

Low-field magnetic resonance imaging using multiplicative regularization

de Leeuw den Bouter, Merel; van Gijzen, Martin; Remis, Rob

DOI

[10.1016/j.mri.2020.10.001](https://doi.org/10.1016/j.mri.2020.10.001)

Publication date

2021

Document Version

Final published version

Published in

Magnetic Resonance Imaging

Citation (APA)

de Leeuw den Bouter, M., van Gijzen, M., & Remis, R. (2021). Low-field magnetic resonance imaging using multiplicative regularization. *Magnetic Resonance Imaging*, 75, 21-33.
<https://doi.org/10.1016/j.mri.2020.10.001>

Important note


To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Original Contribution

Low-field magnetic resonance imaging using multiplicative regularization

Merel de Leeuw den Bouter^{a,*}, Martin van Gijzen^a, Rob Remis^b^a Department of Numerical Analysis, Delft Institute of Applied Mathematics, Delft University of Technology, Van Mourik Broekmanweg 6, 2628 XE Delft, The Netherlands^b Department of Microelectronics, Circuits and Systems Group, Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands

ARTICLE INFO

Keywords:

Magnetic resonance imaging
 Low-field MRI
 Image reconstruction
 Halbach array
 Multiplicative regularization

ABSTRACT

In this paper we present a magnetic resonance imaging (MRI) technique that is based on multiplicative regularization. Instead of adding a regularizing objective function to a data fidelity term, we multiply by such a regularizing function. By following this approach, no regularization parameter needs to be determined for each new data set that is acquired. Reconstructions are obtained by iteratively updating the images using short-term conjugate gradient-type update formulas and Polak-Ribière update directions. We show that the algorithm can be used as an image reconstruction algorithm and as a denoising algorithm. We illustrate the performance of the algorithm on two-dimensional simulated low-field MR data that is corrupted by noise and on three-dimensional measured data obtained from a low-field MR scanner. Our reconstruction results show that the algorithm effectively suppresses noise and produces accurate reconstructions even for low-field MR signals with a low signal-to-noise ratio.

1. Introduction

In magnetic resonance imaging (MRI), the internal structure of the human body is visualized using magnetic fields. To form an image, commercial MR scanners employ strong magnetic background fields with field strengths ranging from 1.5 T to 7 T. Superconducting magnets are used to generate such strong background fields, which obviously adds to the cost, size, and infrastructure demands of high-field MR scanners. In fact, the overall costs of present day MR scanners are so high that they are essentially out of reach for low-income and middle-income countries.

Low-field MR scanners, as described in for example [1–3], may provide a solution to this problem. Compared with the commercial scanners mentioned above, a low-field scanner has a much weaker background field (typically in the centi- or millitesla range) and MR signal quality is reduced. However, a low-field scanner does not require any superconducting magnets and construction and maintenance costs are therefore significantly lower (apart from additional cost reductions that may be achieved). Moreover, low-field scanners may be portable as well and enable us to bring these scanners as diagnostic tools to rural areas in developing countries. For a review on low-field MRI, the reader is referred to [4].

A direct consequence of having a lower background field is that low-field MR scanners inevitably yield significantly noisier signals than their high-field counterparts. Moreover, depending on the type of low-

field scanner that is used, the background field may not be (sufficiently) uniform throughout the object and the gradient fields that are used for imaging may not be linear or are only approximately linear within some subdomain of the object or body part that we want to image. Loosely speaking, a uniform background field and linear gradient fields ultimately lead to a Fourier transform representation of the measured MR signals and MR imaging essentially amounts to applying an inverse Fourier transform to the measured data. Deviations from these ideal background and gradient fields lead to imaging artifacts when an inverse Fourier transform is applied, as can be seen in [5]. Even when the background field and gradient fields can be considered constant and linear throughout the object, then still noisy reconstructions are obtained, since the acquired signals usually have a low signal-to-noise ratio (SNR).

There is a vast literature on image reconstruction in MRI. An excellent overview of many relevant techniques is given in [6]. In the case of nonlinear gradients, a model-based image reconstruction as described in [7,8], among others, can be employed instead of the standard inverse Fourier Transform. In this paper, we will make use of this model-based image reconstruction as well. However, other methods exist that correct for nonlinear magnetic fields as a post-processing step, i.e. on the image obtained by applying a standard technique like the inverse Fourier Transform to the measured data. In [9] for example, spherical harmonic deconvolution methods are used to achieve this result.

* Corresponding author.

E-mail address: M.L.deLeeuwdenBouter-1@tudelft.nl (M. de Leeuw den Bouter).

In this paper, we employ an MR imaging technique that, using model-based image reconstruction, addresses the aforementioned issues of imaging using nonlinear magnetic fields and the contamination of the signal by noise. Specifically, we pose our low-field imaging problem as an optimization problem that minimizes a (least-squares) data fidelity term in which nonuniform background fields and nonlinear gradients are taken into account through a generalized signal model. Furthermore, noise effects are suppressed by incorporating a weighted L^2 -norm total variation objective function into the optimization procedure. Total variation regularization penalizes jumps between neighboring pixels. The additive variant is often used in MR imaging to denoise images while still maintaining their edges, see for example [10–13], among many others.

It is customary to include such a regularization term in an additive manner into an optimization framework. However, one of the main drawbacks of such an additive scheme is that an artificial regularization parameter must be chosen that balances the data fidelity and regularization term. While methods for choosing this parameter exist, see for example [14], they are often computationally expensive and do not allow for fast (real-time) imaging. Moreover, typically a regularizing parameter needs to be determined for each new available data set, which leads to even larger computation times in case multiple data sets need to be processed.

Inspired by the success of multiplicative regularization in wave field inversion, see [15–18], for example, we include regularization in a multiplicative manner as well. In the resulting iterative imaging algorithm, the image is then updated using a Polak-Ribière type of conjugate gradient directions, see for example [19]. Two practical advantages of this scheme are that no effective regularization parameter needs to be computed and reconstruction results can be monitored as the scheme progresses. A similar multiplicative regularization approach was applied to image deblurring problems in [20].

