SPARSE GRADIENT IMAGE RECONSTRUCTION DONE FASTER

R. Maleh   A. C. Gilbert   M. J. Strauss

Department of Mathematics
University of Michigan
Ann Arbor, MI 48109

ABSTRACT

In a wide variety of imaging applications (especially medical imaging), we obtain a partial set or subset of the Fourier transform of an image. From these Fourier measurements, we want to reconstruct the entire original image. Convex optimization is a powerful, recent solution to this problem. Unfortunately, convex optimization in its myriad of implementations is computationally expensive and may be impractical for large images or for multiple images. Furthermore, some of these techniques assume that the image has a sparse gradient (i.e., that the gradient of the image consists of a few nonzero pixel values) or that the gradient is highly compressible. In this paper, we demonstrate that we can recover such images with GRADIENTOMP, an efficient algorithm based upon Orthogonal Matching Pursuit (OMP), more effectively than with convex optimization. We compare both the qualitative and quantitative performance of this algorithm to the optimization techniques.

Index Terms— image reconstruction, image edge analysis, algorithms, linear programming, Fourier transforms

1. INTRODUCTION

Because the discrete Fourier transform is invertible, it is a simple matter to recover an image or a signal given all of its Fourier coefficients. We might be led to believe that if a signal consists of d entries, then we *must* have all d Fourier coefficients to perform this inversion. A number of recent results [1] show that this is *not* the case. If the signal or image is sparse in certain domains (e.g., particular orthonormal bases), then we can, with high probability, reconstruct the signal exactly with many fewer Fourier coefficients than originally thought. Several researchers refer to this paradigm as "compressed sensing" or "compressive imaging."

There are a number of applications in which we cannot afford to collect a complete set of Fourier coefficients. For example, medical imaging devices such as CT scanners and MRIs tend to be noisy and uncomfortable for patients. It would be beneficial for patients if we could significantly reduce the number of measurements that these devices calculate to generate a high quality image. Both hyperspectral and regular imaging [2, 4] are beginning to include compressive imaging technologies. In all of these cases, the amount of reconstructed data is overwhelmingly large compared to the small amount of observed measurements. It is imperative that we develop both effective and efficient algorithms for reconstructing images from partial Fourier measurements.

The usual techniques (such as BasisPursuit, or BP) that perform sparse image reconstruction tend to be slow as they rely on convex optimization techniques. For large applications, such as video processing or hyperspectral image processing, there is too much data for this performance to be acceptable. The greedy algorithm Orthogonal Matching Pursuit (OMP) efficiently recovers images that are sparse with respect to the standard Euclidean basis[1]; however, such images are rare. Images that have a sparse gradient, that are constant valued (or nearly so) over large regions separated by edges, arise much more often. The BP algorithm can recover such images by making a small change to the objective function. The OMP algorithm, on the other hand, can not.

We present a new, efficient algorithm known as GRADIENTOMP, which will use OMP as a subroutine to recover images with a sparse gradient. Furthermore, we can use a much faster, provably correct algorithm in place of the OMP subroutine. We show that we can recover sparse gradient images much more efficiently than convex optimization methods and we demonstrate the trade-off between the algorithms in recovering noisy images and images which do not have a sparse gradient (e.g., textures).

2. PRELIMINARIES

Let $X \in \mathbf{C}^{d\times d}$ be an image with a sparse gradient. By sparse gradient, we mean that the image is sparse under the total-variation (TV) operator. More precisely, we have that $\|TV(X)\|_0 = T \ll d^2$ where $\|\cdot\|_0$ returns the number of

The authors are supported by NSF DMS 0354600 and DARPA/ONR N66001-06-1-2011. MJS is also supported by NSF DMS 0510203. ACG is an Alfred P. Sloan Fellow.
Email:{rmaleh,annacg,martinjs}@umich.edu

[1]Given measurements which are dot products of the signal against Gaussian random variables, we can prove that OMP recovers the signal [6].

non-zero pixels of its argument and

$$TV(X)_{i,j} = |\nabla X|_{i,j} = \sqrt{(X_{i+1,j} - X_{i,j})^2 + (X_{i,j+1} - X_{i,j})^2}.$$

Figure 1 is an image of the Shepp-Logan Phantom and its image under the TV operator. Observe that the latter image is sparse while the former is not.

