Published in final edited form as:

IEEE Signal Process Mag. 2020 January ; 37(1): 33–40. doi:10.1109/MSP.2019.2943645.

Optimization Methods for Magnetic Resonance Image Reconstruction: Key Models and Optimization Algorithms

Jeffrey A. Fessler [Fellow, IEEE]

EECS Department, Univ. of Michigan.

Abstract

The development of compressed sensing methods for magnetic resonance (MR) image reconstruction led to an explosion of research on models and optimization algorithms for MR imaging (MRI). Roughly 10 years after such methods first appeared in the MRI literature, the U.S. Food and Drug Administration (FDA) approved certain compressed sensing methods for commercial use, making compressed sensing a clinical success story for MRI. This review paper summarizes several key models and optimization algorithms for MR image reconstruction, including both the type of methods that have FDA approval for clinical use, as well as more recent methods being considered in the research community that use data-adaptive regularizers. Many algorithms have been devised that exploit the structure of the system model and regularizers used in MRI; this paper strives to collect such algorithms in a single survey.

I. Introduction

A. Scope

Although the paper title begins with “optimization methods,” in practice one first defines a model and cost function, and then applies an optimization algorithm. There are several ways to partition the space of models, cost functions and optimization methods for MRI reconstruction, such as: smooth vs non-smooth cost functions, static vs dynamic problems, single-coil vs multiple-coil data. This paper focuses on the *static* reconstruction problem because the dynamic case is rich enough to merit its own survey paper. This paper emphasizes algorithms for multiple-coil data (parallel MRI) because modern systems all have multiple channels and advanced reconstruction methods with under-sampling are most likely to be used for parallel MRI scans. Main families of parallel MRI methods include “SENSE” methods that model the coil sensitivities in the image domain [1], “GRAPPA” methods that model the effect of coil sensitivity in k-space [2], and “calibrationless” methods that use low-rank or joint sparsity properties [3]. This paper considers all three approaches, emphasizing SENSE methods for simplicity.¹

Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

fessler@umich.edu.

¹Jupyter notebooks with code in the open source language Julia that reproduce the figures in this paper are available in the Michigan Image Reconstruction Toolbox (MIRT) at <http://github.com/JeffFessler/MIRT.jl>

B. Measurement model

The signals recorded by the sensors (receive coils) in MR scanners are linear functions of the object's transverse magnetization. That magnetization is a complicated and highly nonlinear function of the RF pulses, gradient waveforms, and tissue properties, governed by the physics of the Bloch equation [4]. Quantifying tissue properties using nonlinear models is a rich topic of its own [5], but we focus here on the problem of reconstructing images of the transverse magnetization from MR measurements.

Ignoring noise, a vector $s \in \mathbb{C}^M$ of signal samples recorded by a MR receive coil is related (typically) to a discretized version $x \in \mathbb{C}^N$ of the transverse magnetization via a linear Fourier relationship:

$$s = Fx, \quad F_{ij} = \exp(-i2\pi \vec{v}_i \cdot \vec{x}_j), \quad \begin{matrix} i = 1, \dots, M \\ j = 1, \dots, N, \end{matrix} \quad (1)$$

where $\vec{v}_i$ denotes the k-space sample location of the i th sample (units cycles/cm) and $\vec{x}_j$ denotes the spatial coordinates of the center of the j th pixel (units cm). In the usual case where the pixel coordinates $\{\vec{x}_j\}$ and k-space sample locations $\{\vec{v}_i\}$ are both on appropriate Cartesian grids, matrix F is square corresponds to the (2D or 3D) discrete Fourier transform (DFT). In this case $F^{-1} = \frac{1}{N}F'$ so reconstructing x from s is simply an inverse fast Fourier transform (FFT), and that approach is used in many clinical MR scans.

The reconstruction problem becomes more interesting when the k-space sample locations are on a non-Cartesian grid [6], when the scan is “accelerated” by recording $M < N$ samples, when non-Fourier effects like magnetic field inhomogeneity are considered [7] and/or when there are multiple receive coils. In parallel MRI, let s_l denote the samples recorded by the l th of L receive coils. Then one replaces the model (1) with

$$s_l = FC_l x, \quad (2)$$

where C_l is a $N \times N$ diagonal matrix containing the coil sensitivity pattern of the l coil on its diagonal. Note that F does *not* depend on l ; all coils see the same k-space sampling pattern. Stacking up the measurements from all coils and accounting for noise yields the following basic forward model in MRI:

$$\begin{bmatrix} y_1 \\ \vdots \\ y_L \end{bmatrix} = y = \underbrace{(I_L \otimes F)C}_{\mathbf{A}} x + \epsilon, \quad C = \begin{bmatrix} C_1 \\ \vdots \\ C_L \end{bmatrix}, \quad (3)$$

where $\mathbf{A} \in \mathbb{C}^{ML \times N}$ denotes the system matrix, $y \in \mathbb{C}^{ML}$ denotes the measured k-space data, and $x \in \mathbb{C}^N$ denotes the latent image. The noise in k-space is well modeled as complex white Gaussian noise. For extensions that consider other physics effects like relaxation and field inhomogeneity, see [8].

The goal in MR image reconstruction is to recover $\mathbf{x}$ from $\mathbf{y}$ using the model (3). All MR image reconstruction problems are under-determined because the magnetization of the underlying object being scanned is a space-limited continuous-space function on $\mathbb{R}^3$, yet only a finite number of samples are recorded. Nevertheless, the convention in MRI is to treat the object as a finite-dimensional vector $\mathbf{x} \in \mathbb{C}^N$ for which $M \approx N$ appropriate Cartesian k-space samples is considered “fully sampled” and any $M < N$ is considered “accelerated.” Sampling pattern design is a topic of ongoing interest, with renewed interest in data-driven methods [9].

The matrix $\mathbf{F}$ in (3) is known prior to the scan, because the k-space sample locations $\{\vec{v}_i\}$ are controlled by the pulse sequence designer. In contrast, the coil sensitivity maps $\{\mathbf{C}_l\}$ depend on the exact configuration of the receive coils for each patient. To use the model (3), one must determine the sensitivity maps from some patient-specific calibration data, *e.g.*, by joint estimation [10], regularization [11], or subspace methods [12].