We apply our imaging method to low-field noise-corrupted simulated data using nonlinear gradient fields and to measured data obtained with the low-field scanner shown in Fig. 1. This scanner consists of 23 rings spaced 22 mm apart. The total length of the scanner is 50.6 cm and the bore has a diameter of 27 cm. Each ring in the array consists of two concentric layers of cubic neodymium boron iron magnets. The magnets have a side length of 12 mm and the total number of magnets is close to 3000. The magnets are arranged in a Halbach configuration such that an approximately uniform background field is realized within the bore of the scanner. This background field is pointing in a transverse direction (from left to right inside the bore in the photograph of Fig. 1) as opposed to the longitudinal direction as is common in commercial high-field scanners. Consequently, the gradient coils of the scanner have to be redesigned as well. A gradient coil especially designed for the low-field scanner that produces an approximately linear gradient in the longitudinal direction inside the bore is also shown in Fig. 1. Further details about the design and realization of the scanner and the gradient coils can be found in [2,21].

For different types of fruit (apple, melon) we reconstruct an image using data collected with the MR scanner described above. Moreover, the effects of a nonlinear frequency encoding gradient are also studied

using noise-corrupted simulated data. We resort to simulated data in this case, since we do not have measured background or gradient fields available. We do stress, however, that if such measured fields are available, then these can be easily incorporated into our imaging scheme. Imaging results will be presented for a regularizing objective function based on the minimization of a weighted total variation term. We demonstrate that the method indeed effectively suppresses noise for a given data set without any extra computations to determine a regularization parameter.

This paper is organized as follows. In Section 2 we introduce the continuous and discrete signal model and discuss the multiplicative regularization approach. Secondly, we present the multiplicatively regularized MR imaging algorithm and show that it can be used as an image reconstruction algorithm or as a denoising algorithm or both. In Section 3 we illustrate the performance of our imaging algorithm on simulated, noise-corrupted data and on measured data obtained with the low-field scanner. Finally, the conclusions can be found in Section 4.

2. Methods

In Section 2.1 we will introduce the signal model and discuss the multiplicative regularization approach. Section 2.2 contains the details on the numerical discretization that is used. After that, in Section 2.3, the multiplicatively regularized MR imaging algorithm is presented.

2.1. Signal modeling and regularization

The starting point of our imaging procedure is the voltage signal $v(t)$ that is measured by a receive coil of the low-field scanner. Assuming that the complete object is excited, we have for this signal the representation

$$v(t) = \int_{\mathbf{r} \in D} e^{-i\Delta\omega_0(\mathbf{r})t} e^{-i2\pi\mathbf{k}(\mathbf{r},t) \cdot \mathbf{r}} X(\mathbf{r}) dV, \quad (1)$$

where $\Delta\omega_0(\mathbf{r}) = \gamma B_0(\mathbf{r}) - \omega_{\text{mod}}$ is the difference between the local Larmor frequency $\gamma B_0(\mathbf{r})$ and the modulation frequency ω_{mod} with γ the proton gyromagnetic ratio. Furthermore, $X(\mathbf{r})$ is the image having the bounded domain D as its support. The image X depends on the proton density, coil sensitivity, etc. and is expressed in $[\text{Vm}^{-3}]$, but its particular form is irrelevant for our imaging purposes. Explicit expressions for X can be found in [22], for example. Finally,

$$\mathbf{k}(\mathbf{r}, t) = \frac{\gamma}{2\pi} \int_{\tau=0}^t \mathbf{G}(\mathbf{r}, \tau) d\tau, \quad (2)$$

where $\mathbf{G}(\mathbf{r}, \tau) = G_x(\mathbf{r}, \tau)\mathbf{i}_x + G_y(\mathbf{r}, \tau)\mathbf{i}_y + G_z(\mathbf{r}, \tau)\mathbf{i}_z$ [Tm^{-1}] is the gradient vector due to the application of the gradient coils of the scanner. The above model takes background field inhomogeneities into account through $\Delta\omega_0(\mathbf{r})$ and nonuniformities in the gradient fields are taken into account through the spatially dependent gradient vector \mathbf{G} . In case the background field can be considered uniform throughout the object, $B_0(\mathbf{r}) = B_0$, $\mathbf{r} \in D$, and the gradient vector does not depend on position, $\mathbf{G}(\mathbf{r}, \tau) = \mathbf{G}(\tau)$, $\mathbf{r} \in D$, our signal model simplifies to



Fig. 1. *Left*: Low-field MR scanner realized at the Leiden University Medical Center, The Netherlands. The scanner consists of 23 double ring cubic neodymium boron iron magnets arranged in a Halbach configuration such that an approximately homogeneous field inside the bore is realized. The total number of magnets is close to 3000. The scanner has a length of 50.6 cm and the bore diameter is 27 cm. *Right*: A gradient coil realized at the Leiden University Medical Center and Delft University of Technology, The Netherlands, that produces a linear gradient along the longitudinal direction of the Halbach scanner shown on the left.

$$v(t) = \tilde{X}[\mathbf{k}(t)] = \int_{\mathbf{r} \in \mathbb{D}} e^{-i2\pi\mathbf{k}(t) \cdot \mathbf{r}} X(\mathbf{r}) dV, \quad (3)$$

where we have taken $\omega_{\text{mod}} = \gamma B_0$ and

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_{\tau=0}^t \mathbf{G}(\tau) d\tau. \quad (4)$$

Clearly, in this case the signal $v(t)$ is a three-dimensional spatial Fourier transform $\tilde{X}[\mathbf{k}(t)]$ of the image $X(\mathbf{r})$.

The data used for imaging are time samples of the voltage signal. Introducing the time instances $t_n = (n-1)\Delta t$ for $n = 1, 2, \dots, N$ with $\Delta t > 0$ and taking into account that the image X has the domain \mathbb{D} as its support, we may write

$$d(t_n) = \int_{\mathbf{r} \in \mathbb{D}_{\text{im}}} e^{-i\Delta\omega_0(\mathbf{r})t_n} e^{-i2\pi\mathbf{k}(\mathbf{r}, t_n) \cdot \mathbf{r}} X(\mathbf{r}) dV, \quad (5)$$

for $n = 1, 2, \dots, N$ and where \mathbb{D}_{im} is the smallest cube that completely surrounds the domain \mathbb{D} such that X vanishes on its boundary. We refer to \mathbb{D}_{im} as the imaging domain or Field of View (FOV). The side length of this domain is denoted by $\ell > 0$ and \mathbb{D}_{im} has a volume $V = \ell^3$.

