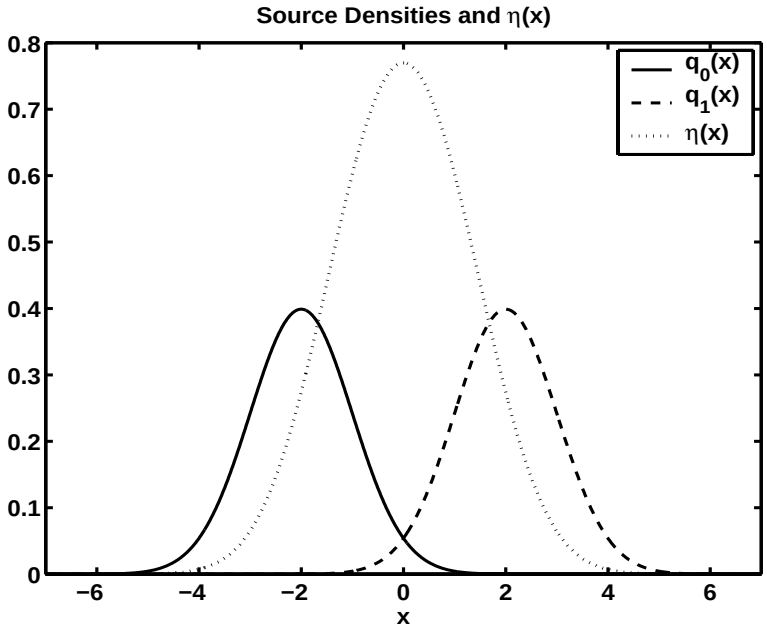

Figure 4: Source densities and $\eta(x)$ for one-dimensional Gaussian example.

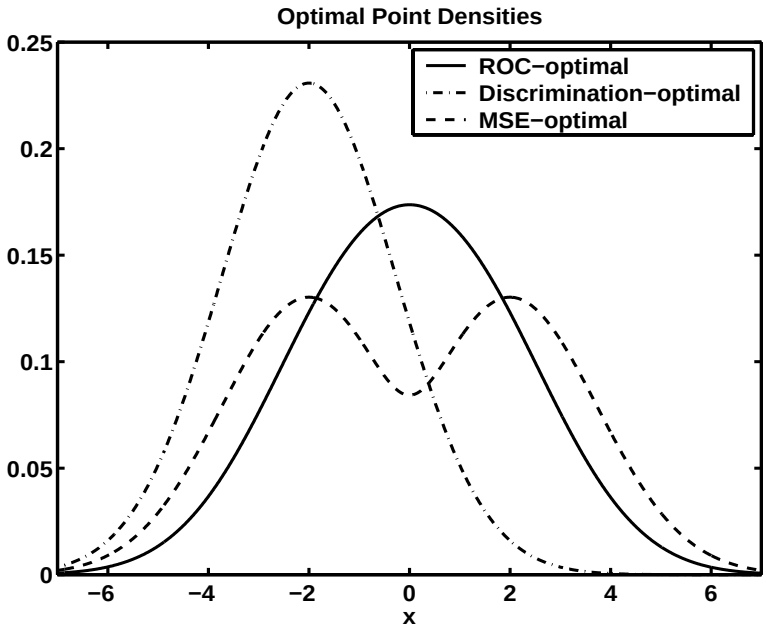

Figure 5: AUC-optimal, discrimination-optimal, and MSE-optimal point densities for one-dimensional Gaussian example.

