An Efficient Region of Interest Acquisition Method for Dynamic Magnetic Resonance Imaging

Wm. Scott Hoge¹, Student Member, IEEE, Eric L. Miller¹, Member, IEEE,
Hanoch Lev-Ari¹, Senior Member, IEEE, Dana H. Brooks¹, Member, IEEE,
W. Clem Karl², Member, IEEE, Lawrence P. Panych³

December 14, 1999

(1) Electrical and Computer Engineering Department, Northeastern University, 360 Huntington Ave., Boston, MA, 02215

(2) Electrical, Computer and Systems Engineering Department, Boston University, 8 St. Mary's St., Boston, MA, 02215

(3) Department of Radiology, Harvard Medical School and Brigham and Women's Hospital, 75 Francis Street, Boston, MA, 02115

Abstract

Motivated by recent work in the area of dynamic magnetic resonance imaging (MRI), we develop a new approach toward the problem of reduced order MRI acquisition. Recent efforts in this field have concentrated on the use of Fourier and Singular Value Decomposition (SVD) methods to obtain low order representations of an entire image plane. We augment this work to the case of imaging an arbitrarily shaped region of interest (ROI) embedded within the full image. After developing a natural error metric for this problem, we show that determining the minimal order required to meet a prescribed error level is in general intractable, but can be solved under certain assumptions. We then develop an optimization approach to the related problem of minimizing the error for a given order. Finally we demonstrate the utility of this approach and its advantages over existing Fourier and SVD methods on a number of MRI images.

EDICS: 3-MRIM

Contact Information:

Eric L. Miller Northeastern University, Room 235 Forsyth, 360 Huntington Ave., Boston, MA 02115, USA Tel: (617) 373-8386, FAX: (617) 373-8627, email: elmiller@cdsp.neu.edu

Acknowledgments:

This research was supported under NSF CAREER Grant MIP-9623721. This work is partially supported by the Army Research Office under Grant ARO DAAG55-97-1-0013, the Air Force Office of Scientific Research under Grant F49620-96-1-0028 and the National Institutes of Health under Grant NINDS 1 R01 NS34189.

1 Introduction

Over the past few decades, magnetic resonance imaging (MRI) has become a powerful tool for imaging soft tissue. MRI provides high resolution images of internal tissue structure through non-invasive means and has been used extensively in medical diagnosis and treatment planning. Traditional MRI techniques use a series of magnetic field gradients and radio-frequency (rf) pulses to encode the position and composition of molecules within a tissue volume. These excitation sequences are used to scan a volume in a sequence of slices, typically by direct sampling of the two-dimensional spatial Fourier domain, or *k-space*, of the slice. An inverse Fourier transform is then used to reconstruct images of the tissue composition within each slice. For a review of MR imaging from a signal processing perspective, see [1].

A promising new application of MRI technology is dynamic MR sequence imaging. In dynamic imaging, one acquires a sequence of images to monitor changes in the tissue [2]. Clinical applications where dynamic MR imaging is of interest include observation of the early flow of contrast agent [3,4], real time monitoring of surgical interventions or thermal treatments [5], and cardiac imaging [6]. In these cases, there is always a trade-off between temporal resolution, spatial resolution, volume coverage and signal-to-noise ratio. For example, the ability to image cardiac activity in real time comes at the expense of limited volume coverage and low spatial resolution [8]. Thus, there is a real need for faster image acquisition techniques.

The physical dynamics of MR imaging constrain the image acquisition time. Typically, one line in k-space is sampled for each excitation sequence. The required acquisition time for an image is proportional to the number of lines sampled in k-space, or equivalently, the number of excitations used. One approach to reduce the acquisition time of a single image is to lower the number of excitations employed and thus obtain a *low order* representation of the underlying image. Thus the problem of reducing acquisition time is equivalent to choosing both an appropriate number of excitations and the particular excitation sequences themselves.

Previous proposals to reduce acquisition time follow one of two general approaches. The first defines the set of excitations using no information about the underlying image. Here, the physical MRI acquisition process suggests using a reduced order Fourier basis set. Such techniques include the Fourier Keyhole method [9], and the multiple region MRI method [10]. Predefined non-Fourier excitations have also been proposed, such as the wavelet techniques in [11,12]. The second approach is to generate excitations based on some known information about the image obtained from a previously acquired scan. For example, a full image scan may be obtained just before the introduction of contrast agent or prior to surgical intervention. Information from this known image can then be used to design excitation sequences to efficiently acquire subsequent image data in the dynamic sequence. One such method is to construct excitation sequences using singular value information from the given image [2].

A key characteristic of both the Fourier and SVD methods is that they rely on image decompositions using basis functions that span the entire image region. For many dynamic MRI problems however, the relevant information is often contained within a small region of interest (ROI) within the scanned volume. Imaging tissue outside the ROI consumes time and resources, and yet provides only extraneous information. If one is able to adequately image just the ROI using a relatively small number of excitations, then the time to acquire the ROI will be correspondingly reduced.

In this work we extend image dependent excitation determination methods to the acquisition of arbitrarily shaped ROIs. As in the full image SVD approach, we assume a known prior image and use it to design appropriate image acquisition sequences. Unlike the SVD, Fourier, or wavelet methods, we are concerned with explicitly selecting a sequence that optimally represents an arbitrary ROI. Our technique is based on the linear system model for the MRI acquisition process developed in [2, 13]. Here the excitation sequences are represented as a set of finite-dimensional *input vectors.* In addition, this model requires a set of corresponding *reconstruction vectors* to build an image from the output data¹.

