


Model-Based Image Reconstruction for MRI

[A review of the use of iterative algorithms]



Magnetic resonance imaging (MRI) is a sophisticated and versatile medical imaging modality. Traditionally, MR images are reconstructed from the raw measurements by a simple inverse two-dimensional (2-D) or three-dimensional (3-D) fast Fourier transform (FFT). However, there are a growing number of MRI applications where a simple inverse FFT is inadequate, e.g., due to non-Cartesian sampling patterns, non-Fourier physical effects, nonlinear magnetic fields, or deliberate under-sampling to reduce scan times. Such considerations have led to increasing interest in methods for model-based image reconstruction in MRI.

INTRODUCTION

The inverse 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 with appropriate density compensation factors [1] 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 article reviews the use of iterative

algorithms for model-based MR image reconstruction. The references give pointers to some recent work but are by no means a comprehensive survey. To see more citations, visit <http://www.eecs.umich.edu/~fessler/>.

MRI BACKGROUND

Any signal processing method aimed at forming images from measurement devices such as MRI scanners must consider the relevant physics. A survey in *IEEE Signal Processing Magazine* [2] and a book written from a signal processing perspective [3] have described MRI physics well. Here we review the physics in a somewhat unconventional way that facilitates describing some of the “non-Fourier” aspects of MRI.

MRI PHYSICS

Standard MRI scanners use a large static magnetic field

$$\vec{B}_0(\vec{r}) = B_0(\vec{r})\vec{k} \quad (1)$$

to induce a net magnetization $\vec{M} = M_x\vec{i} + M_y\vec{j} + M_z\vec{k}$ at each point in space in the body being imaged, where \vec{i} , \vec{j} , and \vec{k} denote the unit vectors along the x , y , and z axes, respectively, and $\vec{r} = (x, y, z)$ denotes 3-D spatial coordinates. Ideally, the static field strength $B_0(\vec{r})$ would be spatially uniform, i.e., a single constant B_0 . In practice, it is never perfectly uniform, due to the unavoidable nonuniformities of all practical coil designs and due to the field strength variations that are induced by the nonuniform magnetic susceptibilities of different tissue types. The electron distributions in different molecules also influence the local magnetic environment experienced by an atom's nucleus, called chemical shift. Some types of MRI scans are robust to such spatial variations of B_0 ; others are sensitive to nonuniformities, necessitating correction methods.

At equilibrium (which is established within a few seconds for a stationary object), the magnetization \vec{M} is aligned with the applied static field and its magnitude is proportional to the product of $B_0(\vec{r})$ and the object-dependent local density of (predominately) hydrogen protons or "spins." This proton density alone is of only modest interest in MRI; in practice one applies time-varying magnetic fields $\vec{B}(\vec{r}, t)$ that induce time-varying changes in the magnetization

$$\vec{M}(\vec{r}, t) = M_x(\vec{r}, t)\vec{i} + M_y(\vec{r}, t)\vec{j} + M_z(\vec{r}, t)\vec{k}. \quad (2)$$

These changes depend on time constants (tissue-dependent relaxation parameters) and other factors, and the goal in MRI is to form images of aspects of this magnetization. By manipulating the applied field $\vec{B}_0(\vec{r}, t)$ appropriately, sometimes in conjunction with injected or inhaled contrast agents, one can examine a multitude of different tissue properties.

An MRI scan consists of one or more alternations between two stages: excitation and readout. During the excitation stage, the applied magnetic field $\vec{B}(\vec{r}, t)$ is designed to tip the magnetization vectors \vec{M} within some slice or slab away from equilibrium, so that they have a component in the transverse plane, i.e., the (x, y) plane. It is convenient to represent this transverse component mathematically using a complex function defined as follows:

$$M(\vec{r}, t) \triangleq M_x(\vec{r}, t) + iM_y(\vec{r}, t), \quad (3)$$

where $i \triangleq \sqrt{-1}$. Note that the field components M_x and M_y are real physical quantities; the "transverse magnetization" $M(\vec{r}, t)$ is complex solely by definition. The excitation process can be quite complicated to model and is beyond the scope of this article. See [2] for an introduction to the role that signal processing plays in the design of excitation pulses and [4] for some recent model-based RF pulse design methods.

During the readout stage, the applied field $\vec{B}(\vec{r}, t)$ is manipulated in ways that help elucidate the transverse magnetization $M(\vec{r}, t)$. For image reconstruction, it is essential to model the effects of the applied field on the

transverse magnetization. The precise relationship is governed by the Bloch equation [2]. For most image reconstruction purposes, it suffices to consider just two aspects of the full relationship: precession and transverse relaxation. The most important equation in MRI is the Larmor relation: $\omega = \gamma|\vec{B}|$, which states that the magnetization precesses (around the axis of the applied field) at a frequency ω that is proportional to the magnitude of the applied field. The constant of proportionality γ is called the gyromagnetic ratio and is about 42.6 MHz/T for hydrogen protons. During a readout, only the longitudinal component of \vec{B} is varied usually, i.e.,

$$\vec{B}(\vec{r}, t) = B_z(\vec{r}, t)\vec{k}, \quad (4)$$

so the magnetization precesses around \vec{k} , i.e., within the transverse plane. This property is why the complex representation (3) is convenient, because precession can be expressed using a complex phase in this form. In general, the applied longitudinal field strength $B_z(\vec{r}, t)$ varies both spatially and temporally, so the Larmor relationship describes the instantaneous frequency at a given spatial location

$$\omega(\vec{r}, t) = \gamma B_z(\vec{r}, t). \quad (5)$$

Without loss of generality, let $t = 0$ be the time when the excitation pulse is completed, and consider some time point $t > 0$ during the readout. The precession of the transverse magnetization between time zero and time t corresponds to a net phase that is the integral of the instantaneous frequency (5), i.e., ideally we would have

$$M(\vec{r}, t) = M(\vec{r}, 0)\exp\left(-i\int_0^t \omega(\vec{r}, t')dt'\right).$$