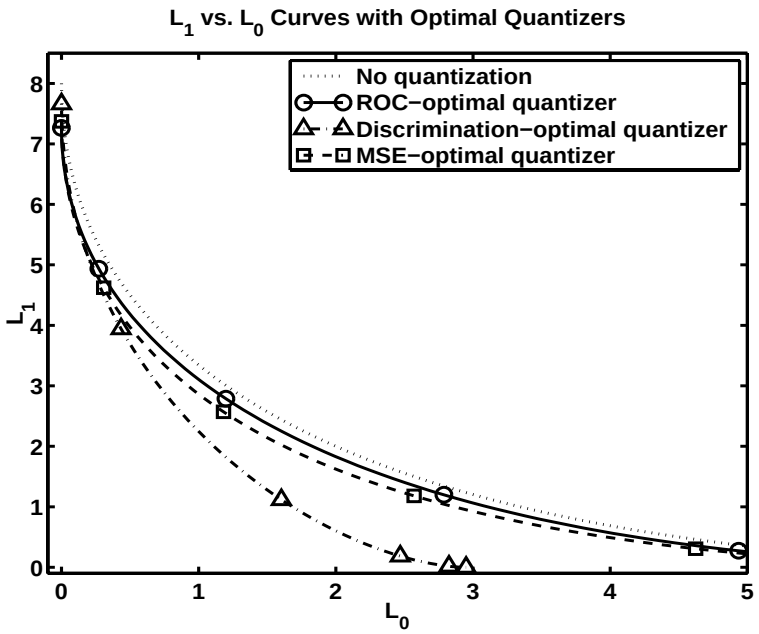

Figure 6: $L_1(L_0)$ curves without quantization and with quantization by AUC-optimal, discrimination-optimal, and MSE-optimal quantizers with $N = 8$ cells for one-dimensional Gaussian example. AUC-optimal quantizer has best performance, on average, while detection-optimal quantizer yields largest value of $L(\bar{q}_0 \| \bar{q}_1)$.

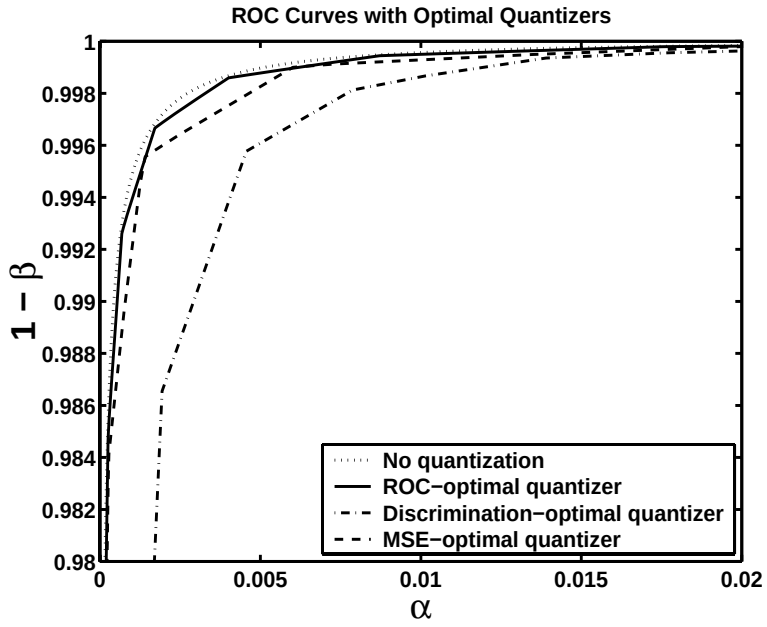

Figure 7: ROC curves with $n = 2$ observations and data quantized by AUC-optimal, discrimination-optimal, and MSE-optimal quantizers with $N = 16$ cells for one-dimensional Gaussian example.