II. Cost functions and algorithms

A. Quadratic problems

When $ML \approx N$, *i.e.*, when the total number of k-space samples acquired across all coils exceeds the number of unknown image pixel values, the linear model (3) is over-determined. If additionally $\mathbf{A}$ is well conditioned, which depends on the sampling pattern and coil sensitivity maps, then it is reasonable to consider an ordinary least-squares estimator

$$\begin{aligned}\hat{\mathbf{x}} &= \arg \min_{\mathbf{x} \in \mathbb{C}^N} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 = (\mathbf{A}'\mathbf{A})^{-1} \mathbf{A}'\mathbf{y} \\ &= \left(\sum_{l=1}^L \mathbf{C}_l' \mathbf{F}' \mathbf{F} \mathbf{C}_l \right)^{-1} \left(\sum_{l=1}^L \mathbf{C}_l' \mathbf{F}' \mathbf{y} \right).\end{aligned}\tag{4}$$

In particular, for fully sampled Cartesian k-space data where $\mathbf{F}^{-1} = \frac{1}{N} \mathbf{F}'$, this least-squares solution simplifies to $\hat{\mathbf{x}} = \left(\sum_{l=1}^L \mathbf{C}_l' \mathbf{C}_l \right)^{-1} \left(\sum_{l=1}^L \mathbf{C}_l' \mathbf{F}^{-1} \mathbf{y} \right)$, which is trivial to implement because each $\mathbf{C}_l$ is diagonal. This is known as the optimal coil combination approach [13]. For regularly under-sampled Cartesian data, where only every n th row of k-space is collected, the matrix $\mathbf{F}'\mathbf{F}$ has a simple block structure with $n \times n$ blocks that facilitates non-iterative block-wise computation known as SENSE reconstruction [1]. This form of least-squares estimation is used widely in clinical MR systems.

B. Regularized least-squares

For under-sampled problems ($ML < N$) the LS solution (4) is not unique. Furthermore, even when $ML \approx N$ often $\mathbf{A}$ is poorly conditioned, particularly for non-Cartesian sampling. Some form of regularization is needed in such cases. Some early MRI reconstruction work used quadratically regularized cost functions leading to optimization problems of the form:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{C}^N} \frac{1}{2} \|\mathbf{Ax} - \mathbf{y}\|_2^2 + \beta \|\mathbf{T}\mathbf{x}\|_2^2, \quad (5)$$

where $\beta > 0$ denotes a regularization parameter and $\mathbf{T}$ denotes a $K \times N$ matrix transform such as finite differences. The **conjugate gradient (CG) algorithm** is well-suited to such quadratic cost functions [6], [7]. The Hessian matrix $\mathbf{A}'\mathbf{A} + \beta\mathbf{T}'\mathbf{T}$ often is approximately Toeplitz [14], so CG with circulant preconditioning is particularly effective. Although the quadratically regularized least-squares cost function (5) is passé in the compressed sensing era, CG is often an inner step for optimizing more complicated cost functions [15].

C. Edge-preserving regularization

The drawback of the quadratically regularized cost function (5) with $\mathbf{T}$ as finite differences is that it blurs image edges. To avoid this blur, one can replace the quadratic regularizer $\|\mathbf{T}\mathbf{x}\|_2^2$ with a nonquadratic function $\psi(\mathbf{T}\mathbf{x})$ where typically ψ is convex and smooth, such as the Huber function, a hyperbola, or the Fair potential function $\psi(z) = \delta^2 (|z|/\delta - \log(1 + |z|/\delta))$, among others as follows:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{C}^N} \Psi(\mathbf{x}) \triangleq \frac{1}{2} \|\mathbf{Ax} - \mathbf{y}\|_2^2 + \beta \psi(\mathbf{T}\mathbf{x}). \quad (6)$$

Such methods have their roots in Bayesian methods based on Markov random fields. The **nonlinear CG algorithm** is an effective optimization method for cost functions with such smooth edge-preserving regularizers. Another appropriate optimization algorithm is the **optimized gradient method (OGM)**, a first-order method having optimal worst-case performance among all first-order algorithms for convex cost functions with Lipschitz continuous gradients [16]. OGM has a convergence rate bound that is twice better than that of **Nesterov's fast gradient method** [17].

Fig. 1 compares two of these methods for the case where $\mathbf{T}$ is finite differences and ψ is the Fair potential with $\delta = 0.1$, which approximates TV fairly closely while being smooth.

D. Sparsity models: synthesis form

Scan time in MRI is proportional to the number of k-space samples recorded. Reducing scan time in MRI can reduce cost, improve patient comfort, and reduce motion artifacts. Reducing the number of k-space samples ML to well below N , necessitates stronger modeling assumptions about $\mathbf{x}$, and sparsity models are prevalent [18] [19]. Two main categories of sparsity models are the synthesis approach and the analysis approach. In a synthesis model, one assumes $\mathbf{x} = \mathbf{B}\mathbf{z}$ for some $N \times K$ matrix $\mathbf{B}$ where coefficient vector $\mathbf{z} \in \mathbb{C}^K$ should be sparse. In an analysis model, one assumes $\mathbf{T}\mathbf{x}$ is sparse, for some $K \times N$ transformation matrix $\mathbf{T}$.

A typical cost function for a synthesis model is

$$\hat{\mathbf{x}} = \mathbf{B}\hat{\mathbf{z}}, \quad \hat{\mathbf{z}} = \arg \min_{\mathbf{z} \in \mathbb{C}^K} \frac{1}{2} \|\mathbf{A}\mathbf{B}\mathbf{z} - \mathbf{y}\|_2^2 + \beta \|\mathbf{z}\|_1, \quad (7)$$

where the 1-norm is a convex relaxation of the ℓ_0 counting measure that encourages $\mathbf{z}$ to be sparse. Typically $\mathbf{B}$ is a wide matrix (often called an over-complete dictionary) so that one can represent $\mathbf{x}$ well using only a fraction of the columns of $\mathbf{B}$. The classical approach for (7) is the **iterative soft thresholding algorithm** (ISTA) [20], also known as the **proximal gradient method** (PGM) [21] and proximal forward-backward splitting [22], having the simple form

$$\mathbf{z}_{k+1} = \text{soft}(\mathbf{z}_k - \mathbf{D}^{-1} \mathbf{B}' \mathbf{A}' (\mathbf{A} \mathbf{B} \mathbf{z}_k - \mathbf{y}), \beta/d), \quad (8)$$

where the soft thresholding function is defined by $\text{soft}(z, c) = \text{sign}(z) \max(|z| - c, 0)$ and $\mathbf{D} = \text{diag}\{\mathbf{d}\}$ is any positive definite diagonal matrix such that $\mathbf{D} - \mathbf{B}' \mathbf{A}' \mathbf{A} \mathbf{B}$ is positive semidefinite [23].

