Multi-Resolution VQ: Parameter Meaning and Choice *

Diego Dugatkin
California Institute of Technology
Department of Electrical Engineering
Pasadena, California. 91125. USA
diego@caltech.edu

Michelle Effros
California Institute of Technology
Department of Electrical Engineering
Pasadena, California. 91125. USA
effros@caltech.edu

Abstract

In multi-resolution source coding, a single code is used to give an embedded data description that may be decoded at a variety of rates. Recent work in practical multi-resolution coding treats the optimal design of fixed- and variable-rate tree-structured vector quantizers for multi-resolution coding. In that work, the codes are optimized for a designer-specified priority schedule over the system rates, distortions, or slopes. The method relies on a collection of parameters, which may be difficult to choose. This paper explores the meaning and choice of the multi-resolution source coding parameters.

1. Introduction

Multi-resolution source codes, also called progressive transmission or successive refinement codes, are compression algorithms in which simple, low-rate source descriptions are embedded in more complex, high-rate descriptions. In a multi-resolution source code, a single code is used to describe a source at a variety of rates. Decoding only the initial segment of the coded bit stream leads to a low-resolution data reconstruction, decoding a greater portion leads to a higher-resolution reconstruction, and so on. Use of multi-resolution codes allows flexibility in systems catering to a variety of users or requiring a variety of rates since one multi-resolution code or description can do the job of many single-resolution codes or descriptions. The result is greater complexity and storage efficiencies. Examples of systems that can benefit from multi-resolution codes include mobile communications, web-based communications, and applications where users uncertain of their precision needs may benefit from the ability to stop data transmission when the desired data reconstruction accuracy is achieved.

Interest in multi-resolution or progressive transmission source coding has inspired an enormous amount of research into practical multi-resolution source coding algorithms. We here focus on [4], which describes a design algorithm and encoder for optimal multi-resolution vector quantizers (MRVQ) based on the tree-structured vector quantization (TSVQ) algorithm (see, for example, [10]). The algorithm generalizes the algorithms of [2, 11] from scalar to vector quantizers. MRVQ relies on a Lagrangian optimization matched to a priority schedule over the multi-resolution code's rates, distortions, or slopes. Since [4] gives little insight on choosing the Lagrangian parameters, we here explore their meaning and choice, demonstrating algorithmic performance when parameters are appropriately chosen.

2. Multi-Resolution Vector Quantization

The MRVQ algorithm replaces the greedy encoder of TSVQ with an optimal encoder. The multi-resolution encoder chooses, from among all possible paths through a tree-structured codebook, the path that minimizes the weighted performance measure $\sum_{\ell=1}^{L}[\alpha_\ell D_\ell+\beta_\ell R_\ell]$, where $\{D_\ell\}$ and $\{R_\ell\}$ are the distortions and rates respectively associated with each resolution of an L-resolution code. The code is designed using an iterative descent technique functionally equivalent to the generalized Lloyd algorithm. As a result of the above definition of the optimal encoder, this iterative descent technique yields a non-greedy design algorithm that simultaneously optimizes all levels of a tree-structured codebook [4].

To design an MRVQ, one must choose the encoder's Lagrangian parameters (α^L, β^L). As discussed in [4], these parameters should match a user-defined priority schedule $p_1, p_2, \ldots, p_L$ on the code's L resolutions. This priority schedule may be derived as a function of a set of priorities or probabilities over the code's resolutions, and may be posed in terms of code rates, distortions, or rate-distortion trade-

*This material is based upon work partially supported by the Oringer Fellowship, by NSF Grant No. MIP-9501977, and by donations through the Intel 2000 Program.

offs (slopes). For example, consider a system in which users decode at rate r_ℓ with probability p_ℓ. Then the goal is to design the system that yields the lowest expected distortion $\sum_\ell p_\ell D_\ell$ subject to a collection of constraints $R_\ell \le r_\ell$ on the rates, or equivalently to minimize $\sum_\ell [p_\ell D_\ell + \beta_\ell R_\ell]$, where $\{\beta_\ell\}$ are the Lagrangian parameters associated with our collections of rate constraints. Thus in theory, the system parameters are simply chosen to match the given design criteria.