Our formulation of the ROI problem is motivated by the framework presented in [13]. When the ROI is in fact the whole image, it is well known that truncating the SVD results in an optimal low order representation of the original image where optimality is measured in a Frobenius norm sense [14]. Thus, we start our formulation of the arbitrary ROI problem by specifying a Frobenius norm cost function relating the error in representing the ROI in the known image to a collection of excitation and reconstruction vectors. Using this cost function, we present a method for choosing these vectors.

From a practical perspective, it seems natural to determine both the minimal number of vectors and the vectors themselves such that the error does not exceed a prescribed threshold. Since image acquisition time is proportional to the number of excitations, we optimize speed for a given level of error by minimizing the order. Unfortunately, this *minimal order* problem is exceptionally difficult to solve. In § 3 we explain why this is the case and discuss some restrictive situations where we can say something concrete about the solution to the problem or properties the solution must possess.

Given the difficulty of the minimal order problem, we next examine the case where we fix the approximation order and seek a *minimal error* solution. This formulation yields a non-linear optimization problem which we solve using a coordinate descent algorithm. A key component of this approach is the *a priori* specification of the number of vectors to be used. We employ upper bounds obtained through analysis of the minimal order problem to guide this choice. Our simulation results indicate that for convex ROIs the upper bound for zero error is especially tight.

Before continuing, we make two observations. First, by solving the minimum error problem for every possible number of excitations, one can obtain a solution to the more difficult minimal order

¹We note that both Fourier and SVD methods also require reconstruction vectors. As explained in § 2.1, in these cases, the reconstruction vectors are trivially specified once the excitation sequences are determined.

problem. Second, an assumption underlying both this work and the SVD method of [13] is that tailoring the excitation and reconstruction vectors using the known, previously acquired image will be useful for the subsequent problem of tracking dynamic changes. For the SVD method, this has been shown to hold for some important applications in [2,15]. The application of these methods to an arbitrarily specified ROI is an area in which we are currently working.

The remainder of this paper is organized as follows: In § 2, we review the linear MRI model used in this paper and present a cost function formulation for finding the excitation vectors for an arbitrarily shaped ROI. The *minimal order* problem is discussed in § 3. We approach the more tractable *minimal error* problem in § 4 and present a coordinate descent algorithm to provide numerical solutions. Simulation results are presented in § 5 with conclusions and future efforts described in § 6.

2 Background

2.1 A Linear Model for MR Imaging

The magnetic resonance imaging process can be described by a linear response model if the acquisition uses rf encoding with a low flip angle excitation [13]. For a given image A and input excitation sequence vector x, the output of a single scan, y, is equal to the matrix-vector product, y = Ax. Here x represents time samples of the rf excitation waveform envelope and y represents time samples of the received rf signal envelope. The matrix A then represents the information of interest in the slice being sampled via the representation detailed in [13]².

Generally, one collects r such scans together to form the matrix equation Y = AX, where A is an $M \times N$ matrix, $Y = [y_1 \ y_2 \ \cdots \ y_r]$ and $X = [x_1 \ x_2 \ \cdots \ x_r]$. To reconstruct the image, one must unwrap the effects of the input vectors to expose the underlying image. This is typically

²As originally presented in [13], x, y and A were frequency domain representations of the input, output, and image. However, appropriate use of the inverse Fourier transform allows for the spatial/temporal interpretation of these quantities which is employed in the work here.

accomplished by applying an $r \times N$ reconstruction matrix L to the output matrix to produce $\hat{A} = YL^* = AXL^*$ where * indicates complex-conjugate (Hermitian) transpose.

A number of MRI modalities can be represented by this model. In the case of the Fourier Keyhole (FK) method [9], the matrix X captures the lower spatial frequency components of kspace. With L = X, \hat{A} contains r low frequency components of the two dimensional spatial Fourier transform of the actual image. For SVD methods [2,16], the vectors in X are chosen from the right singular vectors of the real image matrix A and the acquired image is again reconstructed with L = X.

A key advantage of these Fourier and SVD based methods is that one can easily generate a reduced order reconstruction of A. That is, useful image information is recoverable even if one uses r < N excitation and reconstruction vectors. From our perspective, the primary disadvantage of these methods is that the vectors in X and L are global in nature, while in many cases the desired information in A is restricted to a limited region. As we demonstrate in § 5, methods that derive input and reconstruction vectors solely from global orthogonal basis sets are not well suited for arbitrarily shaped limited region problems.

2.2 Problem Formulation

Given the linear response model above, the problem approached in this paper is to build on current MRI techniques to acquire and represent only certain elements of the image matrix A. In particular, we adopt the outer-product machinery, XL^T , suggested by the previous methods, but choose X and L to reconstruct a specified but arbitrarily shaped region of interest within the image matrix. The elements of interest are described through an $M \times N$ selection matrix matrix S, with elements $s_{ij} \in \{0,1\}^3$. The ROI is designated as the region of A corresponding to the non-zero elements of S.

³We use selection matrices with binary elements in this paper. However, the results can be extended to weighted selection matrices by using selection elements in the range $0 \le s_{ij} \le 1$.

We identify a set of acquisition and reconstruction vectors through explicit formulation and minimization of the cost function

$$\mathcal{J} = \|S \circ (A - AXL^T)\|_F^2,\tag{1}$$

where A and S are of size $M \times N$, and X and L are of size $N \times r$. The \circ operator denotes an element-by-element (Hadamard, or Schur) product. For an arbitrary matrix B, the Frobenius norm is $||B||_F^2 = \sum_{i,j} b_{ij}^2$. We assume that A and any principle minor of A are full rank.

As discussed in § 1, this cost function immediately suggests two problems which could be posed. On the one hand we can set an error tolerance level and seek a minimal r such that some X and L exist which produce a cost not in excess of that value. We term this the *minimal order* problem and discuss it in § 3. Alternatively, we can fix r and seek an X and L which minimize \mathcal{J} . § 4 is devoted to an analysis and solution of this *minimal error* formulation.

