

An Efficient Algorithm for Compressed MR Imaging using Total Variation and Wavelets

Shiqian Ma
Department of IEOR,
Columbia University
New York, NY 10027, USA
sm2756@columbia.edu

Wotao Yin Yin Zhang
Department of CAAM,
Rice University
Houston, TX 77005, USA
{wotao.yin, yzhang}@rice.edu

Amit Chakraborty
Integrated Data Systems,
Siemens Corporate Research
Princeton, NJ 08540, USA
amit.chakraborty@siemens.com

Abstract

Compressed sensing, an emerging multidisciplinary field involving mathematics, probability, optimization, and signal processing, focuses on reconstructing an unknown signal from a very limited number of samples. Because information such as boundaries of organs is very sparse in most MR images, compressed sensing makes it possible to reconstruct the same MR image from a very limited set of measurements significantly reducing the MRI scan duration. In order to do that however, one has to solve the difficult problem of minimizing nonsmooth functions on large data sets. To handle this, we propose an efficient algorithm that jointly minimizes the ℓ_1 norm, total variation, and a least squares measure, one of the most powerful models for compressive MR imaging. Our algorithm is based upon an iterative operator-splitting framework. The calculations are accelerated by continuation and takes advantage of fast wavelet and Fourier transforms enabling our code to process MR images from actual real life applications. We show that faithful MR images can be reconstructed from a subset that represents a mere 20 percent of the complete set of measurements.

1. Introduction

Imaging plays a very important role today in medical diagnosis. In particular, the use of Magnetic Resonance Imaging (MRI) is crucial for understanding soft tissue changes within the body in a non-invasive manner. Its use of non-ionizing radio frequency emission for image acquisition is considered safe for repeated use in a clinical setting. Consequently, it has become a modality of choice for a variety of medical practitioners, such as neuro-radiologists, cariologists and oncologists. Typically, energy from an RF pulse is directed to a small section of the targeted anatomy at a time. When this happens, the protons within that area are

forced to spin in a certain frequency and get aligned to the direction of the magnet. On turning off the signal, the physical system goes back to its natural state and releases the energy, which is consequently captured and sent for analysis. This is then repeated to acquire information for the entire targeted region which is then subsequently analyzed. This data acquired in the frequency domain (also called k-space) is then inverted through the use of the Inverse Discrete Fourier Transform to arrive at the required image of the anatomy under consideration.

Given the Fourier transformation is a linear mapping, the number of samples necessary to image an anatomical area is directly proportional to its size. Further, given that the frequency domain information is acquired in a primarily sequential manner (even though parallel imaging techniques such as SENSE and GRAPPA are available) the time needed for imaging is proportional to the size of the object scanned. During scanning it is important to restrain the subject from moving, otherwise motion artifacts can severely degrade the quality of the acquired images. Even normal breathing can be a problem. Hence, subjects are often either physically restrained or sedated when large complex anatomical parts such as the brain are imaged.

Clearly the adoption of MRI to new clinical applications and even for accepted procedures can be enhanced by reducing the time needed for imaging. From the above discussion, that would mean acquiring fewer samples in the frequency domain. It seems that this however, would directly violate the long established Nyquist criterion: the amount of acquisition must at least match the amount of information needed for recovery, which means perfect reconstruction would not be possible, and hence the images acquired will demonstrate the so called Gibbs aliasing. *Compressed sensing* allows one to do this without the associated artifacts.

1.1. Compressed Sensing

Compressed sensing [2, 9] is an interesting new area of research which has gained enormous popularity of late due to its ability to reconstruct *perfect signals* from a *limited number of samples* by taking advantage of the sparse nature of the signals in a transform domain. In the case of MR images, being able to reconstruct images with high quality with a *subset* of the samples from the k-space means spending less time in the magnet for the patient. This has big consequences for the patient who is now more receptive to the procedure, and for the care giver and owner of the clinic, whose throughput can now be increased.

Compressed sensing for imaging in general is based upon the fact that most images are compressible. Let the vector \bar{u} represent an image. A compression algorithm such as JPEG2000 compresses the image by finding some dictionary Φ (e.g., Fourier or wavelet bases) such that $\Phi\bar{u} = \bar{x}$ is (approximately) sparse, and saving the locations and values of the nonzero entries of \bar{x} . To recover \bar{u} , one simply uses the same dictionary and lets $\bar{u} = \Phi^{-1}\bar{x}$. For ease of notation, we let $k = \|\bar{x}\|_0$ be the number of nonzeros in \bar{x} and n denote the dimension of \bar{x} .

