

FAST REGULARIZED RECONSTRUCTION OF NON-UNIFORMLY SUBSAMPLED PARALLEL MRI DATA

W. Scott Hoge¹, Misha E. Kilmer², Steven J. Haker¹, Dana H. Brooks³, Walid E. Kyriakos¹

(1) Dept. of Radiology, Brigham and Women's Hospital, 75 Francis St., Boston, MA

(2) Dept. of Mathematics, Tufts University, Medford, MA

(3) ECE Dept., Northeastern University, 360 Huntington Ave., Boston, MA

ABSTRACT

Parallel MR imaging is an effective approach to reduce MR image acquisition time. Non-uniform subsampling allows one to tailor the subsampling scheme for improved image quality at high acceleration factors. However, non-uniform subsampling precludes fast reconstruction schemes such as SENSE, and is more likely to require a regularized solution than reconstruction of uniformly subsampled data demands. This means that one needs to choose a good regularization parameter, typically requiring multiple expensive system solves. Here, we present an efficient LSQR-Hybrid algorithm which simultaneously addresses the need for rapid regularization parameter selection and fast reconstruction. This algorithm can reconstruct non-uniformly subsampled parallel MRI data, with automatic regularization and good image quality, in a time competitive with Cartesian SENSE.

1. INTRODUCTION

Parallel MR imaging (pMRI) extends traditional Fourier encoding with spatial encoding, through the use of multiple receiver coils, to improve the efficiency of MR image acquisition. It has been shown to be a very effective approach to reduce MR image acquisition time, increase spatial resolution, or provide some balance of both.

The choice of subsampling pattern used to accelerate the image acquisition is often closely linked to the reconstruction algorithm. For example, the popular SENSE approach [1] typically employs uniform subsampling. However, higher quality images at the same acceleration factor may be possible using non-uniform subsampling. SENSE-like algorithms capable of reconstructing non-uniformly subsampled data typically carry a significantly higher computational burden because non-uniform subsampling destroys the regular structure that SENSE exploits for fast reconstructions [2]. To reconstruct non-uniformly subsampled data in a reasonable time, iterative algorithms such as Conjugate Gradient (CG) descent [3] are often employed [4], although it has been difficult to match the reconstruction time provided by Cartesian SENSE.

In addition, the linear systems of equations associated with the reconstruction of non-uniformly subsampled data typically have very poor conditioning. One approach to address this problem is to employ a damped-least-squares (DLS), or Tikhonov, regularization of the pMRI reconstruction problem, as in [5]. A common approach to determine the best value for the DLS regularization parameter is to employ the L-curve technique [6], which has been effective in SENSE reconstructions [7]. However, straightforward computation of the L-curve requires repeated solution of the linear system, one for each candidate regularization parameter.

In this paper we address both pMRI reconstruction time and regularization parameter selection. Recently, Kilmer and O'Leary proposed an efficient method to rapidly find both the L-curve and multiple solutions for iterative least squares solvers of DLS problems [8]. This method exploits the fact that the Krylov subspace formed by iterative solvers is in fact independent of the regularization parameter for regularization with an identity matrix.

In this work, we have expanded the LSQR-Hybrid algorithm described in [8] to handle reconstruction of non-uniformly sampled, complex-valued parallel MRI data. An additional improvement here is the use of fast matrix-vector products enabled by the particular structure of the pMRI reconstruction problem, similar to that described in [4]. We achieve a reduction in reconstruction time over naive implementations by around two orders of magnitude: one from the LSQR-Hybrid L-curve generation, and the other from the fast matrix-vector products. Together these improvements enable high-quality reconstruction of non-uniformly sampled pMRI data in a time comparable to uniformly-sampled SENSE.

