Tradeoffs and complexities in model-based MR image reconstruction

Jeffrey A. Fessler *
EECS Dept., The University of Michigan
fessler@umich.edu

December 6, 2007

Abstract

Model-based methods for image reconstruction in magnetic resonance imaging (MRI) have seen increasing interest recently. This syllabus summarizes some of the complexities and tradeoffs that arise in model-based image reconstruction methods.

1 Introduction

The inverse fast Fourier transform (FFT) has served the MR community very well as the conventional image reconstruction method for k-space data with full Cartesian sampling. And for well sampled non-Cartesian data, the gridding method [1] with appropriate density compensation factors [2] is fast and effective. But when only under-sampled data is available, or when non-Fourier physical effects like field inhomogeneity are important, then gridding/FFT methods for image reconstruction are suboptimal, and iterative algorithms based on appropriate models can improve image quality, at the price of increased computation. This paper synposizes some of the issues that arise when using iterative algorithms for model-based MR image reconstruction. The references give pointers to some recent work but are by no means a comprehensive survey.

2 Signal Model

We begin by reviewing a typical approach to model-based image reconstruction. Because parallel imaging is of considerable interest, we consider the general case of L receive coils. A standard single receive coil is a simple special case. For simplicity we focus on static imaging, though generalizations to dynamic imaging is of particular interest.

Based on the solution to the Bloch equation, a reasonable model for the (demodulated) received signal of the lth receive coil is

$$s_l(t) = \int f(\vec{r})\, c_l(\vec{r})\, \mathrm{e}^{-R_2^*(\vec{r})t}\, \mathrm{e}^{-\imath\, \omega(\vec{r})\, t}\, \mathrm{e}^{-\imath 2\pi \vec{k}(t)\cdot\vec{r}}\, \mathrm{d}\vec{r}, \tag{1}$$

where $\vec{r}$ denotes spatial position (in 2D or 3D), $c_l(\vec{r})$ denotes the sensitivity of the lth receive coil, $R_2^*(\vec{r})$ denotes the relaxation map of the object, $\omega(\vec{r})$ denotes the off-resonance frequency map (field map), $\vec{k}(t)$ denotes the k-space trajectory of the scan and $f(\vec{r})$ denotes the (unknown) transverse magnetization of the object that we wish to reconstruct from the data.

MR scan data consists of noisy samples of the above signal:

$$y_{li} = s_l(t_i) + \varepsilon_{li}, \qquad i = 1, \ldots, n_{\mathrm{d}}, \quad l = 1, \ldots, L, \tag{2}$$

where y_{li} denotes the ith sample of the lth coil's signal and $\boldsymbol{\varepsilon}_{li}$ denotes complex white gaussian noise, and n_{d} denotes the number of k-space samples.

Typically, the goal of MR image reconstruction is to estimate $f(\vec{r})$ from the measurement vector $\boldsymbol{y} = (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_L)$, where $\boldsymbol{y}_l = (y_{l1}, \ldots, y_{l,n_{\mathrm{d}}})$. This is an ill-posed problem because the given measurements $\boldsymbol{y}$ are discrete whereas the object magnetization $f(\vec{r})$ is an unknown continuous-space function. To proceed, we parameterize the object $f(\vec{r})$ to facilitate parametric estimation using a "finite series expansion" as follows:

$$f(\vec{r}) = \sum_{j=1}^{N} f_j\, b(\vec{r} - \vec{r}_j), \tag{3}$$

*Supported in part by NIH grants EB002683 and DA15410.

where $b(\cdot)$ denotes the object basis function, $\vec{r}_j$ denotes the center of the jth basis function translate, and N is the number of parameters. For simplicity, hereafter we use rect basis functions $b(\vec{r}) = \mathrm{rect}(\vec{r}/\Delta)$, *i.e.*, square pixels of dimension Δ, so N is the number of pixels, or voxels in 3D scans. Many other possible basis function choices can be considered, all of which are imperfect because the true object is not parametric, but nevertheless reasonable basis functions can be useful.

