Non-Iterative Regularized reconstruction Algorithm for Non-CartesiAn MRI: NIRVANA

Satyananda Kashyap, Zhili Yang, & Mathews Jacob*

Department of Electrical & Computer Engineering
Department of Biomedical Engineering, University of Rochester, Rochester, NY 14627.

Abstract

We introduce a novel non-iterative algorithm for the fast and accurate reconstruction of non-uniformly sampled MRI data. The proposed scheme derives the reconstructed image as the non-uniform inverse Fourier transform of a compensated dataset. We derive each sample in the compensated dataset as a weighted linear combination of a few measured k-space samples. The specific k-space samples and the weights involved in the linear combination are derived such that the reconstruction error is minimized. The computational complexity of the proposed scheme is comparable to that of gridding. At the same time, it provides significantly improved accuracy and is considerably more robust to noise and undersampling. The advantages of the proposed scheme makes it ideally suited for the fast reconstruction of large multi-dimensional datasets, which routinely arise in applications such as f-MRI and MR spectroscopy. The comparisons with state of the art algorithms on numerical phantoms and MRI data clearly demonstrate the performance improvement.

Key words: Magnetic resonance imaging, matching pursuits, Tikhonov regularization, non-Cartesian sampling.

1. Introduction

Non-Cartesian sampling schemes continue to play very important roles in several fast MR imaging applications. These schemes are preferred due to their improved sampling efficiency, robustness to under sampling, and insensitivity to motion. The widely used algorithm to recover the image is termed as gridding, which reconstructs the images by weighting the non-uniform k-space samples by appropriate density compensation factors (DCF), followed by the the evaluation of the non-uniform Fourier transform. Most of the gridding methods rely on the non-uniform fast Fourier transform (NUFFT) approximation [1, 2] to accelerate the evaluation of the inverse Fourier transform [3, 4]. Several authors have focussed on the design of DCF using analytic [5] and convolutional formulations [6, 7]; all of these methods compute the DCF as the inverse of the sampling density in k-space to compensate for the varying density.

The use of regularized least-squares (RLS) schemes was recently advocated by several authors for the reconstruction of non-uniformly sampled data [2]; Regularization provides a tradeoff between bias and variance, which enables to suppress the artifacts and reduce noise amplification, especially when k-space is sampled at sub-Nyquist rate. The iterative conjugate-gradient algorithm is often used to realize the RLS reconstruction. The main limitation of this scheme is its computational complexity, especially while dealing with large multi-dimensional datasets. For example, a typical functional MRI dataset has around 1000 images. The reconstruction of each of these images using CG will require several hours. It is thus desirable to have algorithms that are computationally as efficient as gridding, while providing more accurate reconstructions.

The RLS solution to $\mathbf{Ax} = \mathbf{b}$ can be mathematically expressed as $\mathbf{x}^* = \mathbf{A^H Q b}$, where $\mathbf{Q} = \left(\mathbf{AA^H} + \lambda\mathbf{I}\right)^{-1}$ is termed as the compensation matrix. Here, $\mathbf{A}$ and $\mathbf{A}^H$ correspond to the non-uniform Fourier transform (NFT) and the inverse non-uniform Fourier transform (INFT) respectively. These operators are often approximated using NUFFT schemes, when we refer them as NUFFT and INUFFT respectively. Here, $\mathbf{b}$ and $\mathbf{x}$ are the vectors of k-space measurements and the unknown image, respectively. Sedarat et. al. have recently shown that the gridding scheme is essentially an approximation of the RLS method with $\lambda = 0$, where the compensation matrix is approximated by a diagonal matrix [8]; the diagonal entries of the matrix are the density compensation factors.

The main focus of this work is to derive a fast algorithm that closely approximates the RLS reconstruction scheme. If the trajectory is oversampled, the derivation of the compensation matrix is underdetermined. The compensation matrix is often evaluated using the pseudo-inverse, which often results in a dense matrix (most entries being non-zero). The multiplication of the measured data vector $\mathbf{b}$ with a dense compensation matrix will be computationally expensive; each element in the intermediate data set will have to be evaluated as a weighted linear combination of all the measured non-uniform samples, weighted by the corresponding row entries of the compensation matrix. To minimize the complexity, we propose to derive a compen-

sation matrix that results in a computationally efficient algorithm. Specifically, we propose to restrict the number of non-zero elements in each row of the matrix to a small number. This approach results in a computational complexity that is more or less comparable to gridding. We search for the optimal sparsity pattern and the entries of the matrix such that the algorithm closely approximate the RLS scheme.

This paper is inspired by the work on Sedarat et. al., where they derived the optimal DCF to minimize the error in approximating $\mathbf{Q}$ in the least squares sense. [8]. They also generalized this scheme to linearly structured matrices, where the structure is pre-specified (eg. Toeplitz, diagonal, Hankel, Banded). The proposed work is fundamentally different from [8]; in contrast to pre-specifying the structure, we determine the optimal structure and the corresponding matrix entries to minimize the approximation error. Our experiments show a significant improvement in optimizing the structure over banded structured matrices. There are also some similarities between the proposed approach and the block uniform resampling scheme (BURS) [9], even though this approach is fundamentally different from the proposed algorithm. The BURS method avoids the use of INUFFT by directly recovering the uniform Fourier samples of the image as a weighted linear combination of the measured non-uniform samples in a small circular neighborhood. While this approach is computationally efficient, it is difficult to use this scheme for regularized reconstruction. The lack of regularization makes this method sensitive to measurement noise and model misfit, which is reported by several authors [10]. Our experiments also show the noise sensitivity of this method, inspite of the use of the truncated singular value decomposition as described in [11].