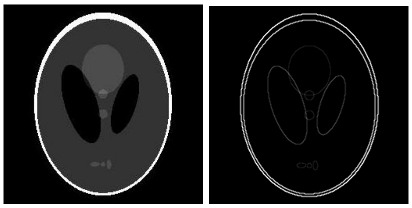

Fig. 1. Shepp-Logan Phantom and its edges.

Suppose that $\Omega \subset \{0, \cdots, d-1\}^2$ is some uniformly random subset of frequencies of size $|\Omega| = N$. Let $\mathcal{F}_\Omega : \mathbf{C}^{d \times d} \to \mathbf{C}^N$ denote the Fourier transform operator restricted to Ω and let $\mathcal{F}_\Omega X \in \mathbf{C}^N$ represent the given partial Fourier measurements of X. The image reconstruction problem is to find X given $\mathcal{F}_\Omega X$.

One approach to the problem is to solve the convex optimization problem

$$X_R = \underset{Y}{\operatorname{argmin}} ||TV(Y)||_1 \quad \text{s.t.} \quad \mathcal{F}_\Omega Y = \mathcal{F}_\Omega X. \tag{2.1}$$

In [1], the authors prove if $T = ||TV(X)||_0 \leq C(\log d)^{-1} N$, then, with high probability, the solution X_R to the convex program of Equation 2.1 is unique and equal to X. We refer to this procedure as TV-minimization.

There are two challenges in devising an efficient algorithm to recover $TV(X)$ from $\mathcal{F}_\Omega X$. First, we cannot directly obtain the Fourier coefficients of the image $TV(X)$ from the Fourier coefficients of the original image $\mathcal{F}_\Omega X$. Second, if we do reconstruct the image $TV(X)$, we cannot easily recover X as $TV(\cdot)$ is not an invertible operator.

The solution is to find a different sparsifying operator Ψ which fulfills the following three conditions.

1. Ψ must be invertible in a numerically stable way.
2. We must be able to express $\mathcal{F}_\Omega \Psi X$ in terms of our observed Fourier coefficients $\mathcal{F}_\Omega X$; i.e., there exists an operator Σ such that $\Sigma \mathcal{F}_\Omega X = \mathcal{F}_\Omega \Psi X$.
3. We must have $||\Psi X||_0 \leq K ||TV(X)||_0$ for some constant K, i.e. X must be at least asymptotically as sparse under Ψ as it is under the total-variation operator.

The first two conditions address the above challenges while the last condition guarantees that reconstring ΨX requires about the same amount of data as reconstructing $TV(X)$.

If we let Ψ be a directional derivative in some direction, then we satisfy the above criteria and can devise a rather naive algorithm which performs poorly. In the next section, we introduce a more sophisticated approach which relies on two orthogonal directional derivatives.

3. GRADIENT ORTHOGONAL MATCHING PURSUIT

Let $(D_v X)_{i,j} = X_{i,j} - X_{i-1,j}$ and $(D_h X)_{i,j} = X_{i,j} - X_{i,j-1}$ be two discrete derivatives in the vertical and horizontal directions, respectively. Observe that we can express these derivatives as linear functions of our partial Fourier measurements,

$$(D_v X) = (1 - e^{-2\pi i k_1/d}) \mathcal{F}_\Omega X$$
$$(D_h X) = (1 - e^{-2\pi i k_2/d}) \mathcal{F}_\Omega X$$

where $(k_1, k_2) \in \Omega$ and we perform the operations elementwise. Furthermore, these derivatives are invertible. Let

$$(D_v^{-1} X)_{i,j} = \sum_{k=1}^{i} X_{k,j}$$

denote the discrete anti-derivative in the vertical direction. We define $D_h^{-1} X$ similarly. Observe that if we know both discrete directional derivatives $D_v X$ and $D_h X$, we can recover X by solving a simple system of differential equations. First, we integrate horizontally and find