Substituting the basis expansion (3) into the signal model (1) and simplifying leads to the model

$$s_l(t_i) = \sum_{j=1}^{N} a_{ij} f_j \tag{4}$$

where the elements $\{a_{lij}\}$ of the system matrix $\boldsymbol{A}_l$ associated with the lth coil are given by

$$a_{lij} = \int b(\vec{r} - \vec{r}_j)\, c_l(\vec{r})\, \mathrm{e}^{-\,z(\vec{r})\,t_i}\, \mathrm{e}^{-\imath 2\pi \vec{k}(t_i)\cdot\vec{r}}\, \mathrm{d}\vec{r}, \tag{5}$$

where we define the "rate map" by combining the relaxation and field maps:

$$z(\vec{r}) \triangleq R_2^*(\vec{r}) + \imath\, \omega(\vec{r})\,.$$

(Often this rate map is assumed to be zero, *i.e.*, relaxation and off resonance are ignored.) We can combine (2) and (4) in matrix-vector form as follows:

$$\boldsymbol{y}_l = \boldsymbol{A}_l \boldsymbol{f} + \boldsymbol{\varepsilon}_l,$$

where $\boldsymbol{f} = (f_1, \ldots, f_N)$ is the vector of parameters (pixel values) that we wish to estimate from the data $\boldsymbol{y}$. Equivalently, stacking up all L vectors and defining the $n_{\mathrm{d}} L \times N$ matrix $\boldsymbol{A} = (\boldsymbol{A}_1, \ldots, \boldsymbol{A}_L)$ yields the linear model $\boldsymbol{y} = \boldsymbol{A}\, f + \boldsymbol{\varepsilon}$. At first glance this linear model appears amenable to a variety of iterative solution methods. However, the first challenge that arises is that the elements of $\boldsymbol{A}$ are quite complicated in the form above, and not evidently amenable to fast computation.

To simplify further, we also expand the coil sensitivity maps $\{c_l(\vec{r})\}$ and the rate map $z(\vec{r})$ using the same rectangular basis functions:

$$\begin{aligned} c_l(\vec{r}) &= \sum_{j=1}^{N} b(\vec{r} - \vec{r}_j)\, c_{lj} \\ z(\vec{r}) &= \sum_{j=1}^{N} b(\vec{r} - \vec{r}_j)\, z_j. \end{aligned} \tag{6}$$

Here, z_j denotes the rate map value in the jth voxel, and c_{lj} denotes the sensitivity map value for the lth coil in the jth voxel.

Substituting these expansions (or approximations) into (5) and simplifying yields

$$a_{lij} = B\Big(\vec{k}(t_i)\Big)\, \mathrm{e}^{-z_j t_i}\, \mathrm{e}^{-\imath 2\pi \vec{k}(t_i)\cdot\vec{r}_j}\, c_{lj}, \tag{7}$$

where B is the Fourier transform of the basis function $b(\cdot)$.

The above model is a generalization of the approach described in [3] to the case of multiple coils, and is amenable to the fast iterative algorithms described therein, *e.g.*, [4]. In particular, because the noise in MR is gaussian, a natural approach is to estimate $\boldsymbol{f}$ by minimizing a regularized least-squares cost function:

$$\hat{\boldsymbol{f}} = \arg\min_{\boldsymbol{f}} \Psi(\boldsymbol{f}), \qquad \Psi(\boldsymbol{f}) \triangleq \|\boldsymbol{y} - \boldsymbol{A}\boldsymbol{f}\|^2 + \beta\, \mathsf{R}(\boldsymbol{f})\,. \tag{8}$$

The norm $\|\cdot\|$ should include the inverse of the covariance matrix that describes the noise correlation between receive coils [5,6].