Before discretizing the volume integral in the above signal model, it is convenient to first normalize the spatial coordinates. Specifically, writing $\mathbf{r} = \ell\mathbf{r}'$ with $\mathbf{r}' \in \mathbb{D}_{\text{u}}$ and $\mathbb{D}_{\text{u}} = [0, 1]^3$ the unit cube, we have

$$d(t_n) = V \int_{\mathbf{r}' \in \mathbb{D}_{\text{u}}} e^{-i\Delta\omega_0(\mathbf{r}')t_n} e^{-i2\pi\mathbf{k}'(\mathbf{r}', t_n) \cdot \mathbf{r}'} X'(\mathbf{r}') dV', \quad (6)$$

where $\Delta\omega_0(\mathbf{r}') = \Delta\omega_0(\ell\mathbf{r}')$, $\mathbf{k}'(\mathbf{r}', t_n) = \ell\mathbf{k}(\ell\mathbf{r}', t_n)$ is the scaled generalized Fourier vector, and $X'(\mathbf{r}') = X(\ell\mathbf{r}')$. From now on we drop the primes and work with normalized spatial coordinates only.

Subsequently, we discretize the unit cube by subdividing this domain into nonoverlapping voxels, where each voxel has a positive side length Δx , Δy , and Δz in the x -, y -, and z -direction, respectively, and such that $P\Delta x = Q\Delta y = R\Delta z = 1$ with P , Q , and R positive integers. The volume of a voxel is denoted by $V_{\Delta} = \Delta x\Delta y\Delta z$ and the centers of the voxels have coordinates (x_p, y_q, z_r) with

$$\begin{aligned} x_p &= \Delta x/2 + (p-1)\Delta x, \\ y_q &= \Delta y/2 + (q-1)\Delta y, \quad \text{and} \\ z_r &= \Delta z/2 + (r-1)\Delta z, \end{aligned} \quad (7)$$

for $p = 1, 2, \dots, P$, $q = 1, 2, \dots, Q$, and $r = 1, 2, \dots, R$. Having introduced the voxelization of our imaging domain, we discretize the integral in (6) using the composite midpoint rule and obtain

$$V^{-1}\mathbf{d} = \mathbf{A}\mathbf{x}, \quad (8)$$

where

$$\mathbf{d} = [d(t_1), d(t_2), \dots, d(t_N)]^T, \quad (9)$$

and \mathbf{A} is the matrix representation of the integral operator in Eq. (6). In case of uniform background and gradient fields and $N = PQR$, we have $\mathbf{A} = V_{\Delta}\mathbf{F}$, where \mathbf{F} the 3D Discrete Fourier Transform (DFT) matrix that satisfies

$$\mathbf{F}^H\mathbf{F} = V_{\Delta}^{-1}\mathbf{I}, \quad (10)$$

from which it follows that in this case $\mathbf{A}^H = V_{\Delta}\mathbf{F}^H = \mathbf{F}^{-1}$.

In practice, the measured data contains noise and the data $\tilde{\mathbf{d}}$ that we really have available is assumed to be of the form $\tilde{\mathbf{d}} = \mathbf{d} + \mathbf{n}$, where \mathbf{n} is a noise vector. Now in high-field MRI, the noise level is much lower than in low-field MRI and as a first approximation the background field and gradient fields can be taken to be uniform (independent of the position vector), especially compared with a low-field setting. Simply applying an inverse 3D-Fourier transform to the measured signal generally yields images of excellent quality in high-field MRI. In a low-field setting, however, the background and gradient fields may not be uniform within the complete object and, as mentioned earlier, the required signals typically have a much lower SNR. Consequently, even when the background and gradient fields can be considered uniform, a straightforward application of an inverse Fourier transform to the measured data will lead to very noisy reconstructions in general.

To address both of these issues, we pose the imaging reconstruction problem as an optimization problem and minimize an objective function that consists of a data fidelity term describing the mismatch between observed and modeled data and a regularizing term, which suppresses the influence of noise on the reconstructions. In particular, in many regularized optimization methods the objective function that needs to be minimized is of the form

$$F(\mathbf{x}) = F^{\text{data}}(\mathbf{x}) + \lambda F^{\text{reg,add}}(\mathbf{x}), \quad (11)$$

where $F^{\text{data}}(\mathbf{x})$ is the data fidelity term, $F^{\text{reg,add}}(\mathbf{x})$ some regularizing function (often chosen to be a variant of the total variation (TV) operator), and λ is a regularization parameter. The main drawback of using objective functions of the form (11) is that to reconstruct an image, the artificial regularization parameter λ needs to be selected for each new data set that is acquired. Strategies for choosing this parameter exist, of course (such as the L-curve method [23]), but such approaches are computationally demanding and generally require extensive numerical experimentation for each new available data set.

Instead of using an additive approach, we follow [15,16,18,20], for example, and set up an iterative reconstruction algorithm that is based on multiplicative regularization in which at each iteration the objective function is of the form

$$F(\mathbf{x}) = F^{\text{data}}(\mathbf{x})F^{\text{reg,mult}}(\mathbf{x}). \quad (12)$$

Here, $F^{\text{reg,mult}}$ is the regularizing function, which is, in general, of a different form than the regularizer in Eq. (11). Specifically, for multiplicative regularization, it is usually defined such that $F^{\text{reg,mult}}(\mathbf{x}) = 1$ at an optimal point. A major advantage of such a multiplicative approach is that no extra computations are required to determine an effective regularization parameter. Here, we focus on the application of multiplicative regularization to invert low-field MR data. For theoretical properties of multiplicative regularization the reader is referred to [16,24,25], for example.

