

Exact reconstruction of sparse signals via nonconvex minimization

Rick Chartrand
Los Alamos National Laboratory
EDICS: DSP-TFSR

Abstract—Several authors have shown recently that is possible to reconstruct exactly a sparse signal from fewer linear measurements than would be expected from traditional sampling theory. The methods used involve computing the signal of minimum ℓ^1 norm among those having the given measurements. We show that by replacing the ℓ^1 norm with the ℓ^p norm with $p < 1$, exact reconstruction is possible with substantially fewer measurements. We give a theorem in this direction, and many numerical examples, both in one complex dimension, and larger-scale examples in two real dimensions.

I. INTRODUCTION

There has been much recent research (e.g., [1], [2], [3]) on the subject of reconstruction of sparse signals from a limited number of linear measurements. This topic is known in the literature as *compressed sensing* or *compressive sampling*, and has some overlap with *basis pursuit* [4]. There are many relevant results, but we focus here on those typified by the following example, from [1].

We consider an $M \times N$ measurement matrix Φ such that for $x \in \mathbb{C}^N$, $y = \Phi x$ is the vector of Fourier coefficients of x at M randomly chosen frequencies. Suppose that the *sparsity* of x is K ; that is, x has K nonzero elements. This can be stated in terms of the ℓ^0 norm of x : $\|x\|_0 = K$. (This is a standard abuse of terminology: $\|\cdot\|_0$ is not positive homogeneous, yet is referred to as a norm.) Then there is a constant C , not depending on K or N , such that whenever $M > CK \log N$, one can reconstruct x exactly, with very high probability, as the solution of the following optimization problem:

$$\min_u \|u\|_0, \quad \text{subject to } \Phi u = y. \quad (1)$$

(The constant C can depend on how small the probability of failure is to be guaranteed to be; it scales as N^{-m} where m depends on C .) Importantly, this result continues to hold if the ℓ^0 norm is replaced by the ℓ^1 norm, resulting in a convex problem:

$$\min_u \|u\|_1, \quad \text{subject to } \Phi u = y. \quad (2)$$

The problem (1) is NP-hard [5], while the convex problem (2) can be solved efficiently. That the former can be replaced by the latter is the reason for the surge in recent interest.

It should be pointed out that the above is not special to Fourier measurements. Similar results hold [3] if the elements

of Φ consist of random samples from the standard normal distribution, or if each element is 1 or -1 with equal probability. The key is that with high probability, Φ will be a sampling from a basis that is *incoherent* with the standard basis in \mathbb{C}^N . Just as a signal cannot be localized in both time and frequency, there is a lower bound on the combined support size of x and Φx .

It is natural to ask what happens if the ℓ^1 norm is replaced by the ℓ^p norm for some $p \in (0, 1)$. (Recall that for $0 < p < \infty$, $\|u\|_p^p = \sum_{j=1}^N |u_j|^p$. As with $\|\cdot\|_0$, $\|\cdot\|_p$ is not a norm when $0 < p < 1$, though $\|\cdot\|_p^p$ satisfies the triangle inequality and induces a metric.) The resulting optimization problem will not be convex, and is described in the literature as intractable. Be that as it may, in this note we demonstrate that the much simpler task of finding a local minimizer can produce exact reconstruction of sparse signals with many fewer measurements than when $p = 1$. We begin with theoretical results concerning when a global minimizer is guaranteed to be an exact reconstruction. We follow with numerical evidence, a suite of one-dimensional examples and some larger-scale, two-dimensional examples. The examples are designed for direct comparison with results in [1].