The ISTA update (8) applies to the 1-norm in (7). If we replace that 1-norm with some other function $\psi(\mathbf{z})$ then one replaces (8) with the more general PGM update of the form

$$\mathbf{z}_{k+1} = \text{prox}_{\beta\psi}(\mathbf{z}_k - \mathbf{D}^{-1} \mathbf{B}' \mathbf{A}' (\mathbf{A} \mathbf{B} \mathbf{z}_k - \mathbf{y}), \beta/d),$$

where the **proximal operator** of a function f is defined by

$$\text{prox}_f(\mathbf{v}) \triangleq \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 + f(\mathbf{x}). \quad (9)$$

Traditionally $\mathbf{D} = \|\mathbf{B}' \mathbf{A}' \mathbf{A} \mathbf{B}\|_2 \mathbf{I}$, but computing that spectral norm (via the power iteration) requires considerable computation for parallel MRI problems in general. However, for Cartesian sampling, $\mathbf{F}' \mathbf{F} \leq \mathbf{I}$ so it suffices to have $\mathbf{N} \mathbf{B}' \mathbf{C}' \mathbf{C} \mathbf{B} \leq \mathbf{D}$. Often the sensitivity maps are normalized such that $\mathbf{C}' \mathbf{C} = \mathbf{I}$ in which case $\mathbf{N} \mathbf{B}' \mathbf{B} \leq \mathbf{D}$ suffices. If in addition $\mathbf{B}'$ is a Parseval tight frame, then $\mathbf{B}' \mathbf{B} \leq \mathbf{I}$ so using $\mathbf{D} = \mathbf{I}$ is appropriate. For non-Cartesian sampling, or non-normalized sensitivity maps, or general choices of $\mathbf{B}$, finding $\mathbf{D}$ is more complicated [23].

Although ISTA is simple, it has an undesirably slow $O(1/k)$ convergence bound, where k denotes the number of iterations. This limitation was first overcome by the **fast iterative soft thresholding algorithm** (FISTA) [24], also known as the **fast proximal gradient method** (FPGM) that has an $O(1/k^2)$ convergence bound. A recent extension is the **proximal optimized gradient method** (POGM) that has worst-case convergence bound about twice better than that of FISTA/FPGM [25]. Both FISTA and POGM are essentially as simple to implement as (8). Recent MRI studies have shown POGM converging faster than FISTA, as one would expect based on the convergence bounds [26], particularly when combined with adaptive restart [25]. So POGM (with restart) is a recommended method for optimization problems having the form (7). Figure 2 provides POGM pseudo-code for solving composite optimization problems like the MRI synthesis reconstruction model (7).

Fig. 2 shows that POGM converges faster than FISTA and ISTA for minimizing (7).

E. Sparsity models: analysis form

A potential drawback of the synthesis formulation (7) is that $\mathbf{x} \approx \mathbf{B}\mathbf{z}$ may be a more realistic assumption than the strict equality $\mathbf{x} = \mathbf{B}\mathbf{z}$ when $\mathbf{z}$ is sparse. The analysis approach avoids constraining $\hat{\mathbf{x}}$ to lie in any such subspace (or union of subspaces when $\mathbf{B}$ is wide). For an analysis form sparsity model, a typical optimization problem involves a composite cost function consisting of the sum of a smooth term and a non-smooth term:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \beta \|\mathbf{T}\mathbf{x}\|_1, \quad (10)$$

where $\mathbf{T}$ is a sparsifying operator such as a wavelet transform, or finite differences, or both [18]. The expression (10) is general enough to handle combinations of multiple regularizers, such as wavelets and finite differences [19], by stacking the operators in $\mathbf{T}$ and possibly allowing a weight 1-norm. When $\mathbf{T}$ is finite differences, the regularizer is called total variation (TV) [11], and combinations of TV and wavelet transforms are useful [19]. Although the details are proprietary, the FDA-approved method for compressed sensing MRI for at least one manufacturer is related to (10) [27].

When $\mathbf{T}$ is invertible, such as an orthogonal wavelet transform, one rewrites the optimization problem (10) as

$$\hat{\mathbf{x}} = \mathbf{T}^{-1}\hat{\mathbf{z}}, \quad \hat{\mathbf{z}} = \arg \min_{\mathbf{z}} \frac{1}{2} \|\mathbf{A}\mathbf{T}^{-1}\mathbf{z} - \mathbf{y}\|_2^2 + \beta \|\mathbf{z}\|_1,$$

which is simply a special case of (7) with $\mathbf{B} = \mathbf{T}^{-1}$. Typically $\mathbf{B}$ is wide and $\mathbf{T}$ is tall so this simplification is not possible in general.

In the general case (10) where $\mathbf{T}$ is not invertible, the optimization problem is much harder than (7) due to the non-differentiability of the 1-norm with the matrix $\mathbf{T}$. The non-invertible case (with redundant Haar wavelets) is used clinically [27]. The PGM for (10) is

$$\mathbf{x}_{k+1} = \arg \min_{\mathbf{x}} \frac{\mathcal{L}}{2} \|\mathbf{x} - \tilde{\mathbf{x}}_k\|_2^2 + \beta \|\mathbf{T}\mathbf{x}\|_1, \quad (11)$$

where $\tilde{\mathbf{x}}_k \triangleq \mathbf{x}_k - \frac{1}{\mathcal{L}} \mathbf{A}'(\mathbf{A}\mathbf{x}_k - \mathbf{y})$ denotes the usual gradient update and the Lipschitz constant is $\mathcal{L} = \|\mathbf{A}\|_2^2$. Unfortunately there is no simple solution for computing the proximal operator (9) in (11) in general, so inner iterative methods are required, typically involving dual formulations [24]. This challenge makes PGM and FPGM and POGM less attractive for (10) and has led to a vast literature on algorithms for problems like (10), with no consensus on what is best. The difficulty of (11) is the main drawback of analysis regularization, whereas a possible drawback of the synthesis regularization in (7) is that often $K \gg N$ for overcomplete $\mathbf{B}$.