2. Background

The measured non-Cartesian samples, $\hat{\rho}(\mathbf{k}_i)$, are the non-uniform Fourier measurements of the object $\rho(\mathbf{x})$:

$$\hat{\rho}(\mathbf{k}_i) = \int_{\mathbf{x}\in\mathbb{R}^2} \rho(\mathbf{x}) \exp\left(j2\pi\mathbf{k}_i^H\mathbf{x}\right) d\mathbf{x};\ i = 0\ldots, L-1 \quad (1)$$

where $\mathbf{k}_i; i = 0, .., L-1$ are the non-uniform Fourier samples. Assuming square voxels, the discretization of the above integral gives

$$\hat{\rho}(\mathbf{k}_i) = \sum_{\mathbf{m}\in\mathbb{Z}^2} \rho[\mathbf{m}] \exp\left(j2\pi\mathbf{k}_i^H\mathbf{m}\right);\ i = 0\ldots, L-1. \quad (2)$$

where $\rho[\mathbf{m}]$ are the image samples. This summation can be expressed in the matrix form as

$$\mathbf{b} = \mathbf{A}\mathbf{p}, \quad (3)$$

where the entries of the $L \times M$ system matrix $\mathbf{A}$ are the Fourier exponentials. Since this system of equations is often ill posed, the standard approach is to formulate the recovery of $\mathbf{p}$ as a regularized reconstruction (RLS) problem:

$$\mathbf{p}^* = \arg\min_{\mathbf{p}} \|\mathbf{A}\mathbf{p} - \mathbf{b}\|^2 + \lambda\|\mathbf{p}\|^2, \quad (4)$$

The above equation can be solved analytically to obtain

$$\begin{aligned} \mathbf{p}^* &= \left(\mathbf{A}^H\mathbf{A} + \lambda\mathbf{I}\right)^{-1}\mathbf{A}^H\mathbf{b} \\ &= \mathbf{A}^H \underbrace{\left(\mathbf{A}\mathbf{A}^H + \lambda\mathbf{I}\right)^{-1}}_{\mathbf{P}^{-1}}\mathbf{b} \end{aligned} \quad (5)$$

This expression can be interpreted as $\mathbf{p}^* = \mathbf{A}^\mathbf{H}\mathbf{s}$, where each element of the vector $\mathbf{s} = \mathbf{P}^{-1}\mathbf{b}$ is obtained as a weighted linear combination of the non uniform samples $\mathbf{b}$. The weights for the non-uniform samples are specified by the corresponding row of $\mathbf{P}^{-1}$. The major challenge in evaluating (5) is the large size of the system matrix $\mathbf{A}$, which makes the evaluation and storage of $\mathbf{P}$ impossible for problems of practical value. Due to these challenges, the standard approach is to solve (4) using a conjugate gradient optimization algorithm [12]. The main drawback of this approach is its computational complexity, which makes the reconstruction of large multidimensional MRI datasets using this scheme impractical.

3. Materials and Methods

We propose to approximate $\mathbf{P}^{-1}$ by a sparse compensation matrix, which closely approximates the RLS scheme. Specifically, we will derive the appropriate square matrix $\mathbf{Q}$, such that

$$\mathbf{Q}\mathbf{P} \approx \mathbf{I}. \quad (6)$$

3.1. Image reconstruction

Once an appropriate $\mathbf{Q}$ is obtained, the image is evaluated as

$$\mathbf{p}^* = \mathbf{A}^H\mathbf{Q}\mathbf{b}. \quad (7)$$

The computational complexity in the evaluation of $\mathbf{s} = \mathbf{Q}\mathbf{b}$ is dependent on the number of non-zero entries of $\mathbf{Q}$. For example, if the n^{th} row of $\mathbf{Q}$ has only N non-zero entries, the n^{th} element of $\mathbf{s}$ is determined as the weighted linear combination of N samples of $\mathbf{b}$. Once $\mathbf{s}$ is obtained, we reconstruct the image using the INUFFT. The combined scheme is illustrated in Fig. 1.

3.2. Pre-computation of the compensation matrix

We precompute the sparse compensation matrix $\mathbf{Q}$ for a specified trajectory and regularization parameter and store it. It is then used to reconstruct all of the images acquired with the same trajectory and similar noise statistics. This approach is analogous to the pre-computation of the density compensation factors in standard gridding algorithms. This provides us with a simpler reconstruction scheme, at the expense of a more complex one-time pre-computation step.

We denote the class of matrices, whose rows have N non zero elements as N-sparse matrices. To obtain a computationally efficient solution, we formulate the derivation of the approximate matrix $\mathbf{Q}$ as

$$\mathbf{Q} = \arg \min_{\mathbf{Q} \in \text{N-sparse}} \|\mathbf{QP} - \mathbf{I}\|_F^2. \tag{8}$$

We rewrite the above metric as

$$\|\mathbf{QP} - \mathbf{I}\|_F^2 = \|\mathbf{PQ}^H - \mathbf{I}\|_F^2 \tag{9}$$

Here, we used the property that $\mathbf{P}$ is a symmetric matrix. Thanks to the use of the Frobenius norm, this expression can be simplified as

$$\|\mathbf{PQ}^H - \mathbf{I}\|^2 = \sum_{i=0}^{M-1} \|\mathbf{Pq}_i - \boldsymbol{\delta}_i\|^2 \tag{10}$$

where $\mathbf{q}_i$ is the i^{th} row of $\mathbf{Q}$ and $\boldsymbol{\delta}_i$ is the Kroneker delta function, defined as

$$\delta_i[m] = \begin{cases} 1 & \text{if} \quad m = i \\ 0 & \text{else} \ . \end{cases} \tag{11}$$

Each of the terms in the summation in (10) are independent of each other. Hence, we decouple the sparse matrix approximation problem, specified by (8), as a series of simpler sparse vector optimization problems:

$$\mathbf{q}_i = \arg \min_{\mathbf{q}; \mathbf{q} \in \text{N-sparse}} \| \underbrace{\mathbf{Pq}_i}_{\mathbf{p}_i} - \boldsymbol{\delta}_i\|^2, \text{ for } i = 0, \ldots, M-1. \tag{12}$$

If the sparsity pattern of the weights $\mathbf{q}_i$ are known apriori (eg. the N k-space samples that are the closest to $\mathbf{k}_i$), the weights can be determined as a simple least square optimization problem as in (16). However, we find that evaluating the optimal pattern provides an algorithm that is significantly more accurate than relying on the nearest neighbors.