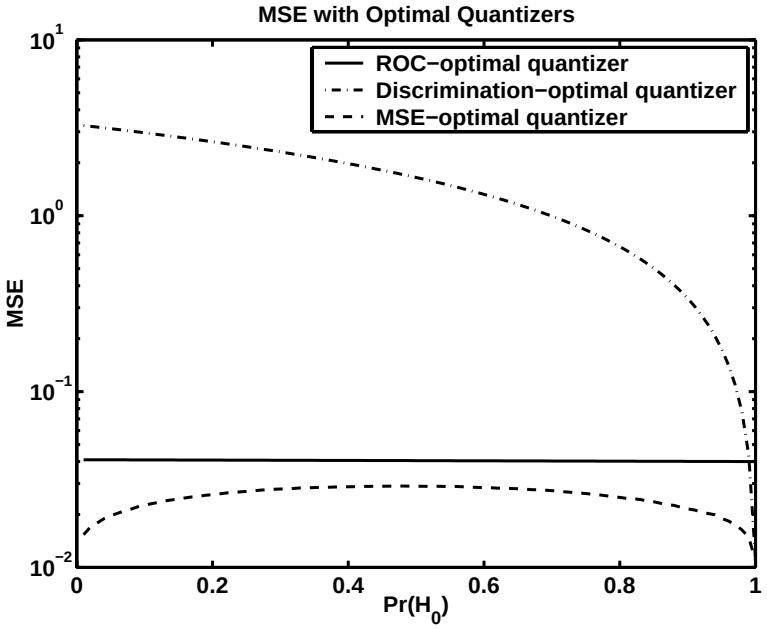

Figure 8: Reconstruction MSE with AUC-optimal, discrimination-optimal, and MSE-optimal quantizers with $N = 16$ cells for one-dimensional Gaussian example.

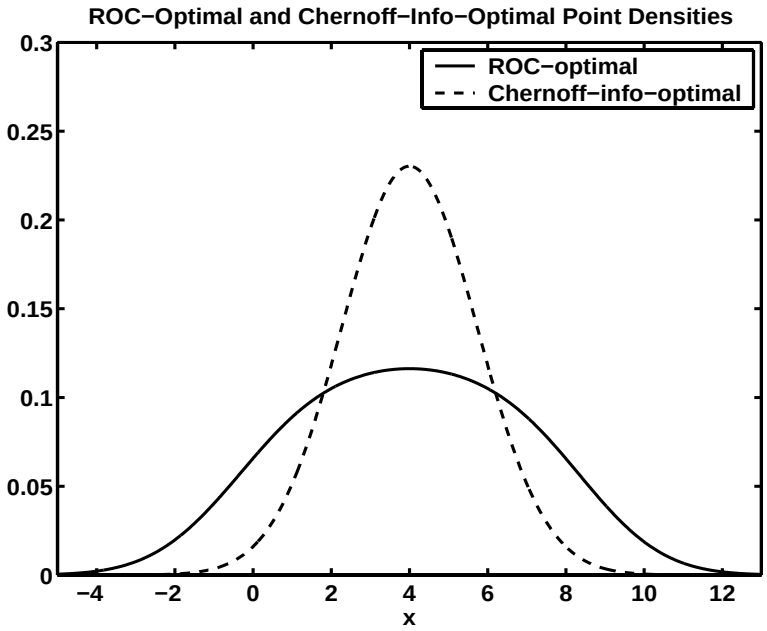

Figure 9: Optimal point densities for ROC area and Chernoff information for one-dimensional Gaussian sources with $\mu_0 = 0$ and $\mu_1 = 8$.

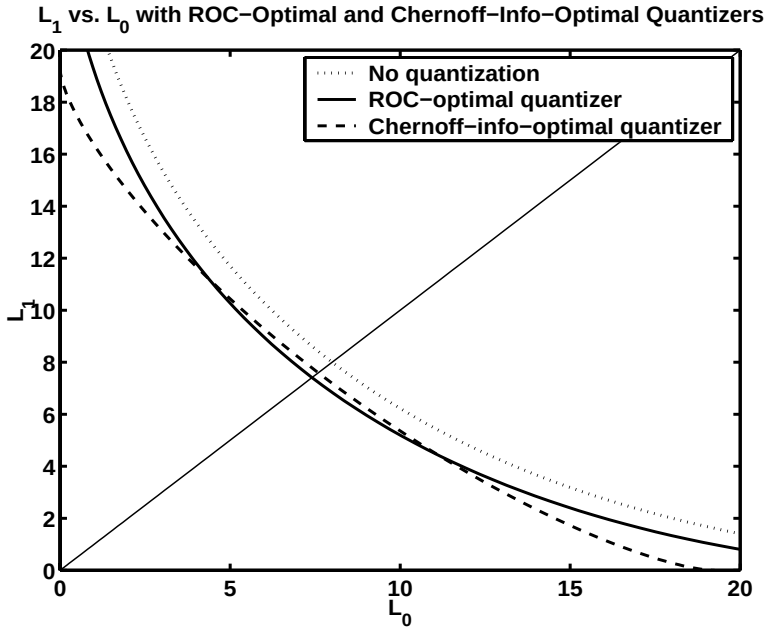

Figure 10: $L_1(L_0)$ curves without quantization and with quantization by AUC-optimal and Chernoff-information-optimal quantizers for one-dimensional Gaussian sources with $N = 8$, $\mu_0 = 0$, and $\mu_1 = 8$.