3 The Minimal Order Problem

It turns out that the general case of the minimal order problem is quite intractable for mathematically precise reasons. To understand why, consider the simpler problem where we ask only for some $Q \equiv XL^T$ such that the cost is zero. We ignore for the moment the requirement that Qbe factorizable into the XL^T form, with X and L of column width r, and seek only the individual elements of Q itself. This formulation belongs to a class of matrix completion problems [17–20].

The best known matrix completion problems in signal processing involve maximum entropy extensions of autocorrelation sequences in which case the matrices possess a Toeplitz structure. Other common problems approach the completion of partially specified Hadamard or symmetric matrices. Solutions to these problems typically make deep use of the intended structural properties of the completed matrix, \hat{A} . The more general problem of choosing an unstructured Q (with or without the factorization constraint) and requiring the cost to be less than some non-zero threshold is much more complex. Other than its known usefulness in extending autocorrelation matrices for spectral estimation [21], no strong results for non-zero costs have been obtained to date.

Despite the difficulty in determining a solution to the general minimal order problem, we have found that there are cases with significant structure that allows us to say a bit more. We present two below which provide some useful insight and results which we use in our approach to the alternate, more tractable, minimal error formulation described in the next section.

3.1 Rectangular ROI, Arbitrary Error Threshold

The first case of interest is when the ROI is rectangular in shape. In this case the optimal solution to the fixed error problem can be found from the SVD of the sub-matrix chosen by S. To begin, let us assume that S takes the form

$$S = \begin{bmatrix} \mathbf{1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$$
(2)

with 1 the matrix of all ones. If the rectangular ROI is not located in the upper left corner, row and column permutations can be performed to arrive at the structure in (2). It is easily shown that \mathcal{J} is permutation invariant and no change in the cost results from these operations. Let A_{11} be the upper left block of $S \circ A$ and define r_{11} to be the number of rows in A_{11} . With these definitions we have

Theorem 1 For rectangular ROIs, the solution to the minimal order, fixed error problem is given by the largest r such that

$$\sum_{i=r+1}^{r_{11}} \sigma_i \le \epsilon$$

where ϵ is the error level, and σ_i is the *i*th singular value of A_{11} with $\sigma_1 > \sigma_2 > \ldots > \sigma_{r_{11}}$. Furthermore, the optimal X and L matrices for this solution can be obtained from the singular vectors of A_{11} .

The proof is given in Appendix A.1.

This theorem has a number of consequences of interest to us in this work. First, if the ROI is rectangular, then the SVD of the ROI (rather than the whole image) will in fact provide an optimal solution to our problem. More importantly, if the ROI is not rectangular, then the SVD of the smallest rectangle covering the ROI represents a sub-optimal solution and provides an upper bound on the solution to the underlying, arbitrary ROI problem. We use this observation in § 4.3 to guide the determination of an appropriate order for the minimum error problem.

3.2 Arbitrarily Specified ROI, Zero Error

This second case concerns arbitrary ROIs and a fixed error of zero. Here we present a suboptimal parameterization of (X, L) which provides the following sufficient condition for an order rsolution to satisfy the zero error requirement:

Theorem 2 For a given selection matrix S, an order r solution of the form

$$X = \begin{bmatrix} I_r \\ \mathbf{0} \end{bmatrix} \quad L = \begin{bmatrix} I_r \\ Q_{12}^T \end{bmatrix} \text{ or } Q_r = XL^T = \begin{bmatrix} I_r & Q_{12} \\ 0 & 0 \end{bmatrix}$$
(3)

will give zero error if $\sum_{i} s_{ij} \leq r$ for each column j of S such that j > r. Here I_r is the $r \times r$ identity matrix, and Q_{12} is an $r \times (N - r)$ sub-matrix of free parameters. The minimum r for this form is found by permuting S such that its columns contain a non-increasing number of non-zero elements.

The proof is given in Appendix A.2.

With the above parameterization of X and L, the solution can be written in closed form. There also exists a simple mechanism to find the minimal order required for zero error for this parameterization. However, there may exist other zero error solutions with an XL^T structure different than the one given in (3) and with a lower order. Thus, this theorem provides an upper bound, r_u , on the order required for zero error. Specifically, this upper bound is the smallest r that satisfies the inequality $\sum_i s_{ij} \leq r$ for all j > r. To summarize, we have presented two results for the minimal order problem. On the one hand, we have a full solution to the problem for a rectangular region and arbitrary error. This result also provides an upper bound on the smallest r required for a given error threshold and arbitrary ROI. Second, we have a sufficient condition for a size r solution to the general ROI, zero error problem. This latter result provides an upper bound on the minimum order required to meet a zero error condition. Given that non-zero error solutions require fewer vectors than zero error solutions, this latter result also provides an upper bound on the order needed to meet any error threshold.

4 The Minimal Error Problem

Given the restrictive nature of the results in the previous section, we now present an alternate formulation to the problem of choosing X and L. Specifically, rather than fixing the error level and minimizing r, we fix r and find some X and L that provide minimum error. Formally, with the cost function given in (1) we seek a solution to

$$(X_{\text{opt}}, L_{\text{opt}}) = \underset{X,L}{\operatorname{argmin}} \mathcal{J}(X, L) = \underset{X,L}{\operatorname{argmin}} \left\| S \circ (A - AXL^T) \right\|_F^2$$
(4)

for a given number of columns r in X and L.

