Compressive Sampling and Lossy Compression

[Do random measurements provide an efficient method of representing sparse signals?]

[Vivek K Goyal,
Alyson K. Fletcher, and
Sundeep Rangan]

Recent results in compressive sampling have shown that sparse signals can be recovered from a small number of random measurements. This property raises the question of whether random measurements can provide an efficient representation of sparse signals in an information-theoretic sense. Through both theoretical and experimental results, we show that encoding a sparse signal through simple scalar quantization of random measurements incurs a significant penalty relative to direct or adaptive encoding of the sparse signal. Information theory provides alternative quantization strategies, but they come at the cost of much greater estimation complexity.

Digital Object Identifier 10.1109/MSP.2007.915001

1053-5888/08/$25.00©2008IEEE

BACKGROUND

SPARSE SIGNALS

Since the 1990s, modeling signals through sparsity has emerged as an important and widely applicable technique in signal processing. Its most well-known success is in image processing, where great advances in compression and estimation have come from modeling images as sparse in a wavelet domain [1].

In this article we use a simple, abstract model for sparse signals. Consider an N-dimensional vector x that can be represented as $x = Vu$, where V is some orthogonal N-by-N matrix and $u \in \mathbb{R}^N$ has only K nonzero entries. In this case, we say that u is *K-sparse* and that x is *K-sparse with respect to V*. The set of positions of nonzeros coefficients in u is called the *sparsity pattern*, and we call $\alpha = K/N$ the *sparsity ratio*.

Knowing that x is K-sparse with respect to a given basis V can be extremely valuable for signal processing. For example, in compression, x can be represented by the K positions and values of the nonzero elements in u, as opposed to the N elements of x. When the sparsity ratio α is small, the compression gain can be significant. Similarly, in estimating x in the presence of noise, one only has to estimate K as opposed to N real parameters.

Another important property of sparse signals has recently been uncovered: they can be recovered in a computationally tractable manner from a relatively small number of random samples. The method, known as *compressive sampling* (sometimes called *compressed sensing* or *compressive sensing*), was developed in [2], [3] and [4] and is detailed in other articles in this issue.

A basic model for compressive sampling is shown in Figure 1. The N-dimensional signal x is assumed to be K-sparse with respect to some orthogonal matrix V. The "sampling" of x is represented as a linear transformation by a matrix Φ yielding a sample vector $y = \Phi x$. Let the size of Φ be M-by-N, so y has M elements; we call each element of y a measurement of x. A decoder must recover the signal x from y knowing V and Φ, but not necessarily the sparsity pattern of the unknown signal u.

Since u is K-sparse, x must belong to one of $\binom{N}{K}$ subspaces in $\mathbb{R}^N$. Similarly, y must belong to one of $\binom{N}{K}$ subspaces in $\mathbb{R}^M$. For almost all Φs with $M \geq K+1$, an exhaustive search through the subspaces can determine which subspace x belongs to and thereby recover the signal's sparsity pattern and values. Therefore, in principle, a K sparse signal can be recovered from as few as $M = K+1$ random samples.

Unfortunately, the exhaustive search described above is not tractable for interesting sizes of problems since the number of subspaces to search, $\binom{N}{K}$, can be enormous; if α is held constant as N is increased, the number of subspaces grows exponentially with N. The remarkable main result of compressive sampling is to exhibit recovery methods that are computationally feasible, numerically stable, and robust against noise while requiring a number of measurements not much larger than K.

SIGNAL RECOVERY WITH COMPRESSIVE SAMPLING

Compressive sampling is based on recovering x via convex optimization. When we observe $y = \Phi x$ and x is sparse with respect to V, we are seeking x consistent with y and such that $V^{-1}x$ has few nonzero entries. To try to minimize the number of nonzero entries directly yields an intractable problem [5]. Instead, solving the optimization problem

$$\text{(LP reconstruction)} \qquad \hat{x}_{\text{LP}} = \underset{x:\, y=\Phi x}{\arg\min} \|V^{-1}x\|_1$$

often gives exactly the desired signal recovery, and there are simple conditions that guarantee exact recovery. Following pioneering work by Logan in the 1960s, Donoho and Stark [6] obtained results that apply, for example, when V is the N-by-N identity matrix and the rows of Φ are taken from the matrix representation of the length-N discrete Fourier transform (DFT). Subsequent works considered randomly selected rows from the DFT matrix [2] and then certain other random matrix ensembles [3], [4]. In this article, we will concentrate on the case when Φ has independent Gaussian entries.

A central question is: How many measurements M are needed for linear program (LP) reconstruction to be successful? Since Φ is random, there is always a chance that reconstruction will fail. We are interested in how M should scale with signal dimension N and sparsity K so that the probability of success approaches one. A result of Donoho and Tanner [7] indicates that $M \sim 2K\log(N/K)$ is a sharp threshold for successful recovery. Compared to the intractable exhaustive search through all possible subspaces, LP recovery requires only a factor $2\log(N/K)$ more measurements.

If measurements are subject to additive Gaussian noise so that $\hat{y} = \Phi x + \eta$ is observed, with $\eta \sim \mathcal{N}(0, \sigma^2)$, then the LP reconstruction should be adjusted to allow slack in the constraint $y = \Phi x$. A typical method for reconstruction is the following convex optimization:

$$\text{(Lasso reconstruction)} \qquad \hat{x}_{\text{Lasso}} = \underset{x}{\arg\min} \left(\|\hat{y} - \Phi x\|_2^2 + \lambda \|V^{-1}x\|_1 \right),$$

where the parameter $\lambda > 0$ trades off data fidelity and reconstruction sparsity. The best choice for λ depends on the variance of the

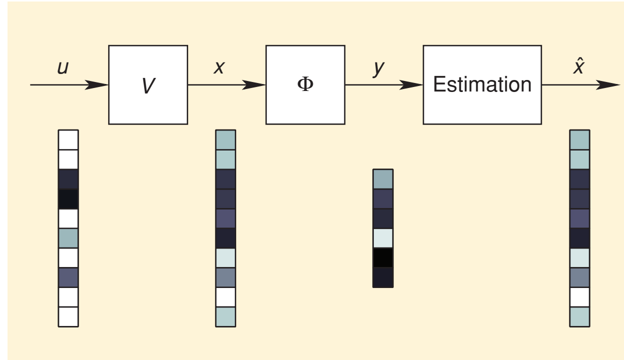

[FIG1] Block diagram representation of compressive sampling. The signal x is sparse with respect to V, meaning that $u = V^{-1}x$ has only a few nonzero entries. $y = \Phi x$ is "compressed" in that it is shorter than x. (White boxes represent zero elements.)

noise and problem size parameters. Wainwright [8] has shown that the scaling $M \sim 2K \log (N - K) + K$ is a sharp threshold for $V^{-1}\hat{x}_{\text{Lasso}}$ to have the correct sparsity pattern with high probability. While this M may be much smaller than N, it is significantly more measurements than required in the noiseless case.