When L is large, finding the appropriate Lagrangian parameters ($\{\beta_\ell\}$ in the above example) may be difficult since these parameters are interdependent. Further, practical design criteria are typically far more vague than in the above example, perhaps stating that most users tend to favor a specific range of rates, but not giving explicit values for either the priority schedule or the rates used. Thus, in practice, choice of (α^L, β^L) may be quite difficult.

3. Parameter Meaning and Choice

For the single-resolution case, $\alpha D + \beta R = \alpha[D + (\beta/\alpha)R]$, where β/α coresponds to λ, the negative slope of the line tangent to the rate-distortion curve $R(D)$ at the target rate and distortion, as noted in [4].

In the two-resolution case $L = 2$, we have four parameters to choose. Each 4-dimensional vector $(\alpha_1, \alpha_2, \beta_1, \beta_2)$ describes the direction of the tangential hyperplane supporting the space of achievable (R_1, D_1, R_2, D_2) points. Since only the relative values of these parameters matter, we will focus on choosing β_1/α_1, β_2/α_2 and α_1/α_2. If we fix R_2, D_2 or R_1, D_1 in this 4-dimensional space, we obtain the 2-dimensional "slices" $D_1 = f_1(R_1)$ and $D2 = f_2(R_2)$ respectively. At the tangent point, the equivalent slopes are $-\beta_1/\alpha_1$ and $-\beta_2/\alpha_2$ respectively, where $\beta_1/\alpha_1 \ge \beta_2/\alpha_2$.

Notice that the ratio α_1/α_2 specifies the priority schedule for resolutions one and two. If the ratio is infinity, the resolution-1 code will effectively be an ECVQ with $\lambda = \beta_1/\alpha_1$, while the resolution-2 code will be the best slope-β_2/α_2 code achievable given the fixed resolution-1 code. For $\alpha_1/\alpha_2 = 0$, the second code will be optimized for slope β_2/α_2 and then the resolution-1 code will be optimized for slope β_1/α_1 subject to the constraint on the resolution-2 code. Between these extremes lies a continuum of possibilities.

This physical interpretation can be extended to the 2L-dimensional vector (α^L, β^L), that describes the angle of the hyperplane tangent to the achievable rate-distortion region at the target L-dimensional rate-distortion vector.

We propose the following approach for estimating (α^L, β^L).

First re-arrange the multi-resolution performing measure: $\sum_{\ell=1}^{L} \alpha_\ell [D_\ell + (\beta_\ell/\alpha_\ell) R_\ell]$. Next, choose α_ℓ/α_k ratios to match the relative priorities of the ℓ^{th} and k^{th} resolutions. Then estimate β_ℓ/α_ℓ for a set of rates covering the range of interest. This can be done using the single-resolution methods of [3]. The next step is to set β_ℓ and β_k to obtain the appropriate ratios β_ℓ/α_ℓ and β_k/α_k, that are equivalent to the desired slopes.

4. Results

We conducted experiments on both synthetic and natural data sets.

The synthetic data follows the distribution described by Equitz and Cover in [7, 8] (and also [1], [6], [9]).

For the natural data experiments, we used a training set of 20 medical brain-scan images, and a test set of 5 brain scans. The training and test sets do not overlap. Figures 3 and 4 show a variety of 9-resolution experiments, in comparison with VQ and TSVQ.

In all the experiments on both synthetic and natural data, the first-resolution codes achieve the same performance of the VQ codes. By choosing a very high priority for the first-resolution, the MRVQ first-resolution is "tied-down" to the VQ curve at that point.

MRVQ is exponentially more complex than TSVQ, since it involves full-search of the best-of-all possible paths of the given tree, at each resolution. MRVQ offers the multi-resolution benefits at the associated cos of doubling the complexity required for full-search VQ.

4.1. Synthetic Data

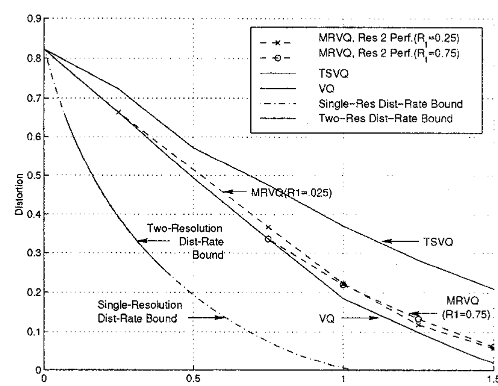

Figure 1. Second-resolution performance of multi-resolution TSVQ (for a first-resolution code fixed at rates 0.25 and 0.75), TSVQ, and VQ rate-distortion performance.