We note that for any given r, the (X, L) pair that minimize \mathcal{J} are not unique. Any solution can be modified by an invertible matrix R of appropriate size via $XL^T = (XR)(R^{-1}L^T) = X_1L_1^T$. In principle, one could think of either seeking an alternate parameterization of the problem which yields a unique solution or of using the extra degrees of freedom in R to achieve other design objectives for X and L. All we desire here are *some* X and L which minimize \mathcal{J} .

Because (4) is quartic in the elements of X and L, a minimum of \mathcal{J} cannot in general be determined in closed form. Thus we pursue a numerical solution to the optimization problem. To start, we observe that all minima of \mathcal{J} must satisfy both of the following equations: $\partial \mathcal{J}/\partial X = 0$, and $\partial \mathcal{J}/\partial L = 0$. Computing these partial derivatives gives

$$\frac{\partial \mathcal{J}}{\partial X} = -2A^T \left[S \circ S \circ (A - AXL^T) \right] L \tag{5}$$

$$\frac{\partial \mathcal{J}}{\partial L} = -2 \left[S \circ S \circ (A - AXL^T) \right]^T AX.$$
(6)

To determine an X and L satisfying (5) and (6) we employ a variant of the Cyclic Coordinate Descent (CCD) algorithm described in [22].

4.1 CCD Algorithm

The CCD algorithm alternately solves each of the two gradient equations, (5) and (6), once in each iteration. For each iteration step, first X is held fixed, and L is found such that $\partial \mathcal{J}/\partial L = 0$. To complete the iteration, L is held fixed, and a corresponding X is found such that $\partial \mathcal{J}/\partial X = 0$. Setting each of the gradient equations above, (5) and (6), equal to zero gives

$$A^{T} (S \circ S \circ A) L = A^{T} \left[S \circ S \circ \left(A X L^{T} \right) \right] L$$

$$\tag{7}$$

$$\left[S \circ S \circ A\right]^{T} (AX) = \left[S \circ S \circ \left(AXL^{T}\right)\right]^{T} (AX).$$
(8)

These equations can be manipulated to yield a system of equations with the form Ba = cthrough the vec{} operator, which stacks the columns of a matrix into a column vector, and the Kronecker product, \otimes [23]. The vectorized versions of (7) and (8) are

$$\operatorname{vec}\{A^{T}(S \circ S \circ A)L\} = \left[(L^{T} \otimes A^{T})\operatorname{diag}\{\operatorname{vec}\{S \circ S\}\}(L \otimes A)\right]\operatorname{vec}\{X\}$$

$$(9)$$

$$\operatorname{vec}\{[S \circ S \circ A]^T (AX)\} = \left[((AX)^T \otimes I_n) \operatorname{diag}\{\operatorname{vec}\{S^T \circ S^T\}\}((AX) \otimes I_n) \right] \operatorname{vec}\{L\} \quad (10)$$

Both of these systems contain a singular matrix of the form $M^T \operatorname{diag} \{\operatorname{vec} \{S \circ S\}\} M$. The elements of S appear as a diagonal matrix embedded in the middle of the matrix product. For

region of interest problems, some elements of S will be zero, thereby causing the overall matrix to be rank deficient even if A is square and full rank. We use the Moore-Penrose pseudo-inverse [14] to solve these systems at each iteration.

Although straightforward, a direct implementation of the coordinate descent algorithm is quite computationally intensive. Equations (9) and (10) both contain a matrix of size $(rN) \times (rN)$ that must be solved via a pseudo-inverse at each iteration. However, as described in Appendix A.3, there exists a significant level of structure in these matrices that can be exploited to reduce the computational requirements for finding the system solution.

4.2 CCD Initialization

The CCD method described in § 4.1 converges to a local minimum of the cost function. Because many local minima may exist on the cost surface, convergence to a "good" minimum is dependent on the initialization point of the algorithm. After experimenting with a number of initialization heuristics, we found an approach that performed particularly well based on the X and L parameterization given in (3) of § 3.2.

Specifically, substituting (3) into (4) leaves a *linear* least squares problem for determining Q_{12} . To determine Q_{12} , we solve a set of normal equations whose structure is similar to that of (8) and (10). We then form the matrix

$$Q_r = \left[\begin{array}{cc} I_r & Q_{12} \\ \mathbf{0} & \mathbf{0} \end{array} \right]$$

and compute its SVD, $Q_r = U\Sigma V^T$. The CCD algorithm is initialized with $X = U\Sigma$ and L = V.

4.3 Choice of approximation order

Solution of the minimum error problem requires a prior specification of the approximation order r, i.e. the number of vectors in X and L. Here we concentrate on selecting the order based on upper bounds of the resulting error. Given the discussion in § 3, it is not surprising that we have

two types of bounds: one error bound based on an SVD argument for the case where we allow the cost to take on some finite, non-zero value; and one order bound based on the restricted forms of X and L in (3).

For those cases where a non-zero error is acceptable, one may use an SVD of the smallest rectangle covering the ROI to find an upper bound on the ROI reconstruction error at a given order. As shown in Theorem 1, the SVD is optimal for reconstruction of a rectangular ROI. If the SVD solution at a given order can provide an acceptable approximation error, we can guarantee that the error resulting from the localized ROI (X, L) solution will be no larger.

For arbitrarily shaped regions of interest, Theorem 2 of § 3.2 provides an upper bound on the minimum order for zero error, which we denote as r_0 . According to this theorem, after permuting S and then comparing each successive column index to the number of non-zero elements in that column, the upper bound, r_u , is the smallest r that satisfies the inequality $\sum_i s_{ij} \leq r$ for all j > r.

5 Examples

In this section, we present results from the application of our method to a few representative examples. For the images in this section, the region of interest is shown with a standard intensity map, while the region outside the ROI is shown with an inverse intensity map. That is, outside the ROI, pixels of high intensity are shown darker than pixels of low intensity.