3 Challenges

3.1 Regularization

The next challenge in iterative reconstruction for MR imaging is choosing the regularizer $\mathsf{R}(\boldsymbol{f})$. If this term is not included, then the image estimate $\hat{\boldsymbol{f}}$ will be suffer from noise and artifacts for under-sampled and/or non-Cartesian data, because this inverse problem is ill-conditioned. The approach for iterative reconstruction that has been adopted in commercial PET scanners is to use an unregularized algorithm [7], initialize it with a uniform image, stop iterating when just as the image gets unacceptly noisy, and then perhaps apply a bit of post-filtering to reduce the noise. One can adopt a similar approach for MR imaging [8]. However, introducing regularization can ensure that the iterative algorithm converges to a stable image, and can enforce prior information that improves image quality particularly for under-sampled data.

The simplest choice is Tikhonov regularization $\mathsf{R}(\boldsymbol{f}) = \|\boldsymbol{f}\|^2$ or $\mathsf{R}(\boldsymbol{f}) = \left\|\boldsymbol{f} - \bar{\boldsymbol{f}}\right\|^2$, where $\bar{\boldsymbol{f}}$ is some prior or reference image (possibly zero). The disadvantage of this choice is that it biases the estimate towards the reference image $\bar{\boldsymbol{f}}$. In particular, if the reference image is zero, then all pixel values in $\hat{\boldsymbol{f}}$ are diminished towards zero, possibly reducing contrast.

Another choice is a quadratic roughness penalty function, which in 1D would be written

$$\mathsf{R}(\boldsymbol{f}) = \sum_{j=2}^{N} (f_j - f_{j-1})^2.$$

This choices biases the reconstruction towards a smooth image where neighboring pixel values are similar. It is convenient for minimization, [9] but it has the drawback of smoothing image edges, particularly if the regularization parameter β in (8) is too large.

More recently, total variation methods have been investigated for MR image reconstruction [10]. In 1D, these methods replace the squared differences between neighboring pixels above with differences:

$$\mathsf{R}(\boldsymbol{f}) = \sum_{j=2}^{N} |f_j - f_{j-1}|.$$

The advantage of this type of regularization is that it biases the reconstructed image towards a piecewise smooth image, instead of a globally smooth image, thereby better preserving image edges. One disadvantage is that it is harder minimize and can lead to the appearance of “blocky” texture in images [11].

3.2 Regularization parameter selection

Another challenge is selection of the regularization parameter β. For quadratic regularization, there is a well developed theory for choosing β in terms of the desired spatial resolution properties of the reconstructed image [12, 13]. This theory extends readily to MR imaging with reasonably well sampled trajectories (or parallel imaging with reasonable acceleration factors) for which the point spread function (PSF) of the reconstructed image is relatively close to a Kronecker impulse so that simple measures like full width at half maximum (FWHM) are reasonable resolution metrics. The extensions to MR have been investigated [14]. For highly under-sampled trajectories the PSF can have “heavy tails” due to aliasing effects, and more investigation is needed to extend the above methods to MR applications.

For nonquadratic regularization such as the total variation method, the analysis in [12,13] is inapplicable so one must resort to other methods for choosing β. Statisticians often use cross validation [15,16] for choosing regularization parameters, with a goal of finding the parameter that minimizes the mean-squared error (MSE) between $\hat{\boldsymbol{f}}$ and the unknown $\boldsymbol{f}$. However, MSE is the sum of variance and bias squared, and where bias is related to spatial resolution and artifacts, and it is unclear whether an equal weighting of noise variance and bias (squared) is optimal from an image quality perspective in medical imaging.

Another method for choosing β is the “L-curve” method [17, 18]. This method is expensive because it requires evaluating $\hat{\boldsymbol{f}}$ for several values of β, and it has some theoretical deficiencies [19].

In summary, choosing β remains a nontrivial issue in most ill-posed imaging problems including MRI.

3.3 Within-voxel gradients

The model (6) treats the field inhomogeneity within each voxel as being a constant, ignoring within-voxel gradients of the off-resonance map. However, these gradients can be significant in functional magnetic resonance imaging (fMRI) based on the BOLD effect [20]. Accurate reconstruction of signals near air-tissue interfaces requires compensation for these within-voxel gradients, which complicates the reconstruction method [21–23].