II. RESTRICTED ISOMETRY CONSTANTS

We begin with a theoretical result concerning the circumstances under which the solution to the problem described in the introduction,

$$\min_u \|u\|_p, \quad \text{subject to } \Phi u = \Phi x, \quad (3)$$

equals exactly x . For this we need the concept of an approximate S -restricted isometry, as introduced in [6]. For $T \subset \{1, \dots, N\}$, let Φ_T be the matrix consisting of the columns φ_j of Φ for $j \in T$. For each number S , define the S -restricted isometry constant of Φ to be the smallest $\delta_S \geq 0$ such that for all subsets T with $|T| \leq S$ and all $c \in \mathbb{C}^{|T|}$,

$$(1 - \delta_S) \|c\|_2^2 \leq \|\Phi_T c\|_2^2 \leq (1 + \delta_S) \|c\|_2^2. \quad (4)$$

Then we have the following sufficient condition for exact reconstruction.

Theorem 1: Let Φ be an $M \times N$ matrix. Let $x \in \mathbb{C}^N$ and let $K = \|x\|_0$ be the size of the support of x . Let $p \in (0, 1]$, $b > 1$, and $a = b^{p/(2-p)}$. Suppose that Φ satisfies

$$\delta_{aK} + b\delta_{(a+1)K} < b - 1. \quad (5)$$

Then the unique minimizer of (3) is exactly x .

Address: Theoretical Division, T-7, MS B284, Los Alamos National Laboratory, Los Alamos, NM 87544, USA. Phone: +1 +505 667-8093. Fax: +1 +505 665-5757. E-mail: rickc@lanl.gov.

This was stated by the author in [7] in the particular case of $b = 3$ and for real-valued matrices, but the same proof works *mutatis mutandis*. The theorem generalizes the corresponding result of Candes et al. for $p = 1$ and $b = 3$ in [8]. For a given Φ , the restricted isometry condition (5) will hold for larger values of K when $p < 1$. It is argued in [9] that the condition of the theorem holds when $M > CK \log N$ for a constant C , and in [1] that this relation is the best possible. The theorem suggests, however, that ℓ^p minimization with $p < 1$ can reduce the value of the constant.

The weakest possible restricted isometry condition that is sufficient to guarantee exact reconstruction is given by Candes and Tao [6], namely that the NP-hard problem (1) gives exact reconstruction if $\delta_{2K} < 1$. We point out that this is the limiting case of the previous theorem.

Corollary 2: Given Φ and x as in the Theorem, suppose $\delta_{2K+1} < 1$. Then there is $p > 0$ such that the unique minimizer of (3) is exactly x .

Proof: Choose b large enough that $\delta_{2K+1} < (b-1)/(b+1)$. Then choose $p > 0$ small enough that $a = b^{p/(2-p)} < 1 + 1/K$. Then

$$\delta_{aK} + b\delta_{(a+1)K} \leq (1+b)\delta_{(a+1)K} \leq (1+b)\delta_{2K+1} < (b-1). \quad (6)$$

The result now follows from the Theorem. ■

On the one hand, the value of these results depends on the ability to compute a global minimizer of a nonconvex functional. On the other hand, any local optimization method can do so if initialized by a point sufficiently close to the global optimum. While we provide no guarantees, the numerical results below strongly suggest that the least-squares solution is often or even always sufficiently close, if there are enough measurements.

III. NUMERICAL EXAMPLES

We proceed with a numerical investigation of the circumstances under which computing a local minimizer of (3) gives exactly x . As in the introduction, Φ will be a matrix having the action of evaluation of the Fourier transform at a randomly chosen collection of frequencies.

For our experiments, we fixed a signal length of $N = 512$. As in [1], the numbers of measurements M were 8, 16, ..., 128. For each value of K , the sparsities K are each multiple of $M/16$ up to M , rounded to the nearest integer. The values of p were 1, 0.95, 0.75, and 0.5. (Smaller values of p can be used, but only with more careful supervision of parameters than that afforded by the automated approach described below.) The experiment was repeated 10 times for each value of M and K . For each repetition, a spike-train signal $x = \sum_{k=1}^K z_k e_k$ was constructed by randomly choosing K elements of the standard basis $\{e_1, e_2, \dots, e_N\}$ of \mathbb{C}^N , then randomly choosing the real and imaginary parts of each z_k from the standard normal distribution. Conceptually, the matrix Φ consisted of a random selection of M rows from the discrete Fourier transform (DFT) matrix. In actual fact, instead of multiplication by a matrix Φ , the fast Fourier transform would be computed and the corresponding M elements selected.