Sparse vector optimization problems similar to (12) are well studied in signal processing literature [13, 14]. We propose to use the matching pursuit algorithm to derive the optimal weights $(\mathbf{q}_i; i = 0, .., M-1)$, mainly due to its simplicity and computational efficiency [14]. This greedy forward selection algorithm proceeds sequentially by adding entries to the vector, starting with the initialization $\mathbf{q}_i = 0$. At each iteration, the algorithm searches over all the non-zero entries of $\mathbf{q}_i$ to determine the one that gives the biggest reduction in the residual error energy: $\|\mathbf{r}_i\|^2 = \|\mathbf{Pq}_i - \boldsymbol{\delta}_i\|^2$. This is often done by evaluating the element with the maximum magnitude of the error gradient:

$$\nabla \left(\|\mathbf{e}_i\|^2 \right) = 2\mathbf{P}^H \left(\mathbf{Pq}_i - \boldsymbol{\delta}_i \right). \tag{13}$$

The above function is often termed as the proxy. Once the support is updated, we recompute the entries of $\mathbf{q}$ on the specified support using a least squares minimization. The sequential algorithm is terminated once the desired number of elements are added to the vector. The steps of the matching pursuit algorithm to determine a specified row of the $\mathbf{Q}$ matrix are described below:

Start with $\mathbf{q}_i = \mathbf{0}$.

1. Determine the proxy as

$$\mathbf{y} = \mathbf{P}^H \underbrace{\left(\mathbf{Pq}_i - \delta_i \right)}_{\mathbf{r}_i}.$$

   For large images, the computation and storage of $\mathbf{P} = \left(\mathbf{AA}^H + \lambda \mathbf{I} \right)$ is not possible, even on modern computers. Hence, we propose to compute the proxy using "on the fly computations":

$$\mathbf{Pq}_i = \text{NUFFT} \left(\text{INUFFT} \left(\mathbf{q}_i \right) \right) + \lambda \mathbf{q}_i \tag{14}$$

   and $\mathbf{P}^H\mathbf{r}$ as

$$\mathbf{P}^H\mathbf{r} = \text{NUFFT} \left(\text{INUFFT} \left(\mathbf{r} \right) \right) + \lambda \mathbf{r} \tag{15}$$

   These simplifications allow us to search over all non-uniform samples to minimize the error. This inturn enables us to obtain significant reductions in the approximation error.
2. The index corresponding to the maximum value of $\mathbf{y}$ is used to update the support T as $T = \{T, \arg\max_i |\mathbf{y}_i|\}$.
3. Update the weight vector $\mathbf{q}_{T,i}$, assuming the support to be T. We perform this update as

$$\mathbf{q}_{T,i} = \arg \min_{\mathbf{q}_{T,i}} \|\mathbf{P}_T \ \mathbf{q}_{T,i} - \boldsymbol{\delta}_i\|^2. \tag{16}$$

   Here, $\mathbf{P}_T$ is a sub-matrix of $\mathbf{P}$ obtained by picking the rows of $\mathbf{P}$ specified by the support T. $\mathbf{q}_{T,i}$ is the subvector of $\mathbf{q}_i$, obtained by picking the corresponding entries of $\mathbf{q}$. Again, since it is often not feasible to compute and store the matrix $\mathbf{P}$, we compute the i^{th} row of $\mathbf{P}_T$ as

$$\mathbf{P}_{T_i} \quad = \quad \mathbf{P}\boldsymbol{\delta}_{T_i} = \text{NUFFT} \left(\text{INUFFT} \right) \left(\boldsymbol{\delta}_{T_i} \right) + \lambda \boldsymbol{\delta}_{T_i} \tag{17}$$

   The above equation is implemented by setting the T_i^{th} non-uniform sample as one and then computing its INUFFT and NUFFT. We determine $\mathbf{q}_i$ as $\mathbf{q}_i = \mathbf{P}_T^{\#}\delta_i$, where $\mathbf{P}^{\#}$ indicates the pseudo-inverse of $\mathbf{P}$.
4. Repeat the steps (1), (2), and (3) until the number of samples in $\mathbf{q}_i$ is equal to N.

We observe that the first point that is added to $\mathbf{q}_i$ is the i^{th} sample itself. Thus, the $\mathbf{Q}$ matrix is diagonal if $N = 1$. Hence, the optimal gridding algorithm [8] is a special case of the proposed algorithm.

4. Results

We first focus on illustrating and optimizing the key aspects of the algorithm in Section 4.1. We determine the computational advantages and the improved accuracy of the proposed algorithm by comparing it to the existing methods in Section 4.2.

4.1. Illustration and optimization of the algorithm

The key aspect of the proposed algorithm is the derivation of the optimal sparsity pattern using matching pursuits. This approach is drastically different from the intuitive approach of selecting the immediate neighbors of the sampling location; we denote this approach as the nearest neighbor (NN) scheme. To determine the utility in optimizing the sparsity pattern, we now compare the optimized pattern with the NN method. The sparsity pattern of the $\mathbf{Q}$ matrices corresponding to the spiral trajectory with six interleaves and $N = 25$ are shown in Fig.2. Fig.2.(a) corresponds to the NN pattern, while (b) corresponds to the optimized sparsity pattern. Since the neighbors of any specified point fall on different interleaves, the NN matrix has a banded structure as seen from Fig. 2.(a). Fig. 3 shows the relation between the sparsity pattern of the matrix and the k-space sampling locations involved in the weighted linear combination. Note that the optimal patterns are drastically different from the NN cases. Towards the k-space center, where the sampling density is high, the optimal samples are more or less aligned along straight lines that intersect at the specified point. However the optimal samples are more clustered for samples in the edge of k-space.

We determine the utility of the optimized sampling patterns in minimizing the reconstruction errors in Fig. 4. Here, we reconstruct a 32x32 Shepp Logan phantom using different values of N assuming NN and the optimal sampling patterns. The reconstruction errors are plotted in 4. Note that this optimization gives a significant decrease in reconstruction errors, thus demonstrating the benefit in using optimal sampling patterns. It is also interesting to see that the proposed scheme with $N = 25$ gives almost a factor of eight decrease in errors over $N = 1$, which is the optimal gridding reconstruction [8]. We observe that the performance improvement saturates around $N = 20$. Hence, we limit our range of supports between N = 15 to N = 25 in all the experiments in this paper.

