

ESTIMATION ERROR BOUNDS FOR DENOISING BY SPARSE APPROXIMATION

Alyson K. Fletcher

University of California, Berkeley
alyson@eecs.berkeley.edu

Kannan Ramchandran

University of California, Berkeley
kannanr@eecs.berkeley.edu

ABSTRACT

If a signal is known to have a sparse representation with respect to a given frame, the signal can be estimated from a noise-corrupted observation of the signal by finding the best sparse approximation to the observation. The ability to remove noise in this manner depends on the frame being designed to efficiently represent the signal while it *inefficiently* represents the noise. This paper gives bounds to show how inefficiently white Gaussian noise is represented by sparse linear combinations of frame vectors. Combined with knowledge of the approximation efficiency of a given family of frames for a given signal class, this work leads to a better understanding of the merits of frame denoising.

1. DENOISING BY SPARSE APPROXIMATION

Consider the problem of estimating an unknown signal $x \in \mathbb{R}^N$ from the noisy observation $y = x + d$ where $d \in \mathbb{R}^N$ has the $\mathcal{N}(0, \sigma^2 I_N)$ distribution. If x is known to lie in a given K -dimensional subspace of \mathbb{R}^N , the situation can immediately be improved by projecting y to the given subspace; since the noise distribution is spherically symmetric this leaves only K/N fraction of the original noise. Further information about the distribution of x could be exploited to remove even more noise.

In many scenarios in which one attempts to estimate x from y , to know with certainty the K -dimensional subspace that contains x is too much to ask. However, it may be reasonable to know a set of K -dimensional subspaces such that the union contains or can closely approximate x . In other words, a basis may be known such that a K -term nonlinear approximation is very accurate [1]. (Wavelet bases have this property for piecewise smooth signals.) At high SNR (small σ^2), the selection of the K -dimensional subspace is unperturbed by d . The approximation error is approximately $K\sigma^2 + f_N(K)$, where $f_N(K)$ is a decreasing function of K that represents the error of the optimal K -term nonlinear approximation of x .

Now suppose $U = \{u_1, \dots, u_M\} \subset \mathbb{R}^N$ is chosen so that the class of signals of interest can be sparsely repre-

sented. I.e.,

$$x \approx \sum_{i=1}^K \alpha_i u_{k_i} \quad (1)$$

for some scalars $\{\alpha_i\}$ gives approximation error $f_M(K)$ where K/N is small. This paper analyzes the denoising of y by finding the best K -term approximation with respect to U .

As M is increased, the approximation quality in (1) can improve, i.e., $f_M(K)$ should decrease, or the required value of K can decrease. This seems to justify large values of M so as to reduce $K\sigma^2 + f_M(K)$. However, increasing M can be problematic. If M is too large, then the subspace chosen based on y will likely not contain x .

Fortunately, there are many classes of signals for which M need not grow too quickly as a function of N to get good sparse approximations. The design of sets U for this problem is essentially the same as designing dictionaries for matching pursuit [2]. Examples for audio with good computational properties are given in [3]. See also [4] for video compression and [5] for an iterative design procedure.

The main result of this paper, Theorem 2, bounds the approximation error as a function of M , N and K when the signal has an exact K -term approximation. This approximation error bound can be balanced against the sparse approximation performance of a particular family of frames to assess whether redundancy is helpful for denoising.

1.1. Related work

Denoising by finding a sparse approximation is similar to the concept of denoising by compression popularized by Saito [6] and Natarajan [7]. More recent work in this area includes [8, 9, 10]. None of these works use representations with respect to frames. To put the present work into a similar framework would require a “rate” penalty for redundancy. Instead, the only penalty for redundancy comes from choosing a subspace that does not contain the true signal (“overfitting” or “fitting the noise”).

2. ERROR ANALYSIS

The set of vectors comprising the dictionary is denoted U . Typically the overcomplete case where M , the cardinality of U , is greater than N , the dimension of the signal, is of interest. A spanning set of cardinality N or higher is called a *frame*. In the following analysis, it is also possible for the dictionary to be undercomplete, i.e. $M < N$. The model for the signal x is that x is known to be some linear combination of K of the M frame vectors. Which of the K vectors to select and the K coefficients of the linear combination must be estimated. Furthermore, it is always assumed that $K \leq N$ because otherwise the existence of a K -term approximation for x is senseless.