The entire process of compressed sensing consists of three steps: encoding, sensing, and decoding. In the first step, \bar{u} is encoded into a *smaller* vector $b = R\bar{u}$ of a size $m < n$ by a linear transform R . Clearly, b contains less information than \bar{u} , so it is a *compression* of \bar{u} . Since $\bar{u} = \Phi^{-1}\bar{x}$, $b = A\bar{x}$ (for $A = R\Phi^{-1}$) is also a compression of \bar{x} . In many applications of compressed sensing, the linear transform R is not calculated by a computer but obtained by certain physical or digital means. In MRI, for example, R represents a partial discrete Fourier transform, where it is “partial” because $R\bar{u}$ only gives the Fourier coefficients corresponding to an incomplete set of frequencies. Notice that since \bar{u} is unknown during this step, R can only be chosen independently (non-adaptively) of \bar{u} . In the second step, b is acquired (by coils in an MRI scanner) and sent to a computer. The third step is to recover \bar{x} (and thus, \bar{u}) from b . Since \bar{x} is sparse, it can be found as the sparsest solution to the underdetermined equations $Ax = b$ unless there exists another *even sparser* solution to these equations. This gives rise to the ℓ_0 -problem: $\min_x \{\|x\|_0 : Ax = b\}$. However, this problem is NP-hard [20], impractical for nearly all real applications, so it is more realistic to solve a computationally tractable alternative such as the ℓ_1 -problem

$$\min_x \{\|x\|_1 : Ax = b\}, \quad (1)$$

which has also been known to yield sparse solutions under some conditions (see [10, 11, 12] for explains). Ideally, we would like to take the least possible number of measurements, that is, m being equal to k . However, we must pay a price for not knowing the locations of the nonzeros in \bar{x}

(there are n choose k possibilities!) while still asking for perfect reconstructions of most sparse \bar{x} . It was shown in [2, 23] that, when R is Gaussian random and partial Fourier (as in the MRI case), (1) can recover \bar{x} (technically, with a high probability) from b of a size $m = O(k \log(n/k))$ and $O(k \log(n)^4)$, respectively. These sample sizes m are larger than k but still much smaller than n . Finally, once \bar{x} is recovered, \bar{u} becomes available through $\bar{u} = \Phi^{-1}\bar{x}$.

When b is contaminated by noise, the constraint $Ax = b$ in (1) must be relaxed, resulting in either the problem

$$\min_x \{\|x\|_1 : \|Ax - b\|_2^2 \leq \sigma^2\} \quad (2)$$

or its Lagrangian version

$$\min_x \mu \|x\|_1 + \frac{1}{2} \|Ax - b\|_2^2, \quad (3)$$

where σ and μ are parameters. Problems (2) and (3) are equivalent in the sense that solving one will determine the parameter in the other such that both give the same solution [21].

1.2. Compressed MR Imaging

For recovering MR images from undersampled Fourier measurements b , one could solve (3) with

$$A = R\Phi^{-1} \quad (4)$$

for R being a partial Fourier transform and Φ being a wavelet transform; however, to get the most out of b , we also exploit the fact that MR images of organs are expected to demonstrate piecewise continuous behavior, i.e. the different anatomical structures are supposed to show uniform characteristics. Such a signal has a small total variation [24], which is defined discretely as $TV(u) := \sum_{ij} ((\nabla_1 u_{ij})^2 + (\nabla_2 u_{ij})^2)^{1/2}$, where ∇_1 and ∇_2 denote the forward finite difference operators on the first and second coordinates, respectively. To capture this we consider minimizing the total variation of $u = \Phi^{-1}x$ in addition, giving rise to the problem

$$\min_x F(x) := \alpha TV(\Phi^{-1}x) + \beta \|x\|_1 + \frac{1}{2} \|Ax - b\|_2^2 \quad (5)$$