4.2. Comparison with existing methods

We compare the proposed scheme with four existing reconstruction techniques:

1. Gridding with optimal least-squares DCF [8]; we will also refer this scheme as the optimal gridding scheme,
2. Gridding with the DCF evaluated using [6]; we will also refer this scheme as the standard gridding scheme,
3. BURS algorithm [9], and
4. Regularized least square (RLS) algorithm [2].

We use (i) the Shepp-Logan numerical phantom, assuming interleaved spiral trajectories and (ii) experimental data acquired from a grid phantom on a 3T Siemens Trio scanner to validate the scheme. In the numerical simulations, we compare the reconstructions to the ground truth as well as the results obtained from the RLS reconstruction. We quantify the performance of the reconstructions using the signal to error ratio metric

$$\mathrm{SE} = -10 \log \left(\frac{\|\mathbf{x}_{\mathrm{rec}} - \mathbf{x}\|^2}{\|\mathbf{x}\|^2} \right). \tag{18}$$

Here, $\mathbf{x}_{\mathrm{rec}}$ is the reconstructed image and $\mathbf{x}$ is the image we are comparing against. We refer to SE1 as the comparison with the ground truth and SE2 as the comparison with the RLS reconstruction. Note that both of the above trajectories fail to sample the corners of k-space, resulting in information loss. Since it is impossible for any algorithm to provide perfect reconstruction from the measured k-space data, we expect the comparisons with the RLS algorithm to be more realistic. Since the ground truth is not available for the experimental data, we rely only on the comparisons to the RLS scheme.

The RLS reconstructions are often evaluated using the iterative conjugate gradients algorithm, which takes significantly more time than the non-iterative algorithms. Specifically, each iteration of the CG algorithm requires two NUFFTs and one INUFFT; a NUFFT-INUFFT pair is needed is to compute the gradient, while one NUFFT is used to compute the optimal step size. In contrast, all of the non-iterative schemes use only a single INUFFT. Since the evaluation of the NUFFT is the most computationally expensive component, the computation time of CG algorithm with 20 iterations is roughly sixty times of that of the non-iterative schemes. We also report the execution times of the MATLAB implementations of the algorithm on an iMac with a 2.4 GHz Intel Core 2 Duo processor and 2 GB of RAM. We quantify the reduction in computational complexity in terms of the acceleration factor (A), which is defined as the ratio of the time taken by the proposed algorithm to the time taken by CG algorithm to converge. While these measures are dependent on the specific implementations, this figure provides a feel for the practical benefits obtained from the use of the proposed scheme.

4.2.1. Validation using the numerical phantom

To perform the numerical simulations, we evaluate the exact Fourier transform of the numerical phantom on the sampling locations specified by the sparsely sampled interleaved spiral trajectories with two interleaves. These experiments enable us to determine the accuracy of the algorithms and their robustness to under-sampling. We also evaluate the robustness of the algorithms by testing them using noise corrupted data. [11]

We study the robustness of the algorithms to under-sampling by comparing the reconstructions of 128x128 phantom sampled on a spiral trajectory with two interleaves. (see Fig. 5). We set $N = 25$ for the proposed scheme, while we assume $\lambda = 0.5$ for the RLS and the proposed methods. As expected, the reconstructions using the proposed and the RLS schemes are more or less similar and provide acceptable image quality. In contrast, the BURS scheme exhibit significant structured artifacts. This may

be attributed to the lack of regularization in this framework and the larger spacing between the k-space samples. It is also seen that the optimal gridding scheme, which is a special case of our approach, gives a 4 dB improvement over standard gridding.

The comparisons in the presence of noise (input SNR of 16 dB) are shown in Fig. 6. We added white Gaussian noise to the simulated k-space measurements, evaluated on the spiral trajectory with two interleaves. We observe that the BURS scheme is very sensitive to measured noise, even with the truncated SVD implementation. This is consistent with what is reported in the literature [10]. We assume $\lambda = 0.5$ for the RLS and the proposed schemes; the value of this parameter was determined using the L-curve method [15]. It is seen that the regularized methods are considerably more robust to measurement noise.

The above experiments show that the gridding methods are surprisingly more robust to undersampling and measurement noise than the BURS scheme. One reason may be fewer degrees of freedom in choosing the $\mathbf{Q}$ matrix, making this scheme significantly more constrained than BURS. This may act as an implicit regularization, resulting in improved robustness. In contrast, the explicit regularization in RLS and the proposed scheme enables them to be robust to noise and undersampling, while providing significantly more accurate reconstructions than gridding.

4.2.2. Validation using the MRI phantom

We show the reconstructions of a grid phantom, acquired using a Nyquist sampled polar trajectory in Fig. 7. We used a polar trajectory with 128 spokes and 256 samples/spoke to sample the phantom. We used a regularization parameter of $\lambda = 1.0$ for the RLS and the proposed method which was determined using the L-curve. For the proposed method, we chose a support size of $N = 25$. We observe that the BURS has the lowest S/E, which may be attributed to its sensitivity to undersampling and model mismatch (due to eddy currents and field-map effects). It is seen that the standard gridding reconstructions exhibit some ringing and intensity variations. The errors with optimal gridding are more subtle visually. However, the proposed scheme provides a 3 dB improvement in SNR over the optimal gridding results. The precomputation of the sparse $\mathbf{Q}$ matrix in the proposed scheme takes approximately 90-100 minutes on an iMac with a 2.4 GHz Intel Core 2 Duo processor and 2 GB of RAM.

All of the above experiments demonstrate that the proposed scheme is capable of closely approximating the RLS algorithm with a significantly reduced computational complexity. Unfortunately, it is not straightforward to generalize the proposed scheme to Tikhonov regularization schemes with gradient based penalties. The proposed scheme can be generalized to reconstruct parallel MRI data.