[TABLE 1] PERFORMANCE SUMMARY: DISTORTIONS FOR SEVERAL SCENARIOS WHEN N IS LARGE WITH $\alpha = K/N$ HELD CONSTANT. RATE R AND DISTORTION D ARE BOTH NORMALIZED BY K. J REPRESENTS THE SPARSITY PATTERN OF u. THE BOXED RED ENTRY IS A HEURISTIC ANALYSIS OF THE COMPRESSIVE SAMPLING CASE. $H(\cdot)$ REPRESENTS THE BINARY ENTROPY FUNCTION AND THE ROTATIONAL LOSS R^* SATISFIES $R^* = O(\log R)$.

		ENCODER	
		(SPARSIFYING BASIS) USES V	(RANDOM MEASUREMENTS) USES Φ
DECODER	KNOWS J A PRIORI	$c2^{-2R}$	$c2^{-2(R-R^*)}$
	IS TOLD J	$c2^{-2(R-H(\alpha)/\alpha)}$	$c2^{-2(R-H(\alpha)/\alpha-R^*)}$
	INFERS J		$c\delta(\log N)2^{-2\delta R}$

COMPRESSIVE SAMPLING AS SOURCE CODING

In the remainder of this article, we will be concerned with the tradeoff between quality of approximation and the number of bits of storage for a signal x that is K-sparse with respect to orthonormal basis V. An immediate distinction from the "Background" section is that the currency in which we denominate the cost of a representation is bits rather than real coefficients.

In any compression involving scalar quantization, the choice of coordinates is key. Traditionally, signals to be compressed are modeled as jointly Gaussian vectors. These vectors can be visualized as lying in an ellipsoid, since this is the shape of the level curves of their probability density [see Figure 2(a)]. Source coding theory for jointly Gaussian vectors suggests to choose orthogonal coordinates aligned with the principal axes of the ellipsoid (the Karhunen–Loève basis) and then allocate bits to the dimensions based on their variances. This gives a *coding gain* relative to arbitrary coordinates [9]. For high-quality (low distortion) coding, the coding gain is a constant number of bits per dimension that depends on the eccentricity of the ellipse.

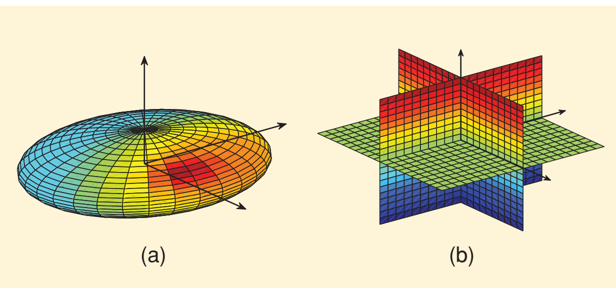

[FIG2] (a) Depiction of Gaussian random vectors as an ellipsoid. Classical rate-distortion theory and transform coding results are for this sort of source, which serves as a good model for discrete cosine transform (DCT) coefficients of an image or MDCT coefficients of audio. (b) Depiction of two sparse signals in $\mathbb{R}^3$, which form a union of three subspaces. This serves as a good conceptual model for wavelet coefficients of images.

Sparse signal models are geometrically quite different than jointly Gaussian vector models. Instead of being visualized as ellipses, they yield unions of subspaces [see Figure 2(b)]. A natural encoding method for a signal x that is K-sparse with respect to V is to identify the subspace containing x and then quantize within the subspace, spending a number of bits proportional to K. Note that doing this requires the encoder to know V and that there is a cost to communicating the subspace index, denoted J, that will be detailed later. With all the proper accounting, when $K \ll N$, the savings is more dramatic than just a constant number of bits.

Following the compressive sampling framework one obtains a rather different way to compress x: quantize the measurements $y = \Phi x$, with Φ and V known to the decoder. Since Φ spreads the energy of the signal uniformly across the measurements, each measurement should be allocated the same number of bits. The decoder should estimate x as well as it can; we will not limit the computational capability of the decoder.

How well will compressive sampling work? It depends both on how much it matters to use the best basis (V) rather than a set of random vectors (Φ) and how much the quantization of y affects the ability of the decoder to infer the correct subspace. We separate these issues, and our

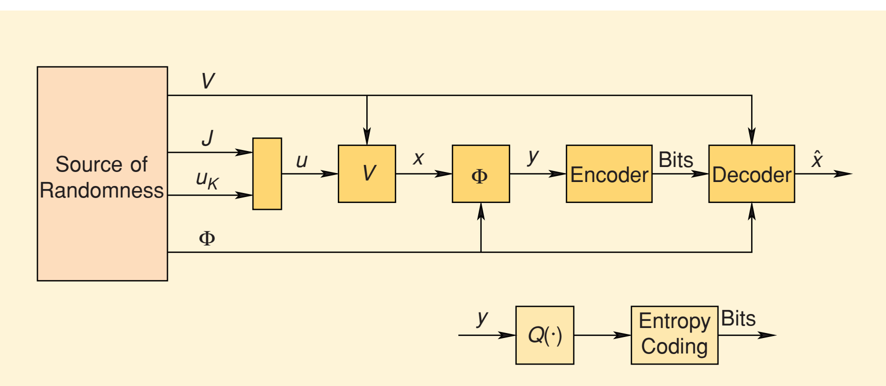

[FIG3] Block diagram representation of the compressive sampling scenario analyzed information theoretically. V is a random orthogonal matrix, u is a K-sparse vector with $\mathcal{N}(0,1)$ nonzero entries, and Φ is a Gaussian measurement matrix. More specifically, the sparsity pattern of u is represented by J and the nonzero entries are denoted u_K. In the initial analysis, the encoding of $y = \Phi x$ is by scalar quantization and scalar entropy coding.

results are previewed and summarized in Table 1. We will derive the results in blue and then the result in red, which requires much more explanation. But first we establish the setting more concretely.