Figures 1 and 2 show results on a synthetically generated data set. The example chosen uses the source alphabet $A =$

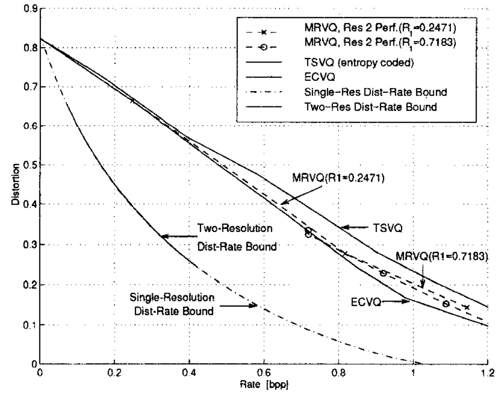

Figure 2. Second-resolution performance of multi-resolution TSVQ (for a first-resolution code fixed at rates 0.2471 and 0.7183), TSVQ, and ECVQ rate-distortion performance.

$\{1, 2, 3\}$, with symbol probabilities $\mu = \{(1-p)/2, p, (1-p)/2\}$, and absolute difference distortion measure $\rho(x, \hat{x}) = |x - \hat{x}|$, introduced by Gerrish [9], and used by Equitz and Cover in their original proof that the optimal performance of a two-resolution source code does not everywhere coincide with the (single-resolution) rate-distortion bound [7, 8] (see also [1, 5, 6]). The plots include both theoretical bounds and results achieved through code design.

In each figure, the lowest curve shows the distortion-rate bound for the given source. Just above (and almost indistinguishable from the the distortion-rate curve) is a curve showing the optimal second-resolution performance theoretically achievable using a first-resolution code whose performance lies exactly on the distortion-rate bound at rates lower than 0.069 bpp and higher than 0.405 bpp. The optimal second-resolution performance and the distortion-rate curve overlap exactly except in the section plotted with a solid line, where the two curves separate.

A variety of codes were designed for the given source and their performances are compared both with each other and with the theoretical results in the same figures. All codes considered are vector quantizers of dimension 4. The lower curve corresponds to the performance of a sequence of VQs; the higher curve shows the multi-resolution performance of a single TSVQ. The curves in between give performance results for MRVQ. Each MRVQ curve shows the range of possible second-resolution rate-distortion performances achievable given that its first-resolution is a fixed, optimal VQ. Notice that the achieved coding results again demonstrate a small degradation in performance associated with going from single-resolution coding to multi-resolution coding. However, the MRVQ offers much better performance than TSVQ, for the range of rates studied here.

For higher dimensions, the experimental results should move closer to the bounds predicted by the theory. Just as vector quantizers are optimal block single-resolution source codes, multi-resolution vector quantizers are asymptotically optimal multi-resolution source codes [5, 6].

4.2. Natural Data

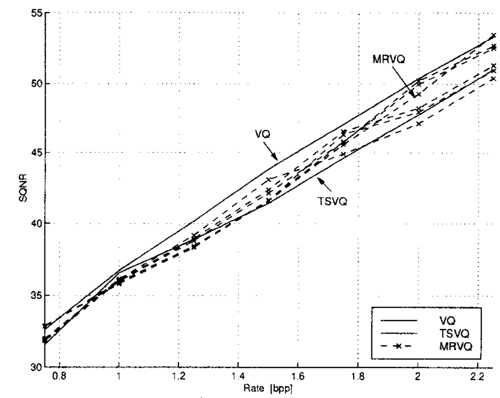

Figure 3. SQNR vs. rate results for fixed-rate MRVQ using the weighted performance measure (dashed lines), non-embedded VQ (upper curve) and TSVQ (lower curve).