5. Conclusions

We introduced a fast and non-iterative algorithm for the accurate reconstruction of non-Cartesian MRI data. The algorithm obtains the reconstructed image as the non-uniform inverse Fourier transform of a compensated dataset, whose samples are obtained as linear combinations of a few of the measured samples. The reconstruction scheme has computational complexity that is comparable to gridding, while it provides reconstructions that are comparable to the regularized least square schemes. This scheme is ideally suited for MRI applications such as MR spectroscopy and f-MRI, where large volumes of data have to be reconstructed.

References

[1] M. Jacob, Optimized least-square non-uniform fast fourier transform (OLS-NUFFT), IEEE Transactions on Signal Processing (2009).

[2] S. Matej, J. Fessler, I. Kazantsev, Iterative tomographic image reconstruction using Fourier-based forward and back-projectors, IEEE Transactions on Medical Imaging 23 (2004) 401 – 412.

[3] J. O'sullivan, A fast sinc function gridding algorithm for fourier inversion in computer tomography., IEEE Transactions on Medical Imaging 4 (1985) 200.

[4] J. Jackson, C. Meyer, D. Nishimura, A. Macovski, Selection of a convolution function for fourier inversion using gridding, IEEE Transactions on Medical Imaging 10 (1991) 473–478.

[5] R. Hoge, R. Kwan, G. Pike, Density compensation functions for spiral MRI, Magnetic Resonance in Medicine 38 (1997) 117–128.

[6] K. Johnson, J. Pipe, Convolution kernel design and efficient algorithm for sampling density correction, Magnetic Resonance in Medicine 61 (2009) 439–447.

[7] M. Bydder, A. Samsonov, J. Du, Evaluation of optimal density weighting for regridding, Magnetic resonance imaging 25 (2007) 695–702.

[8] H. Sedarat, D. Nishimura, On the optimality of the gridding reconstruction algorithm, IEEE Transactions on Medical Imaging 19 (2000) 306–317.

[9] D. Rosenfeld, An optimal and efficient new gridding algorithm using singular value decomposition, Magnetic Resonance in Medicine 40 (1998) 14–23.

[10] H. Moriguchi, M. Wendt, J. Duerk, Applying the uniform resampling(URS) algorithm to a lissajous trajectory: Fast image reconstruction with optimal gridding, Magnetic Resonance in Medicine 44 (2000) 766–781.

[11] H. Moriguchi, J. Duerk, Modified block uniform resampling (burs) algorithm using truncated singular value decomposition: fast accurate gridding with noise and artifact reduction, Magnetic Resonance in Medicine 46 (2001) 1189–1201.

[12] J. Shewchuk, An introduction to the conjugate gradient method without the agonizing pain, 1994.

[13] J. Tropp, A. Gilbert, Signal recovery from partial information via orthogonal matching pursuit, 2005.

[14] S. Mallat, Z. Zhang, Matching pursuits with time-frequency dictionaries, IEEE Transactions on signal processing 41 (1993) 3397–3415.

[15] P. Hansen, The L-curve and its use in the numerical treatment of inverse problems (2000).

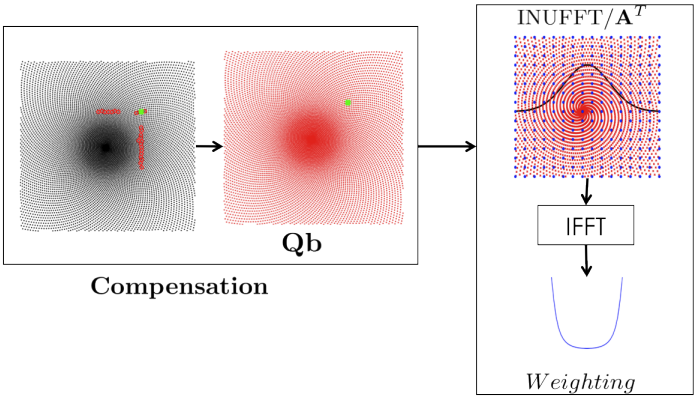

Figure 1: Outline of the proposed reconstruction scheme: The reconstruction involves two steps: (a) compensation, and (b) INUFFT. The compensation step involves the derivation of each sample in the intermediate non-uniform dataset: $\mathbf{s} = \mathbf{Qb}$. Each element of $\mathbf{s}$ is derived as a weighted linear combination of the entries of $\mathbf{b}$. The specific entries and their weights are determined by the sparsity pattern and the entries of the corresponding row of $\mathbf{Q}$. The reconstructed image is obtained as the INUFFT of the compensated dataset $\mathbf{s}$. The INUFFT consists of three steps: interpolation from the non uniform to the uniform/Cartesian grid using an interpolation kernel, IFFT, and multiplying by a weighting function to compensate for the finitely supported interpolation kernel.

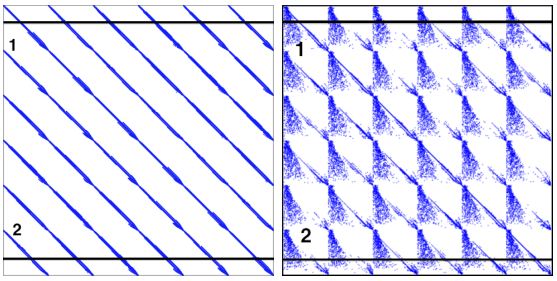

(a) NN (b) Proposed

Figure 2: Sparsity patterns of the compensation matrices for a spiral trajectory with six interleaves: (a) indicates the sparsity pattern corresponding to the nearest neighbor selection, while the pattern derived using the proposed scheme (matching pursuits) is shown in (b). Since the neighbors of a specified point falls on different interleaves, the NN matrix has a banded structure. Note that the pattern determined using the proposed scheme is very different. The flexibility in selecting the sparsity pattern results in lower errors, as seen in Fig. 4. We show the relation between the sparsity pattern of $\mathbf{Q}$ and the samples involved in the weighted linear combination (for two specific rows, indicated by 1 and 2 in the above figures) in Fig. 3.

