QUANTIZATION IN TASK-DRIVEN SENSING AND DISTRIBUTED PROCESSING

Robert M. Gray

Electrical Engineering Department, Stanford University, Stanford, CA 94305

ABSTRACT

Quantization is the mapping of continuous quantities into discrete quantities, an operation far more general and flexible than the ubiquitous example of analog-to-digital conversion of scalar amplitude values. By appropriate choice of distortion measures and transmission constraints, quantization can incorporate signal processing such as statistical classification, estimation, and modeling. We here survey several approaches to incorporating such tasks into the quantization and possible extensions to distributed signal processing.

1. INTRODUCTION

Quantization is widely thought of as A/D conversion and its traditional role in sensing and signal processing has been to convert raw analog scalar measurements into a digital approximation. If the quantized signal is to be used for signal processing with the goal of making decisions, producing estimates, or constructing models of the original observed behavior, the performance of the signal processing can depend critically on the quality of the quantized reproduction of the original signal. Very high bitrates may be needed to provide sufficiently good signal approximations for the subsequent signal processing to perform well. Rigorous analysis of such systems is complicated by the fact that optimizing a quantizer with respect to mean squared error (MSE) does not easily relate to the resulting performance measures such as average Bayes risk in subsequent signal processing. Signal processing is often designed assuming it is operating on the original signal, when in fact it is operating on a quantized approximation. These issues assume added importance when many distributed quantizers are involved in a distributed sensor network. Not only can the approximation errors accumulate, the quantization itself can become inefficient in terms of required bit rates and energy budgets if the individual quantizers are designed without consideration of the overall goals of the system.

A variety of techniques have been developed in the information theory and statistical signal processing literature to handle task-driven quantization and distributed quantization, but both fields remain relatively underdeveloped in comparison with traditional models and little has been done to consider the two aspects together. Existing analyses often assume models which are doubtful fits to some fundamental aspects of real systems.

We survey and discuss several extensions and variations of traditional quantization relevant to task-driven quantization and existing and possible extensions to distributed signal processing.

2. QUANTIZATION FUNDAMENTALS

A *quantizer* consists of an *encoder* α which maps an input space A into in index set $\mathcal{I}$ and a *decoder* β which maps each the index set $\mathcal{I}$ into an output or reproduction space $\hat{A}$. Typically A consists of the real line (scalar quantization) or k-dimensional Euclidean space (vector quantization), but more generally it can be any well-behaved function space. The decoder is equivalent to the range space $\beta(\alpha(A)) = \mathcal{C}$, the *reproduction codebook*, and the encoder is equivalent to a partition $\mathcal{S} = \{S_i\}$, $S_i = \{x : \alpha(x) = i\}$ of the input space A. To quantify the performance of a codebook requires (at least) two cost functions: a *distortion measure* $d(x, y) \geq 0$, which quantifies the cost of reproducing an input x as y, and an *instantaneous rate* $r(i) \geq 0$, which assigns a transmission or memory cost to every index i. By far the most common distortion measure is simple MSE. The instantaneous rate measure can be thought of as the number of bits or nats required to store or losslessly communicate the index. The usual cases are the fixed-rate case where $r(i) = \ln |\mathcal{C}|$ nats, the log of the number of reproduction codewords, and the variable-rate case, where $r(i) = \ell(i)$, where ℓ is a *length function* satisfying the Kraft inequality $\sum_i e^{-\ell(i)} \leq 1$. In general, one might consider a Lagrangian combination $r(i) = (1-\eta)\ell(i) + \eta \ln |\mathcal{C}|$ in order to assign costs both to the bits for transmission and the memory required to hold the codebook.