In practice, microscopic variations in the magnetic field cause the spins within a given voxel to become out of phase over time. So the transverse magnetization vector's magnitude decreases approximately exponentially with a time constant T_2^* . Accounting for this decay, an accurate model for the temporal evolution of the transverse magnetization during a readout is

$$M(\vec{r}, t) = f(\vec{r})e^{-t/T_2^*(\vec{r})}\exp\left(-i\gamma\int_0^t B_z(\vec{r}, t')dt'\right), \quad (6)$$

where $f(\vec{r}) \triangleq M(\vec{r}, 0)$ denotes the object's transverse magnetization immediately after excitation. A typical goal in MRI is to form an image of $f(\vec{r})$. The properties of $f(\vec{r})$ depend not only on spin density, but also on the type of excitation used. Note that for simplicity of exposition, we focus here on the case where the object is static so that $f(\vec{r})$ is not a function of time t . Generalizations to dynamic imaging are very active research areas in MR image reconstruction.

The relaxation factor T_2^* varies spatially, and often is on the order of 10 ms. This relatively rapid decay is a significant limitation in MRI. If T_2^* were longer, then a signal excitation stage

NOTIONS OF SPARSITY HAVE DEEP ROOTS IN STATISTICAL SIGNAL PROCESSING.

followed by a (lengthy) readout stage could be sufficient to form a high-resolution image of $f(\vec{r})$. In practice, the rapid decay limits how much spatial information can be recorded in a single readout stage, so such “single shot” imaging, such as echo-planar imaging (EPI) [5], provides only modest spatial resolution. Therefore, high-resolution imaging uses multiple alternations between excitation stages and readout stages, each with different variations of the applied field $B_z(\vec{r}, t)$.

DATA ACQUISITION: THE MR SIGNAL

By Faraday’s law, the time-varying magnetization $M(\vec{r}, t)$ will induce an electromotive force (emf) in a nearby coil. The emf will be proportional to the volume integral of the time derivative of the magnetization $M(\vec{r}, t)$ multiplied by the coil response pattern $c(\vec{r})$. The resulting electrical potential $v(t)$ across the receive coil is

$$v(t) = \text{real}\left(\int c(\vec{r}) \frac{d}{dt} M(\vec{r}, t) d\vec{r}\right), \quad (7)$$

where $\text{real}(\cdot)$ denotes the real part of a complex number. The coil response $c(\vec{r})$ generally decreases with distance from the coil. If uncorrected, this nonuniformity causes spatial variations in signal strength that can be a challenge for image processing methods like segmentation algorithms. Numerous correction methods have been developed.

Because the time constant T_2^* is on the order of milliseconds whereas the phase variations in (6) are many MHz, it is very reasonable to use a narrow-band approximation when evaluating the time derivative of $M(\vec{r}, t)$ as needed in (7). The time derivative of a narrow-band signal is well approximated by a constant scaling factor $d/dt M(\vec{r}, t) \approx c_0 M(\vec{r}, t)$. We absorb this constant into the coil response pattern and rewrite (7) as

$$v(t) = \text{real}\left(\int c(\vec{r}) M(\vec{r}, t) d\vec{r}\right). \quad (8)$$

The receive coil’s signal is amplified and demodulated using some center frequency ω_0 . Ideally, one would use $\omega_0 = \gamma B_0$ if the static magnetic field had uniform strength B_0 . Usually quadrature demodulation is used, yielding separate in-phase $I(t)$ and quadrature $Q(t)$ baseband signals. In the literature, the demodulated “MR signal” $s(t)$ is defined (implicitly) as

$$s(t) \triangleq I(t) + iQ(t) = \text{lowpass}(e^{i\omega_0 t} v(t)) = e^{i\omega_0 t} \int c(\vec{r}) M(\vec{r}, t) d\vec{r}, \quad (9)$$

where the low-pass operation selects the baseband component of the demodulated signal. This complex analog signal is just a mathematical definition; in practice, the $I(t)$ and $Q(t)$ signals are each sampled and digitized yielding two digital signals. (One can use two separate analog-to-digital (A/D) converters, or a single A/D converter running at twice the normal rate to avoid I/Q

imbalance.) Digitally, these two signals can be combined and stored as complex values, i.e., we record samples

$$I(m\Delta_T) + iQ(m\Delta_T), \quad m = 1, \dots, n_d,$$

where Δ_T denotes the sampling rate (typically around $1 \mu\text{s}$) and n_d denotes the number of recorded samples, typically 64–512 for a given readout stage. Again, the physical quantities are real, but complex quantities are defined in terms of those physical quantities for convenience. (In some systems, digital demodulation is used, but the modeling remains identical.)

SIGNAL MODEL

To improve signal-to-noise ratio and reduce acquisition times, the use of multiple receive coils has become increasingly popular in MRI. Although originally called phased array imaging [6], a term that resonates with other signal processing applications involving multiple receivers, today the use of multiple receive coils in MRI is usually called parallel imaging [7].