where α and β are two positive parameters. This model was previously used by He *et al.* [15] and Lustig *et al.* [18]. However, because $TV(\Phi^{-1}x)$ and $\|x\|_1$ are both nonsmooth in x , problem (5) is much more difficult to solve than any of those with a single nonsmooth term such as (3) and a total variation regularization problem. We emphasize that in variational imaging, it is relatively easy to propose models than to solve them efficiently. For example, introducing a fourth-order regularization term such as the one used in [19] to the objective function in (5) will certainly improve

the quality of reconstructed images (enhancing the recovery of small-scale details); however, the resulting problem is too difficult to solve for real world MR images. Computation is the bottleneck that made (5) impractical in the past, and they still prevent models more complicated than (5) from being proposed and tested. Our contribution in this paper is an efficient algorithm solving (5) on full-size MR images. This algorithm can also be modified to solve similar models that use different regularization functions.

Finally, we briefly review some recent methods for compressed MR imaging to conclude this introduction. [18] is a comprehensive paper that review the steps and requirements of compressed MR imaging. It exploits the sparsity of MR images in a spatial frequency domain (e.g., by a wavelet transform) and uses ℓ_1 -minimization (1) to recover the image. The methods by Jung *et al.* [16] and Ye *et al.* [26] solve a minimization problem involving the ℓ_p -quasinorm, for $p < 1$, using the algorithm FOCUSS [22]. Minimizing an ℓ_p -quasinorm for $0 \leq p < 1$ may achieve a higher compression ratio [6, 7] than $p = 1$ but since the objective function is nonconvex, the ℓ_p -algorithms, including FOCUSS [22] and the recent ones by Candès *et al.* [3] and by Chartrand and Yin [7], do not always give global minima and are also slower. Similarly, Trzasko *et al.* [25] used the homotopic nonconvex ℓ_0 -minimization. Block *et al.* [1] used total variation as a constraint, and the work by He *et al.* [15] and Chang *et al.* [5] is based on solving (5) by PDE methods, the speeds of which are not impressive. We were not able to reproduce these algorithms and compare them with ours by the time this paper was submitted; however, the reader can easily judge the efficiency of a method by its achieved compression ratio (the lower the better), its maximal input image size, and its speed, which are usually reported in each of the papers/reports.

2. The Main Algorithm

We first derive the optimality conditions of (5), based upon which we then propose a fixed-point iteration.

2.1. Optimality Conditions

For ease of notation, we let $u \in \mathbb{R}^{n_1 \times n_2}$ denote a 2D MR image of $n_1 \times n_2$ pixels, $L = (\nabla_1, \nabla_2) : \mathbb{R}^{n_1 \times n_2} \rightarrow \mathbb{R}^{n_1 \times n_2} \times \mathbb{R}^{n_1 \times n_2}$ denote the discrete finite difference operators along the first and second coordinates subject to appropriate boundary conditions, its suboperator $L_{ij}u = (\nabla_1 u_{ij}, \nabla_2 u_{ij})$, $f(\cdot) = \|\cdot\|_2 : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that $TV(u) = \sum_{ij} f(L_{ij}(u))$, $g(\cdot) = \|\cdot\|_1$, $h(\cdot) = \frac{1}{2}\|A \cdot - b\|_2^2$, and $\Psi = \Phi^{-1}$, which equals Φ^* , the adjoint operator of Φ , for any orthonormal transform Φ . Using the notation, (5) can be rewritten as

$$\min_x E(x) = \alpha \sum_{ij} f(L(\Psi x)_{ij}) + \beta g(x) + h(x). \quad (6)$$

Since all terms are convex in (6) and $\alpha, \beta > 0$, the objective function is convex; hence, the first-order optimality condition of (6) is

$$\mathbf{0} \in \partial E(x^*),$$

where $\partial E(x^*)$ is the subdifferential (i.e., the set of subgradients) of $E(\cdot)$ at x^* . We can apply the general property:

$$y \in \partial f(x) \Leftrightarrow x \in \partial f^*(y), \quad (7)$$

for a convex function f and its convex conjugate

$$f^*(y) := \sup_x \{ \langle y, x \rangle - f(x) \} \quad (8)$$

and get the equivalent condition: x^* is optimal if and only if there exists an auxiliary variable $y^* = (y_{ij}^*)$, where $y_{ij} \in \mathbb{R}^2$, such that

$$\mathbf{0} \in \alpha \Phi \sum_{ij} L_{ij}^* y_{ij}^* + \beta \partial g(x^*) + \nabla_x h(x^*), \quad (9)$$

$$L_{ij} \Psi x^* \in \partial f^*(y_{ij}^*), \quad (10)$$

where L_{ij}^* is the adjoint operator of L_{ij} . The reader does not need to be familiar with convex conjugate because f^* can be eliminated as we show below.