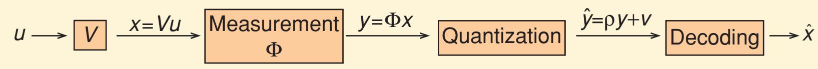

[FIG4] Source coding of *x* with additive noise representation for quantization.

MODELING ASSUMPTIONS

To reflect the concept that the orthonormal basis V is not used in the sensor/encoder, we model V as random and available only at the estimator/decoder. It is chosen uniformly at random from the set of orthogonal matrices. The source vector x is also random; to model it as K-sparse with respect to V, we let $x = Vu$ where $u \in \mathbb{R}^N$ has K nonzero entries in positions chosen uniformly at random. As depicted in Figure 3, we denote the nonzero entries of u by $u_K \in \mathbb{R}^K$ and let the discrete random variable J represent the sparsity pattern. Note that both V and Φ can be considered side information available at the decoder but not at the encoder.

Let the components of u_K be independent and Gaussian $\mathcal{N}(0, 1)$. Observe that $\mathrm{E}[\|u\|^2] = K$, and since V is orthogonal we also have $\mathrm{E}[\|x\|^2] = K$. For the measurement matrix Φ, let the entries be independent $\mathcal{N}(0, 1/K)$ and independent of V and u. This normalization makes the entries of y each have unit variance.

Let us now establish some notation to describe scalar quantization. When scalar y_i is quantized to yield $\widehat{y}_i$, it is convenient to define the relative quantization error $\beta = \mathrm{E}[|y_i - \widehat{y}_i|^2]/\mathrm{E}[|y_i|^2]$ and then further define $\rho = 1 - \beta$ and $v_i = \widehat{y}_i - \rho y_i$. These definitions yield a gain-plus-noise notation $\widehat{y}_i = \rho y_i + v_i$, where

$$\sigma_v^2 = \mathrm{E}[|v_i|^2] = \beta(1-\beta)\mathrm{E}[|y_i|^2], \tag{1}$$

to describe the effect of quantization. Quantizers with optimal (centroid) decoders result in v being uncorrelated with y [10, Lemma 5.1]; other precise justifications are also possible [11].

In subsequent analyses, we will want to relate β to the rate (number of bits) of the quantizer. The exact value of β depends not only on the rate R but also on the distribution of y_i and the particular quantization method. However, the scaling of β with R is as 2^{-2R} under many different scenarios (see "Quantizer Performance and Quantization Error"). We will write

$$\beta = c\, 2^{-2R} \tag{2}$$

without repeatedly specifying the constant $c \geq 1$.

With the established notation, the overall quantizer output vector can be written as

$$\widehat{y} = \rho\Phi V u + v = Au + v, \tag{3}$$

where $A = \rho\Phi V$. The overall source coding and decoding process, with the gain-plus-noise representation for quantization, is depicted in Figure 4. Our use of (3) is to enable easy analysis of linear estimation of x from $\widehat{y}$.

QUANTIZER PERFORMANCE AND QUANTIZATION ERROR

A quantity that takes on uncountably many values—like a real number—cannot have an exact digital representation. Thus digital processing always involves quantized values. The relationships between the number of bits in a representation (rate R), the accuracy of a representation (distortion D), and properties of quantization error are central to this article and are developed in this sidebar.

The simplest form of quantization—*uniform scalar quantization*—is to round $x \in \mathbb{R}$ to the nearest multiple of some fixed resolution parameter Δ to obtain quantized version $\widehat{x}$. For this type of quantizer, rate and distortion can be easily related through the step size Δ. Suppose x has a smooth distribution over an interval of length C. Then the quantizer produces about C/Δ intervals, which can be indexed with $R \approx \log_2(C/\Delta)$ b. The error $x - \widehat{x}$ is approximately uniformly distributed over $[-\Delta/2, \Delta/2]$, so the mean-squared error is $D = \mathrm{E}[(x-\widehat{x})^2] \approx (1/12)\Delta^2$. Eliminating Δ, we obtain $D \approx (1/12)C^2 2^{-2R}$.

The 2^{-2R} dependence on rate is fundamental for compression with respect to MSE distortion. For any distribution of x, the best possible distortion as a function of rate (obtained with high-dimensional vector quantization [25]) satisfies

$$(2\pi e)^{-1} 2^{2h} 2^{-2R} \leq D(R) \leq \sigma^2 2^{-2R},$$

where h and σ^2 are the differential entropy and variance of x. Also, under high resolution assumptions and with entropy coding, $D(R) \approx (1/12)2^{2h}2^{-2R}$ performance is obtained with uniform scalar quantization, which for a Gaussian random variable is $D(R) \approx (1/6)\pi e \sigma^2 2^{-2R}$. Covering all of these variations together, we write the performance as $D(R) = c\sigma^2 2^{-2R}$ without specifying the constant c.

More subtle is to understand the quantization error $e = x - \widehat{x}$. With uniform scalar quantization, e is in the interval $[-\Delta/2, \Delta/2]$, and it is convenient to think of it as a uniform random variable over this interval, independent of x. This is merely a convenient fiction, since $\widehat{x}$ is a deterministic function of x. In fact, as long as quantizers are regular and estimation procedures use linear combinations of many quantized values, second-order statistics (which are well understood [11]) are sufficient for understanding estimation performance. When x is Gaussian, a rather counterintuitive model where e is Gaussian and independent of x can be justified precisely: optimal quantization of a large block of samples is described by the *optimal test channel*, which is additive Gaussian [28].

ANALYSES

Since the sparsity level K is the inherent number of degrees of freedom in the signal, we will let there be KR bits available for the encoding of x and also normalize the distortion by $K: D = (1/K)\mathrm{E}[\|x - \widehat{x}\|^2]$. Where applicable, the number of measurements M is a design parameter that can optimized to give the best distortion-rate tradeoff. In particular, increasing M gives better conditioning of certain matrices, but it reduces the number of quantization bits per measurement.