We adopt a simple computational approach for computing local minimizers of problem (3), namely gradient descent with projection. More sophisticated algorithms are possible; for example, the FOCUSS approach of Rao and Kreutz-Delgado [10] (though this would not be feasible for the examples of the next section, where a gradient operator needs to be incorporated). We are not presenting a new algorithm, merely examining the potential for ℓ^p minimization to improve signal reconstruction in the compressed sensing framework. Similarly, we always begin our iteration with the easily-computed minimum- ℓ^2 norm fit to the data, as this is the most sensible from the perspective of signal reconstruction. Changing the initial iterate would be expected to change the computed local minimizer.

Our iteration is of the form

$$d_n = -|u_n|^{p-2}u_n, \quad u_{n+1} = P_{\Phi u = \Phi x}(u_n + t_n d_n), \quad (7)$$

where the absolute value and arithmetic operations involving u_n are componentwise. We thus compute the steepest descent direction d_n at u_n of the ℓ^p norm, or, more precisely, of $\|\cdot\|_p^p/p$. We then take a step in the direction d_n with a stepsize factor t_n , then project orthogonally onto the affine constraint space $\Phi u = \Phi x$. The projection onto $\Phi u = \Phi x$ was performed every iteration, by computing the DFT of u , setting the previously-chosen M frequency coefficients to those of x , then computing the inverse DFT. The stepsize t_n was chosen every 100 iterations by an exact line search; in other words, it is chosen to give the minimum value of $\|P_{\Phi u = \Phi x}(u_n + t_n d_n)\|_p$. The same t_n was then used until the next 100 iterations had been computed. To avoid division by zero, $|u_n|$ was replaced by $\sqrt{|u_n|^2 + \epsilon^2}$, where ϵ was decremented by 5% every 100 iterations, beginning with 0.01. As mentioned above, the iteration was begun with the minimum-energy solution, in other words the solution u_0 to (3) with $p = 2$. This is simply the projection under $P_{\Phi u = \Phi x}$ of the zero vector. The iteration was run for 40000 iterations, by which point the ℓ^p norm had generally ceased to decrease significantly. The final value of ϵ was thusly 1.23×10^{-11} . The elapsed time was approximately 30 seconds on a 1.3-MHz laptop, running in MATLAB.

For each repetition, the reconstruction was deemed exact if the obtained minimizer u^* satisfied $\|u^* - x\|_\infty < 10^{-6}$. In each of several checked instances when this condition was satisfied, further iteration would produce a solution within less than 10^{-13} of x , an exact reconstruction by any reasonable numerical measure. The percentage of exact reconstructions was recorded for each K and M . It should be noted that in the case of exact reconstruction, there is no way to determine whether the minimizer is global. However, Theorem 1 suggests that it may be, especially for values of K and M for which exact reconstruction was always observed. (Directly checking the condition of the theorem is not feasible.)

A pair of examples are in Figures 1 and 2. The sparsity is $K = 16$, and $M = 48$ measurements were made. The signal and measurement frequencies are the same for both examples. What differs is that $p = 1$ in the first case and $p = 0.5$ in the second case. When $p = 1$, the minimizer u^* differs substantially from the original signal x . The maximum discrepancy is $\|u^* - x\|_\infty = 1.08$, and the relative error is

$\|u^* - x\|_2 / \|x\|_2 = 0.407$. The ℓ^1 norm of the minimizer is indeed lower than that of the signal, $\|u^*\|_1 = 22.7$ versus $\|x\|_1 = 23.4$. In the $p = 0.5$ case, the reconstruction is exact: $\|u^* - x\|_\infty = 7.61 \times 10^{-7}$ after 40000 iterations. Another 50000 iterations gave $\|u^* - x\|_\infty = 3.33 \times 10^{-14}$.