Although it is not easy to directly solve the equations (9)-(10), we can apply the operator splitting method [17] to them with two scalars $\tau_1, \tau_2 > 0$, obtaining:

$$(9) \Leftrightarrow \mathbf{0} \in \tau_1 \beta \partial g(x^*) + x^* - s, \quad (11)$$

$$s = x^* - \tau_1 \left(\alpha \Phi \sum_{ij} L_{ij}^* y_{ij}^* + \nabla_x h(x^*) \right), \quad (12)$$

where $\nabla_x h(x^*) = A^*(Ax^* - b) = \Phi R^*(R\Psi x^* - b)$ and

$$(10) \Leftrightarrow \mathbf{0} \in \tau_2 \partial f^*(y_{ij}^*) + y_{ij}^* - t_{ij}, \quad (13)$$

$$t_{ij} = y_{ij}^* + \tau_2 L_{ij} \Psi x^*. \quad (14)$$

Now, all (11)-(14) are easy to compute! Given x^* and y^* , (12) and (14) compute s and t , respectively, in a straightforward way. On the other hand, given s and t , (11) and (13) uniquely determine x^* and y^* , respectively, because (11) is the optimality condition of the strictly convex problem

$$\min_x \tau_1 \beta \|x\|_1 + \frac{1}{2} \|x - s\|_2^2, \quad (15)$$

and (13) is equivalent to (again, due to (7))

$$\mathbf{0} \in \tau_2 y_{ij}^* + \partial f(y_{ij}^* - t_{ij}) \quad (16)$$

and thus is the optimality condition of

$$\min_y \frac{\tau_2}{2} \|y_{ij}\|_2^2 + \|y_{ij} - t_{ij}\|_2. \quad (17)$$

Both (15) and (17) have closed-form solutions (proved in Theorem 2.1 below) to yield x^* and y^* , respectively:

$$x^*(s) = \text{sign}(s) \max\{0, |s| - \tau_1\beta\} \quad (18)$$

$$y_{ij}^*(t_{ij}) = \min\left\{\frac{1}{\tau_2}, \|t_{ij}\|_2\right\} \frac{t_{ij}}{\|t_{ij}\|_2}, \quad (19)$$

where all operations in (18) are done component-wise, and $\mathbf{0}/0$ is defined to be $\mathbf{0}$ in (19).

Therefore we propose to solve (11)-(14) (hence, the original problem (6)) using the fixed-point iterations:

$$\text{Step a. } s^{(k+1)} \leftarrow (12) \text{ for } (x^*, y^*) = (x^{(k)}, y^{(k)}),$$

$$\text{Step b. } t^{(k+1)} \leftarrow (14) \text{ for } (x^*, y^*) = (x^{(k)}, y^{(k)}),$$

$$\text{Step c. } x^{(k+1)} \leftarrow (18) \text{ for } s = s^{(k)},$$

$$\text{Step d. } y^{(k+1)} \leftarrow (19) \text{ for } t = t^{(k)},$$

for $k = 0, 1, \dots$, starting from a set of initial points $x^{(0)}, y^{(0)}, s^{(0)}, t^{(0)}$.

2.2. Per-Iteration Computation

First, we justify Steps c and d.

Theorem 2.1. *The solutions of problems (15) and (17) are given uniquely by (18) and (19), respectively.*

Proof. First, it has been well known (see [4], for example) that the unique solution of (15) is soft-thresholding or shrinkage:

$$x^*(s) = \begin{cases} s - \tau_1\beta, & s > \tau_1\beta, \\ 0, & -\tau_1\beta \leq s \leq \tau_1\beta, \\ s + \tau_1\beta, & s < -\tau_1\beta, \end{cases} \quad (20)$$

which is precisely (18).