Before analyzing the compressive sampling scenario (Figure 3), we consider some simpler alternatives, yielding the blue entries in Table 1.

SIGNAL IN A KNOWN SUBSPACE

If the sparsifying basis V and subspace J are fixed and known to both encoder and decoder, the communication of x can be accomplished by sending quantized versions of the nonzero entries of $V^{-1}x$. Each of the K nonzero entries has unit variance and is allotted R b, so $D(R) = c2^{-2R}$ performance is obtained, as given by the first entry in Table 1.

ADAPTIVE ENCODING WITH COMMUNICATION OF J

Now suppose that V is known to both encoder and decoder, but the subspace index J is random, uniformly selected from the $\binom{N}{K}$ possibilities. A natural *adaptive* approach is to spend $\log_2\binom{N}{K}$ bits to communicate J and the remaining available bits to quantize the nonzero entries of $V^{-1}x$. Defining $R_0 = (1/K)\log_2\binom{N}{K}$, the encoder has $KR - KR_0$ b for the K nonzero entries of $V^{-1}x$ and thus attains performance

$$D_{\text{adaptive}} = c2^{-2(R-R_0)}, \qquad R \geq R_0. \tag{4}$$

When K and N are large with the ratio $\alpha = K/N$ held constant, $\log_2\binom{N}{K} \approx NH(\alpha)$ where $H(p) = -p\log_2 p - (1-p)\log_2(1-p)$ is the *binary entropy function* [12, p. 530]. Thus $R_0 \approx H(\alpha)/\alpha$, giving a second entry in Table 1.

If R does not exceed R_0, then the derivation above does not make sense, and even if R exceeds R_0 by a small amount, it may not pay to communicate J. A *direct* approach is to simply quantize each component of x with KR/N b. Since the components of x have variance K/N, performance of $\mathrm{E}[(x_i - \widehat{x}_i)^2] \leq c(K/N)2^{-2KR/N}$ can be obtained, yielding overall performance

$$D_{\text{direct}}(R) = c2^{-2KR/N}. \tag{5}$$

By choosing the better between (4) and (5) for a given rate, one obtains a simple baseline for the performance using V at the encoder. A convexification by time sharing could also be applied, and more sophisticated techniques are presented in [13].

LOSS FROM RANDOM MEASUREMENTS

Now let us try to understand in isolation the effect of observing x only through Φx. The encoder sends a quantized version of $y = \Phi x$, and the decoder knows V and the sparsity pattern J.

From (3), the decoder has $\widehat{y} = \rho\Phi Vu + v$ and knows which K elements of u are nonzero. The performance of a linear estimate of the form $\widehat{x} = F(J)\widehat{y}$ will depend on the singular values of the M-by-K matrix formed by the K relevant columns of ΦV. (One should expect a small improvement—roughly a multiplication of the distortion by K/M—from the use of a nonlinear estimate that exploits boundedness of quantization noise [14], [15]. The dependence on ΦV is roughly unchanged [16].) Using elementary results from random matrix theory, one can find how the distortion varies with M and R. (The distortion does not depend on N because the zero components of u are known.) The analysis given in [17] shows that for moderate to high R, the distortion is minimized when $K/M \approx 1 - ((2\ln 2)R)^{-1}$. Choosing the number of measurements accordingly gives performance

$$D_J(R) \approx 2(\ln 2)eR \cdot c2^{-2R} = c2^{-2(R-R^*)} \tag{6}$$

where $R^* = (1/2)\log_2(2(\ln 2)eR)$, giving the final blue entry in Table 1. Comparing to $c2^{-2R}$, we see that having access only to random measurements induces a significant performance loss.

One interpretation of this analysis is that the coding rate has effectively been reduced by R^* b per degree of freedom. Since R^* grows sublinearly with R, the situation is not too bad—at least the performance does not degrade with increasing K or N. The analysis when J is not known at the decoder—i.e., it must be inferred from $\widehat{y}$—reveals a much worse situation.

LOSS FROM SPARSITY RECOVERY

As we have mentioned before, compressive sampling is motivated by the idea that the sparsity pattern J can be detected, through a computationally tractable convex optimization, with a "small" number of measurements M. However, the number of measurements required depends on the noise level. We saw $M \sim 2K\log(N-K) + K$ scaling is required by lasso reconstruction; if the noise is from quantization and we are trying to code with KR total bits, this scaling leads to a vanishing number of bits per measurement.

Unfortunately, the problem is more fundamental than suboptimality of lasso decoding. We will show that trying to code with KR total bits makes reliable recovery of the sparsity pattern impossible as the signal dimension N increases. In this analysis, we assume the sparsity ratio $\alpha = K/N$ is held constant as the problems scale, and we see that no number of measurements M can give good performance.

To see why the sparsity pattern cannot be recovered, consider the problem of estimating the sparsity pattern of u from the noisy measurement y in (3). Let $E_{\text{signal}} = \mathrm{E}[\|Au\|^2]$ and $E_{\text{noise}} = \mathrm{E}[\|v\|^2]$ be the signal and noise energies, respectively, and define the signal-to-noise ratio (SNR) as $\mathrm{SNR} = E_{\text{signal}} / E_{\text{noise}}$. The number of measurements M required to recover the sparsity pattern of u from y can be bounded below with the following theorem.

THEOREM 1

Consider any estimator for recovering the sparsity pattern of a K-sparse vector u from measurements y of the form (3), where v

is a white Gaussian vector uncorrelated with y. Let P_{error} be the probability of misdetecting the sparsity pattern, averaged over the realizations of the random matrix A and noise v. Suppose M, K, and $N-K$ approach infinity with

$$M < \frac{K}{\text{SNR}}\left[(1-\epsilon)\log(N-K) - 1\right] \tag{7}$$

for some $\epsilon > 0$. Then $P_{\text{error}} \to 1$, i.e., the estimator will asymptotically always fail.