Let $c_l(\vec{r})$ denote the sensitivity (response pattern) of the l th coil, for $l = 1, \dots, L$, where L denotes the number of coils. Let $s_l(t)$ denote the demodulated “MR signal” associated with the l th coil, defined as in (9). Substituting (6) into (9) and simplifying yields the following general forward model for the MR signal associated with the l th coil

$$s_l(t) = \int c_l(\vec{r}) f(\vec{r}) e^{-i\omega_0 t} e^{-i\phi(\vec{r}, t)} d\vec{r}, \quad (10)$$

where the space- and time-varying phase is

$$\phi(\vec{r}, t) \triangleq \int_0^t (\gamma B_z(\vec{r}, t') - \omega_0) dt'. \quad (11)$$

In practice, multiple such signals are recorded, one for each excitation/readout pair (“shots”). For simplicity of notation, we consider “single shot” imaging; the extension to multiple shots is conceptually straightforward but notationally cumbersome. Note that the phase variations (11) are common to all receive coils; only the coil response patterns $\{c_l(\vec{r})\}$ differ between coils.

MEASUREMENT MODEL

The recorded measurements in a MR scan consist of noisy samples of the MR signal (10)

$$y_{li} = s_l(t_i) + \varepsilon_{li}, \quad i = 1, \dots, n_d, \quad l = 1, \dots, L, \quad (12)$$

where y_{li} denotes the i th sample of the l th coil’s signal at time t_i and n_d denotes the number of time samples. Usually the t_i values are equally spaced, and often there are one or more time values where the signal is particularly strong due to alignment of the magnetization’s phases; these values are called echo times. The measurement errors ε_{li} are very well modeled by

additive, complex, zero-mean, temporally white Gaussian noise [8]. However, there can be coupling of the noise values between different coils for the same time points, i.e.,

$$\text{Cov}\{\varepsilon_{li}, \varepsilon_{kj}\} = \sum_{lk} \delta[i - j], \quad (13)$$

where δ denotes the Kronecker impulse, and the $L \times L$ matrix Σ characterizes the noise covariance between coils [7].

LINEAR RECONSTRUCTION PROBLEM

Using the measurement model (12) and the signal model (10), the “typical” image reconstruction problem in MRI is to estimate the object $f(\vec{r})$ from the measurement vector $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_L)$, where $\mathbf{y}_l = (y_{l1}, \dots, y_{ln_l})$. (All vectors are column vectors here.) We first consider model-based image reconstruction for this “basic” linear formulation. Because parallel imaging is of considerable interest, we continue to consider the general case of L receive coils. A standard single receive coil is a simple special case.

This is an ill-posed problem because the given measurements \mathbf{y} are discrete whereas the object $f(\vec{r})$ is an unknown continuous-space function. To facilitate parametric estimation, we approximate the object $f(\vec{r})$ using a “finite series expansion” as follows:

$$f(\vec{r}) = \sum_{j=1}^N f_j b(\vec{r} - \vec{r}_j), \quad (14)$$

where $b(\cdot)$ denotes the object basis function, \vec{r}_j denotes the center of the j th translated basis function, and N is the number of parameters. Such approximations are classic in the tomographic image reconstruction literature [9] and are slowly taking root in the MR community. Minimum \mathcal{L}_2 norm methods can postpone the discretization (14) until the final step of displaying the image, but it is unclear if this approach provides image quality benefits that outweigh its computational requirements. For simplicity, hereafter we use rect basis functions $b(\vec{r}) = \text{rect}(\vec{r}/\Delta)$, i.e., square pixels of dimension Δ , so N is the number of pixels, or voxels in 3-D scans. Many other possible basis function choices can be considered, all of which are imperfect because the true object never satisfies the parametric model (14) exactly. Nevertheless simple basis functions can provide useful approximations.

Substituting the basis expansion (14) into the signal model (10) and simplifying leads to the discrete forward model

$$s_l(t_i) = \sum_{j=1}^N a_{lij} f_j, \quad (15)$$

where the elements $\{a_{lij}\}$ of the system matrix A_l associated with the l th coil are given by

$$a_{lij} = \int b(\vec{r} - \vec{r}_j) c_l(\vec{r}) e^{-t_i/T_2^*(\vec{r})} e^{-i\phi(\vec{r}, t_i)} d\vec{r}. \quad (16)$$

MODEL-BASED METHODS THAT ACCOUNT FOR THOSE [PHYSICAL] EFFECTS ARE PROVING TO BE BENEFICIAL FOR IMPROVING IMAGE QUALITY.

In practice the basis functions are usually highly localized (e.g., voxels), so “center of voxel” approximations like the following are nearly always used, often implicitly

$$a_{lij} \approx c_l(\vec{r}_j) e^{-t_i/T_2^*(\vec{r}_j)} e^{-i\phi(\vec{r}_j, t_i)}. \quad (17)$$

For exceptions, see [10].

Typically the decay due to T_2^* is ignored, or it is assumed implicitly that the total readout time $t_{n_d} - t_1$ is small relative to T_2^* in which case one can make the approximation $e^{-t_i/T_2^*(\vec{r})} \approx e^{-t_i/T_2^*(\vec{r}_j)}$. Under this approximation, we can absorb the T_2^* -weighting effect of $e^{-t_i/T_2^*(\vec{r})}$ into the unknown image $f(\vec{r})$.

