

Encoding the ℓ_p Ball from Limited Measurements

Emmanuel Candès and Justin Romberg*
Applied and Computational Mathematics,
Caltech, Pasadena, CA, 91125

Abstract

We address the problem of encoding signals which are *sparse*, i.e. signals that are concentrated on a set of small support. Mathematically, such signals are modeled as elements in the ℓ_p ball for some $p \leq 1$. We describe a strategy for encoding elements of the ℓ_p ball which is *universal* in that 1) the encoding procedure is completely generic, and does not depend on p (the sparsity of the signal), and 2) it achieves near-optimal minimax performance simultaneously for all $p < 1$. What makes our coding procedure unique is that it requires only a limited number of nonadaptive measurements of the underlying sparse signal; we show that near-optimal performance can be obtained with a number of measurements that is roughly proportional to the number of bits used by the encoder. We end by briefly discussing these results in the context of image compression.

1 Introduction

1.1 Signal recovery from incomplete measurements

A series of recent results [3–6, 11] have shown that digital signals can be recovered from surprisingly few *non-adaptive* linear measurements. Suppose that a finite signal $x_0 \in \mathbb{R}^N$ is *sparse* in that there are very few significant components of x_0 . Mathematically, we model such signals as belonging to an ℓ_p ball $B_{p,N}$ with $p \leq 1$:

$$B_{p,N} = \{x \in \mathbb{R}^N : \|x\|_{\ell_p} \leq 1\}$$

where

$$\|x\|_{\ell_p} = \left(\sum_{i=1}^N |x_i|^p \right)^{1/p}.$$

*Contact author. Email: jrom@acm.caltech.edu, Phone: (626) 395-5760.

Now suppose that instead of observing x_0 directly, we instead are allowed to make a small number $K \ll N$ of linear measurements

$$y_k = \langle x_0, \phi_k \rangle, \quad k = 1, \dots, K \quad \text{or} \quad y = \Phi x_0 \quad (1)$$

against test functions $\phi_k \in \mathbb{R}^N$. (The $K \times N$ matrix Φ , whose rows are the test functions ϕ_k , is called the *measurement ensemble*). The central results of [4, 5] state that if Φ obeys a *uniform uncertainty principle*, we can recover something very close to the original signal by solving the convex optimization problem

$$(P_1) \quad \min \|x\|_{\ell_1} \quad \text{such that} \quad \Phi x = y.$$

The uniform uncertainty principle essentially states that small subsets of the columns of Φ for a *restricted isometry*. To make this precise, assume that each column of Φ has unit norm, let $T \subset \{1, \dots, N\}$ be a subset of indices, and let Φ_T be the submatrix formed by extracting the columns of Φ corresponding to the indices in T . We say that Φ obeys a uniform uncertainty principle for sets of size S if there exists a $\delta_S < 1$ such that

$$(1 - \delta_S) \cdot \|c\|_{\ell_2} \leq \|\Phi_T c\|_{\ell_2} \leq (1 + \delta_S) \cdot \|c\|_{\ell_2}$$

for all coefficient sequences c supported on T , and all sets T with fewer than S elements: $|T| \leq S$.

Given measurements y , call \hat{x} the solution to (P1) above. If for a given S , the restricted isometry constants obey

$$3\delta_{4S} + \delta_{3S} < 2, \quad (2)$$

the recovery error can be bounded by

$$\|\hat{x} - x_0\|_{\ell_2} \leq C \cdot S^{-1/2} \cdot \|x_{0,S} - x_0\|_{\ell_1}, \quad (3)$$

where $x_{0,S}$ is the best S -term approximation of x_0 formed by taking the S largest values of x_0 and setting the rest to zero. For $x_0 \in B_{p,N}$, the right hand side of (3) goes to zero quickly as S increases¹:

$$S^{-1/2} \cdot \|x_{0,S} - x_0\|_{\ell_1} \leq C \cdot S^{-\alpha} \quad (4)$$

where $\alpha = 1/p - 1/2$. Note that the $S^{-\alpha}$ in (4) is the same rate at which the nonlinear approximation error is guaranteed to go to zero [9]; (P1) is able to match the accuracy of the *best* adaptive approximation even though it has access to only a small portion of the signal.