The main ideas of a proof of Theorem 1 are given in "Proof Sketch for Theorem 1." Under certain assumptions, the quantization error v in our problem will be asymptotically Gaussian, so we can apply the bound (see "Quantizer Performance and Quantization Error"). The theorem shows that to attain any non-vanishing probability of success, we need the scaling

$$M \geq \frac{K}{\text{SNR}}\left[(1-\epsilon)\log(N-K) - 1\right]. \tag{8}$$

Now, using the normalization assumptions described above, the expression $\rho = 1-\beta$, and σ_v^2 given in (1), it can be shown that the signal and noise energies are given by $E_{\text{signal}} = M(1-\beta)^2$ and $E_{\text{noise}} = M\beta(1-\beta)$. Therefore, the SNR is

$$\text{SNR} = (1-\beta)/\beta. \tag{9}$$

Now, let $\delta = K/M$ be the "measurement ratio," i.e., the ratio of degrees of freedom in the unknown signal to number of measurements. From (2), $\beta \geq 2^{-2\delta R}$ for any quantizer, and therefore, from (9), $\text{SNR} \leq 2^{2\delta R} - 1$. Substituting this bound for the SNR into (8), we see that for the probability of error to vanish (or even become a value less than one) will require

$$\frac{2^{2\delta R}}{(1-\epsilon)\delta} + 1 > \log(N-K). \tag{10}$$

Notice that, for any fixed R, the left hand side of (10) is bounded above uniformly over all $\delta \in (0, 1]$. However, if the sparsity ratio $\alpha = K/N$ is fixed and $N \to \infty$, then $\log(N-K) \to \infty$. Consequently, the bound (10) is impossible to satisfy. We conclude that: *for a fixed rate R and sparsity ratio α, as $N \to \infty$, there is no number of measurements M that can guarantee reliable sparsity recovery. In fact, the probability of detecting the sparsity pattern correctly approaches zero*. This conclusion applies not just to compressive sampling with basis pursuit or matching pursuit detection, but even to exhaustive search methods.

How bad is this result for compressive sampling? We have shown that exact sparsity recovery is fundamentally impossible when the total number of bits scales linearly with the degrees of freedom of the signal and the quantization is regular. However, exact sparsity recovery may not be necessary for good performance. What if the decoder can detect, say, 90% of the elements in the sparsity pattern correctly? One might think that the resulting distortion might still be small.

Unfortunately, when we translate the best known error bounds for reconstruction from nonadaptively encoded undersampled data, we do not even obtain distortion that approaches

PROOF SKETCH FOR THEOREM 1

Since the vector $u \in \mathbb{R}^N$ has K nonzero components, Au belongs to one of the $\binom{N}{K}$ subspaces, each subspace being spanned by K of the N columns of $A \in \mathbb{R}^{M\times N}$. Let $\mathcal{V}$ be the set of all such subspaces and let $V_0 \in \mathcal{V}$ be the "true" K-dimensional subspace—the one that contains Au. The detector with the minimum probability of error would search over all the subspaces for the one with the maximum energy of the received noisy vector y. For the estimator to detect the correct subspace, the true subspace must have the maximum energy. That is,

$$\|P_{V_0}y\|^2 \geq \|P_V y\|^2, \qquad \forall\, V \in \mathcal{V}, \tag{13}$$

where P_S denotes the projection operator onto the subspace S. We can show (7) from (13) as follows.

The true subspace is spanned by K columns of A, which we will denote by $a_1, \ldots, a_K$. Since V_0 contains Au, it must contain the entire signal energy E_{signal}. It also contains a fraction $\delta = K/M$ of the noise energy, E_{noise}. So the average energy in the subspace V_0 is $\|P_{V_0}y\|^2 = E_{\text{signal}} + \delta E_{\text{noise}}$. Although this expression is technically only true in expectation, it is asymptotically exact for large M. So here and in the remainder of the proof, we omit the expectations in the formulas.

Now remove the vector a_1, and let V_1 be the subspace spanned by the remaining $K-1$ vectors $\{a_j\}_{j=2}^{K}$. Since the vectors a_j are i.i.d. and spherically symmetrically distributed, the energy of y in V_1 relative to the energy in V_0 is given by

$$\begin{aligned}\|P_{V_0}y\|^2 - \|P_{V_1}y\|^2 &= \frac{1-\delta}{K}\|P_{V_0}y\|^2 \\ &= \frac{1-\delta}{K}\left[E_{\text{signal}} + \delta E_{\text{noise}}\right].\end{aligned} \tag{14}$$

Now let a_j, $j = K+1, \ldots, N$ be the remaining $N-K$ columns of A. The $M-K$ dimensional subspace $V_0^{\perp} \subseteq V_1^{\perp}$ contains a fraction $1-K/M$ of the noise energy E_{noise}. Each column a_j is independent of the signal in $V_0^{\perp}$. When M is large, it can be shown that adding one of the columns will add a random amount of energy described by $(1-K/M)E_{\text{noise}}u_j^2/(M-K) = E_{\text{noise}}u_j^2/M$, where u_j is an $\mathcal{N}(0,1)$ Gaussian random variable. Let V be the subspace spanned by V_1 and the vector a_j with the maximum energy. The new subspace V is spanned by K columns of A, so $V \in \mathcal{V}$. Also, the energy in V will be

$$\|P_V y\|^2 - \|P_{V_1}y\|^2 = \frac{1}{M}E_{\text{noise}} \max_{j=K+1,\ldots,N} u_j^2.$$

For any $\epsilon > 0$, it can be shown that

$$\Pr\left(\max_{j=K+1,\ldots,N} u_j^2 > (1-\epsilon)\log(N-K)\right) \to 1,$$

as $N-K \to \infty$. Therefore, as $N-K \to \infty$, with high probability,

$$\|P_V y\|^2 - \|P_{V_1}y\|^2 > \frac{1}{M}(1-\epsilon)E_{\text{noise}}\log(N-K). \tag{15}$$

Combining (13), (14), and (15) shows (7).