1) Approximate methods: One popular “work around” option is to “round the corner” of the 1-norm, making smooth approximations like $|z| \approx \sqrt{|z|^2 + \epsilon}$. This approximation is simply the hyperbola function that has a long history in the edge-preserving regularization literature. All of the gradient-based algorithms mentioned for edge-preserving regularization above are suitable candidates when a smooth function replaces the 1-norm. Smooth functions can shrink values towards zero, but their proximal operators never have a thresholding effect that induces sparsity by setting many values exactly to zero. Whether a thresholding effect is truly essential is an open question. Hereafter we focus on methods that tackle the 1-norm directly without any such approximations.

2) Variable splitting methods: Variable splitting methods replace (10) with an exactly equivalent constrained minimization problem involving an auxiliary variable such as $\mathbf{z} = \mathbf{T}\mathbf{x}$, *e.g.*,

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \min_{\mathbf{z}: \mathbf{z} = \mathbf{T}\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \beta \|\mathbf{z}\|_1. \quad (12)$$

This approach underlies the **split Bregman algorithm** [28], various **augmented Lagrangian methods** [29], and the **alternating direction multiplier method** (ADMM) [30]. The augmented Lagrangian for (12) is

$$\begin{aligned} L(\mathbf{x}, \mathbf{z}; \boldsymbol{\gamma}, \mu) &= \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \beta \|\mathbf{z}\|_1 \\ &+ \text{real}\{\langle \boldsymbol{\gamma}, \mathbf{T}\mathbf{x} - \mathbf{z} \rangle\} + \frac{\mu}{2} \|\mathbf{T}\mathbf{x} - \mathbf{z}\|_2^2, \end{aligned}$$

where $\boldsymbol{\gamma} \in \mathbb{C}^K$ denotes the vector of Lagrange multipliers and $\mu > 0$ is an AL penalty parameter that affects the convergence rate but not the final image. Defining the scaled dual variable $\boldsymbol{\eta} \triangleq 1/\mu \boldsymbol{\gamma}$ and completing the square leads to the following scaled augmented Lagrangian:

$$\begin{aligned} L(\mathbf{x}, \mathbf{z}; \boldsymbol{\eta}, \mu) &= \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \beta \|\mathbf{z}\|_1 \\ &+ \frac{\mu}{2} \left(\|\mathbf{T}\mathbf{x} - \mathbf{z} + \boldsymbol{\eta}\|_2^2 - \|\boldsymbol{\eta}\|_2^2 \right). \end{aligned}$$

An augmented Lagrangian approach alternates between descent updates of the primal variables $\mathbf{x}, \mathbf{z}$ and an ascent update of the scaled dual variable $\boldsymbol{\eta}$. The $\mathbf{z}$ update is simply soft thresholding:

$$\mathbf{z}_{k+1} = \text{soft}(\mathbf{T}\mathbf{x}_k + \boldsymbol{\eta}_k, \beta/\mu).$$

The $\mathbf{x}$ update minimizes a quadratic function:

$$\mathbf{x}_{k+1} = (\mathbf{A}'\mathbf{A} + \mu\mathbf{T}'\mathbf{T})^{-1}(\mathbf{A}'\mathbf{y} + \mu\mathbf{T}'(\mathbf{z}_{k+1} + \boldsymbol{\eta}_k)).$$

A few CG iterations is a natural choice for approximating the $\mathbf{x}$ update. Finally the $\boldsymbol{\eta}$ update is

$$\boldsymbol{\eta}_{k+1} = \boldsymbol{\eta}_k + (T\mathbf{x}_{k+1} - \mathbf{z}_{k+1}).$$

The unit step size here ensures dual feasibility [31]. A drawback of variable splitting methods is the need to select the parameter μ . Adaptive methods have been proposed to help with this tuning [31].

One could apply ADMM to the synthesis regularized problem (7), though again it would require parameter tuning that is unnecessary with POGM.

The conventional variable split in (12) ignores the specific structure of the MRI system matrix $\mathbf{A}$ in (3). Important properties of $\mathbf{A}$ include the fact that $\mathbf{F}'\mathbf{F}$ is circulant (for Cartesian sampling) or Toeplitz (for non-Cartesian sampling) and that each coil sensitivity matrix $\mathbf{C}_j$ is diagonal. In contrast, the Gram matrix $\mathbf{A}'\mathbf{A}$ for parallel MRI is harder to precondition, though possible [32]. An alternative splitting that simplifies the updates is [29]:

$$\begin{aligned} \arg \min_{\mathbf{x} \in \mathbb{C}^N} \min_{\mathbf{u} \in \mathbb{C}^{NL}, \mathbf{z} \in \mathbb{C}^K, \mathbf{v} \in \mathbb{C}^N} & \frac{1}{2} \|\mathbf{F}_L \mathbf{u} - \mathbf{y}\|_2^2 + \beta \|\mathbf{z}\|_1 \\ \text{sub. to } & \mathbf{u} = \mathbf{C}\mathbf{x}, \mathbf{z} = \mathbf{T}\mathbf{v}, \mathbf{v} = \mathbf{x}, \end{aligned} \quad (13)$$

where $\mathbf{F}_L \triangleq \mathbf{I}_L \otimes \mathbf{F}$. With this splitting, the $\mathbf{z}$ update again is simply soft thresholding, and the $\mathbf{x}$ update involves the diagonal matrix $\mathbf{C}'\mathbf{C}$ which is trivial. The $\mathbf{v}$ update involves the matrix $\mathbf{T}'\mathbf{T}$ that is circulant for periodic boundary conditions or is very well suited to a circulant preconditioner otherwise, using simple FFT operations. The $\mathbf{u}$ update involves the matrix $\mathbf{F}_L'\mathbf{F}_L$ that is circulant or Toeplitz. This approach exploits the structure of $\mathbf{A}$ to simplify the updates; the primary drawback is that it requires selecting even more AL penalty parameters; condition number criteria can be helpful [29]. Another splitting with fewer auxiliary variables leads to an inner update step that requires solving denoising problems similar to (11).