Second, we prove that (19) uniquely solves (17) by showing that, in each of the two cases: $\|t_{ij}\|_2 \leq 1/\tau_2$ and $\|t_{ij}\|_2 > 1/\tau_2$, (19) uniquely satisfies the first-order optimality condition of (17):

$$\mathbf{0} \in \tau_2 y_{ij}^* + \partial \|y_{ij}^* - t_{ij}\|_2. \quad (21)$$

where

$$\partial \|z\|_2 = \begin{cases} \{z/\|z\|_2\}, & z \neq \mathbf{0}, \\ \{w : \|w\|_2 \leq 1\}, & z = \mathbf{0}. \end{cases} \quad (22)$$

If $\|t_{ij}\|_2 \leq 1/\tau_2$, then simple calculations give $\|y_{ij}^* - t_{ij}\|_2 = 0$; hence, $y_{ij}^* = t_{ij}$, which is given by (19). If $\|t_{ij}\|_2 > 1/\tau_2$, then $y_{ij}^* = t_{ij}$ does not satisfy (21), so $\|y_{ij}^* - t_{ij}\|_2 \neq 0$; this, together with (21), yields $y_{ij}^* = t_{ij}/(\tau_2\|t_{ij}\|_2)$, which is also given by (19). \square

Clearly, (18) and (19) can be computed in times linear in the size of x , i.e., in $O(n_1 \times n_2)$. Therefore Steps c and d are very cheap to compute.

Next, let us study the computation of Steps a and b. All finite difference operators L_{ij} and their adjoint L_{ij}^* can be applied in a total of $O(n_1 \times n_2)$ time, so they are cheaper than the wavelet transform Φ and Fourier transform R , as well as their inverse transforms, which have typically superlinear complexities. In view of (12) and (14), both steps involve the computation $\Psi x^{(k)}$ so only one such computation is needed. In addition, only one Φ is needed in (12) since the last two terms in (12) can be combined. Therefore, the total amount of computation in Steps (12) and (14) for each k is dominated by one forward and one inverse transforms for wavelet and Fourier transforms each.

In our implementation, we store $s^{(k)}, t^{(k)}, x^{(k)}$, and $y^{(k)}$ in memory for current k . However, neither R nor Φ is explicitly expressed; all matrix-vector multiplications involving them are computed by MATLAB's implementation of the corresponding fast transforms. Therefore, our code only requires a few times as much memory as the MR image itself.

2.3. Convergence and a Continuation Strategy

The 4-step iterations are based on splitting the terms in the optimality conditions (9) and (10) into two parts, the forward one (Steps a and b) and the backward one (Steps c and d), each of which is very easy to compute. After embedding our variables and operators into appropriate spaces (for some technical reasons), one can apply the results of Combettes and Wajs [8] to show that $x^{(k)}$ generated by the iterations converge to a global solution as long as the step sizes τ_1 and τ_2 are small enough.

Note the objective function of (6) is convex but not strictly convex. In rare cases, (6) has more than one solution. When this happens, which one of the solutions is the limit of $x^{(k)}$ depends on the initial point. Also, we can avoid computing the theoretical upper bounds of τ_1 and τ_2 , which are used in Combettes and Wajs's proof of convergence, by employing a line search strategy, i.e., choosing τ_1 and τ_2 adaptively to ensure enough descent in the objective at each iteration.

Further, the convergence of the 4-step iterations can be significantly accelerated by adopting a continuation strategy introduced by Hale *et al.* [13] for the simpler problem (3). They applied a splitting-based algorithm, which they call FPC, to (3) and let the penalty parameter β vary with k , starting from an initial large value and gradually decreasing to the given value. They showed that (3) is easier to solve with a larger β , and their continuation algorithm is faster because an approximate solution of (3) corresponding to a large β serves as a good starting point for the problem corresponding to the next and smaller β .

2.4. Extensions to Higher-Dimensional Spaces

Although the underlying images are assumed to be two dimensional in the discussions above, it is straightforward to extend the theories and methods to images in three or higher dimensions. Specifically, one only needs to replace L , R , and Φ by to the higher-dimensional versions of the finite difference, Fourier transform, and wavelet transform operators.