zero as the rate is increased with K, M, and N fixed. [Remember that without undersampling, one can at least obtain the performance (5).] For example, Candès and Tao [18] prove that an estimator similar to the lasso estimator attains a distortion

$$\frac{1}{K}\|x-\widehat{x}\|^2 \le c_1 \frac{K}{M}(\log N)\sigma^2, \quad (11)$$

with large probability, from M measurements with noise variance σ^2, provided that the number of measurements is adequate. There is a constant $\delta \in (0, 1)$ such that $M = K/\delta$ is sufficient for (11) to hold with probability approaching one as N is increased with K/N held constant; but for any finite N, there is a nonzero probability of failure. Spreading RK bits amongst the measurements and relating the number of bits to the quantization noise variance gives

$$D = \frac{1}{K}\mathrm{E}[\|x-\widehat{x}\|^2] \le c_2\delta(\log N)2^{-2\delta R} + D_{\mathrm{err}}, \quad (12)$$

where D_{err} is the distortion due to the failure event. (Haupt and Nowak [19] consider optimal estimators and obtain a bound similar to (12) in that it has a term that is constant with respect to the noise variance. See also [20] for related results.) Thus if D_{err} is negligible, the distortion will decrease exponentially in the rate, but with an exponent reduced by a factor δ. However, as N increases to infinity, the distortion bound increases and is not useful.

NUMERICAL SIMULATION

To get some idea of the possible performance, we performed the following numerical experiment. We fixed the signal dimensions to $N = 100$ and $K = 10$, so the signal has a sparsity of $\alpha = K/N = 0.1$. We varied the quantization rate R from 4 to 12 b per degree of freedom, which spans low to high rate since $(1/K)\log_2\binom{N}{K} \approx 4.4$. The resulting simulated performance of compressive sampling is shown in Figure 5. The performance of direct quantization [$D_{\mathrm{direct}}(R)$ from (5)] and baseline quantization with time sharing [see (4) and (5)] are shown for comparison.

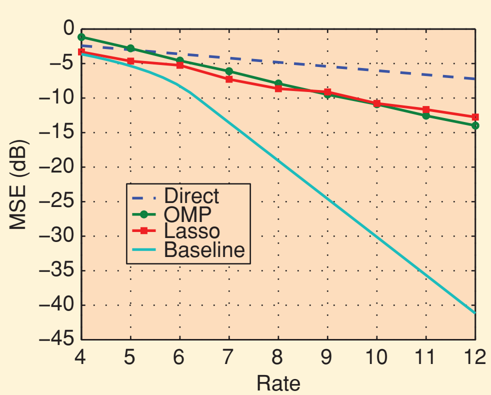

[FIG5] Rate-distortion performance of compressive sampling using reconstruction via OMP and lasso. At each rate, the number of measurements M is optimized to minimize the distortion. Also plotted are the theoretical distortion curves for direct and baseline quantization. In all simulations $(K, N) = (10, 100)$.

The distortion of compressive sampling was simulated as follows: For both lasso and orthogonal matching pursuit (OMP) reconstruction and for integer rates R, the number of measurements M was varied from K to N in steps of ten. At each value of M, the distortion was estimated by averaging 500 Monte Carlo trials with random encoder matrices Φ and quantization noise vectors v. To give the best-case performance of compressive sampling, the distortion was taken to be the minimum distortion over the tested values of M and, for lasso, over several values of the regularization parameter λ. The optimal M is not necessarily the minimum M to guarantee sparsity recovery. Instead, optimizing M trades off errors in the sparsity pattern against errors in the estimated values for the components. The optimal value does not result in small probability of subspace misdetection. More extensive sets of simulations consistent with these are presented in [21].

From (4), the distortion with adaptive quantization decreases exponentially with the rate R through the multiplicative factor 2^{-2R}. This appears in Figure 5 as a decrease in distortion of approximately 6 dB/b. In contrast, simple direct quantization achieves a distortion given by (5), which in this case is only 0.6 dB/b. Thus, there is potentially a large gap between direct quantization that does not exploit the sparsity and adaptive quantization that does.

Both compressive sampling methods, lasso and OMP, are able to perform slightly better than simple direct quantization, achieving approximately 1.4–1.6 dB/b. (A finer analysis that allows computation of the largest possible δ in (12) might predict this slope.) Thus, compressive sampling is able to exploit the sparsity to some degree and narrow the gap between linear and adaptive quantization. However, neither algorithm is able to come close to the performance of the baseline encoder that can use adaptive quantization. Indeed, comparing to the baseline quantization, there a multiplicative rate penalty in this simulation of approximately a factor of four. This is large by source coding standards, and we can conclude that compressive sampling does not achieve performance similar to adaptive quantization.

INFORMATION THEORY TO THE RESCUE?

We have thus far used information theory to provide context and analysis tools. It has shown us that compressing sparse signals by scalar quantization of random measurements incurs a significant penalty. Can information theory also suggest alternatives to compressive sampling? In fact, it does provide techniques that would give much better performance for source coding, but the complexity of decoding algorithms becomes even higher.

Let us return to Figure 3 and interpret it as a communication problem where x is to be reproduced approximately and the number of bits that can be used is limited. We would like to extract *source coding with side information* and *distributed source coding* problems from this setup. This will lead to results much more positive than those developed above.

In developing the baseline quantization method, we discussed how an encoder that knows V can recover J and u_K from x and thus send J exactly and u_K approximately. Compressive sampling is to apply when the encoder does not know (or want to use) the sparsifying basis V. In this case, an information theorist would say that we have a problem of lossy source coding of x with side information V available at the decoder—an instance of the *Wyner-Ziv problem* [22]. In contrast to the analogous lossless coding problem (see "Slepian-Wolf Coding"), the unavailability of the side information at the encoder does in general hurt the best possible performance. Specifically, let $L(D)$ denote the rate loss (increased rate because V is unavailable) to achieve distortion D. Then there are upper bounds to $L(D)$ that depend only on the source alphabet, the way distortion is measured, and the value of the distortion—not on the distribution of the source or side information [23]. For the scenario of interest to us [(continuous-valued source and mean-squared error (MSE) distortion)], $L(D) \leq 0.5$ b for all D. The techniques to achieve this are complicated, but note the *constant additive rate penalty* is in dramatic contrast to Figure 5.