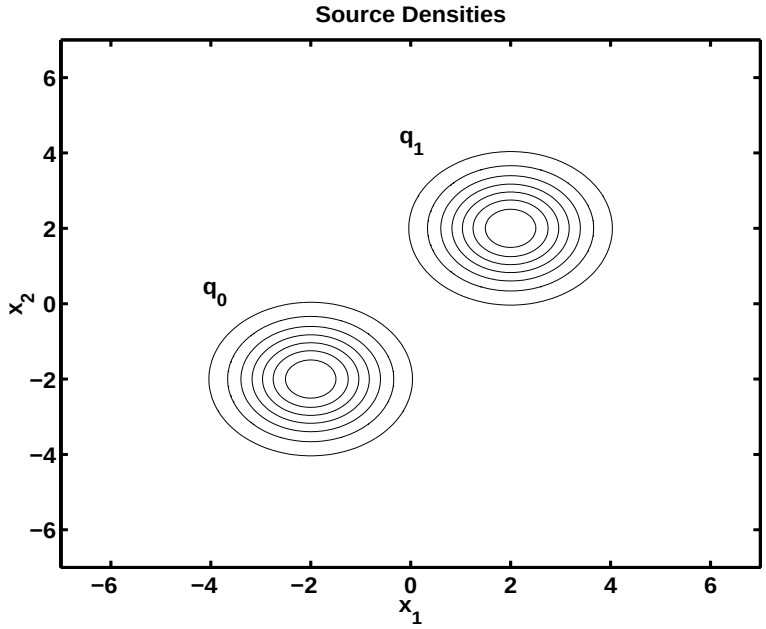

Figure 11: Source densities for two-dimensional uncorrelated Gaussian example.

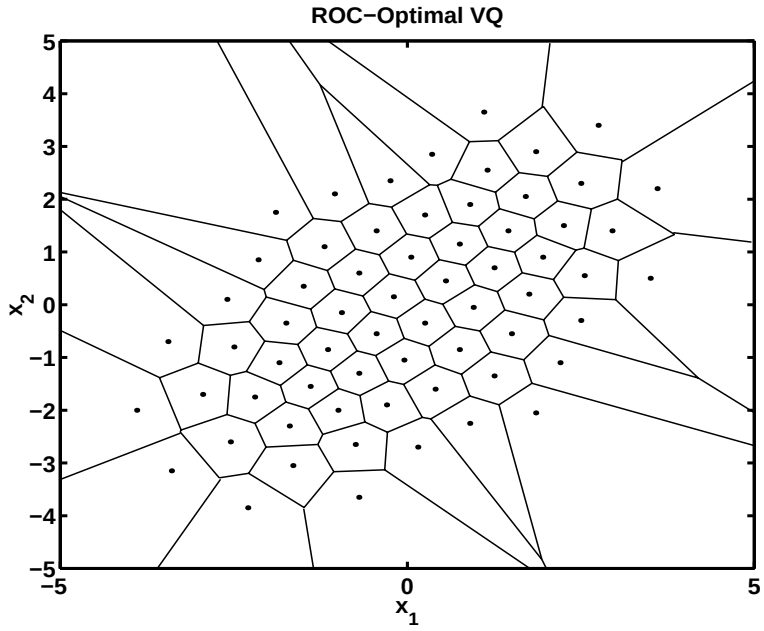

Figure 12: AUC-optimal 64-cell vector quantizer for two-dimensional uncorrelated Gaussian example.