One could use any of several techniques to form the estimate \hat{x} from y . Without introducing a probability distribution for x , the maximum likelihood estimate over the noise d is considered. This yields

$$\hat{x}_{ML}(y) = \operatorname{argmin}_{x \in X} \|y - x\|, \quad (2)$$

where X is the set of all x spanned by K of the M frame vectors in U . In this scenario, the error of the ML estimate provides a *lower* bound on the average error achievable by any other estimator. Many methods used in practice attempt to emulate the ML estimator.

It will later be convenient to see how the ML estimate can be described in terms of projections. Given a frame U , let \mathcal{P}_K be the set of all projections onto spaces spanned by K of the M frame vectors in U . Then \mathcal{P}_K has $\binom{M}{K}$ elements, and the ML estimate is given by

$$\hat{x}_{ML} = \hat{P}y, \quad (3)$$

where

$$\hat{P} = \operatorname{argmax}_{P \in \mathcal{P}_K} \|Py\|.$$

Now the error of the ML estimator is analyzed. The error will generally depend on the true signal x . Define the conditional per-component MSE

$$E(x) = \frac{1}{N} \mathbf{E} (\|x - \hat{x}_{ML}\|^2 \mid x).$$

When $M = K$, there is no issue of selecting the subspace on which to project. In this case the ML estimate \hat{x}_{ML} is the projection of the noisy signal onto a fixed K -dimensional subspace which is equivalent to standard least squares (LS) denoising. The estimation error then does not depend on x . The error reduces by a factor K/N ,

$$E(x) = \sigma^2 \frac{K}{N} \quad \forall x.$$

For $M > K$, recall that \hat{x}_{ML} is the projection of y to the K -dimensional subspace—among the $\binom{M}{K}$ possible ones—with the most energy. When x and d are independent, there

is no reason for this subspace to have less than K/N fraction of the noise energy. It may happen, however, that the selection of the subspace is “wrong,” resulting in estimation error higher than $(K/N)\sigma^2$. This is improbable when the noise energy is very low, but it becomes more likely as M increases, which is the typical problem in model overfitting. The first theorem formalizes the intuitive fact that with high input SNR, the subspace selection becomes independent of the noise. Thus frame denoising is able to achieve the same MSE as if the correct subspace were known.

Theorem 1 Consider the frame denoising problem and maximum likelihood estimator described above. Suppose the true signal x belongs to only one of the subspaces spanned by K of the M frame vectors. Then as $\sigma \rightarrow 0$,

$$E(x) = \sigma^2 \frac{K}{N}. \quad (4)$$

To prove the theorem, observe that for σ sufficiently small, the probability of projecting to the wrong subspace is small. Therefore, with high probability, the ML estimator will project to a fixed K -dimensional subspace containing x . The estimation error is then similar to the case when $M = K$.

The second theorem addresses the following basic question: How does estimation error increase with M ?

Theorem 2 Consider the frame denoising problem and maximum likelihood estimator described above. Then for large N and K/N small,

$$E(x) \leq E(0) \leq \sigma^2 \left[1 - \left(\frac{M}{K}\right)^{-2/N} \left(1 - \frac{K}{N}\right) \right]. \quad (5)$$

Theorem 2 provides a simple expression relating the error to the dictionary size M , the dimension of the signal N , and the dimension of the signal model K . Further comments:

- It is intuitive for $E(x)$ to be bounded by $E(0)$ because for low SNR the subspace selection captures as much of the noise as possible.
- When $M = K$, the bound in (5) reduces to

$$E(0) \leq K\sigma^2/N,$$

which is tight.

- For $M \rightarrow \infty$, $E(0) \rightarrow \sigma^2$. That is, with sufficiently high M , ML estimator may not provide any denoising.
- Define a sparsity measure $\alpha = K/N$ and a redundancy factor $\rho = M/N$. Then for large N , the bound

reduces to

$$E(0) \leq \sigma^2 \left[1 - \left(\frac{\alpha}{\rho} \right)^\alpha \left(1 - \frac{\alpha}{\rho} \right)^{\rho-\alpha} (1 - \alpha) \right]. \quad (6)$$

Equation (6) shows that the MSE increases with ρ^α . With a good signal model, α is small. Therefore, the MSE is not greatly increased with high redundancy ρ .

3. LOW SNR BOUND ANALYSIS

The bound in Theorem 2 is not, in general, tight: the inequality in the theorem provides only an upper bound on the error. The actual error will depend on the specific frame.

To evaluate the tightness of the bound, we compare to the actual error obtained using random frames. Specifically, we first fix some value for N . Then, for various values of M , a random frame is generated by generating a random $M \times N$ matrix A and then finding U , the $M \times N$ matrix whose columns are the M dominant orthonormal eigenvectors of AA' . The matrix U can be found by a singular value decomposition of A . The M frame vectors are then given by the M rows of U .