$$(D_v^{-1} D_v X)_{i,j} = X_{i,j} + H_{i,j}$$

where H is an arbitrary image that is constant in the vertical direction. Next, we solve for H by observing that its horizontal derivative is

$$(D_h H)_{i,j} = (D_h D_v^{-1} D_v X)_{i,j} - (D_h X)_{i,j}$$

and that we can invert the right-hand side of the above expression by computing its anti-derivative in the horizontal direction. Once we know H, we subtract it from $(D_v^{-1} D_v X)$ to estimate X. Assuming periodic boundary conditions and that the set of frequencies Ω contains the set $\{(k_1, 0), (0, k_2)\}$ with $k_1, k_2 \in \{0, \ldots, d-1\}$, then this process can be carried out easily and stably in the frequency domain.

The main idea of the algorithm is to reconstruct $D_v X$ and $D_h X$ independently using OMP, then to solve for X as outlined above.

Algorithm 1. GRADIENTOMP

INPUT:

- *A set of Fourier measurements $\mathcal{F}_\Omega X \in \mathbb{C}^N$ including the frequencies:* $\{(k_1, 0), (0, k_2) | k_1, k_2 \in \{0, \cdots, d-1\}\}$.
- *The numbers T_v and T_h of non-zero entries in $D_v X$ and $D_h X$ respectively.*

OUTPUT:

- *The reconstructed image X_R.*

PROCEDURE:

1. *Initialize the residuals*

$$r_{v,0} = (1 - e^{-2\pi i k_1/d}) \mathcal{F}_\Omega X$$
$$r_{h,0} = (1 - e^{-2\pi i k_2/d}) \mathcal{F}_\Omega X$$

where $(k_1, k_2) \in \Omega$.

2. *Perform OMP twice to recover $D_v X$ and $D_h X$ from $r_{v,0}$ with T_v iterations and $r_{h,0}$ with T_h iterations, respectively.*
3. *Solve for X_R using the equation*

$$X_R = D_v^{-1}[D_v X] + D_h^{-1}[D_h X] - D_h^{-1} D_v^{-1} D_h [D_v X]$$

where all items in brackets were estimated via OMP.

4. TV-MINIMIZATION VS. GRADIENTOMP

Although a great deal of work has been done in the area, no metric truly captures the details and errors that are apparent to the eye. As a result, we have both objective and subjective results comparing the two primary sparse gradient image reconstruction algorithms: TV-Minimization and GRADIENTOMP. We demonstrate that for images with sparse gradients, the GRADIENTOMP algorithm performs better than TV-minimization both in reconstruction error and runtime and that both algorithms' performance suffers on non-sparse images.

For the objective comparison, we used a 32×32 Shepp-Logan phantom. For various values of N (the number of given Fourier coefficients), we chose several uniform randomly generated frequency subsets Ω and compared the two algorithms using the following three criteria: ℓ_1 reconstruction error, probability of exact reconstruction, and runtime in seconds. For TV-Minimization, we used L1Magic [5]. The results of our experiment are shown in Figure 2. We can see that the ℓ_1 reconstruction error decays more rapidly with the GRADIENTOMP algorithm than with TV-minimization. Also, GRADIENTOMP's probability of exact reconstruction converges to one faster as a function of N. Finally, as expected, GRADIENTOMP has a much faster runtime.

For a subjective comparison, we compare the reconstructions of various images pictorially. We begin with a 512×512 Shepp-Logan phantom with 15% of its Fourier coefficients given. The various reconstructions can be seen in Figure 3. In this case, the TV-Minimization algorithm performed poorly whereas GRADIENTOMP quickly gave us a near perfect reconstruction on a large image using a small fraction of Fourier coefficients. The backprojection image refers to setting all unknown Fourier coefficients to zero and then applying an inverse Fourier transform.