Compressive sampling not only allows side information V to be available only at the decoder, it also allows the components of the measurement vector y to be encoded separately. The way to interpret this information theoretically is to consider $y_1, y_2, \ldots, y_M$ as distributed sources whose joint distribution depends on side information (V, Φ) available at the decoder. Imposing a constraint of distributed encoding of y (while allowing joint decoding) generally creates a degradation of the best possible performance. (Again, there is no performance penalty in the lossless case; see "Slepian-Wolf Coding.") Let us sketch a particular strategy that is not necessarily optimal but exhibits only a small additive rate penalty. This is inspired by [23] and [24].

Suppose that each of M distributed encoders performs scalar quantization of its own y_i to yield $q(y_i)$. Before this seemed to immediately get us in trouble (recall our interpretation of Theorem 1), but now we will do further encoding. The quantized values give us a lossless distributed compression problem with side information (V, Φ) available at the decoder. Using Slepian-Wolf coding, a total rate arbitrarily close to $H(q(y))$ can be achieved. The remaining question is how the rate and distortion relate.

For the sake of analysis, let us assume that the encoder and decoder share some randomness Z so that the scalar quantization above can be subtractively dithered (see, e.g., [25]). Then following the analysis in [24] and [26], encoding the quantized samples $q(y)$ at rate $H(q(y) \mid V, Z)$ is within 0.755 b of the conditional rate-distortion bound for source x given V. Thus the combination of universal dithered quantization with Slepian-Wolf coding gives a method of distributed coding with only a constant additive rate penalty. These methods inspired by information theory depend on coding across independent signal acquisition instances, and they generally incur large decoding complexity.

Let us finally interpret the "quantization plus Slepian-Wolf" approach described above when limited to a single instance. Suppose the y_i s are separately quantized as described above. The main negative result of this article indicates that ideal separate entropy coding of each $q(y_i)$ is not nearly enough to get to good performance. The rate must be reduced by replacing an ordinary entropy code with one that collapses some distinct quantized values to the same index. The hope has to be that in the joint decoding of $q(y)$, the dependence between components will save the day. This is equivalent to saying that the quantizers in use are not regular [25], much like multiple description quantizers [27]. This approach is developed and simulated in [21].

CONCLUSIONS—WHITHER COMPRESSIVE SAMPLING?

To an information theorist, "compression" is the efficient representation of data with bits. In this article, we have looked at compressive sampling from this perspective, to see if random measurements of sparse signals provide an efficient method of representing sparse signals.

The source coding performance depends sharply on how the random measurements are encoded into bits. Using familiar forms of quantization (*regular quantizers*; see [25]) even very weak forms of universality are precluded. One would want to

SLEPIAN-WOLF CODING

When two related quantities are to be compressed, there is generally an advantage to doing the compression jointly. What does "jointly" mean? On its face, "jointly" would seem to mean that the quantities are inseparably mapped to a bit string. However, Slepian and Wolf [29] remarkably showed that it can be good enough for the decoding to be "joint"—the encoding can be separate.

To understand the result precisely, suppose a sequence of independent replicas $(X_1^{(1)}, X_2^{(1)}), (X_1^{(2)}, X_2^{(2)}), \ldots$, of the pair of jointly distributed discrete random variables (X_1, X_2) is to be compressed. The minimum possible rate is $H(X_1, X_2)$, the joint entropy of X_1 and X_2. The normal way to approach this minimum rate is to treat (X_1, X_2) as a single discrete random variable (over an alphabet that is the Cartesian product of the alphabets of X_1 and X_2) and apply an entropy code to this random variable. This requires an encoder that operates on X_1 and X_2 together. The main result of [29] indicates that this total rate can be approached with encoders that see X_1 and X_2 separately as long as the decoding is joint. The recovery of the X_k s is perfect (or has vanishing error probability) without requiring any excess total rate (or arbitrarily small excess rate): $R_1 + R_2 = H(X_1, X_2)$. The individual rates need only satisfy $R_1 \geq H(X_1 \mid X_2)$ and $R_2 \geq H(X_2 \mid X_1)$.

As a very simple example, suppose X_1 has any distribution on the integers; and $X_2 - X_1$ equals zero or one with equal probability, independent of X_1. Then (X_1, X_2) has precisely one more bit of information than X_1 alone. The optimal total rate $R_1 + R_2 = H(X_1) + 1$ can be achieved by having Encoder 1 compress X_1 as if communicating X_1 were the only goal and having Encoder 2 send only the parity of X_2.

Slepian-Wolf coding can be extended to any number of correlated sources with no "penalty" in the rate [30, Thm. 14.4.2]. Also, simpler than Slepian-Wolf coding is for one of the sources (say, X_2) to be available to the decoder but not to the encoder. Then a rate of $R_1 = H(X_1 \mid X_2)$ is sufficient to allow the decoder to recover X_1, even though the encoding of X_1 is done without knowledge of X_2. The main text uses these results to give information-theoretic bounds for encoding of random measurements.

spend a number of bits proportional to the number of degrees of freedom of the sparse signal, but this does not lead to good performance. In this case, we can conclude analytically that recovery of the sparsity pattern is asymptotically impossible. Furthermore, simulations show that the MSE performance is far from optimal.

Information theory provides alternatives based on universal versions of distributed lossless coding (Slepian-Wolf coding) and entropy-coded dithered quantization. These information-theoretic constructions indicate that it is reasonable to ask for good performance with merely linear scaling of the number of bits with the sparsity of the signal. However, practical implementation of such schemes remains an open problem.

It is important to keep our mainly negative results in proper context. We have shown that compressive sampling combined with ordinary quantization is a bad compression technique, but our results say nothing about whether compressive sampling is an effective initial step in data acquisition. A good analogy within the realm of signal acquisition is oversampling in analog-to-digital conversion (ADC). Since MSE distortion in oversampled ADC drops only polynomially (not exponentially) with the oversampling factor, high oversampling alone—without other processing—leads to poor rate-distortion performance. Nevertheless, oversampling is ubiquitous. Similarly, compressive sampling is useful in contexts where sampling itself is very expensive, but the subsequent storage and communication of quantized samples is less constricted.

ACKNOWLEDGMENTS

The authors thank Sourav Dey, Lav Varshney, Claudio Weidmann, Joseph Yeh, and an anonymous reviewer for thoughtful comments that helped us improve the article.