With each random frame, we find the best $K = 1$ dimensional approximation, \hat{d} , to a random N -dimensional vector, d , whose components are i.i.d., Gaussian with zero-mean and variance $\sigma^2 = 1$. Using 100 random frames, we estimate the average value $\mathbf{E}\|\hat{d}\|^2$. Fig. 1 plots the ratio

$$-10 \log_{10} \frac{\mathbf{E}\|\hat{d}\|^2}{N\sigma^2} \quad (7)$$

as a function of the frame sizes M for signal lengths $N = 10$ and $N = 100$. In the frame denoising problem above, the quantity in (7) would represent the noise reduction if the true signal is $x = 0$. Also plotted in Fig. 1 is the theoretical lower bound from Theorem 2. Fig. 2 shows the bound (6) that holds for large N .

4. PROOF OF THEOREM 2

Let $J = \binom{M}{K}$. There are at most J possible projections in \mathcal{P}_K . Index the projections by P_j $j = 1, \dots, J$, and let R_j be the corresponding range spaces.

We need only to consider $x = 0$. Therefore $y = d$ and

$$\hat{x}_{ML} = P_T d, \text{ where } T = \underset{j}{\operatorname{argmax}} \|P_j d\|. \quad (8)$$

Here, T is a random variable index of the space where d has the most energy. Now, for any value of T , P_T is a projection so,

$$\|P_T d\|^2 = \|d\|^2 - \|d - P_T d\|^2.$$

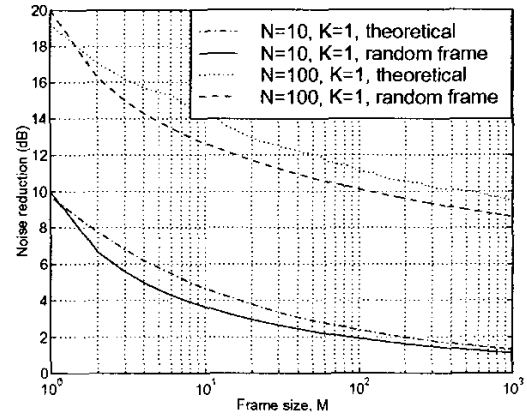


Fig. 1. Comparison of theoretical noise reduction lower bound to actual noise reduction with random frame.

Using this fact along with the fact that $\mathbf{E}\|d\|^2 = N\sigma^2$, we have

$$\begin{aligned} E(0) &= \frac{1}{N} \mathbf{E}\|\hat{x}_{ML} - x\|^2 = \frac{1}{N} \mathbf{E}\|P_T d\|^2 \\ &= \frac{1}{N} \mathbf{E}[\|d\|^2 - \|d - P_T d\|^2] = \sigma^2 - \frac{1}{N} D \end{aligned}$$

where

$$D = \mathbf{E}\|d - P_T d\|^2. \quad (9)$$

Therefore, to bound $E(0)$ above, we need to bound D below. We estimate a lower bound on D with the following rate-distortion argument.

Fix any $n > 0$. For each $j = 1, \dots, J$, let Q_j be an optimal n -bit quantizer of $P_T d$ conditional on $T = j$. Define the quantizer Q by

$$Q(d) = Q_T(P_T(d)).$$

That is, Q projects d to the K -dimensional space with the greatest energy, and then quantizes the projection of d within that space.

Now, for each j , the quantizer Q_j quantizes points in R_j , the range space of P_j . We can assume that for all $z \in R_j$, $Q_j(z) \in R_j$. Therefore, for all d , $P_j(d) - Q_j(P_j(d)) \in R_j$, and hence $P_j(d) - Q_j(P_j(d))$ is orthogonal to $d - P_T(d)$. Therefore,

$$\|d - Q(d)\|^2 = \|d - P_T(d)\|^2 + \|P_T(d) - Q_T(P_T(d))\|^2.$$

Using the definition of D in (9),

$$D = \mathbf{E}\|d - Q(d)\|^2 - \mathbf{E}\|P_T(d) - Q_T(P_T(d))\|^2. \quad (10)$$

We now bound the two terms on the right hand side of (10).