As a first step we introduce the standard 2-norm data misfit objective function as

$$F^{\text{data}}(\mathbf{x}) = \frac{\|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2}{\|\mathbf{b}\|_2^2}, \quad (13)$$

where $\mathbf{b} = V^{-1}\tilde{\mathbf{d}}$ is the scaled and noisy data vector. Subsequently, we set up our iterative scheme and assume that at the k th iteration we have some approximation \mathbf{x}_{k-1} of the image available. The next iterate \mathbf{x}_k is now constructed by minimizing the objective function

$$F_{k-1}(\mathbf{x}) = F^{\text{data}}(\mathbf{x})F_{k-1}^{\text{TV}}(\mathbf{x}), \quad (14)$$

where $F_{k-1}^{\text{TV}}(\mathbf{x})$ is the discretized counterpart of the weighted L^2 -norm total variation functional

$$\mathcal{F}_{k-1}^{\text{TV}}(X) = \int_{\mathbf{r} \in \mathbb{D}_{\text{u}}} \frac{|\nabla X|^2 + \delta_{k-1}^2}{|\nabla X_{k-1}|^2 + \delta_{k-1}^2} dV. \quad (15)$$

Here, X and X_{k-1} are the continuous counterparts of \mathbf{x} and \mathbf{x}_{k-1} , respectively, and δ_{k-1}^2 is given by

$$\delta_{k-1}^2 = F^{\text{data}}(\mathbf{x}_{k-1})^2 \int_{\mathbf{r} \in \mathbb{D}_{\text{u}}} |\nabla X_{k-1}|^2 dV. \quad (16)$$

In our algorithm, we only work with this particular choice for δ_{k-1}^2 , but other choices are possible as well, see [15,17]. In each iteration, Eq. (14) is minimized. This choice of $F^{\text{TV}}(\mathbf{x})$ promotes solutions that are piecewise constant. Note that if this process converges, we have $\lim_{k \rightarrow \infty} \mathbf{x}_{k-1} = \mathbf{x}$ and hence $\lim_{k \rightarrow \infty} \nabla \mathbf{x}_{k-1} = \nabla \mathbf{x}$, which means that $\lim_{k \rightarrow \infty} F_{k-1}^{\text{TV}}(\mathbf{x}) = 1$.

For ill-posed problems, as encountered in wavefield imaging and contrast source inversion (CSI), for example, the total objective function typically shows a convergence behavior as described in [15,16]. Specifically, in CSI the data functional is typically large, at the start of the iterative process, giving a large weight to the regularizing functional. As the iterative process progresses, the regularizing function will start

to approach a value of 1 and the focus will shift towards minimizing the data functional, as illustrated in [26]. However, here we consider imaging problems which are well-posed (Fourier transform) or at least less ill-posed than in CSL, which means that, given a reasonable initial guess, the data functional will be small. Therefore, we cannot expect the same behavior, except that if the process converges, the regularizing functional will converge to 1.

2.2. Numerical discretization

To arrive at the discretized counterpart $F_{k-1}(\mathbf{x})$ of $\mathcal{F}_{k-1}(X)$, we use the weighting function

$$w_{k-1}(\mathbf{r}) = \frac{1}{|\nabla X_{k-1}|^2 + \delta_{k-1}^2} \quad (17)$$

and write Eq. (15) as

$$\mathcal{F}_{k-1}^{\text{TV}}(X) = \delta_{k-1}^2 \int_{\mathbf{r} \in \mathbb{D}_U} w_{k-1}(\mathbf{r}) dV + \int_{\mathbf{r} \in \mathbb{D}_U} w_{k-1}(\mathbf{r}) |\nabla X|^2(\mathbf{r}) dV, \quad (18)$$

where $\nabla = \left(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z} \right)$ is the nabla operator. Discretizing the integrals in (18) using the composite midpoint rule, we obtain

$$\begin{aligned} \mathcal{F}_{k-1}^{\text{TV}}(X) \approx & \delta_{k-1}^2 V_\Delta \sum_{p=1}^P \sum_{q=1}^Q \sum_{r=1}^R w_{k-1}(x_p, y_q, z_r) \\ & + V_\Delta \sum_{p=1}^P \sum_{q=1}^Q \sum_{r=1}^R w_{k-1}(x_p, y_q, z_r) |\nabla X|^2(x_p, y_q, z_r). \end{aligned} \quad (19)$$

The partial derivatives in the gradient operator on the right-hand side of Eq. (19) are approximated by two-point forward or backward difference formulas. For example, for the partial derivative with respect to x we either use the forward or backward difference formulas

$$\begin{aligned} \frac{\partial X}{\partial x}(x_p, y_q, z_r) & \approx \frac{X(x_{p+1}, y_q, z_r) - X(x_p, y_q, z_r)}{\Delta x} \quad \text{or} \\ \frac{\partial X}{\partial x}(x_p, y_q, z_r) & \approx \frac{X(x_p, y_q, z_r) - X(x_{p-1}, y_q, z_r)}{\Delta x} \end{aligned} \quad (20)$$

and similar forward or backward difference formulas are used for the partial derivatives with respect to the y - and z -coordinate. The difference formulas are used for $p = 1, 2, \dots, P, q = 1, 2, \dots, Q$, and $r = 1, 2, \dots, R$ along with the homogeneous boundary conditions

$$X(x_p, y_q, z_r) = 0 \quad (21)$$

whenever $p = 0, p = P + 1, q = 0, q = Q + 1, r = 0$, or $r = R + 1$ in Eq. (21). Through the latter equation we implement the boundary condition that the image X vanishes at the boundary of the imaging domain.

Moreover, we use the forward and backward differencing matrices in the x, y and z directions, defined as

$$\begin{aligned} \mathbf{D}_{\xi;f} &= \frac{1}{\Delta \xi} \begin{pmatrix} -1 & 1 & & & \\ & -1 & 1 & & \\ & & \ddots & \ddots & \\ & & & -1 & 1 \\ & & & & -1 \end{pmatrix} \text{ and} \\ \mathbf{D}_{\xi;b} &= \frac{1}{\Delta \xi} \begin{pmatrix} 1 & & & & \\ -1 & 1 & & & \\ & \ddots & \ddots & & \\ & & & -1 & 1 \\ & & & & -1 \end{pmatrix}, \end{aligned} \quad (22)$$

respectively, with $\xi = \{x, y, z\}$. The dimensions of the matrices are $P \times P$ for $\mathbf{D}_{x;f}$ and $\mathbf{D}_{x;b}$, $Q \times Q$ for $\mathbf{D}_{y;f}$ and $\mathbf{D}_{y;b}$ and $R \times R$ for $\mathbf{D}_{z;f}$ and $\mathbf{D}_{z;b}$.