Two extended examples are given here. In the first, we compare our method to other low order approximation techniques currently used in MR imaging. In the second example, we compare solutions found for three related ROIs applied to one image. To quantitatively compare the ROI approximations in terms of an average error per pixel, we used a normalized value of the cost function. This was calculated as the ratio between the square root of the cost function for a given solution, \mathcal{J} , and the number of non-zero elements in the selection matrix S, or equivalently,

Error per Pixel (EpP) =
$$\sqrt{\frac{\mathcal{J}}{\|S\|_F^2}} = \frac{\|S \circ (A - \hat{A})\|_F}{\|S\|_F}$$
.

5.1 Example 1: Comparison with previous methods

In the following example we compare the ROI reconstruction results given by three different methods: Singular Value Decomposition (SVD); Fourier Keyhole (FK); and the method proposed in this paper, Cyclic Coordinate Descent (CCD). For the FK reconstructions, only the r lowest spatial frequency components of the smallest sub-matrix of A containing the ROI were used in the reconstruction. For the SVD reconstruction, only the right singular vectors of the smallest sub-matrix of A containing the ROI corresponding to the r largest singular values were used. The CCD reconstruction vectors were found as described in § 4.1.



Arc S in 94 x 54 rectangle

Figure 1: Original MR Image and ROI. The region outside the ROI is indicated with an inverse intensity map.

Figure 1 shows an MRI scan of a human head along the sagittal plane. The ROI selection



Figure 2: Permuted selection matrix for Figure 1 and geometric determination of r_u .

matrix is contained within a 94 × 54 pixel rectangle. Thus to achieve zero error using either the SVD or FK techniques, 54 input vectors would be needed. Given the sparseness of the ROI and Theorem 2, we expect the order of the zero error CCD solution, r_0 , to be much lower than this. The permuted selection matrix used to determine the upper bound r_u is shown in the right panel of Figure 2. The non-zero element count for each column and a marker for the upper bound r_u is shown in the right panel. The upper bound on r_0 is determined as per § 4.3 and is found to be $r_u = 43$. We note that if the upper bound is tight, using 43 input vectors to re-scan the ROI will still give a decrease in the acquisition time of about 20% compared to the orthogonal techniques, with perfect accuracy in the ROI.

Figure 3 compares the error curves for the three different methods of low order approximation over a range of solution orders. The right panel shows the upper range of orders with the average error per pixel value plotted on a log scale to show greater detail near zero error. The figure shows that at a given order, the CCD solution provides lower reconstruction error than either the FK or the SVD methods. Furthermore, Figure 3 shows that for a given error tolerance, a CCD reconstruction of the ROI is available at a much lower order than either the SVD or Fourier approximation methods provide. For instance, if the number of input vectors is fixed at 10, the CCD solution has an average pixel error that is one half that given by the SVD. Conversely, for a given error per pixel of 10, the ROI can be acquired in less than half the time using the CCD method. We found similar results for the many image examples we examined with a general ROI specified.



Figure 3: Estimation cost curves for SVD, Fourier Keyhole, and Cyclic Coordinate Descent solutions for Figure 1 ROI. Right panel shows logarithmic scale for better detail near zero error per pixel.

Comparison of the three ROI reconstructions methods, (SVD, FK, and CCD), are given below for two specific orders, r = 10 and r = 25.

Figure 4 shows the order 10 ROI reconstructions. The absolute difference in pixel values for the same methods and order are shown in Figure 5. It is clear from both figures that the solution found by the CCD method has significantly less pixel error than either the FK or SVD methods. The CCD solution also shows a more even distribution of the error across the ROI, and greater structural information in the ROI than either of the global orthogonal approximation methods.

The order 25 reconstructions are shown in Figure 6, with the absolute error illustrated in



Figure 4: Comparison of order r = 10 ROI reconstructions for (a) Cyclic Coordinate Descent $[\hat{A}_{CCD}]$, (b) SVD $[\hat{A}_{SVD}]$, and (c) Fourier keyhole $[\hat{A}_{FK}]$ methods



Figure 5: Pixel value difference comparisons of order r = 25 ROI reconstructions for (a) Cyclic Coordinate Descent $|A - \hat{A}_{CCD}|$, (b) SVD $|A - \hat{A}_{SVD}|$, and (c) Fourier keyhole $|A - \hat{A}_{FK}|$ methods

Figure 7. While all three approximations now show structural detail in the ROI, there is still an order of magnitude difference in the error per pixel measure. This is confirmed visually in the absolute difference illustrations shown in Figure 7. Here, negligible error is shown for the CCD solution, while significant error still occurs in the other two.



Figure 6: Comparison of order r = 25 ROI reconstructions for (a) Cyclic Coordinate Descent $[\hat{A}_{CCD}]$, (b) SVD $[\hat{A}_{SVD}]$, and (c) Fourier keyhole $[\hat{A}_{FK}]$ methods

5.2 Example 2: ROI comparison

In this example, we compare three separate ROIs for the same image. The ROIs used are a square, a circle, and a "horse-shoe" shape, each contained within a 75×75 pixel region. The image and the selection matrices are shown in Figure 8. Note that the square ROI is in fact the smallest rectangle that covers both of the other two ROIs.