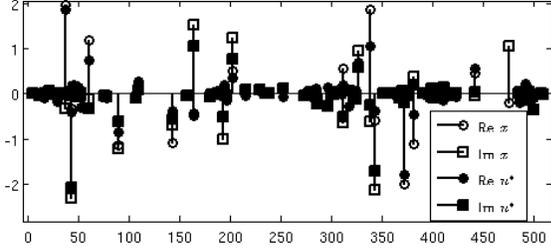


Fig. 1. The real and imaginary parts of a signal x (open markers) and the solution u^* (filled markers) to problem (2). The reconstruction is not even approximately x .

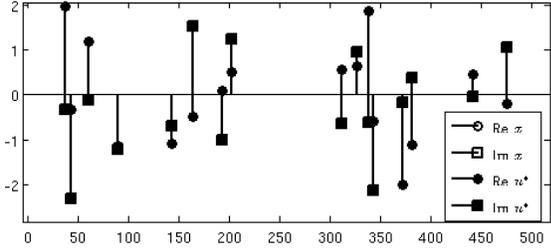


Fig. 2. Same as the previous figure, except u^* is the solution to problem (3) with $p = 0.5$. The reconstruction is exactly x .

Figure 3 shows the results of the experiments, in the format of [1]. The number of measurements M is on the vertical axis and the ratio K/M of sparsity to number of measurements is on the horizontal axis. (Only values of K/M up to $3/4$ are displayed.) The intensity gives the observed frequency of exact reconstruction, with white being 100% and black being 0%. Improvement over the $p = 1$ case is clear even for $p = 0.95$. For $p = 0.5$, exact reconstruction is obtained with a sparsity-to-measurement ratio nearly double that required when $p = 1$.

An additional experiment was conducted to provide an example of the result of varying the number of measurements with a fixed sparsity. We used $K = 16$ and values of M ranging from 24 to 80, for the same values of p used above. The results are in Figure 4. It is remarkable that even a value of p only slightly less than 1 gives exact reconstruction for significantly fewer measurements (12% fewer, in this example). Decreasing p further decreases the required number of measurements, but it appears that the smaller the value of p , the less improvement is seen for a given decrease in p .

A. Two-dimensional example

We present a larger-scale, two-dimensional example, designed as in [1], where the signal x is the 256×256 Shepp-Logan phantom (Figure 5(a)). Here, it is the gradient of the signal that is sparse, with $\|\nabla x\|_0 = 3627$ (out of 65536

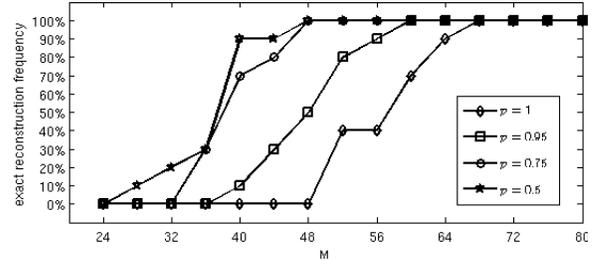


Fig. 4. Observed probabilities of exact reconstruction for a signal of sparsity $K = 16$, for different numbers of measurements M and four values of p . Compared with $p = 1$, substantially fewer measurements are required for exact reconstruction when $p = 0.95$. Decreasing p further decreases the required value of M , but by less and less as p gets smaller.

pixels). The compressed-sensing theory applies in this setting as well, as is well known [1]. The problem (3) becomes:

$$\min_u \|\nabla u\|_p, \quad \text{subject to } \Phi u = \Phi x. \quad (8)$$

As in [1], the Fourier coefficients measured by Φ were not at random frequencies, but instead along radial lines in frequency space (as in Figure 5(b)). By the Fourier slice theorem, sampling along a radial line in frequency space is equivalent to sampling (the Fourier transform of) the Radon transform along the angle determined by the line. The experiment is equivalent to attempting to reconstruct the phantom from a limited number of radiographic projections (or views).