With a simple modification, the recovery procedure can be made stable to the corruption of the measurements y . Suppose now that we observe $y = \Phi x_0 + e$,

¹Throughout the entirety of this paper, we will forgo explicit calculation of constants which do not depend on the quantities of interest. These constants will be denoted everywhere as C .

where e is an unknown perturbation whose size can be bounded $\|e\|_{\ell_2} \leq h$. In place of (P_1) , we solve the relaxed problem

$$(P_2) \quad \min \|x\|_{\ell_1} \quad \text{such that} \quad \|\Phi x - y\|_{\ell_2} \leq h.$$

Call \hat{x} the solution to (P_2) . In [4], it was shown that if the restricted isometry constants for sets of size S obey (2), then the recovery error can be bounded by

$$\|\hat{x} - x_0\|_{\ell_2} \leq C \cdot (h + S^{-1/2} \cdot \|x_{0,S} - x_0\|_{\ell_1}) \quad (5)$$

where C is a small constant. Thus the error in the recovery is on the same order as the larger of the approximation error and the size of the measurement error.

Finally, it should be mentioned how to construct measurement ensembles with isometry constants as in (2) for “useful” set sizes S . In fact, doing so is not difficult at all; in some sense, almost any measurement ensemble will do. More precisely, if we create a Φ by first taking each entry as a independent Gaussian random variable with zero mean and unit variance, and then renormalize the columns so that they have unit length, then (2) is obeyed for [5,6]

$$S \leq C \cdot \frac{K}{\log(N/K)} \quad (6)$$

with overwhelming probability (see also [11] for a different formulation). We will refer to such a Φ as a Gaussian measurement ensemble.

Combining equations (4) and (6), we see how the (noiseless) reconstruction error goes to zero as the number of measurements increases:

$$\|\hat{x} - x_0\|_{\ell_2} \leq C \cdot \left(\frac{K}{\log(N/K)} \right)^{-\alpha}. \quad (7)$$

Let us compare, for a moment, the two processes. To form the best S -term approximation $x_{0,S}$, we must first measure the entire (all N points of the) discrete signal, and then cherry-pick the S most important. In doing so, the nonlinear approximation error goes to zero as $S^{-\alpha}$. What (7) tells us is that we can get an approximation of essentially the same quality from far fewer *nonadaptive* measurements. To within a log factor, the rate at which the reconstruction error of the recovered (via (P_1)) signal approaches zero as the number of measurements increases is the same as the nonlinear approximation rate. Using this scheme, the number of measurements required depends more on the inherent complexity of the signal, rather than its resolution.

Below, we will show that the stability of (P_2) allows us to translate these approximation results into rate-distortion results for a straightforward encoding/decoding scheme.

1.2 Encoding the ℓ_p ball

Before presenting our encoding algorithm, we will quickly review some classical results about the fundamental limits of encoding elements of the ℓ_p ball.

The optimal (in the minimax sense) way to complete the task of encoding elements from a known compact subset \mathcal{C} of \mathbb{R}^N was formulated in simple terms by Kolmogorov [12]. We start with the definition² of an ϵ -net: a finite subset $\mathcal{N} \subset \mathcal{C}$ such that for every point $c \in \mathcal{C}$ there is a point $u \in \mathcal{N}$ $\|c - u\|_{\ell_2} \leq \epsilon$. An ϵ -net \mathcal{N} implicitly specifies a *code* for elements in \mathcal{C} with distortion ϵ . Given an element $c \in \mathcal{C}$, the encoder “sends” the index to the closest point u in \mathcal{N} . No matter which point c was chosen, the difference between the original point and the codeword is bounded by ϵ . The number of bits used to index the u is of course $\lceil \log |\mathcal{N}| \rceil$.

The question naturally arises of how well we can do with a fixed number of bits. That is, given a bit budget R , what is smallest distortion which we can guarantee over the entire set of interest \mathcal{C} ? This quantity is called the *Kolmogorov entropy* [12]:

$$\mathcal{E}_R(\mathcal{C}) = \inf\{\epsilon' > 0 : \text{there exists an } \epsilon'\text{-net of size } 2^R \text{ that covers } \mathcal{C}\}.$$

By definition, $\mathcal{E}_R(\mathcal{C})$ is a bound on the minimax distortion-rate curve for all encoding algorithms, and will be the benchmark for how we judge the encoder/decoder pair proposed in Section 2.

We are particularly interested in the set $\mathcal{C} = B_{p,N}$ (sets of sparse signals). In [13,15], a lower bound on the Kolmogorov entropy of $B_{p,N}$ is calculated:

$$\mathcal{E}_R(B_{p,N}) \geq C \cdot \left(\frac{R}{\log(N/R) + 1} \right)^{-\alpha} \quad (8)$$

for the range $\log N \leq R \leq N$. (Bounds for larger R exist, but are not interesting from a theoretical standpoint in that they are straightforward to achieve. We are analyzing the regime where we use fewer bits than there are components of the signal x_0 —e.g. less than 1 bit-per-pixel in image coding.) If our encoding algorithm achieves something similar to (8), we will consider it efficient.