Combining (12) and (15) in matrix-vector form yields

$$\mathbf{y}_l = A_l \mathbf{f} + \varepsilon_l,$$

where $\mathbf{f} = (f_1, \dots, f_N)$ is the vector of parameters (pixel values) that we wish to estimate from the data \mathbf{y} . Stacking up all L measurement vectors as $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_L)$ and defining the $(n_d L) \times N$ system matrix $A = (A_1, \dots, A_L)$ yields the linear model

$$\mathbf{y} = A \mathbf{f} + \varepsilon. \quad (18)$$

At first glance this linear model appears amenable to a variety of iterative solution methods. However, a significant challenge that arises is that in general the elements of A can be quite complicated in the form above, yet A is too large to store for typical problem sizes. Most iterative algorithms require matrix-vector multiplication by A and its transpose; there are fast algorithms for these operations (without storing A explicitly) in many special cases of interest [10], [11].

Thus far we have allowed the phase function $\phi(\vec{r}_j, t_i)$ to be quite general, without the traditional focus on “Fourier encoding.” Recently there has been interest in investigating nonlinear magnetic field variations $B_z(\vec{r}, t)$ in (4), and reconstruction algorithms have been proposed that use much of the generality in (16) [12], [13]. These are currently specialized research topics, so we now focus on the more common case of linear field gradients.

FOURIER ENCODING

In typical MR scanners, the longitudinal component of the applied field $B_z(\vec{r}, t)$ in (4) consists of three components

$$B_z(\vec{r}, t) = B_0 + \Delta B_0(\vec{r}) + \vec{G}(t) \cdot \vec{r}. \quad (19)$$

The constant B_0 denotes the advertised field strength of the main static field. The function $\Delta B_0(\vec{r})$ denotes the spatial deviations of the field strength from this nominal value. This function is often called a field map, and in general, it is unknown, but it can be estimated by suitable types of

acquisitions and data processing methods [14]. The field gradients $\vec{G}(t) = \vec{G}_x(t)\vec{j} + \vec{G}_y(t)\vec{j} + \vec{G}_z(t)\vec{j}$ consist of three user-controlled functions that are the historical key to providing spatial information in standard MR imaging. Many different types of MR scans are possible by changing $\vec{G}(t)$.

Substituting (19) into (11) using $\omega_0 \triangleq \gamma B_0$ and simplifying yields

$$\phi(\vec{r}, t) = \int_0^t \gamma \Delta B_0(\vec{r}) + \gamma \vec{G}(t) \cdot \vec{r} dt$$

or equivalently

$$e^{-i\phi(\vec{r}, t)} = e^{-i\Delta\omega_0(\vec{r})t} e^{-i2\pi\vec{k}(t) \cdot \vec{r}}, \quad (20)$$

where $\Delta\omega_0(\vec{r}) \triangleq \gamma \Delta B_0(\vec{r})$ denotes the off-resonance frequency and the k-space trajectory is defined by

$$\vec{k}(t) \triangleq \frac{1}{2\pi} \int_0^t \gamma \vec{G}(t) dt. \quad (21)$$

Usually the phase accrual $e^{-i\Delta\omega_0(\vec{r})t}$ due to off resonance is undesirable and can distort reconstructed images if ignored. Therefore some image reconstruction methods, particularly in fMRI, account for its effects [10]. In some cases, the map $\Delta\omega_0(\vec{r})$ is found from a separate ‘‘prescan,’’ in other cases it is estimated jointly with f [15]. In chemical shift imaging, e.g., to separate fat and water components, the term $\Delta\omega_0(\vec{r})$ includes both useful information about the chemical shift effect as well as the undesirable variations due to field inhomogeneity [16].

For the linear field gradients (19), substituting (20) into (17) yields simpler expressions for the system matrix

$$a_{ij} \approx c_i(\vec{r}_j) e^{-z(\vec{r}_j)t_i} e^{-i2\pi\vec{k}(t_i) \cdot \vec{r}_j}, \quad (22)$$

where we define the ‘‘rate map’’ $z(\vec{r})$ by combining the relaxation and field maps

$$z(\vec{r}) \triangleq 1/T_2^*(\vec{r}) + i\Delta\omega_0(\vec{r}). \quad (23)$$

When this rate map is assumed to be zero, i.e., if relaxation and off resonance are ignored, then a_{ij} is the product of a Fourier encoding matrix having elements $e^{-i2\pi\vec{k}(t_i) \cdot \vec{r}_j}$ with a diagonal sensitivity encoding matrix having elements $c_i(\vec{r}_j)$.

If the k-space sample locations $\vec{k}(t_i)$ lie on an appropriate subset of a Cartesian grid, then FFT operations provide efficient multiplication by A and its transpose. If non-Cartesian k-space sampling is used, then a nonuniform FFT (NUFFT) is needed [17].

When $z(\vec{r})$ in (22) is nonzero, then the elements (22) no longer correspond to a standard Fourier transform. Approximations are needed to provide fast computation of matrix-vector products. In particular, often one can approximate the exponentials in (17) using an additively separable form

$$e^{-z(\vec{r}_j)t_i} \approx \sum_k b_{ik} c_{kj}$$

for various choices for the basis functions b_{ik} and coefficients c_{kj} [11]. With this type of approximation, we can rewrite matrix-vector multiplication as follows:

$$[A_i f]_i \approx \sum_k b_{ik} \sum_{j=1}^N (c_{kj} c_i(\vec{r}_j) f_j) e^{-i2\pi\vec{k}(t_i) \cdot \vec{r}_j}.$$

The inner sum is simply a FFT or NUFFT so this approach is relatively fast. Free software for this is available [18].