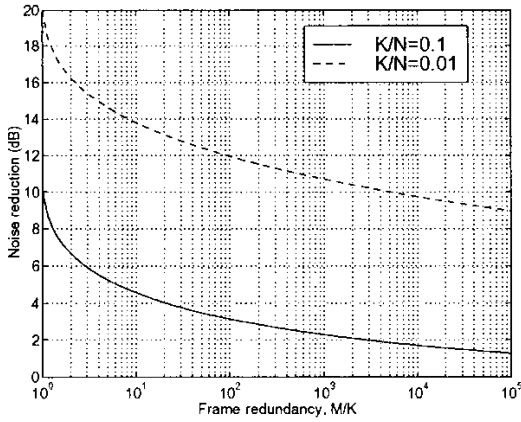


Fig. 2. Theoretical lower bound on the zero-signal noise reduction for large values of N

For the first term, the quantized point $Q(d)$ can be described by $\log_2 J$ bits to quantize the index T along with n bits for the point $Q_T(P_T(d))$. Therefore, $Q(d)$ can be described by a total of $n + \log_2 J$ bits. The term $\mathbf{E}\|d - Q(d)\|^2$ is the average distortion of the quantizer Q on the source d . Since d is an N -dimensional Gaussian variable with variance $\sigma^2 I$, the distortion is bounded below by the distortion-rate bound,

$$\mathbf{E}\|d - Q(d)\|^2 \geq N\sigma^2 2^{-2(n+\log_2 J)/N}. \quad (11)$$

For the second term on the right hand side of (10), let

$$\sigma_j^2 = \mathbf{E}(\|P_j(d)\|^2 | T = j).$$

Now, in general, the distribution of $P_j(d)$ conditional on $T = j$ is *not* Gaussian. However, the distortion achievable for any distribution is always less than or equal to the minimum distortion for a Gaussian white source with the same total variance. Consequently, for every j , the quantizer Q_j can attain a distortion for a K -dimensional Gaussian white source with total variance σ_j^2 . Therefore,

$$\mathbf{E}(\|P_j(d) - Q_j(P_j(d))\|^2 | T = j) \leq \sigma_j^2 2^{-2n/K}. \quad (12)$$

Also, observe that

$$\mathbf{E}\sigma_T^2 = \mathbf{E}\|P_T(d)\|^2 = \mathbf{E}\|d\|^2 - \mathbf{E}\|d - P_T(d)\|^2 = N\sigma^2 - D. \quad (13)$$

Substituting (13) in (12),

$$\mathbf{E}(\|P_T(d) - Q_T(P_T(d))\|^2) \leq (N\sigma^2 - D)2^{-2n/K}. \quad (14)$$

Then, substituting (11) and (14) in (10) we obtain

$$D \geq N\sigma^2 \left(\frac{2^{-2(n+\log_2 J)/N} - 2^{-2n/K}}{1 + 2^{-2n/K}} \right).$$

This bound must be true for all n , so it can be maximized over n . This maximization, under the assumption that N is large and K/N is small, gives the desired result.

5. REFERENCES

- [1] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies. Data compression and harmonic analysis. *IEEE Trans. Inform. Th.*, 44(6):2435–2476, October 1998.
- [2] S. G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Trans. Signal Proc.*, 41(12):3397–3415, December 1993.
- [3] M. M. Goodwin. *Adaptive Signal Models: Theory, Algorithms and Audio Applications*. Kluwer Acad. Pub., 1998.
- [4] F. Moschetti, L. Granai, P. Vandergheynst, and P. Frossard. New dictionary and fast atom searching method for matching pursuit representation of displaced frame difference. In *Proc. IEEE Int. Conf. Image Proc.*, volume 3, pages 685–688, Rochester, NY, September 2002.
- [5] K. Engan, S. O. Aase, and J. H. Husøy. Designing frames for matching pursuit algorithms. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Proc.*, volume 3, pages 1817–1820, Seattle, WA, May 1998.
- [6] N. Saito. Simultaneous noise suppression and signal compression using a library of orthonormal bases and the minimum description length criterion. In E. Foufoula-Georgiou and P. Kumar, editors, *Wavelets in Geophysics*, pages 299–324. Academic Press, San Diego, CA, 1994.
- [7] B. K. Natarajan. Filtering random noise from deterministic signals via data compression. *IEEE Trans. Signal Proc.*, 43(11):2595–2605, November 1995.
- [8] H. Krim, D. Tucker, S. Mallat, and D. Donoho. On denoising and best signal representation. *IEEE Trans. Inform. Th.*, 45(7):2225–2238, November 1999.
- [9] S. G. Chang, B. Yu, and M. Vetterli. Adaptive wavelet thresholding for image denoising and compression. *IEEE Trans. Image Proc.*, 9(9):1532–1546, September 2000.
- [10] J. Liu and P. Moulin. Complexity-regularized image denoising. *IEEE Trans. Image Proc.*, 10(6):841–851, June 2001.