3. Numerical Experiments

3.1. Selection of Sampling Matrix

In compressed MR imaging, the sampling matrix A is given by $A = R\Phi^{-1}$, where Φ is the wavelet transform and R is a partial Fourier transform. Assume that an MR image has n pixels. In our algorithm, R consists of m rows of the $n \times n$ matrix corresponding to the full 2D discrete Fourier transform, where $m \ll n$. (Recall that neither A nor R is stored in memory.) The m selected rows specify the selected frequencies at which the measurements in b are collected. The smaller the m , the lesser the amount of time required for an MR scanner to acquire b . The *sampling ratio* is defined to be m/n . In MR imaging, one has certain freedom to select the rows (however, practical constraints may affect the selections but they are out of the scope of this paper), and we select the rows in our experiments in the following manner. In the k-space, we chose to sample more points near the bottom left and bottom right corners, fewer points near the center. Because of the symmetry of the 2D Fourier transform, we mask the upper half space. Following these guidelines, we randomly created sampling matrices. Figure 1 highlights the positions of the selected frequencies (in one of the several experiments) in the k-space. We found that this kind of selection allowed us to recover MR images from a much smaller number of samples than a uniformly random selection. In practice, the set of frequencies as well as the sampling speed, in an MRI scan are constrained by physical and physiological limits [18], so our sampling strategy is idealized.

3.2. Experiment Setup

We developed our 2D code, which we call TVCMRI (Total Variation ℓ_1 Compressed MR Imaging), in MATLAB based upon the code FPC [14] by Hale *et al.*, and applied it on 2D real MR images.

All of our experiments were carried out in MATLAB v7.3 on a laptop with a 1.66GHz Intel Core Duo T2300E processor and 2GB memory.

We let the (final) regularization parameters $\bar{\alpha} = 1 \times 10^{-3}$ and $\bar{\beta} = 3.5 \times 10^{-2}$ in the underlying model (5) while, for the continuation procedure in the code, we chose the initial regularization parameters $\alpha_0 = \bar{\alpha}/(\eta_\alpha^3)$ and $\beta_0 =$




Figure 1. The positions of samplings (highlighted in white) with sampling ratio 21%.

$\max\{\eta_\beta \|A^T b\|_\infty, \bar{\beta}\}$ where the rate of reduction in α and β are $\eta_\alpha = 0.25$ and $\eta_\beta = 0.25$, respectively. The maximal allowed number of iterations for each value of α and β was set to 200.

For each original MR image $\bar{u} \in \mathbb{R}^n$ of n pixels, the observation data b was synthesized as

$$b = A\bar{u} + \mathbf{n}, \quad (23)$$

where \mathbf{n} is Gaussian white noise generated by $\sigma \times \text{randn}(m, 1)$ in MATLAB, and $A \in \mathbb{R}^{m \times n}$ is the sampling matrix. A and b were given to our code as data, and u was the unknown.

3.3. Recovered MR Images

We tested our code on three different 2D MR images: a 210×210 brain image, a 220×220 chest image, and a 924×208 full body image. The noise level σ is set to 0.01 in all the test problems. The original and recovered images are shown in Figures 2-4. The relative errors,

$$rerr = \frac{\|u - \bar{u}\|_2}{\|\bar{u}\|_2}, \quad (24)$$

signal to noise ratios (SNRs) and running times are given in Table 1, where \bar{u} and u are original and recovered images, respectively.

The algorithm is very efficient in recovering the images. For the brain and chest images, TVCMRI can usually recover the image within around 10 seconds. Even for the big full body image, the algorithm still can recover the image within 30 seconds. Although the input images have different levels of complexity, the recovery qualities are consistent across these images. Specifically, for each of them, a




Figure 2. (a) is the original Brain image. (b), (c) and (d) are the recovered images at the sampling ratios of 38.65%, 21.67% and 8.66%, respectively.




Figure 3. (a) is the original Chest image. (b), (c) and (d) are the recovered images at the sampling ratios of 38.50%, 21.58% and 8.73%, respectively.

sampling ratio of 38% is always sufficient for reconstructing a faithful image; 21% yields clean images without significant artifacts; yet a very low sampling ratio of 8.7% still gives acceptable results with obvious artifacts. These results can certainly be improved if we allow more number of iterations in our code. Moreover, slight under-sampling causes minor artifacts but the reconstruction still doesn't fail; this is not the case in classical compressed sensing for sparse signals, where insufficient number of measurements often yield completely incorrect results.




Figure 4. (a) is the original MR Full body image. (b), (c) and (d) are the recovered images at the sampling ratios of 38.38%, 21.49%, and 8.32%, respectively.