2. THEORY

Parallel MR imaging methods reconstruct an image of the excited spin distribution from down-sampled k -space data acquired using multiple coils. The signal acquired in each coil, $l = 1, 2, \dots, L$, can be modeled by the following equation:

$$s_l(\mathbf{k}) = \int_V W_l(\mathbf{r})\rho(\mathbf{r})e^{j2\pi\mathbf{k}\cdot\mathbf{r}} d\mathbf{r}. \quad (1)$$

Support for this research provided in part by NIH U41 RR019703-01A2 (Jolesz, PI) and GE Medical Systems

Here, $\rho(\mathbf{r})$ is the excited spin density function throughout the volume V (weighted by various imaging factors such as spin relaxation), \mathbf{r} is a vector describing the spatial position within the FOV, $W_l(\mathbf{r})$ is the spatial sensitivity of coil l at \mathbf{r} , and \mathbf{k} is a reciprocal spatial (wave number) vector determined by the gradients employed between RF excitation and data acquisition. Note that the data is sampled in k -space, and reducing the number of acquired k -space lines reduces the image acquisition time. The solution of this linear system gives an estimate image of the spin distribution computed from the sampled data.

2.1. The SPACE RIP pMRI method

In the SPACE RIP method, for the case of 2D imaging with phase-encoding along the k_y direction used here, Kyriakos et al. [9] take a 1D Fourier transform along k_x to recast Eq. (1) as

$$s_l(k_y, x) = \sum_{n=1}^N \rho(n, x) W_l(n, x) e^{j2\pi k_y n}. \quad (2)$$

for each column, x , in the image. Using the Golub and Van Loan colon notation [3], one can compactly recast Eq. (2) as

$$s_l(:, x) = E^H \text{diag}\{W_l(:, x)\} \rho(:, x) \quad (3)$$

where all of the acquired data points associated with a particular column x have been collected into a vector of length M , the number of phase encodes employed. The elements of the M -by- N matrix E^H are the exponential terms in Eq. (2), $E(n, k_y) = e^{-jk_y n}$, where each row of E^H corresponds to a particular phase encode k_y . We note here that the rows of E^H correspond to the mutually orthogonal rows of a discrete Fourier transform operator. However, due to subsampling, the columns of E^H are not generally orthogonal.

Recognizing that the underlying spin density terms $\rho(:, x)$ are common in each of the coils, one can form the SPACE RIP linear system presented in [9] by concatenating the respective equations (3) for each acquisition coil:

$$\begin{bmatrix} s_1(:, x) \\ s_2(:, x) \\ \vdots \\ s_L(:, x) \end{bmatrix} = \begin{bmatrix} E^H \text{diag}\{W_1(:, x)\} \\ E^H \text{diag}\{W_2(:, x)\} \\ \vdots \\ E^H \text{diag}\{W_L(:, x)\} \end{bmatrix} \rho(:, x) \quad (4)$$

$$s_x = P_x \rho_x$$

This formulation exposes the particular structure of the SPACE RIP linear system matrix that we exploit for very fast matrix-vector computations.

2.2. Exploiting the SPACE RIP linear system structure

Iterative algorithms such as CG or LSQR require repeated calculations of matrix-vector products, here $P_x r$, and projections $P_x^H P_x r$. In the SPACE RIP formulation, and similar to the

approach in [4], these two products can be replaced with a sequence of more efficient operations through use of the FFT. To facilitate a compact description of this procedure, we define the following: a vector a to hold the indices corresponding to the acquired subsampled phase-encode steps; a subsampling operator D composed of the a rows of the identity matrix, $D = I(a, :)$; and the matrix W whose L columns hold the coil sensitivity data for the coils, $W = [W_1(:, x) \ W_2(:, x) \ \cdots \ W_L(:, x)]$. We denote the unitary Fourier transform along the i^{th} dimension of a matrix as $\mathcal{F}_i\{\cdot\}$.