For the square ROI, Theorem 1 shows that the SVD is in fact the optimal solution. For the other two ROIs however, the CCD algorithm provides a significant improvement over the SVD solution at any order specified. This is demonstrated in Figures 9 and 10 which show the average error per pixel for each ROI reconstruction over a range of approximation orders. Again, the right panel



Figure 7: Pixel value difference comparisons of order r = 25 ROI reconstructions for (a) Cyclic Coordinate Descent $|A - \hat{A}_{CCD}|$, (b) SVD $|A - \hat{A}_{SVD}|$, and (c) Fourier keyhole $|A - \hat{A}_{FK}|$ methods



Figure 8: MRI image of a knee with a square ROI shown. Other ROIs applied to the same 75×75 pixel region are: (a) a circle, and (b) a "horse-shoe" shape.

shows a logarithmic plot for greater detail near zero error. In both figures, the ROI approximation error for the CCD solution is significantly smaller than for the SVD solution. This improvement is more pronounced in the case of Figure 10, and is attributed to the fact that the horse-shoe selection matrix is much more sparse than the circle.



Figure 9: Estimation cost curves for the SVD and CCD reconstructions of Figure 8 with the circular ROI. Right panel shows logarithmic scale for better detail near zero error per pixel.

In addition, the order required for zero error occurs at a much lower order as the ROI becomes more sparse. From Theorem 1, we know that for the square ROI, $r_0 = 75$ is required for zero error in the reconstruction. Theorem 2 provides an upper bound on r_0 for the other two ROIs. As per the discussion in § 3.2, the upper bound for zero error is found to be $r_u = 53$ for the circle ROI and $r_u = 26$ for the horse-shoe ROI. This is confirmed by Figure 11, where the results of the simulations show $r_0 = 75 = r_u$ for the square, $r_0 = 53 = r_u$ for the circle, and $r_0 = 20 < r_u = 26$ for the horse-shoe. Note that with the exception of the horse-shoe ROI, the upper bound on zero error is tight despite the constrained form of (X, L) used to compute r_u .

Reconstructions of order 10 for the circular ROI are shown in Figure 12. Reconstructions of



Figure 10: Estimation cost curves for the SVD and CCD reconstructions of Figure 8 with the horse-shoe ROI. Right panel shows logarithmic scale for better detail near zero error per pixel.



Figure 11: Error per pixel comparison over a range of orders for the CCD reconstruction of three ROIs in Figure 8.

order 5 for the horse-shoe selection matrix are shown in Figure 13. Both figures show the CCD and SVD ROI reconstructions and the absolute value of the approximation error at each pixel. These figures clearly show that the CCD algorithm provides much higher quality ROI reconstructions at low orders over the SVD. This difference in quality performance becomes more pronounced as the ROI grows more sparse in a given rectangular region.



Figure 12: SVD and CCD order r = 10 comparison for the circular ROI in Example 2. Top row shows ROI reconstructions, \hat{A} . Bottom row shows absolute difference in pixel values, $|A - \hat{A}|$.

6 Summary

In this paper, we presented a problem formulation, associated analysis, and a computational method to find a set of excitation and reconstruction vectors for efficient MRI acquisition of an arbitrarily shaped ROI. The solution provides an ROI reconstruction of a given order with minimum error in the least squares sense. We illustrated that solutions found with this method have significantly lower approximation error in the ROI than solutions found through traditional SVD or Fourier based methods. Additionally, we showed that one can find a zero error ROI reconstruction



Figure 13: SVD and CCD order r = 5 comparison for the horse-shoe ROI in Example 2. Top row shows ROI reconstructions, \hat{A} . Bottom row shows absolute difference in pixel values, $|A - \hat{A}|$.

with fewer scans (i.e. with lower approximation order) than traditional methods allow.

A key focus of our future efforts will be the use of the framework developed here in a dynamic setting. Assessing the utility of our optimal vectors for tracking changes in a ROI is clearly of interest. Developing efficient methods for jointly estimating the ROI and adaptively updating our low order excitation sequence as the scene changes is another area of interest.

A Appendix

A.1 Proof of Theorem 1

Theorem 1: For rectangular ROIs, the solution to the minimum order, fixed error problem is given by the largest r such that $\sum_{i=r+1}^{r_{11}} \sigma_i \leq \epsilon$ where ϵ is the error level, and σ_i is the *i* singular value of A_{11} with $\sigma_1 > \sigma_2 > \ldots > \sigma_{r_{11}}$. Furthermore, the optimal X and L matrices for this solution can be obtained from the singular vectors of A_{11} .

Proof: As mentioned in § 3.1, selection matrices with a rectangular ROI can always be permuted

to the form of (2). Such matrices can be described by an outer product of two vectors, $S = s_1 s_2^T$. If the non-zero sub-block **1** is of size $m \times n$, then s_1 and s_2 are vectors with m and n leading ones and (M-m) and (N-n) trailing zeros, respectively. As shown in [23, p.304], a Hadamard product involving such a matrix can be rewritten as a conventional matrix product containing two diagonal matrices. Thus for matrices with a rectangular ROI, the cost function can be written as

$$\mathcal{J} = \|(s_1 s_2^T) \circ (A - AXL^T)\|_F^2 = \|D_1 (A - AXL^T) D_2\|_F^2$$

where D_1 and D_2 are diagonal matrices with s_1 and s_2 along their respective diagonals. This can further be simplified to

$$\mathcal{J} = \left\| (S \circ A) - ZW^T \right\|_F^2$$

where $Z = D_1 A X$ and $W = D_2 L$. The optimal solution for W and Z can be found through the SVD of $(S \circ A)$. The structure of the optimal solution is

$$Z = D_1 A X = \begin{bmatrix} Z_1 \\ 0 \end{bmatrix}$$
$$W^T = L^T D_2 = \begin{bmatrix} L_1^T & 0 \end{bmatrix}$$
$$(S \circ A) = \begin{bmatrix} A_{11} & 0 \\ 0 & 0 \end{bmatrix}$$