3) Primal-dual methods: A key idea behind duality-based methods is the fact:

$$\|\mathbf{T}\mathbf{x}\|_1 = \max_{\mathbf{z} \in \mathbb{C}^K: \|\mathbf{z}\|_\infty \leq 1} \text{real}\{\langle \mathbf{z}, \mathbf{T}\mathbf{x} \rangle\}.$$

Thus the (nonsmooth) analysis regularized problem (10) is equivalent to this constrained problem:

$$\arg \min_{\mathbf{x}} \min_{\mathbf{z} \in \mathcal{Z}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \beta \text{real}\{\langle \mathbf{z}, \mathbf{T}\mathbf{x} \rangle\}, \quad (14)$$

where $\mathcal{Z} \triangleq \{\mathbf{z} \in \mathbb{C}^K: \|\mathbf{z}\|_\infty \leq 1\}$. The **primal-dual methods** typically alternate between updating the primal variable $\mathbf{x}$ and the dual variable $\mathbf{z}$, using more convenient alternatives to

(14) that involve separate multiplication by $\mathbf{A}$ and by $\mathbf{A}'$ without requiring inner CG iterations. These methods provide convergence guarantees and acceleration techniques that lead to $\mathcal{O}(1/k^2)$ rates [33]. A drawback of such methods is they typically require power iterations to find a Lipschitz constant, and, like AL methods, have tuning parameters that affect the practical convergence rates. Finding a simple, convergent, and tuning-free method for the analysis regularized problem (10) remains an important open problem.

F. Patch-based sparsity models

Using (10) with a finite-difference regularizer is essentially equivalent to using patches of size 2×1 . It is plausible that one can regularize better by considering larger patches that provide more context for distinguishing signal from noise. There are two primary modes of patch-based regularization: synthesis models and analysis methods.

A typical synthesis approach attempts to represent each patch using a sparse linear combination of atoms from some signal patch dictionary. Let $\mathbf{P}_p$ denote the $d \times N$ matrix that extracts the p th of P patches (having d pixels) when multiplied by an image vector $\mathbf{x}$. Then the synthesis model is that $\mathbf{P}_p \mathbf{x} \approx \mathbf{D} \mathbf{z}_p$ where $\mathbf{D}$ is a $d \times J$ dictionary, such as the discrete cosine transform (DCT) [34], and $\mathbf{z}_p \in \mathbb{C}^J$ is a sparse coefficient vector for the p th patch. Under this model, a natural regularizer is

$$R(\mathbf{x}) = \min_{\{\mathbf{z}_p\}} \sum_{p=1}^P \frac{1}{2} \|\mathbf{P}_p \mathbf{x} - \mathbf{D} \mathbf{z}_p\|_2^2 + \alpha \|\mathbf{z}_p\|_1. \quad (15)$$

See [34] for an extension to the case of multiple images. The regularizer has an inner minimization over the sparse coefficients $\{\mathbf{z}_p\}$, so the overall problem involves both optimizing the image $\mathbf{x}$ and those coefficients. This structure lends itself to **alternating minimization** algorithms. The work in [34] used ISTA for updating $\mathbf{z}_p$; the results in Fig. 2 suggest that POGM may be beneficial.

A typical analysis approach for patches assumes there is a sparsifying transform $\mathbf{\Omega}$ such that $\mathbf{\Omega} \mathbf{P}_p \mathbf{x}$ tend to be sparse. For example, [35] uses a directional wavelet transform for each patch. Under this model, a natural regularizer is

$$R(\mathbf{x}) = \min_{\{\mathbf{z}_p\}} \sum_{p=1}^P \frac{1}{2} \|\mathbf{\Omega} \mathbf{P}_p \mathbf{x} - \mathbf{z}_p\|_2^2 + \alpha \|\mathbf{z}_p\|_1. \quad (16)$$

Again a double minimization over the image $\mathbf{x}$ and the transform coefficients $\{\mathbf{z}_p\}$ is needed, so **alternating minimization** algorithms are natural. For alternating minimization (block coordinate descent), the update of each $\mathbf{z}_p$ is simply soft thresholding, and the update of $\mathbf{x}$ is a quadratic problem involving $\mathbf{A}' \mathbf{A} + \beta \sum_p \mathbf{P}_p' \mathbf{\Omega}' \mathbf{\Omega} \mathbf{P}_p$. When the transform $\mathbf{\Omega}$ is unitary and the patches are selected with periodic boundary conditions and a stride of one pixel, then this simplifies to $\mathbf{A}' \mathbf{A} + \beta \mathbf{I}$. A few inner iterations of the (preconditioned) CG algorithm is useful for the $\mathbf{x}$ update. Under these assumptions, and using just a single gradient descent update for $\mathbf{x}$, an alternating minimization algorithm for least-squares with regularizer (16) simply alternates between a denoising step and a gradient step:

$$\begin{aligned}\tilde{\mathbf{x}}_k &= \sum_{p=1}^P \mathbf{P}_p' \boldsymbol{\Omega}' \text{soft}(\boldsymbol{\Omega} \mathbf{P}_p \mathbf{x}_k, \alpha) \\ \mathbf{x}_{k+1} &= \mathbf{x}_k - (\mathbf{D} + \beta \mathbf{I})^{-1} (\mathbf{A}'(\mathbf{A} \mathbf{x} - \mathbf{y}) + \beta \tilde{\mathbf{x}}_k).\end{aligned}\quad (17)$$

For this algorithm the cost function is monotonically nonincreasing.

G. Adaptive regularization

The patch dictionary $\mathbf{D}$ in (15) or the sparsifying transform $\boldsymbol{\Omega}$ in (16) can be chosen based on mathematical models like the DCT, or they can be learned from a population of preexisting training data and then used in (15) or (16) for subsequent patients. A third possibility is to adapt $\mathbf{D}$ or $\boldsymbol{\Omega}$ to each specific patient [36]. The “dictionary learning MRI” (DLMRI) approach [36] uses a non-convex regularizer of the following form:

$$R(\mathbf{x}) = \min_{\mathbf{D} \in \mathcal{D}} \min_{\{\mathbf{z}_p\}_{p=1}^P} \sum_{p=1}^P \|\mathbf{P}_p \mathbf{x} - \mathbf{D} \mathbf{z}_p\|_2^2 + \alpha \|\mathbf{z}_p\|_1, \quad (18)$$

where $\mathcal{D}$ is the feasible set of dictionaries (typically constrained so that each atom has unit norm). Now there are three set of variables to optimize: $\mathbf{x}$, $\{\mathbf{z}_p\}$, $\mathbf{D}$, so alternating minimization methods are well suited. The update of the image $\mathbf{x}$ is a quadratic optimization subproblem, the $\mathbf{z}_p$ update is soft thresholding, and the $\mathbf{D}$ update is simple when considering one atom at a time.