With this notation, to compute $P_x r$, one first forms the element-by-element product between r and each vector of coil sensitivity data, $[\text{diag}\{r\}W]$. The inverse FFT along the rows of this product are then computed, and one saves only those elements corresponding to the phase encodes acquired:

$$z_1 = D \mathcal{F}_1^{-1}\{\text{diag}\{r\}W\}. \quad (5)$$

Reshaping this matrix, z_1 , to a vector gives the product $P_x r$.

In the second step, one ‘‘inflates’’ z_1 to form a matrix with non-zero rows with indices in a as the corresponding rows of z_1 and the other rows filled with zeros. The forward FFT is then computed along the rows of this matrix and then element-by-element multiplied by the conjugate of the coil data vectors, and summed along the columns. The i^{th} elements of this vector can be described by

$$[P_x^H P_x r]_i = \sum_j [W^* \circ \mathcal{F}_1\{D^T z_1\}]_{i,j}, \quad (6)$$

where \circ represents an element-by-element product.

For an image with full phase-encode resolution of N , M acquired phase encode lines, and L coils, the matrices are sized as follows: $r : N \times 1$; $W : N \times L$; $D : M \times N$; $z_1 : M \times L$. The base-2 FFT of a vector of length N is of complexity $\mathcal{O}\{N \log_2(N)\}$. Thus, the computation of $P_x r$ in the first step is of complexity $\mathcal{O}\{NL \log_2(N)\}$. In contrast, the standard matrix-vector product is of complexity $\mathcal{O}\{NLM\}$. In typical pMRI reconstruction problems M is typically much larger than $\log_2(N)$ (e.g. for a 3x acceleration of a 256x256 image, $N = 256$, $M = 85$ vs. $\log_2(N) = 8$). This savings is evident in the second step as well, again with a complexity savings of $\mathcal{O}\{N \log_2(N)L\}$ vs. $\mathcal{O}\{NML\}$.

2.3. The LSQR-Hybrid algorithm

The LSQR algorithm solves $\min_{\rho} \|P\rho - s\|_2^2$ directly while the CG algorithm instead solves the associated system of normal equations, $P^H P x = P^H s$. If P is full rank, there is a unique solution to the least squares problem, which is also the solution to the normal equations, and vice versa. The inclusion of damped-least-squares regularization adds a second term to the minimization problem:

$$\min_{\rho} \{\|P\rho - s\|_2^2 + \lambda^2 \|\rho\|_2^2\} \quad (7)$$

where the scalar λ is referred to as the regularization parameter. With non-zero λ , the direct least-squares solution and the solution to the normal equations are identical and unique.

The selection of a good value for λ is non-trivial. To use the L-curve technique [6], one constructs multiple solutions of the reconstruction problem across a range of regularization parameter values, and plots the residual error, $\|P\rho - s\|$, against the solution norm, $\|\rho\|$, in a log-log plot. This function is parametric in λ and typically results in an L-shaped curve. The inflection point on this curve corresponds to a good trade-off between the residual error and noise in the solution, and guides the regularization parameter choice.

Both the CG and LSQR algorithms operate over same underlying *Krylov subspace*, defined as:

$$\mathcal{K}_k = \text{span}\{P^H s, (P^H P)P^H s, \dots, (P^H P)^{k-1}P^H s\}. \quad (8)$$

That is, the solution ρ_k at iteration k lies in the space \mathcal{K}_k , and both algorithms create the same orthonormal basis v_1, v_2, \dots, v_k that spans \mathcal{K}_k . In LSQR, this basis is generated through the following iterative steps,

$$PV_k = U_{k+1}B_k; P^H U_k = V_{k+1}B_k \quad (9)$$

Here, B_k is a $(k+1)$ -by- k bidiagonal matrix, and U_{k+1} is a $(k+1)$ -by- n matrix with orthonormal columns with $u_1 = b/\|b\|$.