Denoting the 3D image array by X , we introduce its vectorized counterpart as $\mathbf{x} = \text{vec}(X)$, where we use lexicographical ordering. Similarly, \mathbf{w}_{k-1} represents the vector form of Eq. (17).

Having introduced these vectors and matrices, we can write Eq. (19) more compactly as

$$\begin{aligned} \mathcal{F}_{k-1}^{\text{TV}}(X) & \approx \delta_{k-1}^2 V_\Delta \mathbf{e}^T \mathcal{W}_{k-1} \mathbf{e} \\ & + V_\Delta (\mathbf{g}_x^H \mathcal{W}_{k-1} \mathbf{g}_x + \mathbf{g}_y^H \mathcal{W}_{k-1} \mathbf{g}_y + \mathbf{g}_z^H \mathcal{W}_{k-1} \mathbf{g}_z) \\ & = \delta_{k-1}^2 V_\Delta \mathbf{e}^T \mathcal{W}_{k-1} \mathbf{e} + V_\Delta \sum_{i=x,y,z} \mathbf{g}_i^H \mathcal{W}_{k-1} \mathbf{g}_i, \end{aligned} \quad (23)$$

where \mathbf{e} is the $PQR \times 1$ vector of ones, $\mathcal{W}_{k-1} = \text{diag}(\mathbf{w}_{k-1})$, and

$$\begin{aligned} \mathbf{g}_x &= (\mathbf{I}_R \otimes \mathbf{I}_Q \otimes \mathbf{D}_x) \mathbf{x}, \\ \mathbf{g}_y &= (\mathbf{I}_R \otimes \mathbf{D}_y \otimes \mathbf{I}_P) \mathbf{x}, \quad \text{and} \\ \mathbf{g}_z &= (\mathbf{D}_z \otimes \mathbf{I}_Q \otimes \mathbf{I}_P) \mathbf{x} \end{aligned} \quad (24)$$

with \mathbf{D}_ξ a forward or backward differencing matrix and $\mathbf{I}_p, \mathbf{I}_q$ and \mathbf{I}_r being the identity matrices of order P, Q and R , respectively, and \otimes denotes the Kronecker product. Finally, substituting the gradient vectors of Eq. (24) into Eq. (23) gives

$$\mathcal{F}_{k-1}^{\text{TV}}(X) \approx F_{k-1}^{\text{TV}}(\mathbf{x}) := \delta_{k-1}^2 V_\Delta \mathbf{e}^T \mathcal{W}_{k-1} \mathbf{e} + V_\Delta \mathbf{x}^H \mathbf{L}_w \mathbf{x}, \quad (25)$$

where \mathbf{L}_w is a three-dimensional weighted approximate Laplacian given by

$$\begin{aligned} \mathbf{L}_w &= (\mathbf{I}_R \otimes \mathbf{I}_Q \otimes \mathbf{D}_x^T) \mathcal{W}_{k-1} (\mathbf{I}_R \otimes \mathbf{I}_Q \otimes \mathbf{D}_x) \\ &+ (\mathbf{I}_R \otimes \mathbf{D}_y^T \otimes \mathbf{I}_P) \mathcal{W}_{k-1} (\mathbf{I}_R \otimes \mathbf{D}_y \otimes \mathbf{I}_P) \\ &+ (\mathbf{D}_z^T \otimes \mathbf{I}_Q \otimes \mathbf{I}_P) \mathcal{W}_{k-1} (\mathbf{D}_z \otimes \mathbf{I}_Q \otimes \mathbf{I}_P). \end{aligned} \quad (26)$$

In practice, forward finite differences or backward finite differences can be used to implement the total variation functional. Another option is to use mixed finite differences, which combines these two.

2.2.1. Mixed finite difference approach

As mentioned before, the partial derivatives in Eq. (19) are approximated by forward or backward two-point finite difference formulas and the differentiation matrices $\mathbf{D}_{x,y,z}$ in the above formulas are either all forward differentiation matrices $\mathbf{D}_{x,y,z;f}$ or backward differentiation matrices $\mathbf{D}_{x,y,z;b}$. Another option is to mix the forward and backward differencing operators. Specifically, introducing the forward and backward x -coordinate gradient vectors as

$$\begin{aligned} \mathbf{g}_{x;f} &= (\mathbf{I}_R \otimes \mathbf{I}_Q \otimes \mathbf{D}_{x;f}) \mathbf{x} \quad \text{and} \\ \mathbf{g}_{x;b} &= (\mathbf{I}_R \otimes \mathbf{I}_Q \otimes \mathbf{D}_{x;b}) \mathbf{x} \end{aligned} \quad (27)$$

with similar definitions for the forward and backward y - and z -coordinate gradient vectors, we approximate $\mathcal{F}_{k-1}^{\text{TV}}(X)$ as

$$\begin{aligned} \mathcal{F}_{k-1}^{\text{TV}}(X) \approx & \delta_{k-1}^2 V_\Delta \mathbf{e}^T \mathcal{W}_{k-1} \mathbf{e} \\ & + \frac{V_\Delta}{2} \sum_{i=x,y,z} \mathbf{g}_{i;b}^H \mathcal{W}_{k-1} \mathbf{g}_{i;b} + \mathbf{g}_{i;f}^H \mathcal{W}_{k-1} \mathbf{g}_{i;f} \end{aligned} \quad (28)$$