The “transform learning MRI” (TLMRI) approach uses a regularizer of this form:

$$R(\mathbf{x}) = \min_{\boldsymbol{\Omega}} \min_{\{\mathbf{z}_p\}_{p=1}^P} \sum_{p=1}^P \|\boldsymbol{\Omega} \mathbf{P}_p \mathbf{x} - \mathbf{z}_p\|_2^2 + \alpha \|\mathbf{z}_p\|_1 + \gamma r(\boldsymbol{\Omega}),$$

where $r(\boldsymbol{\Omega})$ enforces or encourages properties of the sparsifying transform such as orthogonality. Again, alternating minimization methods are well suited; the $\boldsymbol{\Omega}$ update involves (small) SVD operations.

H. Convolutional regularizers

An alternative to patch-based regularization is to use convolutional sparsity models [37]. A convolutional synthesis regularizer replaces (15) with

$$R(\mathbf{x}) = \min_{\{\mathbf{z}_k\}} \frac{1}{2} \left\| \mathbf{x} - \sum_{k=1}^K \mathbf{h}_k * \mathbf{z}_k \right\|_2^2 + \alpha \|\mathbf{z}_k\|_1,$$

where $\{\mathbf{h}_k\}$ is a set of filters learned from training images [37] and $*$ denotes convolution. Again, **alternating minimization** algorithms are a natural choice because the $\mathbf{x}$ update is quadratic and the $\mathbf{z}_k$ update is a sparse coding problem for which proximal methods like POGM are well-suited.

A convolution analysis regularizer replaces (16) with

$$R(\mathbf{x}) = \min_{\{\mathbf{z}_k\}} \sum_{k=1}^K \frac{1}{2} \|\mathbf{h}_k * \mathbf{x} - \mathbf{z}_k\|_2^2 + \alpha \|\mathbf{z}_k\|_1.$$

Again, **alternating minimization** algorithms are effective, where the $\mathbf{z}_k$ update is soft thresholding. One can either learn the filters $\{\mathbf{h}_k\}$ from good quality (*e.g.*, fully sampled) training data, or adapt the filters for each patient by jointly optimizing $\mathbf{x}$, $\{\mathbf{h}_k\}$ and $\{\mathbf{z}_k\}$ using alternating minimization.

I. Other methods

The summation in (17) is a particular type of patch-based denoising of the current image estimate $\mathbf{x}_k$. There are many other denoising methods, some of which have variational formulations well-suited to inverse problems, but many of which do not, such as nonlocal means (NLM) [38] and block-matching 3D (BM3D) [39]. One way to adapt most such denoising methods for image reconstruction is to use a plug-and-play ADMM approach [40] that replaces a denoising step like (17) with a general denoising procedure.

J. Non-SENSE methods

The measurement model (2) and (3) has a single latent image $\mathbf{x}$, viewed by each receive coil. An alternate formulation is to define a latent image for each coil $\mathbf{x}_I \triangleq \mathbf{C}_I \mathbf{x}$ and write the measurement model as $\mathbf{y}_I = \mathbf{F} \mathbf{x}_I + \mathbf{e}_I$. For such formulations, the problem becomes to reconstruct the L images $\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_L]$ from the measurements, while considering relationships between those images. Because multiplication by the smooth sensitivity map $\mathbf{C}_I$ in the image domain corresponds to convolution with a small kernel in the frequency domain, any point in k-space can be approximated by a linear combination of its neighbors in all coil data [2]. This “GRAPPA modeling” leads to an approximate consistency condition $\text{vec}(\mathbf{X}) \approx \mathbf{G} \text{vec}(\mathbf{X})$ where $\mathbf{G}$ is a matrix involving small k-space kernels that are learned from calibration data [2]. This relationship leads to “SPIRiT” optimization problems like:

$$\begin{aligned} \hat{\mathbf{X}} = & \arg \min_{\mathbf{X} \in \mathbb{C}^{N \times L}} \frac{1}{2} \|\mathbf{F} \mathbf{X} - \mathbf{Y}\|_{\text{Frob}}^2 \\ & + \beta_1 \frac{1}{2} \left\| (\mathbf{G} - \mathbf{I}) \text{vec}(\mathbf{X}) \right\|_2^2 + \beta_2 R(\mathbf{X}), \end{aligned}$$

where $\mathbf{Y} = [\mathbf{y}_1 \dots \mathbf{y}_L] \in \mathbb{C}^{M \times L}$ and $R(\mathbf{X})$ is a regularizer that encourages joint sparsity because all of the images $\{\mathbf{x}_I\}$ have edges in the same locations. No sensitivity maps $\mathbf{C}$ are needed for this approach. When $\beta_2 = 0$ the problem is quadratic and CG is well suited. Otherwise, ADMM is convenient for splitting this optimization problem into parts with easier updates. The ESPIRiT approach uses the redundancy in k-space data from multiple coils to estimate sensitivity maps from the eigenvectors of a certain block-Hankel matrix [12]; this approach helps bridge the SENSE and GRAPPA approaches while building on related signal processing tools like subspace estimation and multichannel blind deconvolution.

III. Summary

Although the title of this paper is “optimization methods for...” before selecting an optimization algorithm it is far more important (for under-sampled problems) to first select an appropriate cost function that captures useful prior information about the latent object $\mathbf{x}$. The literature is replete with numerous candidate models, each of which often lead to different optimization methods. Nevertheless, common ingredients arise in most formulations, such as alternating minimization (block coordinate descent) at the outer level, preconditioned CG for inner iterations related to quadratic terms, and soft thresholding or other proximal operators for nonsmooth terms that promote sparsity.

This survey has focused on 1-norm regularizers for simplicity, but (nonconvex) p “norms” with $0 < p < 1$ have also been investigated and appear to be beneficial particularly for very undersampled measurements. This survey considers a single image $\mathbf{x}$ but many MRI scan protocols involve several images with different contrast and it may be useful to reconstruct them jointly, *e.g.*, by considering common sparsity or subspace models.