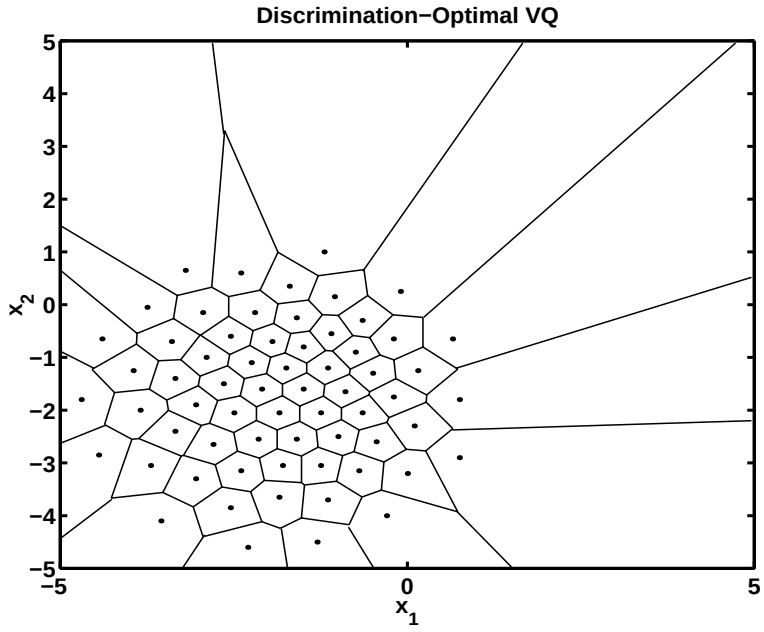

Figure 13: Discrimination-optimal 64-cell vector quantizer for two-dimensional uncorrelated Gaussian example.

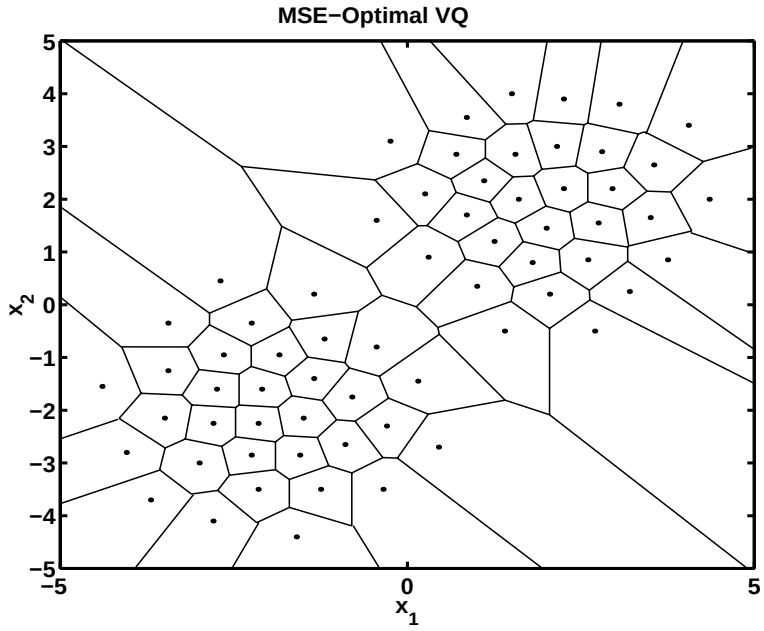

Figure 14: MSE-optimal 64-cell vector quantizer for two-dimensional uncorrelated Gaussian example.

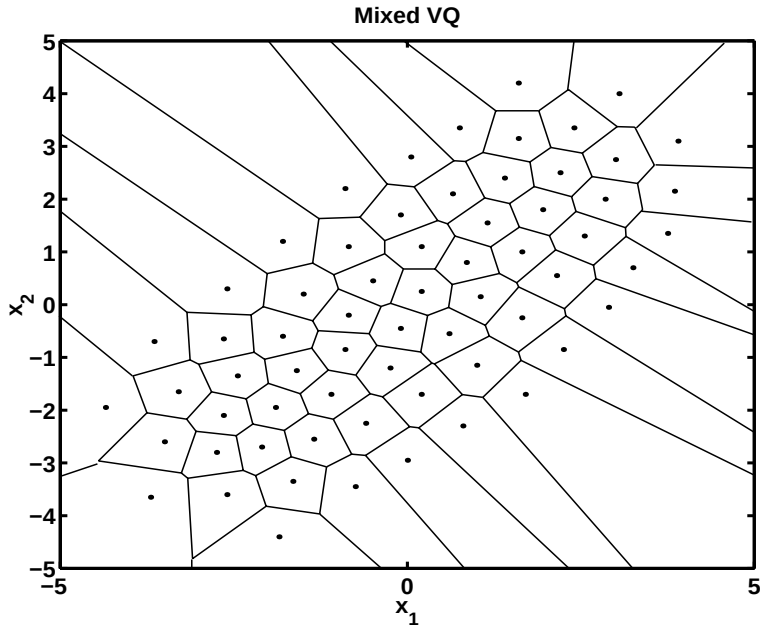

Figure 15: Optimal 64-cell vector quantizer with mixed objective function with $\rho = 1/2$ for two-dimensional uncorrelated Gaussian example.