If X is a random object (variable, vector, process, or field) described by a probability distribution P_X, then the average distortion is the corresponding expectation $E[d(X, \beta(\alpha(X)))]$ and the average rate is the expectation $E[r(\alpha(X))]$. Ideally one would like both of these quantities to be small, but the point of the theory and practice is to find the best tradeoff between the two. For both theory and design, it is useful to combine the cost measures into a single Lagrangian distortion and consider unconstrained minima. Toward this end we focus on

This work was supported by the National Science Foundation under NSF Grant CCR-0073050 and by Hewlett-Packard Laboratories.

the basic optimization problem as being the minimization of the average of the Lagrangian distortion $E[d(X,\beta(\alpha(X)))] + \lambda E[r(\alpha(X))]$ over all quantizers, possibly with added structural constraints, There are three major theoretical approaches to this optimization for random vectors. One is Shannon's rate-distortion theory, which provides lower bounds that are achievable in the asymptopia of large dimension and delay when the random vectors are drawn from a random process with suitable stationarity properties. A second approach is the theory of high rate quantization, which quantifies the optimal tradeoffs between rate and distortion in the asymptopia of large rate and small distortion (but fixed dimension). Lastly, there are nonasymptotic results which provide necessary conditions for quantizers to be optimal for general random objects. These have their origins in Lloyd's original optimality conditions for PCM and are easily summarized as optimality conditions for each component in terms of the other components:

Optimal Encoder $\alpha(x) = \mathrm{argmin}_i\,(d(x,\beta(i)) + \lambda r(i))$.
Optimal Decoder $\beta(i) = \mathrm{argmin}_y\, E(d(X,y|\alpha(X)=i)$.
Optimal Length Function $\ell(i) = -\ln \Pr(\alpha(X)=i)$.
Given a reproduction codebook $\mathcal{C}$ (decoder) or partition (encoder), the Lloyd properties determine the remaining components so optimizing over quantizers is equivalent to optimizing over reproduction codebooks or partitions. Hence we can add one more condition for pruning codebooks.
Codebook Pruning For a given reproduction codebook $\mathcal{C}$ a necessary condition for optimality is that there be no subcodebook $\mathcal{C}' \subset \mathcal{C}$ for which $\mathcal{C}'$ has strictly smaller average Lagrangian distortion than does $\mathcal{C}$, e.g., it must be true that $\Pr(\alpha(X)=i) \neq 0$ for $i \in \alpha(A)$.

Iterative application of these properties yields a Lloyd clustering algorithm for quantizer design.

3. QUANTIZATION FOR ESTIMATION

Suppose that the task is not to reproduce the observed random object X at the receiver, but instead to reproduce an unobserved random object Y which is correlated with X. The goal is now to encode the observed X into an index i, and then for the decoder to decode i as an estimate $\hat{Y}$. This is considered *estimation* or *regression* if Y is continuous and *statistical classification* or *detection* if Y is discrete. The average distortion for the system now takes the form $E[d(Y,\beta(\alpha(X)))]$, which might be a Bayes risk in general or MSE for estimation or probability of error for classification. One can think of this system as having an input Y, a sensor described by a conditional probability distribution $P_{X|Y}$, and an encoder which acts upon the observed X.

In the estimation case this problem is known as remote source coding or coding a noisy source and has a long history (see, e.g., [4]). If one uses MSE as the distortion measure, then the optimal estimator for Y given the true value of X is well known to be the conditional expectation $E[Y|X]$, which

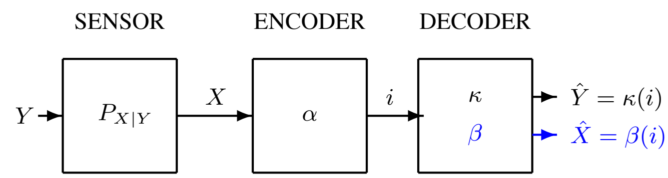

suggests a natural approximation — use the quantizer to encode X so that $\hat{X} = \beta(i)$ is designed to well approximate X, then form the estimator by inserting $\hat{X}$ in place of X in the conditional expectation, that is, $\hat{Y} = E[Y|\hat{X}]$. The technique is simple and intuitively it should work well if the approximation of X by $\hat{X}$ is good, e.g., if the quantizer has very high rate, but few careful theorems exist to justify the approximation and clearly it has problems outside the high rate regime.