3.4 Computation time and preconditioning

To attempt to accelerate algorithm convergence for non-Cartesian MR data, some authors have advocated using a weighted norm in the cost function (8) where the weights are based on the density compensation factors that would have been used for gridding type reconstruction methods. This weighting can be thought of as a type of preconditioning of the equation $\mathsf{E}[\boldsymbol{y}] = \boldsymbol{A}\boldsymbol{f}$. If the unregularized conjugate gradient (CG) algorithm is initialized with a zero image, then the first iterate is equivalent to a conjugate phase (CP) reconstruction. So in this case using density compensation weighting will improve that first iterate. But when regularization is used, initializing with an appropriate (density compensated) CP image and then iterating without any further density compensation can be just as effective in terms of convergence rate. More importantly, weighting properly by the statistics (rather than by the sampling pattern) avoids the noise amplification that results from density compensation [24].

3.5 Modeling error

The model (1) assumes the field map $\omega(\vec{r})$ is known. In practice it must be estimated from noisy MR scans, *e.g.*, [25, 26]. Errors on the field map estimates will propagate into errors the reconstructed image, but the effect is not understood thoroughly.

In addition, object motion that occurs between the field map scans and the scan of interest will lead to an inconsistency between the actual scan data and the assumed model (1) used by the reconstruction algorithm. This possibility has motivated the development of dynamic field mapping methods that estimate the field map separately for each frame in a dynamic study, *e.g.*, [27].

3.6 Prior information

More recently, particularly in dynamic imaging, a variety of methods have been proposed that introduce various forms of prior information to augment k-t space data that is inherently undersampled, *e.g.*, [28–31]. Choosing the right balance between the prior information and the measured data remains a very important area for further investigation.

References

[1] P. J. Beatty, D. G. Nishimura, and J. M. Pauly. Rapid gridding reconstruction with a minimal oversampling ratio. *IEEE Trans. Med. Imag.*, 24(6):799–808, June 2005.

[2] M. Bydder, A. A. Samsonov, and J. Du. Evaluation of optimal density weighting for regridding. *Mag. Res. Im.*, 25(5):695–702, June 2007.

[3] J. A. Fessler, S. Lee, V. T. Olafsson, H. R. Shi, and D. C. Noll. Toeplitz-based iterative image reconstruction for MRI with correction for magnetic field inhomogeneity. *IEEE Trans. Sig. Proc.*, 53(9):3393–402, September 2005.

[4] V. Olafsson, S. Lee, J. A. Fessler, and D. C. Noll. Fast Toeplitz based iterative SENSE reconstruction. In *Proc. Intl. Soc. Mag. Res. Med.*, page 2459, 2006.

[5] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: sensitivity encoding for fast MRI. *Mag. Res. Med.*, 42(5):952–62, November 1999.

[6] K. P. Pruessmann, M. Weiger, P. Börnert, and P. Boesiger. Advances in sensitivity encoding with arbitrary k-space trajectories. *Mag. Res. Med.*, 46(4):638–51, October 2001.

[7] H. M. Hudson and R. S. Larkin. Accelerated image reconstruction using ordered subsets of projection data. *IEEE Trans. Med. Imag.*, 13(4):601–9, December 1994.

[8] P. Qu, K. Zhong, B. Zhang, J. Wang, and G. X. Shen. Convergence behavior of iterative SENSE reconstruction with non-Cartesian trajectories. *Mag. Res. Med.*, 54(4):1040–5, October 2005.

[9] B. P. Sutton, D. C. Noll, and J. A. Fessler. Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities. *IEEE Trans. Med. Imag.*, 22(2):178–88, February 2003.

[10] K. T. Block, M. Uecker, and J. Frahm. Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint. *Mag. Res. Med.*, 57(6):1086–98, June 2007.

[11] A. Raj, G. Singh, R. Zabih, B. Kressler, Y. Wang, N. Schuff, and M. Weiner. Bayesian parallel imaging with edge-preserving priors. *Mag. Res. Med.*, 57(1):8–21, January 2007.