leading to the mixed finite difference Laplacian

$$\begin{aligned} \mathbf{L}_w &= \frac{1}{2} (\mathbf{I}_R \otimes \mathbf{I}_Q \otimes \mathbf{D}_{x;b}^T) \mathcal{W}_{k-1} (\mathbf{I}_R \otimes \mathbf{I}_Q \otimes \mathbf{D}_{x;b}) \\ &+ \frac{1}{2} (\mathbf{I}_R \otimes \mathbf{I}_Q \otimes \mathbf{D}_{x;f}^T) \mathcal{W}_{k-1} (\mathbf{I}_R \otimes \mathbf{I}_Q \otimes \mathbf{D}_{x;f}) \\ &+ \frac{1}{2} (\mathbf{I}_R \otimes \mathbf{D}_{y;b}^T \otimes \mathbf{I}_P) \mathcal{W}_{k-1} (\mathbf{I}_R \otimes \mathbf{D}_{y;b} \otimes \mathbf{I}_P) \\ &+ \frac{1}{2} (\mathbf{I}_R \otimes \mathbf{D}_{y;f}^T \otimes \mathbf{I}_P) \mathcal{W}_{k-1} (\mathbf{I}_R \otimes \mathbf{D}_{y;f} \otimes \mathbf{I}_P) \\ &+ \frac{1}{2} (\mathbf{D}_{z;b}^T \otimes \mathbf{I}_Q \otimes \mathbf{I}_P) \mathcal{W}_{k-1} (\mathbf{D}_{z;b} \otimes \mathbf{I}_Q \otimes \mathbf{I}_P) \\ &+ \frac{1}{2} (\mathbf{D}_{z;f}^T \otimes \mathbf{I}_Q \otimes \mathbf{I}_P) \mathcal{W}_{k-1} (\mathbf{D}_{z;f} \otimes \mathbf{I}_Q \otimes \mathbf{I}_P). \end{aligned} \quad (29)$$

Extensive numerical testing has shown that the mixed finite-difference approach leads to faster convergence than implementations that use forward or backward difference operators only. Therefore, we use this mixed finite-difference approach in our implementation of multiplicative regularization.

2.3. MR imaging using multiplicative regularization

As mentioned above, at the k th step of the algorithm we assume that we have an approximation \mathbf{x}_{k-1} available. We update the image according to the update formula

$$\mathbf{x}_k = \mathbf{x}_{k-1} + \beta_k \mathbf{d}_k, \quad (30)$$

where \mathbf{d}_k is the Polak-Ribière update direction given by [15]

$$\mathbf{d}_k = \mathbf{g}_k + \frac{\text{Re}[\mathbf{g}_k^H (\mathbf{g}_k - \mathbf{g}_{k-1})]}{\|\mathbf{g}_{k-1}\|_2^2} \mathbf{d}_{k-1} \quad (31)$$

with $\mathbf{d}_0 = \mathbf{0}$ and \mathbf{g}_k the gradient of the objective function

$$F_{k-1}(\mathbf{x}) = F^{\text{data}}(\mathbf{x}) F_{k-1}^{\text{TV}}(\mathbf{x}) \quad (32)$$

evaluated at $\mathbf{x} = \mathbf{x}_{k-1}$. Using the product rule, this gradient is given by

$$\mathbf{g}_k = \mathbf{g}_k^{\text{data}} F_{k-1}^{\text{TV}}(\mathbf{x}_{k-1}) + F^{\text{data}}(\mathbf{x}_{k-1}) \mathbf{g}_k^{\text{TV}} = \mathbf{g}_k^{\text{data}} + F^{\text{data}}(\mathbf{x}_{k-1}) \mathbf{g}_k^{\text{TV}}, \quad (33)$$

where we have used $F_{k-1}^{\text{TV}}(\mathbf{x}_{k-1}) = 1$ and $\mathbf{g}_k^{\text{data}}$ is the gradient of F^{data} at $\mathbf{x} = \mathbf{x}_{k-1}$ given by

$$\mathbf{g}_k^{\text{data}} = -2\|\mathbf{b}\|_2^2 \mathbf{A}^H \mathbf{r}_{k-1} \quad (34)$$

with $\mathbf{r}_{k-1} = \mathbf{b} - \mathbf{A}\mathbf{x}_{k-1}$ the data residual. Furthermore, \mathbf{g}_k^{TV} is the gradient of F_{k-1}^{TV} at $\mathbf{x} = \mathbf{x}_{k-1}$ and using Eq. (25) is easily obtained as

$$\mathbf{g}_k^{\text{TV}} = 2V_\Delta \mathbf{L}_w \mathbf{x}_{k-1}. \quad (35)$$

The gradient of F_{k-1} at $\mathbf{x} = \mathbf{x}_{k-1}$ now follows as

$$\mathbf{g}_k = 2[-\|\mathbf{b}\|_2^2 \mathbf{A}^H \mathbf{r}_{k-1} + V_\Delta F^{\text{data}}(\mathbf{x}_{k-1}) \mathbf{L}_w \mathbf{x}_{k-1}]. \quad (36)$$

Note that for position independent background and linear gradient fields we have $\mathbf{A}^H = \mathbf{F}^{-1}$, i.e. \mathbf{A}^H is an inverse Fourier Transform.

Finally, to obtain the update coefficient β_k , we substitute $\mathbf{x} = \mathbf{x}_{k-1} + \beta \mathbf{d}_k$ in the objective function $F_{k-1}(\mathbf{x}) = F^{\text{data}}(\mathbf{x}) F_{k-1}^{\text{TV}}(\mathbf{x})$ and determine the update coefficient by solving the equation

$$\left. \frac{\partial F_{k-1}(\mathbf{x}_{k-1} + \beta \mathbf{d}_k)}{\partial \beta} \right|_{\beta=\beta_k} = 0. \quad (37)$$

Explicitly, $F_{k-1}(\mathbf{x}_{k-1} + \beta \mathbf{d}_k) = F^{\text{data}}(\mathbf{x}_{k-1} + \beta \mathbf{d}_k) F_{k-1}^{\text{TV}}(\mathbf{x}_{k-1} + \beta \mathbf{d}_k)$, where

$$F^{\text{data}}(\mathbf{x}_{k-1} + \beta \mathbf{d}_k) = a_0 + a_1 \beta + a_2 \beta^2 \quad \text{and} \\ F_{k-1}^{\text{TV}}(\mathbf{x}_{k-1} + \beta \mathbf{d}_k) = 1 + b_1 \beta + b_2 \beta^2 \quad (38)$$

with

$$a_0 = F^{\text{data}}(\mathbf{x}_{k-1}), \quad a_1 = -2 \frac{\text{Re}(\mathbf{r}_{k-1}^H \mathbf{A} \mathbf{d}_k)}{\|\mathbf{b}\|_2^2}, \quad \text{and} \quad a_2 = \frac{\|\mathbf{A} \mathbf{d}_k\|_2^2}{\|\mathbf{b}\|_2^2} \quad (39)$$

and

$$b_1 = 2V_\Delta \text{Re}(\mathbf{x}_{k-1}^H \mathbf{L}_w \mathbf{d}_k) \quad \text{and} \quad b_2 = V_\Delta \mathbf{d}_k^H \mathbf{L}_w \mathbf{d}_k. \quad (40)$$