Let the singular value decomposition of the rectangular sub-matrix A_{11} be $A_{11} = U_1 \Sigma_1 V_1^T$. The error at a given approximation r is therefore the sum of the discarded singular values, or equivalently

$$\mathcal{J} = \sum_{i=r+1}^{r_{11}} \sigma_i$$

For the approximation to be less than a given error threshold ϵ , one need only choose r such that $\sum_{i=r+1}^{r_{11}} \sigma_i < \epsilon.$ Returning now to choose an optimal X and L^T , one may use the SVD decomposition to find $X = (D_1 A)^{\dagger} \begin{bmatrix} U_1 \Sigma_1 \\ 0 \end{bmatrix} = \begin{bmatrix} V_1 \\ 0 \end{bmatrix} \text{ and } L^T = \begin{bmatrix} V_1^T & 0 \end{bmatrix}.$

A.2 Proof of Theorem 2

Theorem 2: For a given selection matrix S, an order r solution of the form

$$X = \begin{bmatrix} I_r \\ \mathbf{0} \end{bmatrix} \quad L = \begin{bmatrix} I_r \\ Q_{12}^T \end{bmatrix} \text{ or } Q_r = XL^T = \begin{bmatrix} I_r & Q_{12} \\ 0 & 0 \end{bmatrix}$$

will give zero error if $\sum_{i} s_{ij} \leq r$ for each column j of S such that j > r. Here I_r is the $r \times r$ identity matrix, and Q_{12} is an $r \times (N - r)$ sub-matrix of free parameters. The minimum r for this form is found by permuting S such that the columns contain a non-increasing number of non-zero elements.

Proof: This theorem is shown true by considering that for $\mathcal{J} = 0$, the following equation must hold,

$$S \circ A = S \circ (AXL^T) = S \circ (AQ_r).$$
⁽¹¹⁾

From this, one can recognize that all columns in the Q_r formulation may be treated independently. The first r columns of the Q_r parameterization contain the identity matrix in the upper sub-block, and zeros elsewhere. Thus, the first r columns of the approximation (AQ_r) will be identical to the first r columns of A, satisfying (11) for those columns.

For the remaining columns, indexed from (r + 1) to N, each of the column equations can be rewritten as a system of equations with row size dependent on the number of non-zero elements in the j^{th} column of S. If one constructs the vector α_j to contain the index values of the non-zero elements in the j^{th} column of S, then this column system may be written as

$$A(\alpha_j, j) = A(\alpha_j, 1:r) \ q_j. \tag{12}$$

Here q_j is a length r vector containing the free parameters of the j^{th} column of Q_r . The vector $A(\alpha_j, j)$ is composed of elements from the j^{th} column of A as specified α_j . The matrix $A(\alpha_j, 1:r)$ is composed by taking certain rows as specified by α_j from the first r columns of the original A matrix. The number of rows in each column system depends on the number of ones in the j^{th} column of S. If $\sum_i s_{ij}$ is greater than r, then the system is over-determined and the system can only be solved in an approximate sense. However, if $\sum_i s_{ij}$ is less than or equal to r, then the system is under- or exactly-determined, and with our previous assumptions on A, an exact solution exists. Thus, if r is chosen such that $\sum_i s_{ij} \leq r$ for each of the columns $\{j; j > r\}$, then none of the column systems will be over-determined. Under this condition, collecting each of these column systems together, the order r solution of the form

$$Q_r = \left[\begin{array}{cc} I_r & Q_{12} \\ 0 & 0 \end{array} \right]$$

has a sufficient number of free parameters to ensure that $\mathcal{J} = 0$. If we permute S such that the columns contain a non-increasing number of non-zero elements, we will find the minimum value of r in the above expressions.

We note here that other zero-error solutions exist for a given problem, possibly with order less than the order of the Q_r solution. Thus, Theorem 2 provides an upper bound on the minimum order needed for a zero error solution. This result is used in the algorithm initialization discussion, presented in § 4.2, and verified in § 5.

A.3 Efficient solution of vectorized systems

In the development of the CCD algorithm, systems of linear equations appear which are described by matrix equations that contain Kronecker products. Systems containing Kronecker products tend to be very large and require a substantial amount of memory and processing power to solve. However, the system matrices presented in this paper, (9) and (10), contain a significant level of structure that can be exploited to assist in finding the system solution.

In each of the equations mentioned, the matrix to invert is symmetric, and can be written to contain expressions of the form $R = (B^T \otimes I)$ diag $\{vec\{C\}\}(B \otimes I)$. For example, in (9), B = (AX) and $C = (S \circ S)$. Expressions of this type can be rewritten in block matrix form, with each block containing a diagonal matrix, determined as

$$R_{ij} = (B_{(:,i)}^T \otimes I) \operatorname{diag}\{C\}(B_{(:,j)} \otimes I) = \operatorname{diag}\{C(B_{(:,i)} \circ B_{(:,j)})\}.$$

Collecting the sub-blocks together, we find

$$R = \begin{bmatrix} \operatorname{diag}\{C(B_{(:,1)} \circ B_{(:,1)})\} & \operatorname{diag}\{C(B_{(:,1)} \circ B_{(:,2)})\} & \cdots & \operatorname{diag}\{C(B_{(:,1)} \circ B_{(:,m)})\} \\ \operatorname{diag}\{C(B_{(:,2)} \circ B_{(:,1)})\} & \operatorname{diag}\{C(B_{(:,2)} \circ B_{(:,2)})\} & \cdots & \operatorname{diag}\{C(B_{(:,2)} \circ B_{(:,m)})\} \\ \vdots & \vdots & \ddots & \vdots \\ \operatorname{diag}\{C(B_{(:,m)} \circ B_{(:,1)})\} & \operatorname{diag}\{C(B_{(:,m)} \circ B_{(:,2)})\} & \cdots & \operatorname{diag}\{C(B_{(:,m)} \circ B_{(:,m)})\} \end{bmatrix}.$$

This matrix can then be permuted to form a block diagonal matrix via

$$P_r R P_c = \begin{bmatrix} \Box & & & \\ & \Box & & \\ & & \ddots & \\ & & & \Box \end{bmatrix}$$

To find the pseudo-inverse of this matrix, one may use the SVD of each individual sub-block. By decomposing the matrix in this way, the processing resources required to compute the pseudoinverse can be dramatically reduced. This enables the solutions required in each iteration of the CCD algorithm, \S 4, to be calculated very quickly.