The problem can be recast as a quantization problem with input X by modifying the distortion measure to measure the distortion between a source sample value x and the encoder output index i by $d(x,i) = E[C(Y,\beta(i))|X=x]$, where C is a Bayes risk cost function. The Lloyd algorithm immediately extends to this variation. The encoder picks the index to minimize the modified distortion and the decoder κ uses the centroid with respect to the modified distortion — the optimal nonlinear estimate of Y given the observed index [16]. In [4] it is shown that under certain conditions the optimum strategy is to first form an optimal nonlinear estimate of the unknown Y based on the observed X, and then optimally quantize the estimate using the usual MSE distortion. As a variation on this theme, one can have a decoder that produces both a reproduction of X and an estimate of Y from the received index and combine the two resulting distortions using an additional Lagrangian multiplier which allows one to weight the relative importance of the estimation task and the accuracy of the observed signal reproduction. Once again Lloyd has a natural extension and each decoder can be optimized for the remaining components.

This particular case makes two points that are common to many task-driven quantization systems: First, one needs to know the underlying distributions in order to apply the theory, in particular the conditional distribution $P_{Y|X}$ is required, and hence it must be estimated based on training data. Second, the quantizer is made task-specific by choosing a distortion measure that reflects the overall goal.

4. QUANTIZATION FOR CLASSIFICATION

Detection or classification based on quantized data also has a long history. Early techniques attempted to optimize indirectly by causing the quantizer to maximize an Aly-Silvey distance between the conditional distributions of the quantized data under the different hypotheses [14, 10, 13, 2]. The classifier performance can be directly incorporated as in the estimation task using a Bayes risk, either alone or in a Lagrangian combination with MSE on the observed signal [12]. As in the estimation case, this yields modified Lloyd conditions:

Optimal Classifier Bayes optimal classifier given index
$\operatorname{argmin}_k\{\sum_j C_{j,k}P(Y=j|\alpha(X)=i)\}$.

Optimal Encoder
$\operatorname{argmin}_i\{d(x,\beta(i))+\lambda\sum_{j=1}^M C_{j,\kappa(i)}P(Y=j|X=x)\}$.

5. QUANTIZATION FOR MODELING

In some applications the task is to construct a statistical model for the observed data, which might be used for synthesis or to choose an algorithm for further processing. One can think of quantizing observed random vectors into a codebook of models. One such approach is based on high rate quantization theory, a quantitative measure of the *mismatch* resulting from applying MSE quantization designed for one random vector to another, and the worst case property of Gaussian models for quantization. The development [9, 1] boils down to generating a reproduction codebook of Gaussian models, each consisting of a covariance matrix K_i, mean vector μ_i, and prior probability w_i. Taken together the collection is equivalent to a Gauss mixture model. The resulting distortion between an observed vector x and the encoder output index i is $d(x,i) = -\ln w_i + \frac{1}{2}\ln\left((2\pi)^k|K_i|\right) + \frac{1}{2}(x-\mu_m)^t K_i^{-1}(x-\mu_m)$. The distortion is a minimum discrimination information distortion and has its origins in the Itakura-Saito distortion of speech processing. The key points are that the distortion measure is a weighted quadratic distortion measure and hence has well defined Lloyd centroids and a computable minimum distortion encoder [9, 1] and there is no need to explicitly estimate class conditional probabilities. The Lloyd algorithm can be applied to design the quantizer which directly provides a Gauss mixture model for the observed data which can in turn be used to perform classification or as a component in a subsequent classification algorithm.