Next we test a 128×128 Shepp-Logan phantom with AWGN ($SNR \approx 20$ dB). With medical images, it is typical to assume that the SNR will be worse than 40 dB. The reconstructions are shown in Figure 4. It is interesting to observe that TV-minimization gives a cleaner image but softens the details. As a result, some features such as the small ellipses in the bottom part of the image are missing. However, in the GRADIENTOMP reconstruction, despite the obvious vertical line errors present, all the important features of the original image are present. Thus, we have a clear trade-off between the two algorithms. When TV-minimization begins to fail, it smooths out important details whereas when

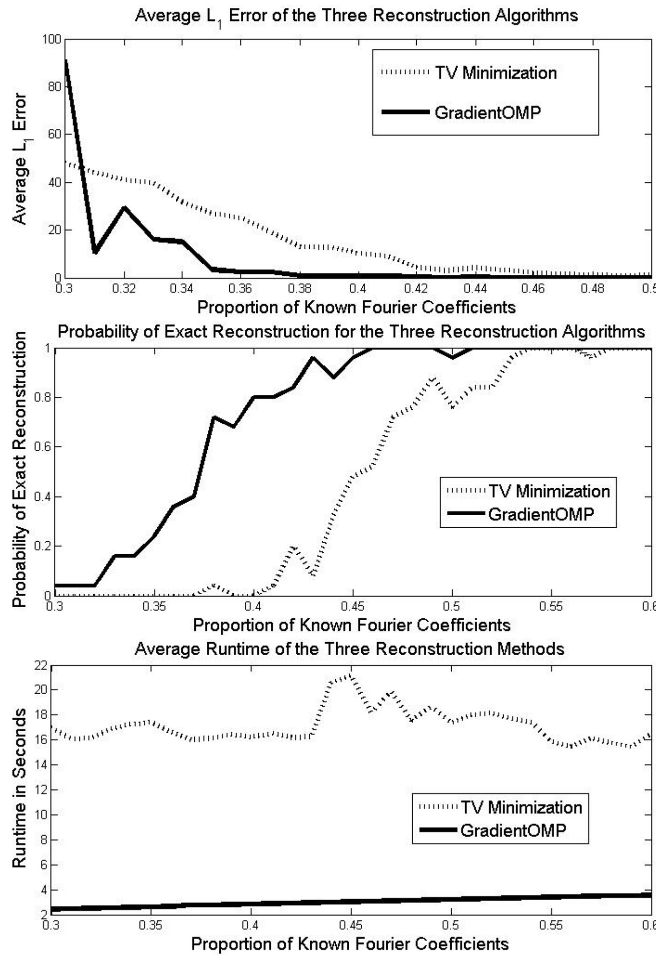

Fig. 2. The ℓ_1 reconstruction errors, probabilities of exact reconstruction, and runtimes of TV-minimization and GRADIENTOMP

GRADIENTOMP begins to fail, it introduces vertical stripes but maintains the important details. The larger the image, the less these stripes impact the appearance of the image. A very surprising result of our experiment is that the GRADIENTOMP reconstruction has a smaller ℓ_1 error than the original noisy image.

Next we compare the performance of both algorithms on images that do not have sparse gradients. The first image is a 128×128 version of Lena. The second image is a 128×128 straw texture whose nature is far from that of a sparse gradient image. The results can be seen in Figure 5. Both TV-minimization and GRADIENTOMP yield poor reconstructions in these extreme cases. The former does seem to perform slightly better, but not without the cost of severe high frequency attenuation.

5. CONCLUSIONS AND FUTURE WORK

The GRADIENTOMP algorithm is a new algorithm which significantly reduces the amount of time needed to recover

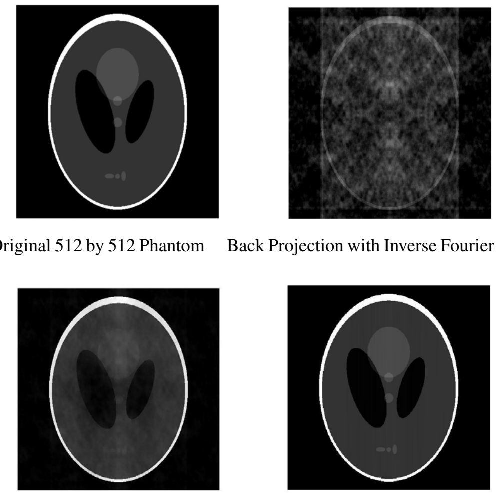

Fig. 3. Comparison of 512 by 512 Shepp-Logan Phantoms with 15% of the Fourier coefficients known.