RECONSTRUCTION COST FUNCTION

Having specified the linear model (18), we now turn to solution methods. Because the noise in MRI measurements is Gaussian, a natural approach is to estimate f by minimizing a regularized least-squares cost function

$$\hat{f} = \arg \min_f \Psi(f), \quad \Psi(f) \triangleq \|y - Af\|^2 + \beta R(f). \quad (24)$$

For a single coil, the noise variance in the k-space data is white (uncorrelated with uniform variance), so the usual Euclidian norm $\|\cdot\|$ is appropriate. For parallel MRI, noise is stationary across time samples (i), but the norm should include the inverse of the $L \times L$ covariance matrix Σ in (13) that describes the noise correlation between receive coils [7].

If the k-space samples lie on an equally spaced grid (Cartesian sampling) with appropriate sample spacings relative to the object field of view, and if the rate map $z(\vec{r})$ is zero (i.e., we ignore relaxation and field inhomogeneity), and if we consider just a single coil ($L = 1$) and treat the sensitivity pattern as uniform, i.e., $c_1(\vec{r}) = 1$, then the system matrix A_i is orthogonal. In this special case, no regularization is needed and $A^{-1} = 1/NA'$ and the solution is simply $\hat{f} = 1/NA'y$, which can be evaluated by an inverse FFT. This is the most common MR image-reconstruction method. However, if any of these conditions do not hold, then typically the system matrix A is not well conditioned, and the unregularized LS solution can lead to undesirable noise amplification. To avoid this problem, some form of regularization is needed.

REGULARIZATION

An open problem in most image reconstruction problems, including MRI, is how to best choose the regularizer $R(f)$. If this term is not included, then the image estimate \hat{f} will 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 positron emission tomography scanners is to use an unregularized algorithm, initialize it with a uniform image, stop iterating just as the image gets unacceptably noisy, and then perhaps apply a bit of post-filtering to reduce the noise. One could adopt a similar approach for MR imaging. 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 $R(f) = \|f\|^2$ or $R(f) = \|f - \bar{f}\|^2$, where \bar{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{f} . In particular, if the reference image is zero, then all pixel values in \hat{f} are diminished towards zero, possibly reducing contrast.

Another choice is a quadratic roughness penalty function, which in one-dimensional (1-D) would be written

$$R(f) = \sum_{j=2}^N |f_j - f_{j-1}|^2. \quad (25)$$

This choice biases the reconstruction towards a smooth image where neighboring pixel values are similar. It is convenient for minimization [10], but it has the drawback of smoothing image edges, particularly if the regularization parameter β in (24) is too large. One can prove that using (25) guarantees that the cost function (24) has a unique minimizer.

More recently, total variation methods have been investigated for MR image reconstruction [19]. In 1-D, these methods replace the squared differences between neighboring pixels above with absolute differences

$$R(f) = \sum_{j=2}^N |f_j - f_{j-1}|. \quad (26)$$

In 2-D continuous space, the analogous functional is

$$\int \|\nabla f\| d\vec{r} = \iint \sqrt{\left| \frac{\partial f(\vec{r})}{\partial x} \right|^2 + \left| \frac{\partial f(\vec{r})}{\partial y} \right|^2} dx dy.$$

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. However it is harder to minimize and can lead to the appearance of “blocky” texture in images. Numerous alternatives of the form

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

for various choices of the “potential function” $\psi(\cdot)$ have been proposed in the imaging literature. Many of these compromise between the quadratic case (25) and the absolute difference case (26), for example the hyperbola

$$\psi(t) = \sqrt{1 + |t/\delta|^2} - 1 \quad (27)$$

is approximately quadratic near zero, which aids noise reduction, yet approximately linear away from zero, which helps preserve edges.

ALGORITHMS

Iterative algorithms are needed to minimize (24). For differentiable regularizers such as (25), the conjugate gradient algorithm is a natural choice [10]. For nondifferentiable regularizers like (26), more sophisticated algorithms are needed and this is an active research area [20].

RECONSTRUCTION CHALLENGES

Although a variety of useful problems can be solved in MRI using the formulation (24), there are numerous challenges that provide research opportunities.

REGULARIZATION PARAMETER SELECTION

A practical challenge with regularized methods is selection of the regularization parameter β in (24). For quadratic regularization, there is a well-developed theory for choosing β in terms of the desired spatial resolution properties of the reconstructed image [21]. This theory extends readily to MR imaging with reasonably well sampled

trajectories (and to 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. 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 (26), the analysis in [21] is inapplicable so one must resort to other methods for choosing β . Statisticians often use cross validation for choosing regularization parameters, with a goal of finding the parameter that minimizes the mean-squared error (MSE) between \hat{f} and the unknown 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. This method is expensive because it requires evaluating \hat{f} for several values of β , and it has some theoretical deficiencies [22].

In summary, choosing β for nonquadratic regularization remains a nontrivial issue in most ill-posed imaging problems including MRI, and remains an active research area [23].

PARTIAL K-SPACE METHODS

If the object $f(\vec{r})$ were real, then its Fourier transform would be Hermitian symmetric so in principle only half of k-space would need to be sampled. In practice, the magnetization (3) is complex due to a variety of physical effects. However, in many cases the phase of $M(\vec{r}, t)$ can be assumed to be a smooth function. This property has led to a variety of partial k-space methods where one samples a bit more than half of k-space, then

RECENTLY IT HAS BECOME VERY POPULAR TO EXPRESS PRIOR INFORMATION IN TERMS OF SOME TYPE OF SPARSITY OF THE OBJECT.

estimates the phase from the central portion of k-space (corresponding to low spatial frequencies), and then uses this estimated phase to reconstruct the entire image [24]. Such methods are used routinely in many types of MR scans.