There are many open problems in optimization that are relevant to MRI. The analysis form regularized problem (10) remains challenging, and further investigation of analysis vs synthesis approaches is needed. There has been considerable recent progress on finding optimal worst-case methods [16], but these optimality results are for very broad classes of cost functions, whereas the cost functions in MRI reconstruction have particular structure. Finding algorithms with optimal complexity (fastest possible convergence) for MRI-type cost functions would be valuable both for clinical practice and for facilitating research.

Finally, the current trend is to use convolutional neural network (CNN) methods to process under-sampled images, or for direct reconstruction, or as denoising operators. The stochastic gradient descent method currently is the universal optimization tool for training CNN models. Many “deep learning” methods for MRI are based on network architectures that are “unrolled” versions of iterative optimization methods like PGM. Thus, familiarity with “classical” optimization methods for MR image reconstruction is important even in the machine learning era.

IV. Acknowledgement

The author thanks many students and postdocs for all they have taught him, Doug Noll for numerous MRI insights, and the reviewers for their detailed comments that improved the paper,

Research supported in part by NIH Grants R01 EB023618, U01 EB026977, and R21 AG061839.

References

- [1]. Pruessmann KP, Weiger M, Scheidegger MB, and Boesiger P, “SENSE: sensitivity encoding for fast MRI,” *Mag. Res. Med*, vol. 42, no. 5, pp. 952–62, 11 1999.
- [2]. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, and Haase A, “Generalized autocalibrating partially parallel acquisitions (GRAPPA),” *Mag. Res. Med*, vol. 47, no. 6, pp. 1202–10, 6 2002.

- [3]. Shin PJ, Larson PEZ, Ohliger MA, Elad M, Pauly JM, Vigneron DB, and Lustig M, "Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion," *Mag. Res. Med.*, vol. 72, no. 4, pp. 959–70, 10 2014.
- [4]. Wright GA, "Magnetic resonance imaging," *IEEE Sig. Proc. Mag.*, vol. 14, no. 1, pp. 56–66, 1 1997.
- [5]. Ma D, Gulani V, Seiberlich N, Liu K, Sunshine JL, Duerk JL, and Griswold MA, "Magnetic resonance fingerprinting," *Nature*, vol. 495, pp. 187–93, 3 2013. [PubMed: 23486058]
- [6]. Pruessmann KP, Weiger M, Boernert P, and Boesiger P, "Advances in sensitivity encoding with arbitrary k-space trajectories," *Mag. Res. Med.*, vol. 46, no. 4, pp. 638–51, 10 2001.
- [7]. Sutton BP, Noll DC, and Fessler JA, "Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities," *IEEE Trans. Med. Imag.*, vol. 22, no. 2, pp. 178–88, 2 2003.
- [8]. Fessler JA, "Model-based image reconstruction for MRI," *IEEE Sig. Proc. Mag.*, vol. 27, no. 4, pp. 81–9, 7 2010, invited submission to special issue on medical imaging.
- [9]. Baldassarre L, Li Y-H, Scarlett J, Gozcu B, Bogunovic I, and Cevher V, "Learning-based compressive subsampling," *IEEE J. Sel. Top. Sig. Proc.*, vol. 10, no. 4, pp. 809–22, 6 2016.
- [10]. Ying L and Sheng J, "Joint image reconstruction and sensitivity estimation in SENSE (JSENSE)," *Mag. Res. Med.*, vol. 57, no. 6, pp. 1196–1202, 6 2007.
- [11]. Block KT, Uecker M, and Frahm J, "Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint," *Mag. Res. Med.*, vol. 57, no. 6, pp. 1086–98, 6 2007.
- [12]. Uecker M, Lai P, Murphy MJ, Virtue P, Elad M, Pauly JM, Vasanawala SS, and Lustig M, "ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA," *Mag. Res. Med.*, vol. 71, no. 3, pp. 990–1001, 3 2014.
- [13]. Roemer PB, Edelstein WA, Hayes CE, Souza SP, and Mueller OM, "The NMR phased array," *Mag. Res. Med.*, vol. 16, no. 2, pp. 192–225, 11 1990.
- [14]. Wajer F and Pruessmann KP, "Major speedup of reconstruction for sensitivity encoding with arbitrary trajectories," in *Proc. Intl. Soc. Mag. Res. Med.*, 2001, p. 767 [Online]. Available: <http://cds.ismrm.org/ismrm-2001/PDF3/0767.pdf>
- [15]. Aelterman J, Luong HQ, Goossens B, Pizurica A, and Philips W, "Augmented Lagrangian based reconstruction of non-uniformly sub-Nyquist sampled MRI data," *Signal Processing*, vol. 91, no. 12, pp. 2731–42, 1 2011.
- [16]. Drori Y, "The exact information-based complexity of smooth convex minimization," *J. Complexity*, vol. 39, pp. 1–16, 4 2017.
- [17]. Nesterov Y, "A method of solving a convex programming problem with convergence rate $O(1/k^2)$," *Soviet Math. Dokl.*, vol. 27, no. 2, pp. 372–76, 1983 [Online]. Available: <http://www.core.ucl.ac.be/~nesterov/Research/Papers/DAN83.pdf>
- [18]. Lustig M, Donoho D, and Pauly JM, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Mag. Res. Med.*, vol. 58, no. 6, pp. 1182–95, 12 2007.
- [19]. Lustig M, Donoho DL, Santos JM, and Pauly JM, "Compressed sensing MRI," *IEEE Sig. Proc. Mag.*, vol. 25, no. 2, pp. 72–82, 3 2008.
- [20]. Daubechies I, Defrise M, and De Mol C, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Comm. Pure Appl. Math.*, vol. 57, no. 11, pp. 1413–57, 11 2004.
- [21]. Combettes PL and Pesquet J-C, "Proximal splitting methods in signal processing," 2011, fixed-Point Algorithms for Inverse Problems in Science and Engineering, Springer, Optimization and Its Applications, pp 185–212.
- [22]. Combettes P and Wajs V, "Signal recovery by proximal forward-backward splitting," *SIAM J. Multi. Mod. Sim.*, vol. 4, no. 4, pp. 1168–200, 2005.
- [23]. Muckley MJ, Noll DC, and Fessler JA, "Fast parallel MR image reconstruction via B1-based, adaptive restart, iterative soft thresholding algorithms (BARISTA)," *IEEE Trans. Med. Imag.*, vol. 34, no. 2, pp. 578–88, 2 2015.
- [24]. Beck A and Teboulle M, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE Trans. Im. Proc.*, vol. 18, no. 11, pp. 2419–34, 11 2009.