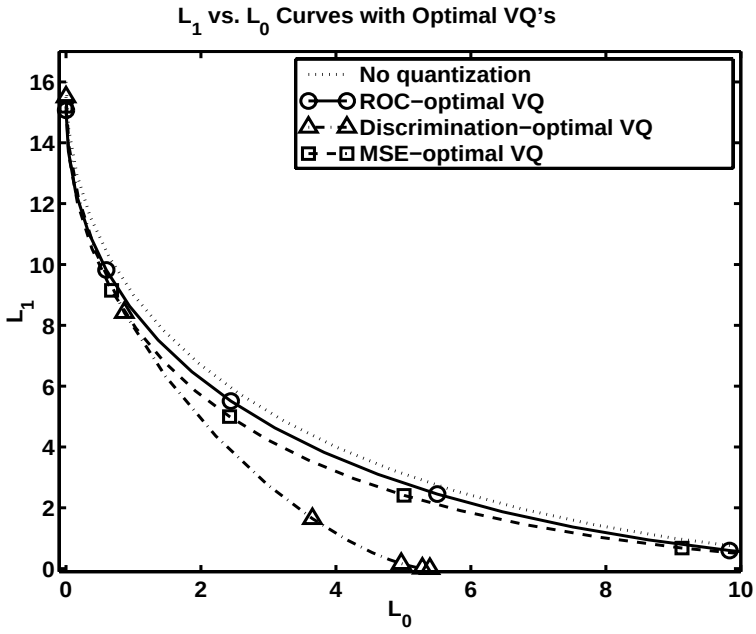

Figure 16: $L_1(L_0)$ curves without quantization and with quantization by AUC-optimal, discrimination-optimal, and MSE-optimal vector quantizers with $N = 64$ cells for two-dimensional uncorrelated Gaussian example. AUC-optimal quantizer has best performance, on average, while detection-optimal quantizer yields largest value of $L(\bar{q}_0 \| \bar{q}_1)$.

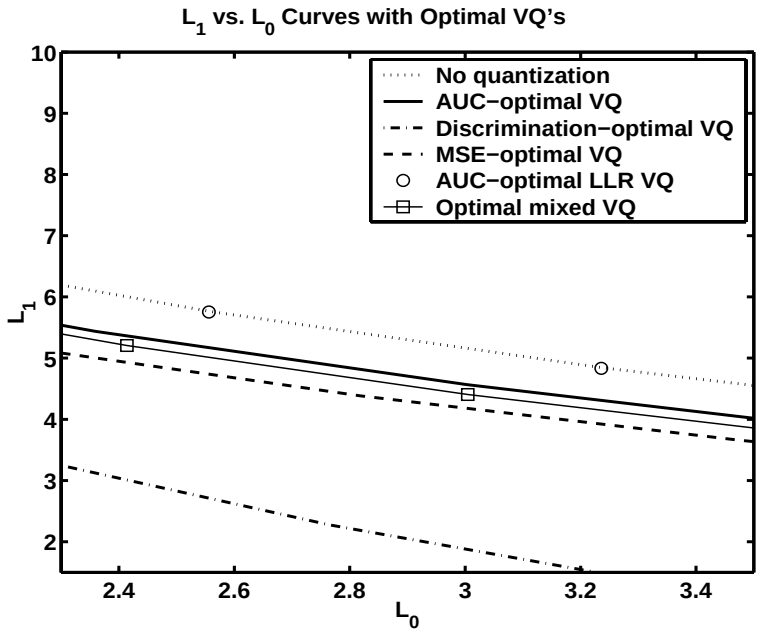

Figure 17: $L_1(L_0)$ curves with several 64-cell vector quantizers for two-dimensional uncorrelated Gaussian example.