6. DISTRIBUTED PROCESSING

Consider a sensor network consisting of many nodes, each of which makes measurements on its environment and is capable of some signal processing and communication. We assume that all nodes make analog measurements, but have communication constraints on what they can transmit and hence must quantize. The nodes might or might not be able to receive information from neighboring nodes and incorporate it into their own signal processing. We also assume that there are central processing units which can receive information from many nodes and are permitted sophisticated signal processing in order to perform some overall task. Much of the literature ignores the quantization component and evaluates distributed estimation and classification assuming perfect measurements. As in the single user case, this provides an approximate approach by simply replacing a measurement by its quantized approximation, but it compounds the problems with such approximations.

Quantization enters into the system in several ways. The most obvious way is the quantization on the measurements at each node prior to the transmission of information to a central facility. Quantization also enters in the selection of node locations, that is, two or three dimensional space is quantized into a finite number of points corresponding to the location of the nodes. Much of the theory for sensor networks involves assumptions that the number of nodes is large, their locations dense, and there is a point density function describing the distribution, the same idea as the quantizer point density function of Lloyd and Gersho [7] for amplitude quantization. Intuitively, one might suspect that there should be a fundamental tradeoff between these two types of quantization. For example, if the rate of the spatial quantization is small and hence there are few nodes, then a higher rate will be required for the amplitude quantization at each node for a given level of performance at the central facility. Conversely, if the spatial quantization rate is high and the nodes are dense, then one would think that the individual nodes could have very low amplitude quantization rates and still provide good performance. An old (1985) and simple result along these lines showed that if a common random variable was viewed plus conditionally independent noise at each node, then binary quantizers suffice to achieve arbitrarily small estimation error at a central facility [8]. This result involved a simple application of the ergodic theorem, but it makes the point that many nodes with low rate simple scalar quantization can combine to produce a good estimator of a common underlying random variable.

In the naive design case, each node might behave as if it alone was sending information to the central facility and the discussion thus far describes how that quantization might be performed in a task-specific way to make best use of the node signal processing and the communication constraints. That approach, however, is clearly inefficient since it takes no advantage of the correlation of inputs and observations at the many nodes. A variety of approaches have evolved under the general headings of multiterminal information theory and distributed compression to treat such systems. In general the nodes might interract with each other, which leads to complicated formulations which to date have resulted in the solution of only a few very simple systems, including the classic Slepian-Wolf and Wyner-Ziv codes. Of particular interest in sensor networks with complexity and communication constraints is the case where the sensors do not receive information from other sensors or do their coding incorporating side information, but they are allowed to design their codes with knowledge of the statistical behavior of the other sensors in communication with a central facility — called “cooperative design, separate encoding” [11]. Extending the model of [8] one can assume a common random object Y and a collection of observations X_m corresponding to the measurements at separate sensors and described by conditional probability distributions $P_{X_m|Y}$. A variety of Shannon coding theorems and specific coding schemes have been developed for such systems under the assumption of conditional independence of the X_m given Y [6, 11, 17, 3, 15]. Such systems are sometimes referred to as the CEO problem [3] or multiple access coding [5]. The assumption of conditionally independent noise is dubious if the nodes are assumed dense unless one assumes the noise is entirely circuit noise in the nodes themselves. Traditional Shannon theoretic results have the drawback of assuming that sequences of measurements can be grouped as large dimensional vectors for compression, adding to complexity and delay. Specific modest complexity coding schemes have been considered, e.g., in [6, 11, 15, 5]. In [15] attention was also paid to channel coding to allow for the well known failure of the source/channel coding separation in networks. As a simple example of cooperation, both the index reuse techniques of [6, 11] and the syndrome codes of [15] make use of the fact that if one has separate MSE quantizers at, say, two nodes, then one can merge quantizer cells at both nodes to produce lower rate quantizers at each node and then disambiguate the pairs of merged cells at the decoder by picking the most likely pair of high rate input cells given the received pair of indexes. Lloyd-algorithms might be able to find good structured and simple codes of this variety.