UNDER-SAMPLED K-SPACE DATA

The need for some type of regularization is essential when the k-space data is under sampled, i.e., when the number of measurements Ln_d is less than the number of unknown voxels N . In MRI, the scan time is roughly proportional to the number of measurements, so collecting fewer samples can reduce scan time, which is particularly desirable in dynamic imaging.

In the broader field of tomographic image reconstruction, there is a long history of using prior information, such as assuming objects are piecewise smooth, to reconstruct images from an under-sampled set of projection views, e.g., [25]. Many of these methods involve cost functions of the form (24) with a suitable system matrix A for the application and appropriate regularizers $R(f)$ that capture prior information about the object.

Recently it has become very popular to express prior information in terms of some type of sparsity of the object. Notions of sparsity have deep roots in statistical signal processing [26]. Sparsity is especially apparent in MR angiography. The moniker of compressed sensing or compressive sampling has become widespread for such techniques, and recently entire sessions at MR conferences have been devoted to this topic [20]. Some compressed sensing formulations ignore the noise in the data. In the presence of noise, a typical formulation is

$$\arg \min_f \|\Psi f\|_1 \text{ s.t. } \|y - Af\|_2 \leq \epsilon,$$

where Ψ transforms the image f into a domain (such as wavelet coefficients) where one postulates that the signal is sparse.

Often this optimization problem is solved using a Lagrange multiplier approach

$$\arg \min_f \|y - Af\|_2^2 + \beta \|\Psi f\|_1,$$

which corresponds to a particular regularizer in (24). Rarely is the ℓ_1 norm implemented exactly; in practice usually a continuously differentiable approximation is used, such as

$$\|v\|_1 \approx \sum_i (\sqrt{|v_i|^2 + \delta^2} - \delta) \quad (28)$$

for some small value of $\delta > 0$. This approximation is equivalent to the hyperbola (27) used frequently for edge-preserving image reconstruction. Nonconvex methods that enforce sparsity even more strongly are also under investigation. In the usual case where A corresponds to an under-sampled discrete Fourier transform (DFT), a variety of algorithms are available that have numerous potential applications in MR [20]. Challenges with

this approach include choosing the sparsifying transform Ψ and regularization parameters β and δ appropriately. Furthermore, when δ is small, the regularizer (28) has very high curvature near zero, which can slow convergence.

NONLINEAR RECONSTRUCTION PROBLEMS

The linear image reconstruction problem (24) is just one of many estimation problems of interest in MRI. Returning to the elements of the system matrix (22), there has been research on estimating essentially every component therein, as summarized below.

FIELD MAP ESTIMATION

For scans with long readout times, the effect of field inhomogeneity $\Delta\omega_0$ in (22) is important. In practice, the field map $\omega(\vec{r})$ is not known a priori but rather it must be estimated from noisy MR scans. One can examine the phase differences between two scans having different echo times to determine $\Delta\omega_0$. If these two scans have

short readouts, then there are simple image-domain methods for estimating $\Delta\omega_0$, which is known as B_0 field mapping [14]. Errors in the field map estimates may cause artifacts in reconstructed images that are based on models like (22).

In addition, object motion that occurs between the field map scans and subsequent scans of interest, e.g., in fMRI, will lead to an inconsistency between the actual scan data and the assumed model (22) 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., [15]. For scans with long readout durations, the appearance of $\Delta\omega_0$ in a complex exponential in (22) makes this a somewhat complicated nonlinear estimation problem.

RELAXATION MAP ESTIMATION

In some MR applications, it is useful to estimate tissue relaxation parameters, particularly T_2 or T_2^* , on a pixel-by-pixel basis. One approach to measuring such relaxation parameters is to acquire a “baseline” scan of the object and then acquire one or more additional scans having different echo times. One then reconstructs images from each of those scans and then performs linear regression on a voxel-by-voxel basis using the logarithm of the image voxel values. This approach can be adequate if the readout durations are sufficiently small. But for acquisitions with long readouts, the effect of time t_i in the $e^{-z(\vec{r})/T_i}$ in (22) should be considered, i.e., we should account for relaxation during the signal readout. This requires methods that estimate the relaxation map directly from the k-space data. These are more challenging nonlinear estimation problems because T_2^* appears in an exponent in (22). Several methods for jointly estimating T_2^* , $\Delta\omega_0$, and $f(\vec{r})$ have been investigated [27].

**THE LINEAR IMAGE RECONSTRUCTION
PROBLEM IS JUST ONE OF MANY
ESTIMATION PROBLEMS OF
INTEREST IN MRI.**

SENSITIVITY MAP ESTIMATION

The coil sensitivity patterns $c_l(\vec{r})$ in (22) also must be determined for parallel imaging based on sensitivity encoding. Normally this is done by acquiring well-sampled data both with local receive coils and with a reference body coil and dividing the two [7]. Acquiring the extra reference data can be inconvenient, so normalizing by the square root of the sum of squares of the local receive coils is also used. A variety of other estimation methods have been proposed, including methods that jointly estimate the sensitivity maps $\{c_l(\vec{r})\}$ and the image $f(\vec{r})$ [28]. Note that if $f(\vec{r})$ were known, then the problem of estimating $c_l(\vec{r})$ would be a linear estimation problem because $c_l(\vec{r})$ appears as a linear scaling in (22). But when both $f(\vec{r})$ and $c_l(\vec{r})$ are to be estimated, the model is bilinear because $f(\vec{r})$ and $c_l(\vec{r})$ appear as a product in (10). This complicates joint estimation.