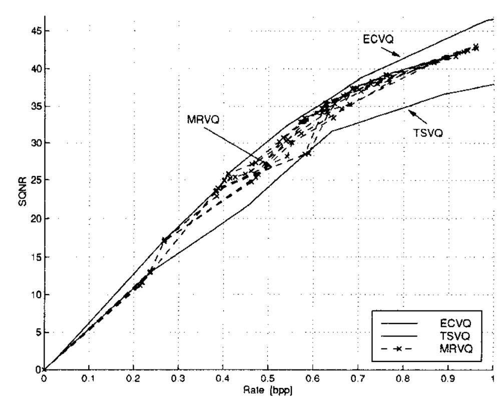

Figure 4. SQNR vs. rate results for variable-rate MRVQ using the weighted performance measure (dashed lines), non-embedded ECVQ (upper curve) and TSVQ (lower curve).

In Figure 3, the lowest curve shows the performance of TSVQ while the highest curve is the performance of a (non-embedded) collection of full search VQs. The lowest curve in Figure 4 shows the results for variable-rate entropy-coded TSVQ, and highest curve describes the performance of ECVQs. In both cases, the curves in between show the performance results for MRVQ, with a range of (α^L, β^L) vectors. Each dashed line shows the rates and signal to quantization noise ratios (SQNRs) associated with the above described code for a single set of (α^L, β^L) values. The (α^L, β^L) sets of values used in these experiments were chosen to fix a desired resolution $\{R_\ell\}$,$\{D_\ell\}$ to be as close as possible to the full search VQ or ECVQ performance by choosing a relatively very high α_ℓ in comparison with the other parameters. An intermediate priority α_k was set to achieve the best possible secondary performance, and all higher rate points were given a very low (but non-zero) remaining priorities. The results show that at the high-priority fixed-resolution point, it effectively achieves the ECVQ performance, and the performance at second-priority resolutions is generally better than TSVQ and less than 1 dB below ECVQ.

All codes use vector dimension 4. The distortion is measured as squared error $\rho(x, \hat{x}) = (x - \hat{x})^2$.

5. Conclusions

We have explored the meaning and choice of the MRVQ parameters, offering an interpretation of them, and showing practical algorithmic performance for both fixed- and variable-rate experiments.

MRVQ performance has been compared to both TSVQ and a (non-embedded) collection of full-search VQs (and ECVQs for variable rate). MRVQ effectively starts to close the gap between TSVQ and ECVQ.

References

[1] T. Berger. *Rate distortion theory. A mathematical basis for data compression.* Prentice Hall, Englewood Cliffs, N.J., 1971.

[2] H. Brunk and N. Farvardin. Fixed-rate successively refinable scalar quantizers. In *Proceedings of the Data Compression Conference*, pages 250–259, Snowbird, UT, Apr. 1996. IEEE.

[3] P. A. Chou, T. Lookabaugh, and R. M. Gray. Entropy-constrained vector quantization. *IEEE Transactions on Acoustics Speech and Signal Processing*, 37(1):31–42, Jan. 1989.

[4] M. Effros. Practical multi-resolution source coding: TSVQ revisited. In *Proceedings of the Data Compression Conference*, pages 53–62, Snowbird, UT, Mar. 1998. IEEE.

[5] M. Effros. Rate-distortion bounds for fixed- and variable-rate multi-resolution source codes. Mar. 1998. Submitted to IEEE Transactions on Information Theory.

[6] M. Effros and D. Dugatkin. Multi-resolution vector quantization. 1998. In preparation.

[7] W. H. R. Equitz. *Succesive refinement of information.* Ph. D. Dissertation, Stanford University, Stanford, CA, 1989.

[8] W. H. R. Equitz and T. M. Cover. Succesive refinement of information. *IEEE Transactions on Information Theory*, 37(2):269–275, Mar. 1991.

[9] A. M. Gerrish. *Estimation of information rates.* Ph. D. Dissertation, Yale University, New Haven, CT, 1963.

[10] A. Gersho and R. M. Gray. *Vector quantization and signal compression.* Kluwer Academic Publishers, Boston, 1992.

[11] H. Jafarkhani, H. Brunk, and N. Farvardin. Entropy-constrained successively refinable scalar quantization. In *Proceedings of the Data Compression Conference*, pages 337–346, Snowbird, UT, Mar. 1997. IEEE.