The key to efficient L-curve generation with the LSQR-Hybrid algorithm is the fact that the V_k and U_k matrices are *not affected* by the value of the regularization parameter λ , since the Krylov subspace does not change with the addition of a scalar multiple of I to $P^H P$. In the LSQR-Hybrid algorithm [8], one implicitly solves a regularized projected system with matrix B_k . Moreover, one can compute the quantities needed for the L-curve for this smaller problem without explicitly computing the solution to that system. It turns out that this L-curve is *identical* to the L-curve for the original system. An additional savings comes from the fact that one does not need to save the V_k in Eq. (9) to obtain a short-term recurrence for ρ . This opportunity to embed an L-curve iteration within each iteration of LSQR leads to significant computational savings over calculating an L-curve using a non-hybrid approach.

3. RESULTS

We implemented and tested in Matlab (The Mathworks, Natick, MA) two algorithms to solve the SPACE RIP pMRI problem: a CG algorithm with fast matrix-vector multiplications but without embedded L-curve computation, and an LSQR algorithm with both improvements, on data acquired using a Fast Spin Echo (TR = 500 ms, TE = 13.8 ms) sequence on a 1.5T GE Signa EXCITE 11.0 MR scanner with an 8 channel bird-cage head coil. The non-uniform subsampling pattern used followed an exponentially weighted distribution, $Z_{ed}(y) = \exp\{-0.87|y|\}$, [2], with an overall acceleration

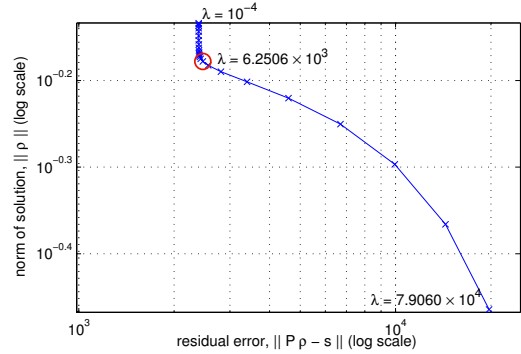


Fig. 1. L-curve for one column in the reconstruction. The circle shows λ at the inflection point of the curve.

(down-sampling) factor of 3x. Coil sensitivity estimates were calculated from 20 central lines of k -space data acquired from the same slice. A Gaussian envelope was applied to this data to limit ringing, which was then normalized in the spatial domain such that the root-sum-of-squares was equal to 1. Note that this number of lines is much lower than typically needed for good coil sensitivity estimation when reconstructing uniformly subsampled data with SENSE.

Regularization parameter selection: The regularization parameter λ was automatically chosen for each column using the embedded L-curve approach for all reconstructions [8]. For the LSQR-Hybrid algorithm, all the required values were obtained with one pass through the iterative solution, while for CG SPACE RIP one set of iterations was computed for each value of λ . Fig. 1 shows the L-curve for one column of the reconstruction. The inflection point of the L-curve is identified by the circle in the figure and corresponds to the highest value of regularization that suppressed noise in the reconstruction without introducing visibly objectionable smoothing and aliasing artifacts.

Image quality: Fig. 2 shows images corresponding to four values of λ . In panel (a), the solution is under-regularized and we see significant noise amplification throughout the image. Panel (b) shows the reconstruction for the value of λ at the inflection point of the L-curve in Fig. 1. Here, the noise in (a) is no longer present while aliasing artifacts are minimal and spatial resolution is reasonable. At higher levels of regularization, as seen in panels (c) and (d), the regularization is excessive and thus visually important modes from acquisition encoding are suppressed. This results in smoothed anatomical detail and increased aliasing artifacts.

Reconstruction time: Table 1 shows a comparison of the relative reconstruction times between the LSQR-Hybrid algorithm and CG with FFT-type matrix-vector calculations but without embedded L-curve calculations. These calculations were performed on a moderate performance 1GHz PIII with 512MB of RAM. Computation times are dependent on many factors (e.g. operating system overhead, coding implementation, memory/cache size, etc.), and the numbers we report should be taken only as an indication of the relative computa-




Fig. 2. Comparison of image reconstructions for four values of λ .