Image	Samp.Ratio	Rel.Err.	SNR	Time(sec)
Brain 210 × 210	38.65%	0.0795	18.74	4.64
	21.67%	0.2200	9.90	6.88
	8.66%	0.4479	3.73	8.97
Chest 220 × 220	38.50%	0.0573	21.19	6.67
	21.58%	0.1537	12.63	8.64
	8.73%	0.3134	6.44	11.47
Full body 924 × 208	38.38%	0.0503	24.12	25.58
	21.49%	0.1513	14.55	19.41
	8.32%	0.4182	5.72	22.55

Table 1. Numerical results for different MR images with different sampling ratios.

3.4. Comparison of recoveries with and without total variation

In this subsection, we demonstrate the advantage of using the total variation in the model (5). We used the FPC code [14] to solve (3) on the same images that we tested with TVCMRI.

We recorded the relative errors and signal to noise ratios (SNRs), which show differences in performance. Figures 5 and 6 give the relative errors and SNRs, respectively, of the recovered brain MR images by TVCMRI (blue curves) and FPC (green curves) from the measures at a sequence of different sampling ratios. From Figure 5 we can see that the relative errors of the recovered images from TVCMRI, which solves (5), are much smaller than those from FPC, which solves (3). In terms of SNR, TVCMRI is also significantly better than FPC. This indicates that model (5) performs better for recovering MR images from undersampled images than (3).




Figure 5. The comparison of relative errors by TVCMRI and FPC with different sampling ratios.




Figure 6. The comparison of signal noise ratios by TVCMRI and FPC with different sampling ratios.

The model (5) gave cleaner images than (3). In Figures 7 and 8 we compare the images produced by these two models, and we also record the relative errors and signal to noise ratios in Table 2. At the sufficient sampling ratio of 38%, both models give faithful images though the ones from FPC have slightly less contrast. However, the difference in image quality becomes larger and visually obvious at a lower sampling ratio of 21%. The advantage of using total variation is obvious.

4. Conclusion

In this paper, based on the compressed sensing theory, we use a joint total variation and ℓ_1 minimization model to recover MR images from a small number of measurements,




Figure 7. TVCMRI versus FPC. (a): TVCMRI result at $m/n = 38.66\%$; (b): TVCMRI result at $m/n = 21.67\%$; (c) FPC result at $m/n = 38.66\%$; (d): FPC result at $m/n = 21.67\%$.




Figure 8. TVCMRI versus FPC. (a): TVCMRI result at $m/n = 38.50\%$; (b): TVCMRI result at $m/n = 21.58\%$; (c) FPC result at $m/n = 38.50\%$; (d): FPC result at $m/n = 21.58\%$.

Image	Method	Samp.Ratio	Rel.Err.	SNR
Brain 210 × 210	TVCMRI	38.66%	0.0795	18.74
	FPC	38.66%	0.1457	13.56
	TVCMRI	21.67%	0.2200	9.90
	FPC	21.67%	0.3491	6.36
Chest 220 × 220	TVCMRI	38.50%	0.0573	21.19
	FPC	38.50%	0.1278	14.58
	TVCMRI	21.58%	0.1537	12.63
	FPC	21.58%	0.2505	8.32

Table 2. Numerical results for different MR images with different sampling ratios.

for which we developed an efficient algorithm. The numer-

ical experiments on the real MR images show that this algorithm can give a faithful recovery in less than a minute even when the sampling ratio is relatively small. By comparing the relative errors and SNRs of the recovered images by this total variation model and one without the total variation term, we show the usage of total variation is critical for restoring MR image from compressed measurements. Our algorithm can still be accelerated by incorporating optimization techniques such as smoothing and more efficient line search. Further, we believe that the algorithm presented in this paper for compressed sensing can be extended to other interesting imaging and vision applications.

Acknowledgements

S. Ma did most of the work during an internship with the Integrated Data Systems Department at Siemens Corporate Research and was supported in part by NSF Grant DMS 06-06712 and ONR Grant N000140310514. The work of W. Yin was supported by an internal faculty research grant from the Dean of Engineering at Rice University. The work of Y. Zhang was supported in part by NSF Grant DMS-0405831.