We used gradient descent with projection as in (7), though now the descent direction is

$$d_n = \nabla \cdot \left((\sqrt{|\nabla u_n|^2 + \epsilon^2})^{p-2} \nabla u_n \right), \quad (9)$$

obtained from the Euler-Lagrange equation for $\min_u \|u\|_p$. Gradients were computed with simple forward differencing, with backward differencing used for divergences, and Neumann boundary conditions. The value of ϵ was initialized to 10, fixed until the ℓ^p norm had ceased to decrease significantly, then decreased by a factor of 10. This process was continued until no further decrease was observed, or until the ℓ^p norm was less than that of the phantom. Reconstruction was deemed exact if the computed minimizer was within 10^{-13} of the phantom at every pixel. The number of iterations required for exact reconstruction varied widely, generally fewer for smaller p , and more when the number of measurements was near minimal. For example, with $p = 0.5$, 85000 iterations were needed with 10 views, 10000 with 18 views, while 452000 were needed for 18 views with $p = 1$. Iterations were computed at a rate of about 1000 per minute.

In [1], exact reconstruction is claimed for 22 views in the $p = 1$ case, for which there are 5481 Fourier coefficients measured, for $M = 10962$ total measurements (counting both the real and imaginary parts). In fact, we find that 18 views suffice (or $M = 9010$; see Figure 5(d)), while 17 views do not ($M = 8514$). For $p = 0.5$, however, 10 views ($M = 5042$) are sufficient for exact reconstruction. Decreasing p further was not found to decrease the number of views required.

IV. CONCLUSIONS

The ability to reconstruct signals from very few measurements is an important development in signal processing. In

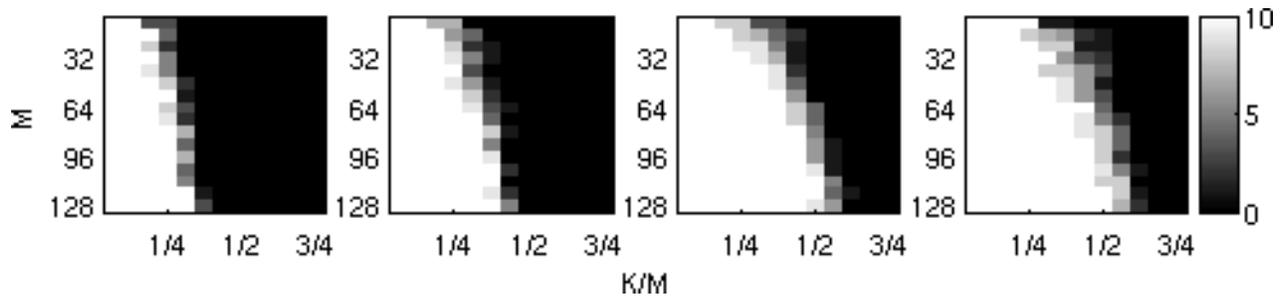


Fig. 3. Observed probabilities of exact reconstruction for different numbers of measurements (M) and sparsity-to-measurements ratio (K/M), for four values of p : $p = 1$ (left), $p = 0.95$ (second from left), $p = 0.75$ (second from right), and $p = 0.5$ (right). Compared with $p = 1$, exact reconstruction is obtained with larger values of K/M even for $p = 0.95$, and almost double the value of K/M for $p = 0.5$.