2 A simple encoding algorithm

The recovery results of Section 1.1 suggest a straightforward encoding/decoding algorithm with the following blueprint:

1. **Measure.** K linear measurements of the form (1) are made using a measurement ensemble known to both the encoder and decoder. We will assume that

²Here and throughout the paper we use the standard Euclidean metric to measure approximation errors. The definitions can of course be extended to general metric spaces.

the restricted isometry constants of the measurement ensemble obey (2) for S as in (6).

2. **Quantize.** Each entry in the measurement vector y is quantized using a uniform scalar quantizer with stepsize Δ . That is, the vector y^q is formed by “rounding” each y_k to the closest value in $\{\iota \cdot \Delta, \iota \in \mathbb{Z}\}$. In the Appendix, it is shown that $|y_k| \leq 1$; as a result, each quantized observations can be indexed with $\approx \log(1/\Delta)$ bits. The total quantization error $\|y - y^q\|_{\ell_2}$ will be less than $\sqrt{K} \cdot \Delta$.
3. **Send.** The quantized measurements y^q are communicated to the receiver.
4. **Decode.** The signal is reconstructed by solving (P_2) with $h = \sqrt{K} \cdot \Delta$.

The central question remains: Can we choose values of K, Δ so that the performance of our codec behaves as (8)? The following calculations show that we can.

From (5), we know that the difference between the original signal x_0 and the signal \hat{x} reconstructed at the decoder is at most

$$\|x_0 - \hat{x}\|_2 \leq C \cdot \sqrt{K} \cdot \Delta + C \cdot \left(\frac{K}{\log(N/K)} \right)^{-\alpha}. \quad (9)$$

The term on the left is the *quantization error*, while the term on the right is the *approximation error*. For the encoding process to be universal, we need the quantization error (as a function of the number of bits used R) to go to zero faster than $R^{-\alpha}$ for all $\alpha > 0$. With stepsize Δ in the uniform quantization scheme outline above, the total number of bits used is

$$R = K \cdot \log \left(\frac{1}{\Delta} \right),$$

meaning that given a bit budget R , we choose

$$\Delta = 2^{-R/K}.$$

To have $\sqrt{K}\Delta = o(R^{-\alpha})$ for all $\alpha > 0$, we a slightly smaller number of observations K than bits,

$$K = \frac{R}{(\log R)^2},$$

(since $2^{-(\log R)^2}$ goes to zero faster than any power of R). Then to write the approximation error in terms of the number of bits spent, we use

$$\frac{K}{\log(N/K)} = \frac{R}{(\log R)^2 \log(N/K)} \leq \frac{R}{(\log R)^2 \log(N/R)},$$

where the last inequality follows from the fact that $R > K$. In the context of our encoding scheme, (5) becomes

$$\begin{aligned} \|x_0 - \hat{x}\|_2 =: \epsilon(R) &\leq C \cdot \left(2^{-(\log R)^2} + \left(\frac{R}{(\log R)^2 \log(N/R)} \right)^{-\alpha} \right) \\ &\leq C_\alpha \cdot \left(\frac{R}{(\log R)^2 (\log(N/R) + 1)} \right)^{-\alpha}, \end{aligned} \quad (10)$$

where the constant in the second expression depends on α (which is unknown to the encoder/decoder pair).

The performance of our simple encoder/decoder pair is within a log term of the optimal asymptotic performance dictated by (8). Note that we were a bit sloppy above in using \log^2 above; the result (10) will hold for $(\log R)^{1+\gamma}$ in the denominator in place of $(\log R)^2$ for all $\gamma > 0$.

3 Applications in imaging

The theory above can also be applied to imaging object of a continuous variable. We will consider the example where we are interested in capturing a photograph-like image f modeled as a piecewise-smooth function on the continuous domain $[0, 1]^2$; we will consider f as a member of the class of functions which are twice continuously differentiable everywhere except along contours which are themselves twice differentiable (as a function of one variable, see [2]). We will measure and subsequently reconstruct f in an appropriate finite dimensional subspace V of $L_2([0, 1]^2)$, chosen such that the distance between f and its projection f_V onto V is on the same order as the distance between f_V and the reconstruction \hat{f}_V from the measurements taken in V .