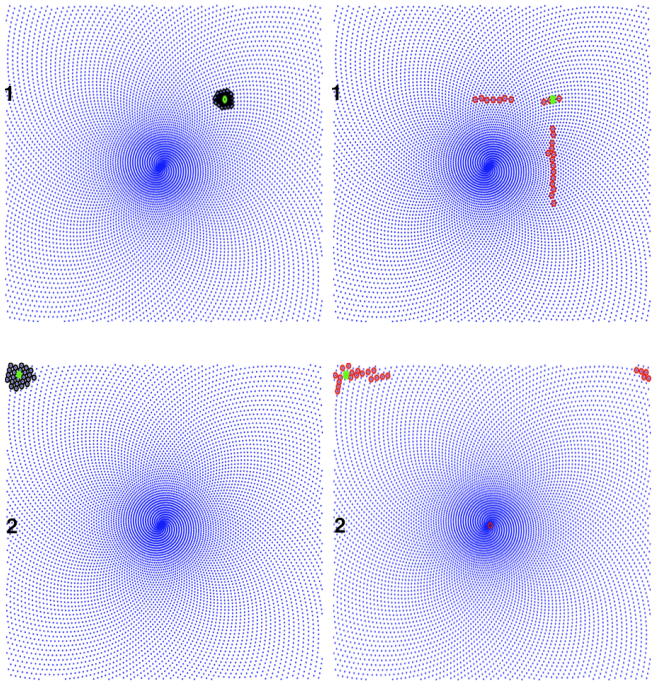

(a) Nearest neighbor (b) Proposed

Figure 3: Relation between the sparsity pattern of the **Q** matrix and the entries involved in the weighted linear combination. We focus on the evaluation of two points in the compensated dataset **s**, one at the centre of k-space and another in the outer k-space. These green points correspond to the rows of the matrix marked in Fig. 2. The non-uniform samples involved in the sparse weighted linear combination are denoted by colored circles (black for NN and red for proposed algorithm), along with the green point. In the nearest neighbor case, each sample in the compensated dataset is derived as a weighted linear combination of its immediate neighbors; since these samples fall on different interleaves, the corresponding **Q** matrix (shown in Fig. 2.(a)) has a banded structure. In contrast, the proposed scheme selects patterns that are very different, thus resulting in significantly lower reconstruction errors. Note that for the sample in outer k-space (shown in the bottom row), the linear combination also involves a sample in from the center of k-space.

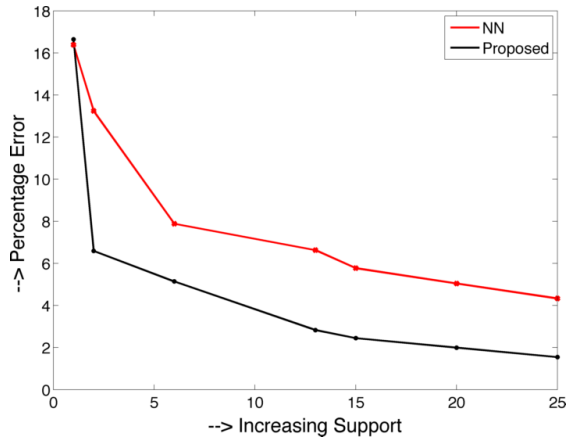

Figure 4: Error with increasing support sizes.We see that for each support size the error due to choosing supports by Nearest Neighbor is more that that for proposed method.