Note that a_2 and b_2 are always positive. With these results, the update coefficient follows from Eq. (37) as the root for which the polynomial $a_0 b_1 + a_1 + 2(a_0 b_2 + a_1 b_1 + a_2) \beta + 3(a_1 b_2 + a_2 b_1) \beta^2 + 4a_2 b_2 \beta^3$ is minimized. The roots can be found analytically, or using a built-in polynomial root-finding algorithm. To summarize, the overall algorithm is as follows:

LOW-FIELD MULTIPLICATIVELY REGULARIZED MR IMAGING

1. Given an initial guess of the low-field MR image \mathbf{x}_0 .
2. For $k = 1, 2, \dots$,
 - a) Compute the gradient vector \mathbf{g}_k as given by Eq. (36);
 - b) Compute the Polak-Ribière update direction \mathbf{d}_k ;
 - c) Compute the update coefficient β_k ;
 - d) Update the low-field MR image: $\mathbf{x}_k = \mathbf{x}_{k-1} + \beta_k \mathbf{d}_k$.

Remark: Consider the case of position independent background and

linear gradient fields. We then have $\mathbf{A}^H = \mathbf{F}^{-1}$ and the gradient vector becomes

$$\mathbf{g}_k(\mathbf{x}_{k-1}) = 2[-\|\mathbf{b}\|_2^2 \mathbf{F}^{-1} \mathbf{r}_{k-1} + V_\Delta F^{\text{data}}(\mathbf{x}_{k-1}) \mathbf{L}_w \mathbf{x}_{k-1}]. \quad (41)$$

As an initial guess, let us take $\mathbf{x}_0 = \mathbf{F}^{-1} \mathbf{b}$. This initial guess is very noisy, but the data is exactly matched and $F^{\text{data}}(\mathbf{x}_0) = 0$. Consequently, $\mathbf{g}(\mathbf{x}_0) = \mathbf{0}$ and the algorithm stops after one iteration with $\mathbf{x}_1 = \mathbf{x}_0$. In this case of homogeneous background and linear gradient fields, which is often assumed in practice, we use our algorithm as a denoising algorithm by setting the gradient at the k th iteration to $\mathbf{g}(\mathbf{x}_{k-1}) = \mathbf{L}_w \mathbf{x}_{k-1}$. We now update in directions determined by the gradient of the TV-term only. Starting with a masked initial guess $\mathbf{x}_0 = \mathbf{M} \mathbf{F}^{-1} \mathbf{b}$, where the mask \mathbf{M} zeroes out any components of $\mathbf{F}^{-1} \mathbf{b}$ for which it is a priori known that the image at the corresponding location vanishes, $F^{\text{data}}(\mathbf{x}_k)$ will generally increase as k increases and if $\lim_{k \rightarrow \infty} F_{k-1}^{\text{TV}}(\mathbf{x}_k) = 1$ then the objective function $F_{k-1}(\mathbf{x})$ will asymptote to $F^{\text{data}}(\mathbf{x}^*)$ as $k \rightarrow \infty$, where $\mathbf{x}^* = \lim_{k \rightarrow \infty} \mathbf{x}_{k-1} = \lim_{k \rightarrow \infty} \mathbf{x}_k$ is the converged image in which noise variations have been minimized (suppressed).

3. Results

In this section we illustrate the performance of our multiplicatively regularized imaging and denoising algorithm. In Section 3.1 we use the algorithm to reconstruct the well-known 2D Shepp-Logan phantom from simulated and noise-corrupted low-field data with a frequency encoding gradient that is nonlinear. The data matrix \mathbf{A} is not equal to a standard DFT matrix in this case and simply applying an inverse DFT to the data would lead to highly distorted reconstructions. We therefore use the algorithm as an iterative reconstruction algorithm to recover the Shepp-Logan phantom from the data.

Subsequently, in Section 3.2 we apply our algorithm to data as measured by the low-field scanner described in [2]. This scanner has been constructed such that the background and gradient fields are approximately constant and linear throughout the imaging domain. We take this into account in our data model and the data matrix \mathbf{A} is assumed to be equal to a scaled three-dimensional DFT matrix in this case (see Eq. (10)). The algorithm is now used as a denoising algorithm and images of an apple and a melon will be presented to demonstrate the performance of the method on measured data for which a Fourier signal representation is assumed to be applicable. In Section 3.3, we compare our multiplicatively regularized 2D and 3D results with images obtained by solving the additively regularized minimization problem

$$\mathbf{x}^{\text{add}} = \arg \min_{\mathbf{x}} \{ \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda \|\mathbf{T}\mathbf{x}\|_1 \}, \quad (42)$$

where the second term is the total variation term, with \mathbf{T} the anisotropic total variation operator, which we define as

$$\mathbf{T}^{2D} = \begin{pmatrix} \mathbf{I}_Q \otimes \mathbf{D}_{x,f} \\ \mathbf{D}_{y,f} \otimes \mathbf{I}_P \end{pmatrix}, \quad \mathbf{T}^{3D} = \begin{pmatrix} \mathbf{I}_R \otimes \mathbf{I}_Q \otimes \mathbf{D}_{x,f} \\ \mathbf{I}_R \otimes \mathbf{D}_{y,f} \otimes \mathbf{I}_P \\ \mathbf{D}_{z,f} \otimes \mathbf{I}_Q \otimes \mathbf{I}_P \end{pmatrix} \quad (43)$$

for 2D and 3D, respectively. We solve this minimization problem using the Alternating Direction Method of Multipliers (ADMM), see for example [27].