References

- G. A. Wright, "Magnetic resonance imaging," *IEEE Signal Processing Mag.*, vol. 14, pp. 56–66, Jan 1997.
- [2] L. P. Panych, C. Oesterle, G. P. Zientara, and J. Henning, "Implementation of a fast gradientecho SVD encoding technique for dynamic imaging," *J Magn Reson Imaging*, vol. 35, pp. 554– 562, 1996.

- [3] J. Tacke, G. Adam, H. Classen, A. Muhler, A. Prescher, and R. W. Gunther, "Dynamic MRI of a hypovascularized liver tumor model," *J Magn Reson Imaging*, vol. 7, pp. 678–82, Jul-Aug 1997.
- [4] F. Wallis and F. J. Gilbert, "Magnetic resonance imaging in oncology: an overview," J R Coll Surg Edinb, vol. 44, pp. 117–25, Apr 1999.
- [5] F. A. Jolesz and S. M. Blumenfeld, "Interventional use of magnetic resonance imaging," Magn Reson Q, vol. 10, pp. 85–96, Jun 1994.
- [6] P. C. Yang, A. B. Kerr, A. C. Liu, D. H. Liang, C. Hardy, C. H. Meyer, A. Macovski, J. M. Pauly, and B. S. Hu, "New real-time interactive cardiac magnetic resonance imaging system complements echocardiography," J Am Coll Cardiol, vol. 32, pp. 2049–56, Dec 1998.
- [7] H. L. Liu, Y. Pu, Y. Liu, L. Nickerson, T. Andrews, P. T. Fox, and J. H. Gao, "Cerebral blood flow measurement by dynamic contrast MRI using singular value decomposition with an adaptive threshold," *Magn Reson Med*, vol. 42, pp. 167–72, Jul 1999.
- [8] A. B. Kerr, J. M. Pauly, B. S. Hu, K. C. Li, C. J. Hardy, C. H. Meyer, A. Macovski, and D. G. Nishimura, "Real-time interactive mri on a conventional scanner," *Magn Reson Med*, vol. 38, pp. 355–67, Sep 1997.
- [9] M. E. Brummer, W. Dixon, B. Gerety, and H. Tuithof, "Composite k-space windows (keyhole techniques) to improve temporal resolution in a dynamic series of images following contrast administration," in *Proc. ISMRM 11th Annual Meeting*, p. 4236, 1992.
- [10] S. K. Nagle and D. Levin, "Multiple region MRI," Magn Reson Med, vol. 41, no. 4, pp. 774–786, 1999.
- [11] D. M. Healy, Jr and J. B. Weaver, "Two applications of wavelet transforms in magnetic resonance imaging," *IEEE Trans. Information Theory*, vol. 38, pp. 840–860, Mar 1992.
- [12] L. P. Panych, P. D. Jakab, and F. A. Jolesz, "An implementation of wavelet encoded MRI," J Magn Reson Imaging, vol. 3, p. 649, 1993.
- [13] L. P. Panych, G. P. Zientara, and F. A. Jolesz, "MR image encoding by spatially selective rf excitation: An analysis using linear response models," *Int J Imaging Syst Technol*, vol. 10, no. 2, pp. 143–150, 1999.
- [14] R. A. Horn and C. R. Johnson, *Matrix Analysis*. New York, NY: Cambridge University Press, 1985.
- [15] G. P. Zientara, L. P. Panych, and F. A. Jolesz, "Applicability and efficiency of near-optimal spatial encoding for dynamically adaptive MRI," *Magn Reson Med*, vol. 39, pp. 204–13, Feb 1998.
- [16] G. P. Zientara, L. P. Panych, and F. A. Jolesz, "Dynamically adaptive MRI with encoding by singular value decomposition," *Magn Reson Med*, vol. 32, no. 2, pp. 268–74, 1994.
- [17] N. Cohen, C. R. Johnson, L. Rodman, and H. J. Woerdeman, "Ranks of completions of partial matrices," in *Oper. Theory: Adv. Appl.* (H. Dym, S. Goldberg, M. A. Kaashoek, and P. Lancaster, eds.), vol. 40 of *The Gohberg anniversary collection*, Vol. I, pp. 165–185, 1989.

- [18] C. R. Johnson and G. T. Whitney, "Minimum rank completions," *Linear and Multilinear Algebra*, vol. 28, pp. 271–273, 1991.
- [19] C. R. Johnson, "Matrix completion problems: a survey," in *Matrix Theory and Applications* (C. R. Johnson, ed.), vol. 40 of *Proc. Sympos. Appl. Math*, (Providence, RI), pp. 171–198, Amer. Math. Soc., 1990.
- [20] C. R. Johnson, B. Kroschel, and H. Wolkowicz, "An interior-point method for approximate positive semidefinite completions," *Comput. Optim. Appl.*, vol. 9, no. 2, pp. 175–190, 1998.
- [21] W. I. Newman, "Extension to the maximum entropy method," IEEE Trans. Image Processing, vol. IT-23, p. 89, Jan 1997.
- [22] D. G. Luenberger, Linear and Non-Linear Programming. Reading, MA: Addison-Wesley, 2nd ed., 1984.
- [23] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*. New York, NY: Cambridge University Press, 1991.