References

- [1] K. T. Block, M. Uecker, and J. Frahm. Undersampled radial MRI with multiple coils. iterative image reconstruction using a total variation constraint. *Magnetic Resonance in Medicine*, 57:1086–1098, 2007. 3
- [2] E. Candès and T. Tao. Near optimal signal recovery from random projections: universal encoding strategies. *IEEE Transactions on Information Theory*, 52(1):5406–5425, 2006. 2
- [3] E. Candès, M. Wakin, and S. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Preprint*, 2007. 3
- [4] A. Chambolle, R. A. DeVore, N.-Y. Lee, and B. J. Lucier. Nonlinear wavelet image processing: variational problems, compression, and noise removal through wavelet shrinkage. *IEEE Transactions on Image Processing*, 7(3):319–335, 1998. 4
- [5] T.-C. Chang, L. He, and T. Fang. MR image reconstruction from sparse radial samples using Bregman iteration. *Proceedings of the 13th Annual Meeting of ISMRM, Seattle*, page 696, 2006. 3
- [6] R. Chartrand. Exact reconstruction of sparse signals via nonconvex minimization. *IEEE Signal Processing Letters*, 14:707–710, 2007. 3
- [7] R. Chartrand and W. Yin. Iteratively reweighted algorithms for compressive sensing. *Submitted to ICASSP'08*, 2007. 3
- [8] P. L. Combettes and V. R. Wajs. Signal recovery by proximal forward-backward splitting. *SIAM Journal on Multi-scale Modeling and Simulation*, 4(4):1168–1200, 2005. 4
- [9] D. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52:1289–1306, 2006. 2
- [10] D. Donoho and M. Elad. Optimally sparse representation in general (nonorthogonal) dictionaries vis ℓ_1 minimization. *Proceedings of the National Academy of Sciences*, 100:2197–2202, 2003. 2
- [11] D. Donoho and X. Huo. Uncertainty principles and ideal atomic decompositions. *IEEE Transactions on Information Theory*, 47:2845–2862, 2001. 2
- [12] J. Fuchs. On sparse representations in arbitrary redundant bases. *IEEE Transactions on Information Theory*, 50. 2
- [13] E. Hale, W. Yin, and Y. Zhang. A fixed-point continuation method for ℓ_1 -regularization with application to compressed sensing. *Rice University CAAM Technical Report TR07-07*, 2007. 4
- [14] E. Hale, W. Yin, and Y. Zhang. FPC: A fixed-point continuation method for ℓ_1 -regularization, <http://www.caam.rice.edu/~optimization/>. 2007. 5, 6
- [15] L. He, T.-C. Chang, S. Osher, T. Fang, and P. Speier. MR image reconstruction by using the iterative refinement method and nonlinear inverse scale space methods. *UCLA CAM Report 06-35*, 2006. 2, 3
- [16] H. Jung, J. Ye, and E. Kim. Improved k-t blask and k-t sense using focuss. *Phys. Med. Biol.*, 52:2101–3226, 2007. 3
- [17] P. L. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16:964–979, 1979. 3
- [18] M. Lustig, D. Donoho, and J. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, in press, 2007. 2, 3, 5
- [19] M. Lysaker, A. Lundervold, and X. C. Tai. Noise removal using fourth-order partial differential equations with applications to medical magnetic resonance images in space and time. *IEEE Transactions on Image Processing*, 12:1579–1590, 2002. 2
- [20] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24:227–234, 1995. 2
- [21] J. Nocedal and S. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006. 2
- [22] B. D. Rao and K. Kreutz-Delgado. An affine scaling methodology for best basis selection. *IEEE Transactions on Signal Processing*, 47:187–200, 1999. 3
- [23] M. Rudelson and R. Vershynin. Geometric approach to error correcting codes and reconstruction of signals. *International Mathematical Research Notices*, 64:4019–4041, 2005. 2
- [24] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992. 2
- [25] J. Trzasko, A. Manduca, and E. Borisch. Highly undersampled magnetic resonance image reconstruction via homotopic ℓ_0 -minimization. *SSP 2007*, 2007. 3
- [26] J. C. Ye, S. Tak, Y. Han, and H. W. Park. Projection reconstruction MR imaging using FOCUSS. *Magnetic Resonance in Medicine*, 57(4):764–775, 2007. 3