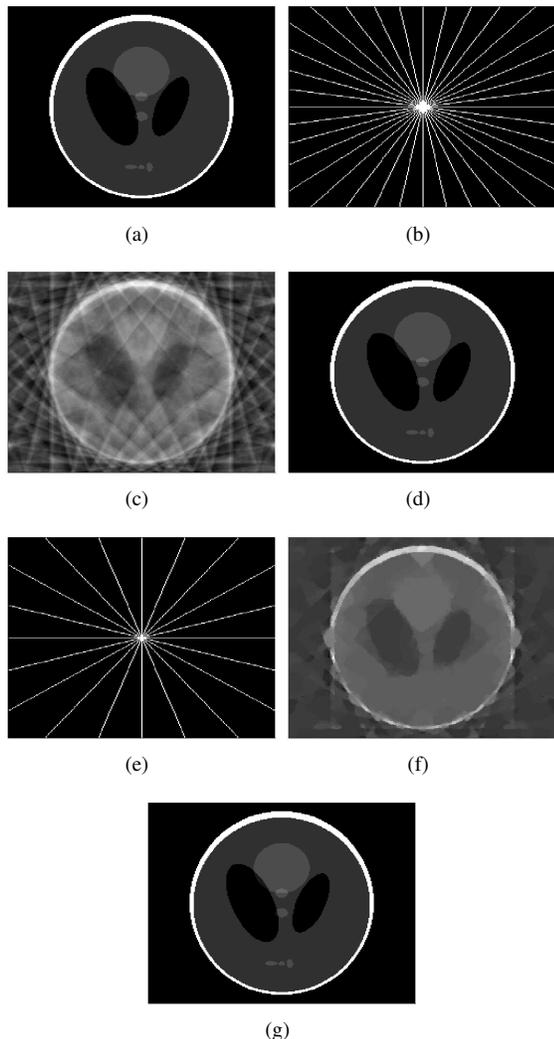


Fig. 5. (a) The 256×256 Shepp-Logan phantom. (b) The white pixels show where the Fourier transform is sampled with 18 views. (c) The minimum-energy or backprojection reconstruction with 18 views is very poor. (d) The reconstruction from 18 views with $p = 1$ is exact. (e) The Fourier transform sample pattern for 10 views. (f) The reconstruction from 10 views with $p = 1$ is poor. (g) The reconstruction from 10 views with $p = 0.5$ is exact.

applications where data acquisition is expensive or difficult, compressed sensing can allow good results to be obtained in a manner that would once have been infeasible. In this note, we have seen that by computing local ℓ^p minimizers with $p < 1$, fewer measurements are required than previously observed. The required reconstruction time is generally longer than with $p = 1$, but much less than with $p = 0$. Moreover, the algorithmic approach in this note is relatively naive; more sophisticated ℓ^p optimization methods should reduce the reconstruction time further. This will increase the number of applications for which the ℓ^p approach is worthwhile.

REFERENCES

- [1] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, 2006.
- [2] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, pp. 1289–1306, 2006.
- [3] E. Candes and T. Tao, "Near optimal signal recovery from random projections: universal encoding strategies?," *IEEE Trans. Inf. Theory*, vol. 52, pp. 5406–5425, 2006.
- [4] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, pp. 33–61, 1998.
- [5] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, vol. 24, pp. 227–234, 1995.
- [6] E. Candes and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, pp. 4203–4215, 2005.
- [7] R. Chartrand, "Nonconvex compressed sensing and error correction," in *International Conference on Acoustics, Speech, and Signal Processing*, IEEE, 2007.
- [8] E. Candes, M. Rudelson, T. Tao, and R. Vershynin, "Error correction via linear programming," in *46th Annual IEEE Symposium on Foundations of Computer Science*, pp. 295–308, 2005.
- [9] E. J. Candes, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, pp. 1207–1223, 2006.
- [10] B. D. Rao and K. Kreutz-Delgado, "An affine scaling methodology for best basis selection," *IEEE Trans. Signal Process.*, vol. 47, pp. 187–200, 1999.