To make this concrete, we will use as our finite dimensional subspace of L_2 a *wavelet scaling space* [14]. The scaling space V_J at scale J has dimension $N = 2^{2J}$, and can be orthogonally decomposed by a collection of N wavelet basis functions $\{\psi_n\}$. The measurement test functions are members of V_J ; the measurements are made by

$$y_k = \int_{t \in [0,1]^2} f(t) \phi_k(t) dt. \quad (11)$$

The distance between a piecewise smooth image and its projection f_{V_J} onto V_J is [14]

$$\|f - f_{V_J}\|_{\ell_2} \leq C \cdot 2^{-J/2} = C \cdot N^{-1/4}, \quad (12)$$

we will discuss how to choose N below. A suitable set of K test functions can be generated from a $K \times N$ Gaussian measurement ensemble Φ simply by setting

$$\phi_k(t) = \sum_n \Phi_{k,n} \psi_n(t),$$

where $\Phi_{k,n}$ is the entry in the k th row and n th column of Φ . Measuring f as in (11) is equivalent to calculating the wavelet coefficients $w_n(f) = \langle f, \psi_n \rangle$, and then using Φ to measure $w = \{w_n\}$ as in the finite case (1). These measurements are quantized as before, and the decoder solves the (finite dimensional) optimization problem

$$\min \|w\|_{\ell_1} \quad \text{such that} \quad \|\Phi w - y\|_{\ell_2} \leq \epsilon \quad (13)$$

and chooses as the continuous-space reconstruction $\hat{f} = \sum_n \hat{w}_n \psi_n$, where \hat{w} is the solution to (13) above. By design, we will have

$$\|\hat{f} - f\|_{L_2} \leq C \cdot N^{-1/4} + \epsilon(R). \quad (14)$$

Now suppose that we would like to code a piecewise smooth image f using R bits. The wavelet coefficients of such images will lie in the ℓ_1 ball in V_J [14]. From (14) and our result (10) in the previous section, we can choose $N = R^2$ and have reconstruction error on the order of

$$\|\hat{f} - f\|_{L_2} \leq C \cdot \left(R^{-1/2} + \left(\frac{R}{\log^3 R} \right)^{-1/2} \right) \leq C' \cdot \left(\frac{R}{\log^3 R} \right)^{-1/2}.$$

Sparser representations for piecewise smooth images exist. We could just as easily represent functions in V_J using *curvelets* [2]. The decoder will solve a slightly different problem, searching for the sparsest set of curvelet coefficients that explain the observations:

$$\min \|Hw\|_{\ell_1} \quad \text{such that} \quad \|\Phi w - y\|_{\ell_2} \leq \epsilon \quad (15)$$

where H maps a sequence of wavelet coefficients (which we are using as the canonical basis for V_J) into the corresponding sequence of curvelet coefficients. The curvelet coefficients of a piecewise smooth image³ lie in the $\ell_{2/3}$ ball [2], and so the reconstruction error will obey

$$\|f - \hat{f}\|_{L_2} \leq C \cdot \left(\frac{R}{\log^3 R} \right)^{-1}. \quad (16)$$

bits to code the measurements. The recovery error (16) matches the best asymptotic coding rates published to date for this class of images (see [2, 16]), **even though the encoder only has access to a limited number of measurements.**

The result (16) is near-optimal, as the Kolmogorov entropy for this class of piecewise smooth images can be bounded below by $C \cdot \epsilon^{-1}$ [7]. Hence even though we are given relatively few measurements of the image, we can still encode it efficiently.

4 Discussion

To put our results in context, let us compare the procedure outlined in Section 2 to that used by a common measurement/encoding device: the digital camera. Roughly speaking, when a picture is taken using a digital camera, an image is formed by taking N measurements (one for each pixel), transformed into a set of N wavelet

³The situation is slightly more complicated here in that the curvelets introduced in [2] are an overcomplete frame, and not an orthobasis; H in (15) is an $N \times N'$ matrix, where $N' = C \cdot N$, $C \approx 4$. Nevertheless, a result such as (5) exists, and we can proceed as before.

coefficients (for example), and then coded, which usually involves discarding many (or even most) of the computed wavelet coefficients. Our procedure, on the other hand, takes about as many measurements (to log factors) as bits it will produce, and requires a bare minimum of computation.

The benefits of our encoding algorithm are manifold. First, the measuring/encoding device is extremely simple; it just needs to take linear measurements of the form (1) and quantize them. The vast majority of the computation is done at the decoder. Second, the procedure is perfectly robust against "packet loss": none of the observations are more important than others, as the number of observations received by the decoder decreases, the reconstruction quality degrades gracefully. Finally, the data received by the decoder can be used in many different ways. As sparser representations for signals/images are discovered through advances in applied harmonic analysis, the reconstruction from the same data will improve.