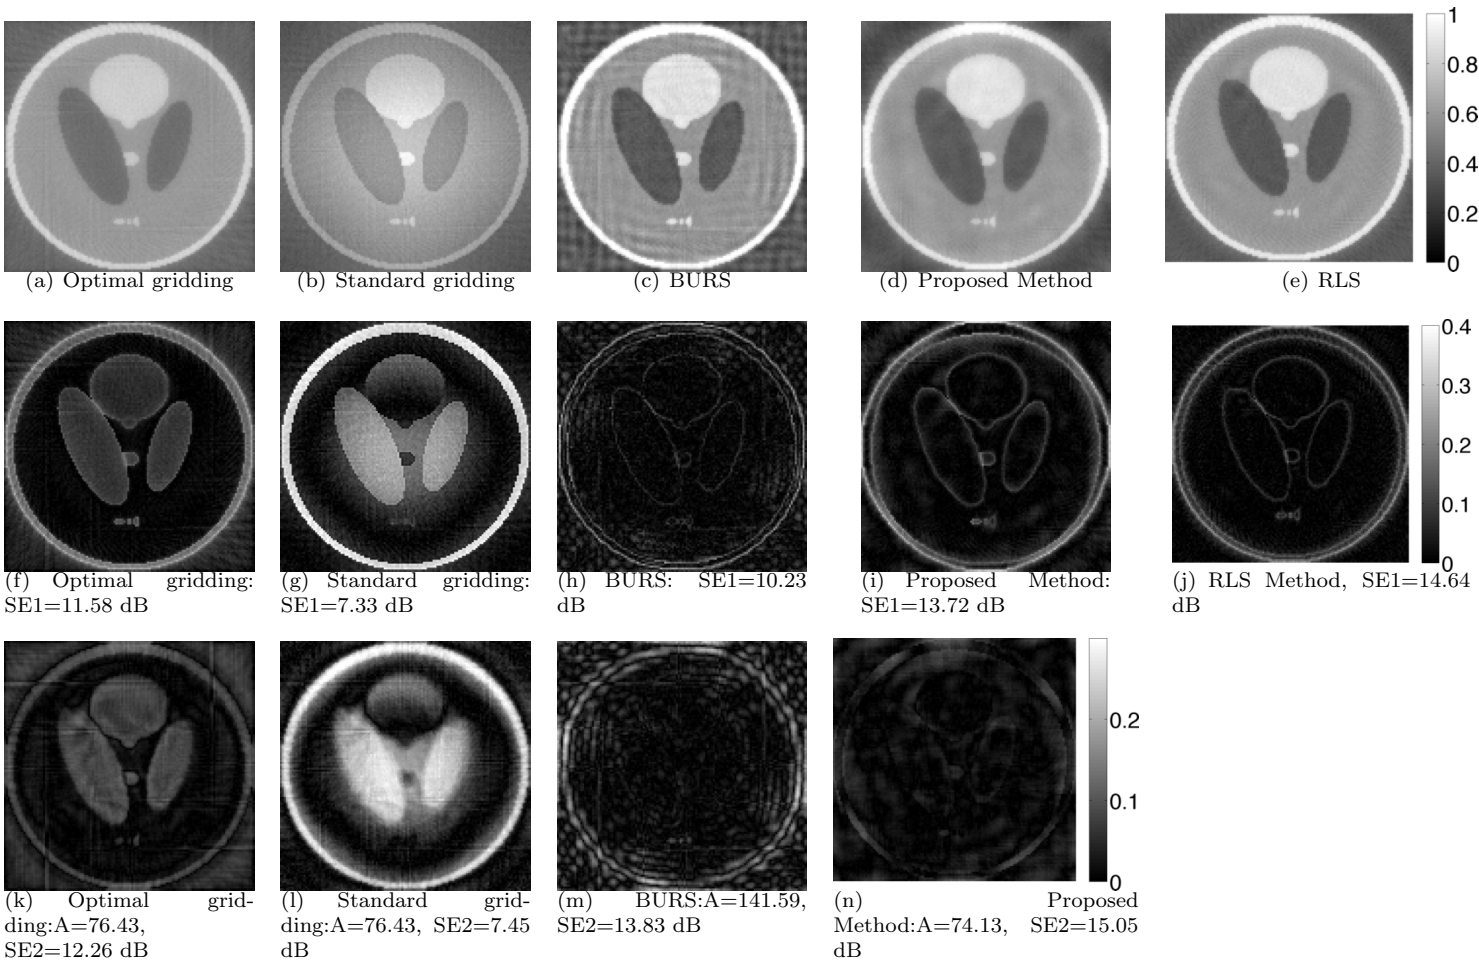

(a) Optimal gridding (b) Standard gridding (c) BURS (d) Proposed Method (e) RLS

(f) Optimal gridding: SE1=11.58 dB (g) Standard gridding: SE1=7.33 dB (h) BURS: SE1=10.23 dB (i) Proposed Method: SE1=13.72 dB (j) RLS Method, SE1=14.64 dB

(k) Optimal gridding:A=76.43, SE2=12.26 dB (l) Standard gridding:A=76.43, SE2=7.45 dB (m) BURS:A=141.59, SE2=13.83 dB (n) Proposed Method:A=74.13, SE2=15.05 dB

Figure 5: Reconstruction of the Shepp Logan phantom from its spiral samples with two interleaves: The reconstructed images are shown in the first row, the comparisons with the original phantom in the second row, and the comparisons with the RLS method is shown in the bottom row. As mentioned in the text we see that the proposed and the RLS reconstructions are comparable. Note that the BURS scheme gives structured artifacts. The results of the optimal gridding scheme is more comparable to the RLS scheme than the standard gridding method. We used $\lambda = 0.5$ for the RLS and the proposed algorithms.

