Nullspace compressed sensing for accelerated imaging

L. J. Grady[1], and J. R. Polimeni[2]

[1]Imaging and Visualization, Siemens Corporate Research, Princeton, NJ, United States, [2]NMR Center, Radiology Department, Massachusetts General Hospital, Charlestown, MA, United States

Introduction

Compressed sensing techniques have recently become very popular for image reconstruction given sparsely sampled data. However, previous methods only approximately enforce the constraint that the reconstructed image has k-space data at specified locations, which requires manual parameter tuning and slower speeds for convergence [1,2,3,4]. In this work we show that it is possible to quickly reconstruct a sparsely sampled image that does not alter the initial k-space data.

Methods

We seek to find a reconstructed image x that minimizes $\|Gx\|_0$ subject to the constraints $Bx = y$, where G is a gradient operator (implemented with forward differences), B is a 2D or 3D Fourier transform and y is a set of measured Fourier coefficients (i.e., k-space data). Since the constrained minimization of the L_o semi-norm is NP-Hard, the L_o semi-norm has been approximated by various functions such as the L_1 norm [2,4] and, more recently a L_2 reweighted norm [1,3]. The reweighted L_2 norm can be viewed as minimization of a robust approximation to the L_0 semi-norm in which the parameter is decreased [1]. We follow this reweighted L_2 approach by approximating $\|Gx\|_0 \approx \sigma_\alpha(Gx) = \Sigma\, 1 - e^{-\alpha(Gx)2}$ where the functions become equal as α approaches infinity, as shown in Figure 1. By taking a derivative of σ, we can find a (local) minimum of x by iteratively solving a weighted least-squares minimization where the weights are given by the function $\sigma(x)$, using the current value of x. The standard approach to enforcing the constraints is to minimize the unconstrained function $Q(x) = \|Gx\|_0 + \lambda\ \|Bx - y\|_2$ with the free parameter λ penalizing discrepancy from the constraints and λ is chosen manually.

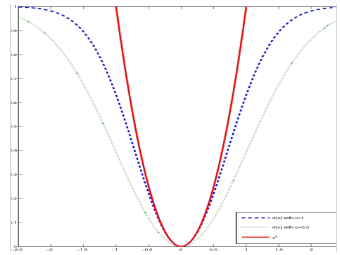

Fig. 1: *Robust function used to estimate L_0 semi-norm for various values of α.*

Penalizing discrepancy from the constraints in this manner is problematic both because there is no guarantee that the constraints will be enforced exactly and because manual selection of λ for each image makes an algorithm difficult to use in practice. We observe that the constraints can be enforced exactly without selection of λ by making a change of variables to employ the *nullspace* of B. Specifically, if matrix C spans the nullspace of B (i.e., $BC = 0$) and we have *some* solution, x_0 such that $Bx_0=y$, then we can change variables to z, where $x = x_0 + Cz$. By using z, we ensure that *every* solution satisfies the constraints and we can simply minimize $Q(x) = \|Gx_0 + GCz\|_0 \approx \sigma_\alpha(Gx_0 + GCz)$. When B represents the discrete Fourier transform sampled at locations in y, then the nullspace C is simply the inverse discrete Fourier transform sampled at every location not represented by y. In other words our new variables, z, are the missing k-space data. An initial solution, x_0, is easy to generate as the inverse Fourier transform of y with zeros at all non-sampled locations. If $Q(x)$ were convex, then any x_0 would produce the global optimum. Unfortunately, since $Q(x)$ is not convex, it is possible that choosing a particularly good x_0 or particularly bad x_0 could affect the final solution and the best choice is a topic for further research. The iterative reweighted least-squares approach is simple to implement using conjugate gradients requiring two FFTs and a sparse matrix-vector multiply at each iteration. Although our experiments were performed in MATLAB (requiring roughly 20s to reconstruct a 257×257 pixel image), the conjugate gradient method could be trivially parallelized on a GPU for a real-time implementation [5].

Results

We tested our results on the 257×257 Shepp-Logan phantom, which contains 1,278 nonzero gradients (of 131,584 total gradients). We first sampled Fourier space radially around the origin (261 samples of 66,049 total) and then sampled Fourier space by taking only every fifth y-axis line. Both sampling methods allowed for a perfect reconstruction of the Shepp-Logan phantom using our method, despite the fact that such a small number of Fourier samples were sampled and there was no parallel acquisition. Figure 2 shows the sampling structure in Fourier space with progressive iterations before achieving the final reconstruction given on the right. Although similar performance has been reported elsewhere in the literature [1,2,3], these methods required manual selection of λ or non-optimal projection onto the constraint set. Consequently, our method produces the same quality results with exact fidelity to the constraints, no need to manually set λ and faster speed.

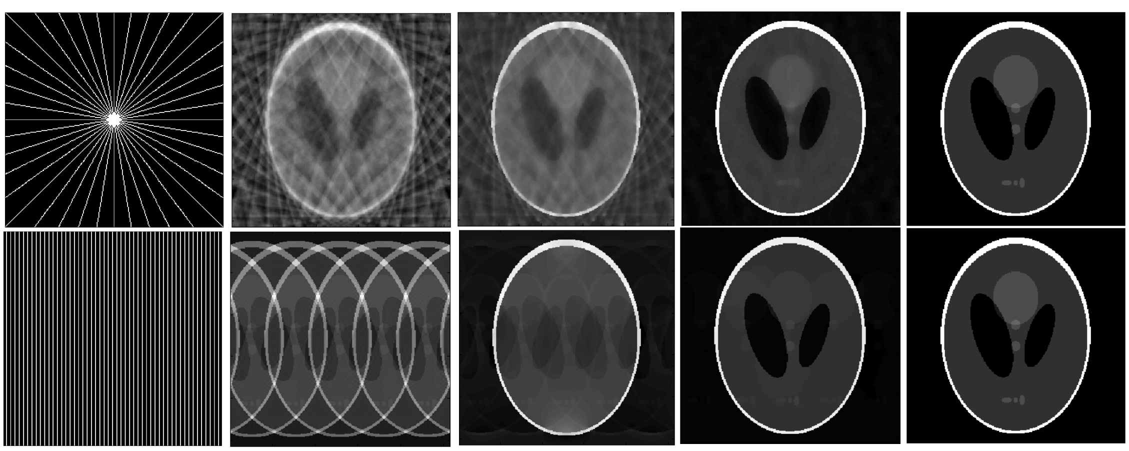

Fig. 2: *Exact reconstruction of the Shepp-Logan phantom using samples in k-space given by the leftmost columns. Subsequent columns are progressive iterations of the algorithm which produce an exact reconstruction.*

Discussion

Previous uses of compressed sensing techniques in image reconstruction loosely enforced the constraints with a manually tuned parameter. In this work, we showed that it is possible to employ popular compressed sensing techniques while enforcing fidelity of the reconstructed image to the known data in k-space space by performing optimization in the nullspace of the constraints. This technique is fast, parallelizable and produces exact reconstructions on the Shepp-Logan phantom for two different sparse samplings. Future work will focus on extending this method to parallel image reconstruction.

References

[1] J. Trzasko et al. Proc. of MICCAI 2007, p. 809-816, [2] E. Candes et al., J. of Fourier Anal. and App., 2008, [3] R. Chartrand, Signal Proc. Let. 10(14) p. 707-710, [4] Donoho, IEEE Trans. Info. Theory 2006, 52, 1289, [5] J. Bolz et al. ACM Trans. on Graphics 2003, 3(22), p. 917-924