All of these approaches attempt to recover the common underlying random object seen by the many sensors. If the problem is really one of classification, however, task-driven quantization can produce lower rate bitstreams at individual nodes and hence ameliorate the problem of efficiency lost by lack of cooperation. This lower bit rate also means that simple channel coding or robust quantizers can more easily provide reliable communication. A central processing facility can then pool the outcomes using well established statistical techniques for an overall decision. For this reason single-user task-specific quantization may play a useful role in distributed sensor node applications where statistical classification is the overall goal.

7. REFERENCES

[1] A. Aiyer, K. Pyun, Y.-Z. Huang, D. B. O'Brien, and R. M. Gray, "Lloyd Clustering of Gauss Mixture Models for Image Compression and Classification," *Signal Processing: Image Communication*, pp. 459–485, 2005.

[2] G. R. Benitz and J. A. Bucklew, "Asymptotically optimal quantizers for detection of i.i.d. data," *IEEE Trans. Inform. Theory*, pp. 316–325, 1989.

[3] T. Berger, Z. Zhang, and H. Viswanathan, "The CEO Problem," *IEEE Transactions on Information Theory*, Vol. 42, pp. 887–902, May 1996.

[4] Y. Ephraim and R.M. Gray, "A unified approach for encoding clean and noisy sources by means of waveform and autoregressive model vector quantization," *IEEE Trans. on Inform. Theory*, pp. 826–834, July 1988.

[5] M. Fleming, Q. Zhao, and M. Effros, "Network vector quantization," *IEEE Trans. on Inform. Theory*, 50(8), pp. 1584–1604, August 2004.

[6] T.J. Flynn and R.M. Gray, "Encoding of correlated observations," *IEEE Transactions on Information Theory*, pp. 773–787, November 1987.

[7] A. Gersho, "Asymptotically optimal block quantization," *IEEE Trans. Inform. Theory.*, pp. 373–380, July 1979.

[8] R.M. Gray, S. Boyd, and T. Lookabaugh, "Low rate distributed quantization of noisy sources," *Proceedings of the 23rd Annual Allerton Conference on Communication, Control, and Computing*, Allerton House, Monticello, Illinois, pp. 354–358, October 1985.

[9] R.M. Gray and T. Linder, "Mismatch in high rate entropy constrained vector quantization," Vol. 49, pp. 1204–1217g, *IEEE Trans. Inform. Theory*, May, 2003.

[10] S. A. Kassam, "Optimum quantization for signal detection," *IEEE Trans. Comm.*, pp. 479–484, May 1977.

[11] M. Longo, T. Lookabaugh, and R. M. Gray, "Quantization for decentralized hypothesis testing under communication constraints," *IEEE Transactions on Information Theory*, pp. 241–255, March 1990.

[12] K.L. Oehler and R.M. Gray, "Combining image compression and classification using vector quantization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 461–473, May 1995.

[13] B. Picinbono and P. Duvaut, "Optimum quantization for detection," *IEEE Trans. Comm.*, pp. 1254–1258, Nov. 1988.

[14] H. V. Poor and J. B. Thomas, "Applications of Ali-Silvey distance measures in the design of generalized quantizers for binary decision systems," *IEEE Trans. Comm.*, pp. 893–900, September 1977.

[15] S. S. Pradhan, J. Kusuma, and K. Ramchandran, K., "Distributed compression in a dense microsensor network," *IEEE Signal Processing Magazine*, pp. 51–60, March 2002.

[16] A. Rao, D. Miller, K. Rose, A. Gersho, "Generalized vector quantization: jointly optimal quantization and estimation" *Proceedings 1995 IEEE International Symposium on Information Theory*, p. 432, Sept. 1995.

[17] H. Yamamoto and K. Itoh, "Source coding theory for multiterminal communication systems with a remote source," *Trans on IECE Japan*, pp. 700–706, Oct 1980.