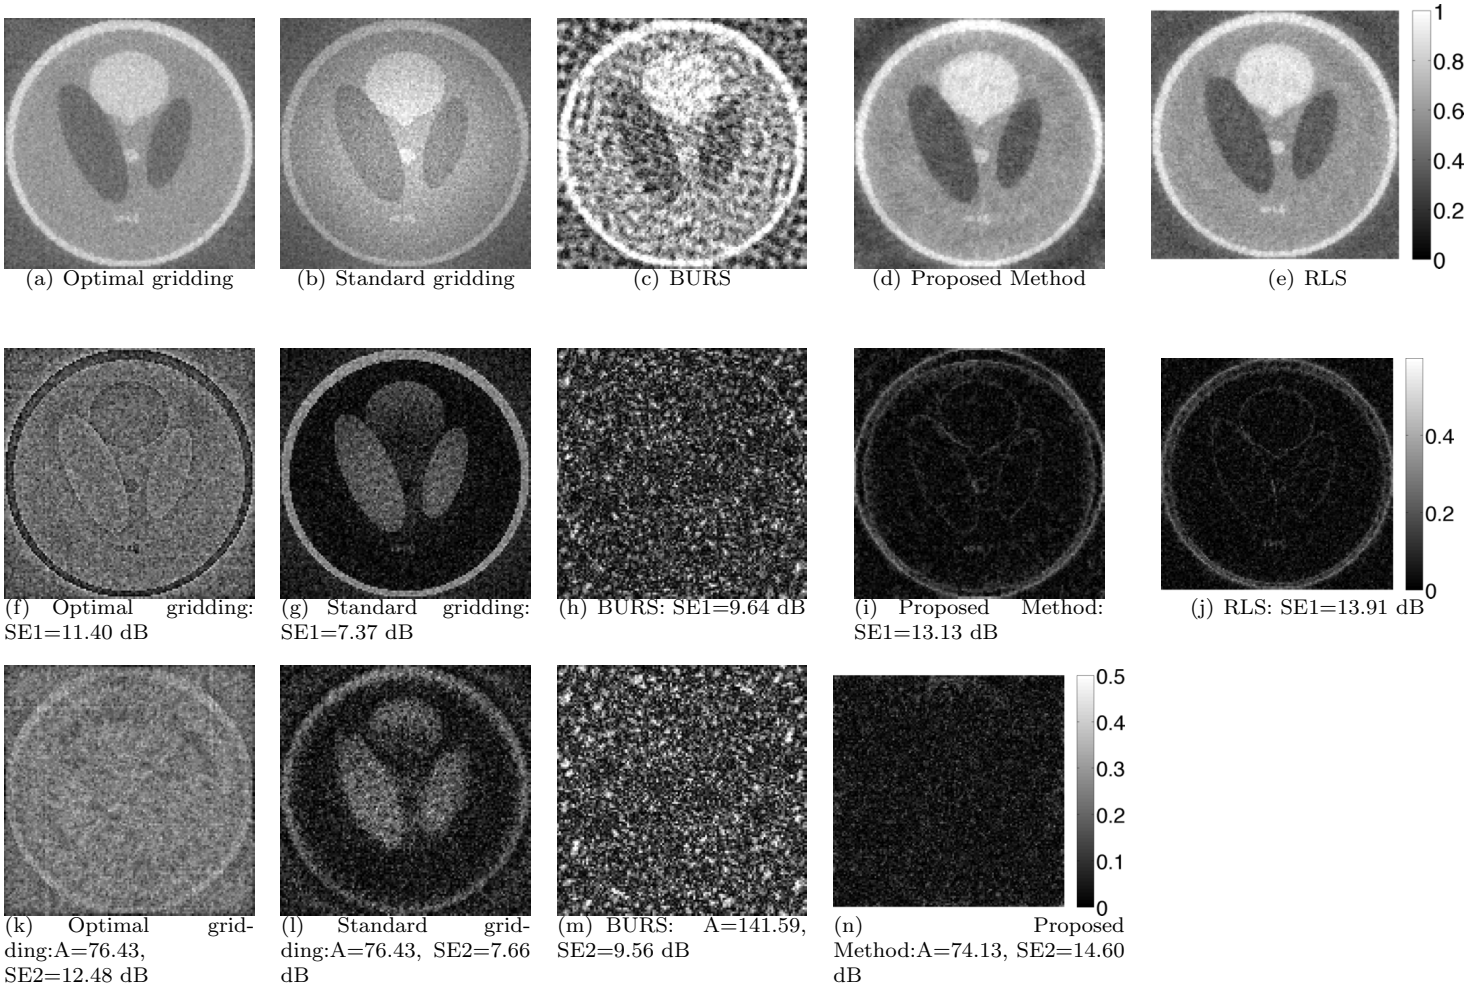

(a) Optimal gridding (b) Standard gridding (c) BURS (d) Proposed Method (e) RLS

(f) Optimal gridding: SE1=11.40 dB (g) Standard gridding: SE1=7.37 dB (h) BURS: SE1=9.64 dB (i) Proposed Method: SE1=13.13 dB (j) RLS: SE1=13.91 dB

(k) Optimal gridding:A=76.43, SE2=12.48 dB (l) Standard gridding:A=76.43, SE2=7.66 dB (m) BURS: A=141.59, SE2=9.56 dB (n) Proposed Method:A=74.13, SE2=14.60 dB

Figure 6: Reconstruction of Shepp-Logan phantom from its noisy Fourier samples on a spiral trajectory with two interleaves: The comparisons of the reconstructions to the original phantom are shown in the second row, while the comparisons with the RLS solution are shown in the last row. White Gaussian noise of 16 dB was added to the computed k-space samples, prior to the reconstructions. It is seen that the BURS is highly sensitive to noise as reported in the literature. The proposed scheme and the RLS scheme are the most robust to noise due to the explicit regularization in the framework. We used $\lambda = 0.5$ was used for RLS and the proposed schemes.

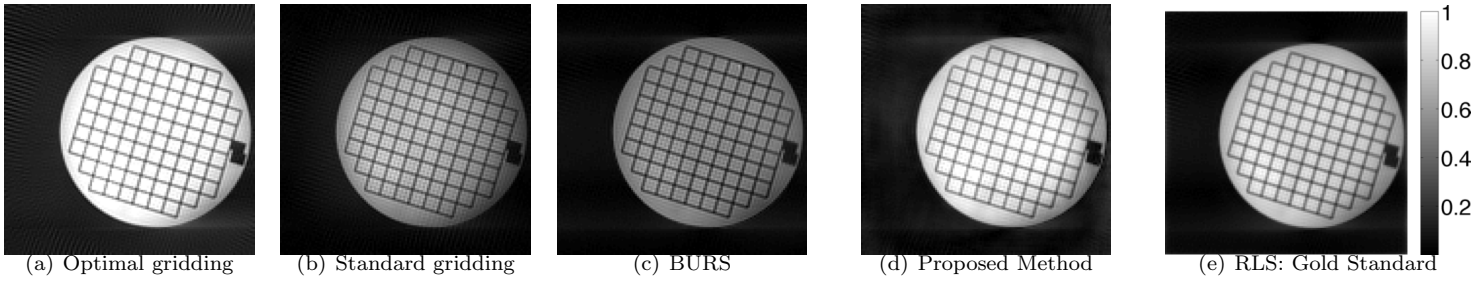

(a) Optimal gridding (b) Standard gridding (c) BURS (d) Proposed Method (e) RLS: Gold Standard

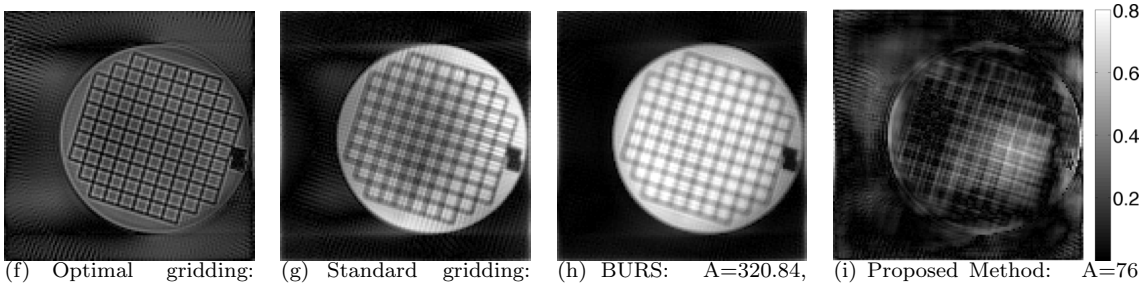

(f) Optimal gridding: A=82.78, SE2=7.49 dB (g) Standard gridding: A=82.78, SE2=3.69 dB (h) BURS: A=320.84, SE2=2.79 dB (i) Proposed Method: A=76 SE2=10.62 dB

Figure 7: Reconstruction of a grid phantom from its polar samples using different methods. The reconstructions obtained from the measured data is shown in the first row, while the errors (in comparison with RLS) are shown in the second row. Note that the reconstructions obtained with the proposed scheme are in close agreement with the RLS reconstructions and has minimal aliasing artifacts. Here, we used $\lambda = 1.0$ for the RLS and the proposed scheme. The reconstructions provided by the BURS scheme and standard gridding scheme exhibit significant artifacts. The performance improves with least square evaluation DCF, although the SNR is around 3 dB lower than the proposed method.