DESPITE OVER THREE DECADES OF MR RESEARCH, THERE REMAIN CHALLENGING AND INTRIGUING PROBLEMS IN MR IMAGE RECONSTRUCTION.

TRAJECTORY MAPPING

The k-space trajectory $\vec{k}(t_i)$, defined as an integral of the gradient waveforms in (21), should be calibrated carefully to ensure that the system model (22) is accurate. In practice, the field gradients induced by the gradient coils in the scanner are not exactly proportional to the waveforms applied to those coils due to eddy currents. Therefore the physical k-space trajectory realized in the system can depart somewhat from the desired k-space trajectory. These differences can degrade the reconstructed image, particularly for non-Cartesian trajectories with long readout durations. Therefore, a variety of techniques have been developed for mapping the actual k-space trajectory experimentally.

WITHIN-VOXEL GRADIENTS

The model (23) 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 [29]. Accurate reconstruction of signals near air-tissue interfaces requires compensation for these within-voxel gradients, which complicates the reconstruction method [30].


EXAMPLE

To illustrate the capabilities of model-based image reconstruction methods for MRI, we simulated k-space data for a four-shot EPI sequence with matrix size 128×128 and $5 \mu\text{s}$ sampling so the readout duration was $27.3 \mu\text{s}$ per shot. The field map $\Delta B_0(\vec{r})$ appears in Figure 2 of [14] and is based on a brain slice above the sinuses and ear canals where susceptibility effects occur. Figure 1 shows the true image used in the simulations and images from three different

reconstruction methods. The “uncorrected” reconstruction simply uses an inverse 2-D FFT, with no consideration of field inhomogeneity. The field inhomogeneity causes spatial distortion in the read-out (vertical) direction (that increases NRMSE dramatically), as well as significant intensity artifacts above the ears and sinuses where the susceptibility effects are largest. The classical conjugate phase reconstruction method, which corresponds to $A'y$ in this single-coil case, reduces the spatial distortion but the intensity artifacts persist. Applying 15 iterations of a conjugate gradient algorithm with a monotonic line search [11] to the cost function (24) with the edge-preserving hyperbola (27) yields the right-most image in Figure 1. This model-based image reconstruction method yields the lowest RMS error, but it requires about 30 times more computation than the noniterative conjugate phase method [11] because each iteration requires multiplication by A and A' . The software that generated this figure is available online [18].

SUMMARY

Image reconstruction is not a single problem in MRI but rather is a wide family of problems depending on what physical



[FIG1] Comparison of model-based image reconstruction with convention methods.

effects are included in the signal model. The most widely studied case, particularly in the signal processing community, is when nearly all physical effects are disregarded and the system model consists solely of sampled of the Fourier transform of the object. This basic model is amenable to familiar signal processing tools and is applicable to many MR scans. But there are also many interesting applications where other physical effects are relevant, and model-based methods that account for those effects are proving to be beneficial for improving image quality. Model-based methods themselves depend on estimates of a variety of model parameters, leading to interesting problems where those parameters are determined either by separate calibration scans or by jointly estimating the image and those parameters. Despite over three decades of MR research, there remain challenging and intriguing problems in MR image reconstruction.

ACKNOWLEDGMENT

This work was supported in part by NIH grants CA87634 and NS058576.

AUTHOR

Jeffrey A. Fessler (fessler@umich.edu) received the B.S.E.E. degree from Purdue University in 1985, the M.S.E.E. degree from Stanford University in 1986, and the M.S. degree in statistics from Stanford University in 1989. From 1985 to 1988, he was a National Science Foundation Graduate Fellow at Stanford, where he earned a Ph.D. degree in electrical engineering in 1990. He has worked at the University of Michigan since then where he is now a professor in the Departments of Electrical Engineering and Computer Science, Radiology, and Biomedical Engineering. He is a Fellow of the IEEE. He received the Francois Erbsmann Award. He is an associate editor for *IEEE Transactions on Medical Imaging* and was an associate editor for *IEEE Transactions on Image Processing* and *IEEE Signal Processing Letters*. He was cochair of the 1997 SPIE Conference on Image Reconstruction and Restoration, technical program cochair of the 2002 IEEE International Symposium on Biomedical Imaging (ISBI), and general chair of ISBI 2007. His research interests are in statistical aspects of imaging problems.

REFERENCES

- [1] M. Bydder, A. A. Samsonov, and J. Du, "Evaluation of optimal density weighting for regridding," *Magn. Reson. Imag.*, vol. 25, no. 5, pp. 695–702, June 2007.
- [2] G. A. Wright, "Magnetic resonance imaging," *IEEE Signal Processing Mag.*, vol. 14, no. 1, pp. 56–66, Jan. 1997.
- [3] Z.-P. Liang and P. C. Lauterber, *Principles of Magnetic Resonance Imaging*. New York: IEEE Press, 2000.
- [4] W. A. Grissom, D. Xu, A. B. Kerr, J. A. Fessler, and D. C. Noll, "Fast large-tip-angle multidimensional and parallel RF pulse design in MRI," *IEEE Trans. Med. Imag.*, vol. 28, no. 10, pp. 1548–1559, Oct. 2009.
- [5] P. Mansfield and I. L. Pykett, "Biological and medical imaging by NMR," *J. Magn. Reson.*, vol. 29, no. 2, pp. 355–373, Feb. 1978.