- [25]. Kim D and Fessler JA, "Adaptive restart of the optimized gradient method for convex optimization," *J. Optim. Theory Appl.*, vol. 178, no. 1, pp. 240–63, 7 2018.
- [26]. El Gueddari L, Lazarus C, Carrie H, Vignaud A, and Ciuciu P, "Self-calibrating nonlinear reconstruction algorithms for variable density sampling and parallel reception MRI," in *Proc. IEEE SAM*, 2018, pp. 415–9.
- [27]. Wetzl J, Forman C, Wintersperger BJ, D'Errico L, Schmidt M, Mailhe B, Maier A, and Stalder AF, "High-resolution dynamic CE-MRA of the thorax enabled by iterative TWIST reconstruction," *Mag. Res. Med.*, vol. 77, no. 2, pp. 833–40, 2 2017.
- [28]. Goldstein T and Osher S, "The split Bregman method for L1-regularized problems," *SIAM J. Imaging Sci.*, vol. 2, no. 2, pp. 323–43, 2009.
- [29]. Ramani S and Fessler JA, "Parallel MR image reconstruction using augmented Lagrangian methods," *IEEE Trans. Med. Imag.*, vol. 30, no. 3, pp. 694–706, 3 2011.
- [30]. Eckstein J and Bertsekas DP, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, no. 1–3, pp. 293–318, 4 1992.
- [31]. Boyd S, Parikh N, Chu E, Peleato B, and Eckstein J, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. & Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2010.
- [32]. Koolstra K, van Gemert J, Boernert P, Webb A, and Remis R, "Accelerating compressed sensing in parallel imaging reconstructions using an efficient circulant preconditioner for Cartesian trajectories," *Mag. Res. Med.*, vol. 81, no. 1, pp. 670–85, 1 2019.
- [33]. Chambolle A and Pock T, "A first-order primal-dual algorithm for convex problems with applications to imaging," *J. Math. Im. Vision*, vol. 40, no. 1, pp. 120–145, 2011.
- [34]. Wang S, Tan S, Gao Y, Liu Q, Ying L, Xiao T, Liu Y, Liu X, Zheng H, and Liang D, "Learning joint-sparse codes for calibration-free parallel MR imaging (LINDBERG)," *IEEE Trans. Med. Imag.*, vol. 37, no. 1, pp. 251–61, 1 2018.
- [35]. Qu X, Guo D, Ning B, Hou Y, Lin Y, Cai S, and Chen Z, "Undersampled MRI reconstruction with patch-based directional wavelets," *Mag. Res. Im.*, vol. 30, no. 7, pp. 964–77, 9 2012.
- [36]. Ravishanker S and Bresler Y, "MR image reconstruction from highly undersampled k-space data by dictionary learning," *IEEE Trans. Med. Imag.*, vol. 30, no. 5, pp. 1028–41, 5 2011.
- [37]. Wohlberg B, "Efficient algorithms for convolutional sparse representations," *IEEE Trans. Im. Proc.*, vol. 25, no. 1, pp. 301–15, 1 2016.
- [38]. Buades A, Coll B, and Morel J, "Image denoising methods. A new nonlocal principle," *SIAM Review*, vol. 52, no. 1, pp. 113–47, 2010.
- [39]. Dabov K, Foi A, Katkovnik V, and Egiazarian K, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Im. Proc.*, vol. 16, no. 8, pp. 2080–95, 8 2007.
- [40]. Chan SH, Wang X, and Elgendy OA, "Plug-and-play ADMM for image restoration: fixed-point convergence and applications," *IEEE Trans. Computational Imaging*, vol. 3, no. 1, pp. 84–98, 3 2017.

Fig. 1.

Comparison of CG and OGM convergence for single-coil MRI reconstruction with edge-preserving regularization (akin to anisotropic TV with corner rounding). From left to right: Top row: k-space sampling pattern where only 34% of the phase-encodes are collected, true image, initial image from zero-filled k-space data, minimizer $\hat{x}$ of (6). (Both CG and OGM converge to the same limit $\hat{x}$.) Bottom row: cost function $\Psi(x_k)$ in (6) and normalized root mean squared error (NRMSE) $\|x_k - x\|_2 / \|x\|_2$ versus iteration k .

Fig. 2.

Comparison of ISTA/PGM, FISTA/FPGM and POGM for single-coil MRI reconstruction with orthogonal discrete wavelet transform sparsity regularizer using the 1-norm. Minimizer $\hat{\mathbf{x}}$ of (7); NRMSE versus iteration k . FISTA requires about 40% more iterations to converge than POGM, consistent with the $2\times$ better worst-case bound of POGM.

Initialize $\mathbf{w}_0 = \mathbf{x}_0$, $\theta_0 = 1$. Then for $k = 1 : N$:

$$\begin{aligned}\theta_k &= \begin{cases} \frac{1}{2} \left(1 + \sqrt{4\theta_{k-1}^2 + 1} \right), & k < N \\ \frac{1}{2} \left(1 + \sqrt{8\theta_{k-1}^2 + 1} \right), & k = N \end{cases} \\ \gamma_k &= \frac{1}{L} \frac{2\theta_{k-1} + \theta_k - 1}{\theta_k} \\ \mathbf{w}_k &= \mathbf{x}_{k-1} - \frac{1}{L} \nabla f(\mathbf{x}_{k-1}) \\ \mathbf{z}_k &= \mathbf{w}_k + \frac{\theta_{k-1} - 1}{\theta_k} (\mathbf{w}_k - \mathbf{w}_{k-1}) + \frac{\theta_{k-1}}{\theta_k} (\mathbf{w}_k - \mathbf{x}_{k-1}) \\ &\quad + \frac{\theta_{k-1} - 1}{L\gamma_{k-1}\theta_k} (\mathbf{z}_{k-1} - \mathbf{x}_{k-1}) \\ \mathbf{x}_k &= \text{prox}_{\gamma_k g}(\mathbf{z}_k) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{z}_k\|_2^2 + \gamma_k g(\mathbf{x})\end{aligned}$$

Fig. 3.

POGM method for minimizing $f(x) + g(x)$ where f is convex with L -Lipschitz smooth gradient and g is convex. See [25] for adaptive restart version.