Number of λ solutions:	1	5	10	15	25	40	50
LSQR recon time (sec):	14	16	20	22	30	37	45
CG recon time (sec):	9	49	97	144	240	417	518

Table 1. Comparison between CG and LSQR reconstruction time for multiple λ L-curve generation.

tional load. Nonetheless, they clearly illustrate the computational savings provided by the LSQR-Hybrid approach when solutions with multiple λ 's are found. We note that for a fixed value of λ , our CG implementation with the FFT-based multiplication is as fast as the comparable SENSE reconstruction of uniformly subsampled data, which required ≈ 10 sec on the same machine at the same acceleration factor.

Thus for automatic selection of the regularization parameter via the L-curve, the LSQR-Hybrid approach shows a clear computational advantage, requiring only minimal additional time for multiple values of λ , unlike the same multiple solutions found through the fast CG method.

4. SUMMARY

In this paper we used two iterative algorithms, CG and LSQR, to reconstruct non-uniform subsampled pMRI data using the SPACE RIP formulation. The particular structure of the SPACE RIP linear system allows the use of FFT operations to compute the required matrix-vector products in either algorithm, which leads to significant savings over standard CG implementations and previous pseudo-inverse approaches.

The linear system matrix associated with non-uniform subsampling typically has a high condition number and the reconstruction algorithm requires regularization. Through use of the LSQR-Hybrid algorithm, we can achieve additional computational savings when the regularization parameter value is automatically selected using the L-curve. (We note that one could also derive a CG-Hybrid algorithm with embedded L-curve calculations, although LSQR-Hybrid is expected to have some technical advantages.) Together, the fast FFT-multiply and the hybrid L-curve computation allows high image quality reconstruction of non-uniformly subsampled parallel MRI data in a time competitive with SENSE reconstruction of uniformly sub-sampled data.

5. REFERENCES

- [1] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: Sensitivity encoding for fast MRI," *Magn. Reson. Med.*, vol. 42, no. 5, pp. 952–62, Nov. 1999.
- [2] W. S. Hoge, D. H. Brooks, B. Madore, and W. E. Kyriakos, "A tour of accelerated parallel MR imaging from a linear systems perspective," *Concepts in MR*, vol. 27A, no. 1, pp. 17–37, 2005.
- [3] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore, MD: Johns Hopkins University Press, 3rd ed., 1996.
- [4] K. P. Pruessmann, M. Weiger, P. Börnert, and P. Boesiger, "Advances in sensitivity encoding with arbitrary k -space trajectories," *Magn. Reson. Med.*, vol. 46, no. 4, pp. 638–651, Oct. 2001.
- [5] L. Ying, D. Xu, and Z.-P. Liang, "On Tikhonov regularization for image reconstruction in parallel MRI," in *Proc Intl Conf of the IEEE EMBS. (San Francisco, CA)*, pp. 1056–1059, Sep. 2004.
- [6] P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems*. SIAM Press, 1998.
- [7] F.-H. Lin, K. K. Kwong, J. W. Belliveau, and L. L. Wald, "Parallel imaging reconstruction using automatic regularization," *Magn. Reson. Med.*, vol. 51, no. 3, pp. 559–567, Mar. 2004.
- [8] M. E. Kilmer and D. P. O'Leary, "Choosing regularization parameters in iterative methods for ill-posed problems," *SIAM J. Matrix Anal. Appl.*, vol. 22, no. 4, pp. 1204–1221, 2001.
- [9] W. E. Kyriakos, L. P. Panych, D. F. Kacher, C.-F. Westin, S. M. Bao, R. V. Mulkern, and F. A. Jolesz, "Sensitivity profiles from an array of coils for encoding and reconstruction in parallel (SPACE RIP)," *Magn. Reson. Med.*, vol. 44, no. 2, pp. 301–308, Aug. 2000.