We will close by briefly mentioning some previous work in universal encoding. In [10], a near-optimal encoding strategy for coding elements from (weak) ℓ_p was introduced. The basic strategy was to measure each coefficient, figure out which ones are important, code their locations and quantize their amplitudes. This strategy, although minimax near-optimal in the same way as ours, is not universal (the quantization the encoder uses depends on the parameter α), and works from a "full set" of measurements.

Several universal encoding strategies for continuous-time signals and images in certain smoothness classes have been developed in recent years [1, 7, 8]. Again, these algorithms work by measuring many wavelet coefficients, coding the important ones, and discarding the rest. The algorithm presented in Section 2 is less wasteful in that it uses (quantizes) everything that it measures.

5 Appendix

We will show that the measurements $y = \Phi x_0$ of a signal $x_0 \in B_{p,N}$, $p \leq 1$, where Φ is a Gaussian measurement ensemble, have energy that can be bounded by a constant. Let Φ_j be the j th column of Φ . Then

$$\begin{aligned} \|y\|_{\ell_2} &= \left\| \sum_{j=1}^N \Phi_j x_0(j) \right\|_{\ell_2} \\ &\leq \sum_j \|\Phi_j x_0(j)\|_{\ell_2} \\ &\leq \sum_j |x_0(j)| \\ &\leq 1 \end{aligned}$$

where the second-to-last inequality follows from the fact that the columns of Φ are normalized, and the last inequality follows from the fact that $x \in B_{p;N} \subset B_{1;N}$.

Of course, we also have that

$$\|y\|_{\infty} \leq 1,$$

a fact which is used in Section 2.

References

- [1] L. Birgé and P. Massart. An adaptive compression algorithm in Besov spaces. *Constr. Approx.*, 16(1):1–36, 2000.
- [2] E. Candès and D. L. Donoho. New tight frames of curvelets and optimal representations of objects with piecewise c^2 singularities. *Comm. on Pure and Applied Math.*, 57:219–266, 2004.
- [3] E. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *submitted to IEEE Trans. Inform. Theory*, June 2004. Available on the ArXiv preprint server: [math.GM/0409186](https://arxiv.org/abs/math.GM/0409186).
- [4] E. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *submitted to Communications on Pure and Applied Mathematics*, March 2005.
- [5] E. Candès and T. Tao. Near-optimal signal recovery from random projections and universal encoding strategies. *submitted to IEEE Trans. Inform. Theory*, November 2004. Available on the ArXiv preprint server: [math.CA/0410542](https://arxiv.org/abs/math.CA/0410542).
- [6] E. J. Candès and T. Tao. Decoding by linear programming. *submitted to IEEE Trans. Inform. Theory*, December 2004.
- [7] A. Cohen, W. Dahmen, I. Daubechies, and R. A. DeVore. Tree approximation and optimal encoding. *Appl. Comp. Harmonic Analysis*, 11(2):192–226, 2001.
- [8] A. Cohen, I. Daubechies, O. G. Guleryuz, and M. T. Orchard. On the importance of combining wavelet-based nonlinear approximation with coding strategies. *IEEE Trans. Inform. Theory*, 48(7):1895–1921, July 2002.
- [9] R. A. DeVore. Nonlinear approximation. *Acta Numerica*, 7:51–150, 1998.
- [10] D. L. Donoho. Unconditional bases and bit-level compression. *Appl. Comp. Harmonic Analysis*, 3(4):388–392, 1996.
- [11] D. L. Donoho. Compressed sensing. *submitted to IEEE Trans. Inform. Theory*, September 2004.

- [12] A. N. Kolmogorov and V. M. Tihomirov. ϵ -entropy and ϵ -capacity of sets in functional spaces. *Translations of the AMS*, 17:277–364, 1961.
- [13] T. Kühn. A lower estimate for entropy numbers. *J. Approx. Theory*, 110:120–124, 2001.
- [14] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, San Diego, second edition, 1999.
- [15] C. Schutt. Entropy numbers of diagonal operators between symmetric Banach spaces. *J. Approx. Theory*, 40:121–8, 1984.
- [16] M. B. Wakin, J. Romberg, H. Choi, and R. G. Baraniuk. Wavelet-domain approximation and compression of piecewise smooth images. *to appear in IEEE Trans. Image Proc.*, 2005.