Fig. 4. Comparison of 128 by 128 noisy Shepp-Logan Phantoms with 40% of the Fourier coefficients known.

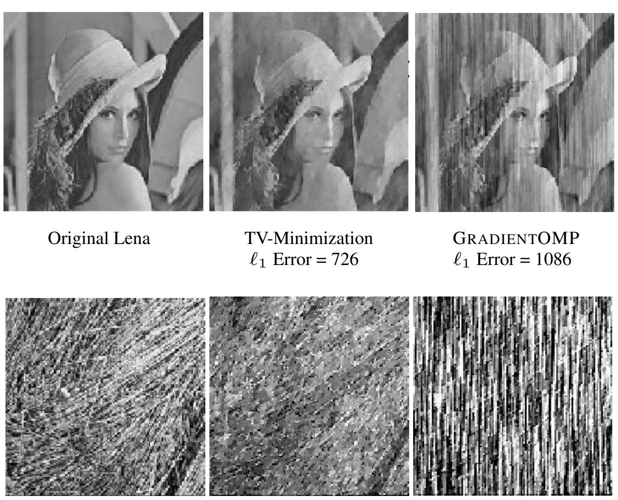

Fig. 5. Reconstructions of non-sparse gradient images with 40% of Fourier coefficients known.

a sparse-gradient image from a partial set of Fourier transform coefficients. It is significantly faster than TV minimization and may be more appropriate for large scale imaging problems such as hyperspectral imaging. We need efficient image reconstruction algorithms for these applications as the amount of data they generate is quite large. In fact, for large enough images, algorithms which run in linear time (in the image size) are not sufficiently fast. The GRADIENTOMP algorithm has the property that the OMP step can be replaced by any sublinear algorithm, such as that in [3], that performs the same action. In addition, rather than running two independent instances of OMP for the vertical and horizontal directions, we could use Simultaneous Orthogonal Matching Pursuit as described in [7]. This algorithm couples the horizontal and vertical directions.

Finally, GRADIENTOMP assumes that we have partial Fourier information. The algorithm relies upon a crucial property of the Fourier transform—it diagonalizes differentiation. There are other orthonormal transforms in which the derivative operator is almost diagonal and there are binary versions of the Fourier transform (e.g., the Hadamard transform). We are currently extending our ideas to these cases.

Compressive imaging presents exciting opportunities for new imaging technologies and algorithmic advances. While our research covers only a small component of this area, it is clear that we need a variety of reconstruction algorithms to match our technological advances.

6. REFERENCES

[1] E. J. Candes, J. Romberg and T. Tao, *Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information.*, IEEE Trans. Inform. Theory **52** (2006), no. 2, pp. 489–509.

[2] Duke Imaging and Spectroscopy Program, Duke University, http://www.disp.duke.edu

[3] A. C. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, M. J. Strauss, *Near-Optimal Sparse Fourier Representations via Sampling*, Proc. of the 2002 ACM Symposium on Theory of Computing STOC, pp. 152–161, 2002.

[4] Rice University DSP Group, *Rice Single-Pixel Camera Project*, http://www.dsp.rice.edu/cscamera.

[5] J. Romberg, *L1-Magic*, http://www.acm.caltech.edu /~l1magic/.

[6] J. Tropp, A. C. Gilbert, *Signal Recovery from Partial Information via Orthogonal Matching Pursuit*, Submitted for publication, April 2005.

[7] J. Tropp, A. C. Gilbert, M. J. Strauss, *Simultaneous Sparse Approximation via Greedy Pursuit*, in Proc. 2005 IEEE International Conference on Acoustics, Speech, and Signal Processing ICASSP, March 2005.