AUTHORS

Vivek K Goyal (vgoyal@mit.edu) received the B.S. degree in mathematics and the B.S.E. degree in electrical engineering from the University of Iowa and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Berkeley. He is currently Esther and Harold E. Edgerton Assistant Professor of Electrical Engineering at the Massachusetts Institute of Technology. He has received the UC-Berkeley Eliahu Jury Award for outstanding achievement in systems, communications, control, or signal processing, the IEEE SPS Magazine Award and an NSF CAREER Award. He is on the IEEE SPS Image and Multiple Dimensional Signal Processing Technical Committee and cochair of the SPIE Wavelets conference series. He is a Senior Member of the IEEE.

Alyson K. Fletcher received the B.S. degree in mathematics from the University of Iowa and the M.A. degree in mathematics and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Berkeley. She is currently a President's Postdoctoral Fellow at the University of California, Berkeley. She is a member of SWE, SIAM, and Sigma Xi. She has been awarded the University of California Eugene L. Lawler Award, the Henry Luce Foundation's Clare Boothe Luce Fellowship, and the Soroptimist Dissertation Fellowship. Her research interests include estimation, image processing, statistical signal processing, sparse approximation, wavelets, and control theory.

Sundeep Rangan received the B.A.Sc. degree in electrical engineering from the University of Waterloo, Canada, and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Berkeley. He was a postdoctoral research fellow at the University of Michigan. He then joined the Wireless Research Center at Bell Laboratories, and in 2000, he co-founded Flarion Technologies with four others. He is currently a director of engineering at Qualcomm Technologies, where he is involved in the development of next generation cellular wireless systems.

REFERENCES

[1] D.L. Donoho, M. Vetterli, R.A. DeVore, and I. Daubechies, "Data compression and harmonic analysis," *IEEE Trans. Inform. Theory*, vol. 44, no. 6, pp. 2435–2476, Oct. 1998.

[2] E.J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[3] E.J. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.

[4] D.L. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[5] B.K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, vol. 24, no. 2, pp. 227–234, Apr. 1995.

[6] D.L. Donoho and P.B. Stark, "Uncertainty principles and signal recovery," *SIAM J. Appl. Math.*, vol. 49, no. 3, pp. 906–931, June 1989.

[7] D.L. Donoho and J. Tanner, "Counting faces of randomly-projected polytopes when the projection radically lowers dimension," submitted for publication.

[8] M.J. Wainwright, "Sharp thresholds for high-dimensional and noisy recovery of sparsity," Univ. California, Berkeley, Dept. of Statistics, Tech. Report #arXiv:math.ST/0605740 v1 30, May 2006.

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

[10] A.K. Fletcher, "A jump linear framework for estimation and robust communication with Markovian source and channel dynamics." Ph.D. dissertation, Dept. of Electrical Eng. Comp. Sci., Univ. California, Berkeley, Nov. 2005.

[11] H. Viswanathan and R. Zamir, "On the whiteness of high-resolution quantization errors," *IEEE Trans. Inform. Theory*, vol. 47, no. 5, pp. 2029–2038, July 2001.

[12] R.G. Gallager, *Information Theory and Reliable Communication*. New York: Wiley, 1968.

[13] C. Weidmann and M. Vetterli, "Rate-distortion analysis of spike processes," in *Proc. IEEE Data Compression Conf.*, Snowbird, Utah, Mar. 1999, pp. 82–91.

[14] V.K. Goyal, M. Vetterli, and N.T. Thao, "Quantized overcomplete expansions in $\mathbb{R}^N$: Analysis, synthesis, and algorithms," *IEEE Trans. Inform. Theory*, vol. 44, no. 1, pp. 16–31, Jan. 1998.

[15] S. Rangan and V.K. Goyal, "Recursive consistent estimation with bounded noise," *IEEE Trans. Inform. Theory*, vol. 47, no. 1, pp. 457–464, Jan. 2001.

[16] V.K. Goyal, J. Kovačević, and J.A. Kelner, "Quantized frame expansions with erasures," *Appl. Comput. Harmon. Anal.*, vol. 10, no. 3, pp. 203–233, May 2001.

[17] A.K. Fletcher, S. Rangan, and V.K. Goyal, "On the rate-distortion performance of compressed sensing," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, Honolulu, HI, vol. 3, Apr. 2007, pp. 885–888.

[18] E.J. Candès and T. Tao, "The Dantzig selector: Statistical estimation when p is much larger than n," submitted for publication.

[19] J. Haupt and R. Nowak, "Signal reconstruction from noisy random projections," *IEEE Trans. Inform. Theory*, vol. 52, no. 9, pp. 4036–4048, Sept. 2006.

[20] N. Meinshausen and B. Yu, "Lasso-type recovery of sparse representations for high-dimensional data," Univ. of California, Berkeley, Dept. of Statistics, Sept. 2007.

[21] R.J. Pai, "Nonadaptive lossy encoding of sparse signals," M.S thesis, Dept. of Electrical Eng. Comp. Sci., MIT, Cambridge, MA, Aug. 2006.

[22] A.D. Wyner and J. Ziv, "The rate-distortion function for source coding with side information at the decoder," *IEEE Trans. Inform. Theory*, vol. IT-22, no. 1, pp. 1–10, Jan. 1976.

[23] R. Zamir, "The rate loss in the Wyner–Ziv problem," *IEEE Trans. Inform. Theory*, vol. 42, no. 6, pp. 2073–2084, Nov. 1996.

[24] J. Ziv, "On universal quantization," *IEEE Trans. Inform. Theory*, vol. IT-31, no. 3, pp. 344–347, May 1985.

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

[26] R. Zamir and M. Feder, "On universal quantization by randomized uniform/lattice quantization," *IEEE Trans. Inform. Theory*, vol. 38, no. 2, pp. 428–436, Mar. 1992.

[27] V.K. Goyal, "Multiple description coding: Compression meets the network," *IEEE Signal Processing Mag.*, vol. 18, no. 5, pp. 74–93, Sept. 2001.

[28] T. Berger, *Rate Distortion Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1971.

[29] D. Slepian and J.K. Wolf, "Noiseless coding of correlated information sources," *IEEE Trans. Inform. Theory*, vol. IT-19, no. 4, pp. 471–480, July 1973.

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