[12] J. A. Fessler and W. L. Rogers. Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs. *IEEE Trans. Im. Proc.*, 5(9):1346–58, September 1996.

[13] J. Qi and R. M. Leahy. Resolution and noise properties of MAP reconstruction for fully 3D PET. *IEEE Trans. Med. Imag.*, 19(5):493–506, May 2000.

[14] V. Olafsson, J. A. Fessler, and D. C. Noll. Spatial resolution analysis of iterative image reconstruction with separate regularization of real and imaginary parts. In *Proc. IEEE Intl. Symp. Biomed. Imag.*, pages 5–8, 2006.

[15] M. Stone. Cross-validation: A review. *Math Oper Stat Ser Stat.*, 9(1):127–139, 1978.

[16] G. H. Golub, M. Heath, and G. Wahba. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–23, May 1979.

[17] P. C. Hansen. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Review*, 34(4):561–580, December 1992.

[18] P. C. Hansen and D. P. O'Leary. The use of the L-curve in the regularization of discrete ill-posed problems. *SIAM J. Sci. Comp.*, 14(6):1487–506, 1993.

[19] C. R. Vogel. Non-convergence of the L-curve regularization parameter selection method. *Inverse Prob.*, 12(4):535–47, August 1996.

[20] D. C. Noll, J. A. Fessler, and B. P. Sutton. Conjugate phase MRI reconstruction with spatially variant sample density correction. *IEEE Trans. Med. Imag.*, 24(3):325–36, March 2005.

[21] B. P. Sutton. *Physics-based reconstruction of magnetic resonance images*. PhD thesis, Univ. of Michigan, Ann Arbor, MI, 48109-2122, Ann Arbor, MI, 2003.

[22] B. P. Sutton, D. C. Noll, and J. A. Fessler. Compensating for within-voxel susceptibility gradients in BOLD fMRI. In *Proc. Intl. Soc. Mag. Res. Med.*, page 349, 2004.

[23] J. A. Fessler and D. C. Noll. Model-based MR image reconstruction with compensation for through-plane field inhomogeneity. In *Proc. IEEE Intl. Symp. Biomed. Imag.*, pages 920–3, 2007. Invited paper (oral) in special session on model-based imaging.

[24] R. E. Gabr, P. Aksit, P. A. Bottomley, A-B. M. Youssef, and Y. M. Kadah. Deconvolution-interpolation gridding (DING): Accurate reconstruction for arbitrary k-space trajectories. *Mag. Res. Med.*, 56(6):1182–91, December 2006.

[25] J. A. Fessler, D. Yeo, and D. C. Noll. Regularized fieldmap estimation in MRI. In *Proc. IEEE Intl. Symp. Biomed. Imag.*, pages 706–9, 2006.

[26] A. Funai, J. A. Fessler, D. T. B. Yeo, and D. C. Noll. Regularized field map estimation in MRI. *IEEE Trans. Med. Imag.*, 2007. Submitted as TMI-2007-0218.

[27] B. P. Sutton, D. C. Noll, and J. A. Fessler. Dynamic field map estimation using a spiral-in / spiral-out acquisition. *Mag. Res. Med.*, 51(6):1194–204, June 2004.

[28] J. Tsao, P. Boesiger, and K. P. Pruessmann. k-t BLAST and k-t SENSE: Dynamic MRI with high frame rate exploiting spatiotemporal correlations. *Mag. Res. Med.*, 50(5):1031–42, November 2003.

[29] H. Jung, J. C. Ye, and E. Y. Kim. Improved k-t BLAST and k-t SENSE using FOCUSS. *Phys. Med. Biol.*, 52(11):3201–26, June 2007.

[30] D. Xu, K. F. King, and Z-P. Liang. Improving k-t SENSE by adaptive regularization. *Mag. Res. Med.*, 57(5):918–30, May 2007.

[31] B. Sharif and Y. Bresler. Affine-corrected PARADISE: Free-breathing patient-adaptive cardiac MRI with sensitivity encoding. In *Proc. IEEE Intl. Symp. Biomed. Imag.*, pages 1076–9, 2007.