- [6] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller, "The NMR phased array," *Magn. Reson. Med.*, vol. 16, no. 2, pp. 192–225, Nov. 1990.
- [7] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: Sensitivity encoding for fast MRI," *Magn. Reson. Med.*, vol. 42, no. 5, pp. 952–962, Nov. 1999.
- [8] A. Macovski, "Noise in MRI," *Magn. Reson. Med.*, vol. 36, no. 3, pp. 494–497, Sept. 1996.
- [9] Y. Censor, "Finite series expansion reconstruction methods," *Proc. IEEE*, vol. 71, no. 3, pp. 409–419, Mar. 1983.
- [10] 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.*, vol. 22, no. 2, pp. 178–188, Feb. 2003.
- [11] 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. Signal Processing*, vol. 53, no. 9, pp. 3393–3402, Sept. 2005.
- [12] J. Hennig, A. M. Welz, G. Schultz, J. Korvink, Z. Liu, O. Speck, and M. Zaitsev, "Parallel imaging in non-bijective, curvilinear magnetic field gradients: A concept study," *Magn. Reson. Mater. Biol. Phys. Med.*, vol. 21, no. 1–2, pp. 5–14, Mar. 2008.
- [13] J. Stockmann, P. Ciris, and R. T. Constable, "Efficient 'O-space' parallel imaging with higher-order encoding gradients and no phase encoding," in *Proc. Int. Soc. Magnetic Resonance in Medicine*, 2009, p. 761.
- [14] A. K. Funai, J. A. Fessler, D. T. B. Yeo, V. T. Olafsson, and D. C. Noll, "Regularized field map estimation in MRI," *IEEE Trans. Med. Imag.*, vol. 27, no. 10, pp. 1484–1494, Oct. 2008.
- [15] B. P. Sutton, D. C. Noll, and J. A. Fessler, "Dynamic field map estimation using a spiral-in/spiral-out acquisition," *Magn. Reson. Med.*, vol. 51, no. 6, pp. 1194–1204, June 2004.
- [16] W. Dixon, "Simple proton spectroscopic imaging," *Radiology*, vol. 153, no. 1, pp. 189–194, Oct. 1984.
- [17] J. A. Fessler and B. P. Sutton, "Nonuniform fast Fourier transforms using min-max interpolation," *IEEE Trans. Signal Processing*, vol. 51, no. 2, pp. 560–574, Feb. 2003.
- [18] J. A. Fessler. MATLAB tomography toolbox, 2004 [Online]. Available: <http://www.eecs.umich.edu/~fessler>
- [19] K. T. Block, M. Uecker, and J. Frahm, "Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint," *Magn. Reson. Med.*, vol. 57, no. 6, pp. 1086–1098, June 2007.
- [20] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Signal Processing Mag.*, vol. 25, no. 2, pp. 72–82, Mar. 2008.
- [21] J. A. Fessler and W. L. Rogers, "Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs," *IEEE Trans. Image Processing*, vol. 5, no. 9, pp. 1346–1358, Sept. 1996.
- [22] C. R. Vogel, "Non-convergence of the L-curve regularization parameter selection method," *Inverse Probl.*, vol. 12, no. 4, pp. 535–547, Aug. 1996.
- [23] S. Ahn and R. M. Leahy, "Analysis of resolution and noise properties of non-quadratically regularized image reconstruction methods for PET," *IEEE Trans. Med. Imag.*, vol. 27, no. 3, pp. 413–424, Mar. 2008.
- [24] D. C. Noll, D. G. Nishimura, and A. Macovski, "Homodyne detection in magnetic resonance imaging," *IEEE Trans. Med. Imag.*, vol. 10, no. 2, pp. 154–163, June 1991.
- [25] A. H. Delaney and Y. Bresler, "Globally convergent edge-preserving regularized reconstruction: An application to limited-angle tomography," *IEEE Trans. Image Processing*, vol. 7, no. 2, pp. 204–221, Feb. 1998.
- [26] G. Harikumar and Y. Bresler, "A new algorithm for computing sparse solutions to linear inverse problems," in *Proc. IEEE Conf. Acoustics, Speech, and Signal Processing*, 1996, vol. 3, pp. 1331–1334.
- [27] V. T. Olafsson, D. C. Noll, and J. A. Fessler, "Fast joint reconstruction of dynamic R_2^* and field maps in functional MRI," *IEEE Trans. Med. Imag.*, vol. 27, no. 9, pp. 1177–1188, Sept. 2008.
- [28] L. Ying and J. Sheng, "Joint image reconstruction and sensitivity estimation in SENSE (JSENSE)," *Magn. Reson. Med.*, vol. 57, no. 6, pp. 1196–1202, June 2007.
- [29] D. C. Noll, J. A. Fessler, and B. P. Sutton, "Conjugate phase MRI reconstruction with spatially variant sample density correction," *IEEE Trans. Med. Imag.*, vol. 24, no. 3, pp. 325–336, Mar. 2005.
- [30] J. A. Fessler and D. C. Noll, "Model-based MR image reconstruction with compensation for through-plane field inhomogeneity," in *Proc. IEEE Int. Symp. Biomedical Imaging*, 2007, pp. 920–923.