The experiments were carried out in MATLAB R2015a on a desktop PC with an Intel(R) Xeon(R) W-2123 CPU (3.60GHz).

3.1. Two-dimensional imaging of simulated noise-corrupted low-field MR data

In low-field MRI, inhomogeneities may be present in the background field and gradients may not be perfectly linear. To investigate the performance of the algorithm as an image reconstruction algorithm in such cases, we consider a two-dimensional low-field reconstruction problem in which field perturbations are taken into account. Many

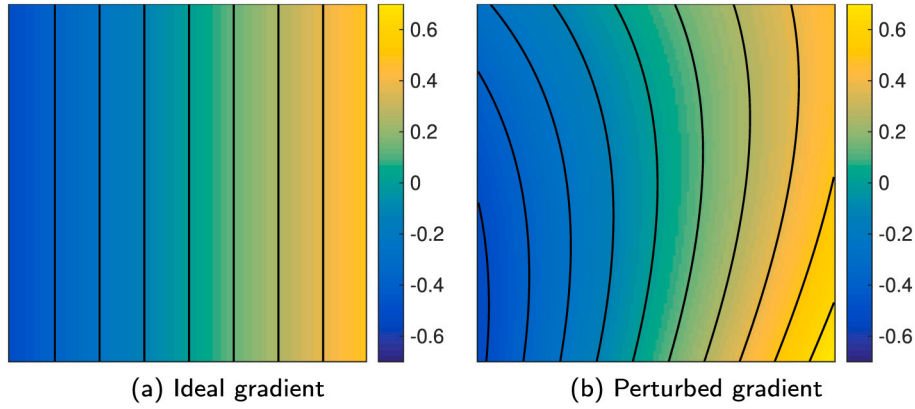


Fig. 2. In an ideal scenario, the frequency encoding gradient is linear (a). In these simulations, however, the gradient is perturbed and nonlinear (b). The black lines are contour lines.

different field perturbations can be considered, of course, but in all cases signals are obtained for which the relationship between the signal and the object is no longer governed by a DFT. To fix the idea, we therefore focus on signal generation in an MR scanner in which the frequency encoding gradient is not perfectly linear. Specifically, we introduce a perturbation of the gradient field that is relative to the gradient strength. To this end, we scale the frequency encoding gradient such that we can describe its strength to be in the range $[-.5, .5]$ and we use \bar{x} and \bar{y} to denote the location in the Field of View (FoV) such that $\bar{x} = -.5$ corresponds to the left and $\bar{x} = .5$ to the right boundary of the FoV, and $\bar{y} = -.5$ and $\bar{y} = .5$ to the lower and upper boundary, respectively. The gradient field is now perturbed by the function $.5\bar{x}\bar{y} + .05\bar{x}^2 + .35\bar{y}^2$. The unperturbed gradient field is shown in Fig. 2a, while the perturbed gradient field is shown in Fig. 2b.

Since the background and (nonlinear) gradient profiles are known, we can use the discretized version of Eq. (5) to obtain the model matrix \mathbf{A} . As an image or model solution, we take the Shepp-Logan phantom of Fig. 3a and we use matrix \mathbf{A} to generate the data. Subsequently, white Gaussian noise with an SNR of 20 is added to the data in the frequency domain. We note that not taking the distorted gradient into account leads to the image shown in Fig. 3b, which is obviously a deformed image of the true object.

Having the noise-corrupted data available, we apply our multiplicative regularization scheme to this data in an attempt to reconstruct the Shepp-Logan model solution. Mixed finite differences and the 2D version of the Laplacian matrix given in Eq. (29) are used to implement the total variation functional. As an initial guess, we use $\mathbf{x}_0 = \alpha \mathbf{A}^H \mathbf{b}$, where α is chosen such that $F^{\text{data}}(\mathbf{x}_0)$ is minimized. This initial guess is depicted in Fig. 3c. Subsequently, we use this \mathbf{x}_0 to define $F_0^{\text{TV}}(\mathbf{x})$, and start iterating. The reconstruction that we obtain after 50 iterations is shown in Fig. 3d. The algorithm's progress is shown in the Appendix, where we can see that the algorithm manages to progressively denoise the image, while maintaining the edges and structures in the phantom.

Fig. 4 shows the value of the objective function $F(\mathbf{x})$ as a function of the iteration number, for the two-dimensional Shepp-Logan image. The values of $F^{\text{TV}}(\mathbf{x})$ and $F^{\text{data}}(\mathbf{x})$ are plotted as well. We note that, as expected, F^{TV} converges to 1, while the other two steadily grow to a larger value that enables this convergence. The increasing value of $F^{\text{data}}(\mathbf{x})$ can be explained by the observation that our matrix \mathbf{A} does not deviate all that much from a Fourier Transform, which means that the term $\|\mathbf{b} - \mathbf{A}\mathbf{x}_0\|^2$ is close to 0. However, since \mathbf{x}_0 is too noisy, this solution does not meet the smoothness requirements of the total variation functional. Instead we iterate towards a solution that minimizes the least-squares term under the constraint that $F^{\text{TV}}(\mathbf{x})$ is equal to 1. One iteration takes 1.79 s, which is relatively long but this can be explained by the fact that we need to explicitly calculate the matrix \mathbf{A} and its transpose, since we cannot rely on the Fast Fourier Transform (FFT) due to the nonlinear gradient.

Additionally, we consider the same Shepp-Logan phantom, but with a decreased SNR of 5. This is more realistic in a low-field MRI setting. The inverse FFT reconstruction, initial guess and the multiplicatively regularized reconstruction are shown in Fig. 5b, c and d, respectively, and for different iteration numbers, the reconstructions are shown in the Appendix. This result was obtained after 50 iterations. We see that the algorithm manages to denoise the image, but some structures are lost. This is not due to the method itself, but because the SNR is simply too low. As will be shown later, an additively regularized method is not able to recover these structures either. The convergence plots are shown in Fig. 6, which showcase the same behavior as in the higher SNR case.

3.2. Three-dimensional imaging of measured data

This section contains our imaging results of an apple and a melon that were scanned using the low-field scanner of [2]. Images of slices through these pieces of fruit will be presented.

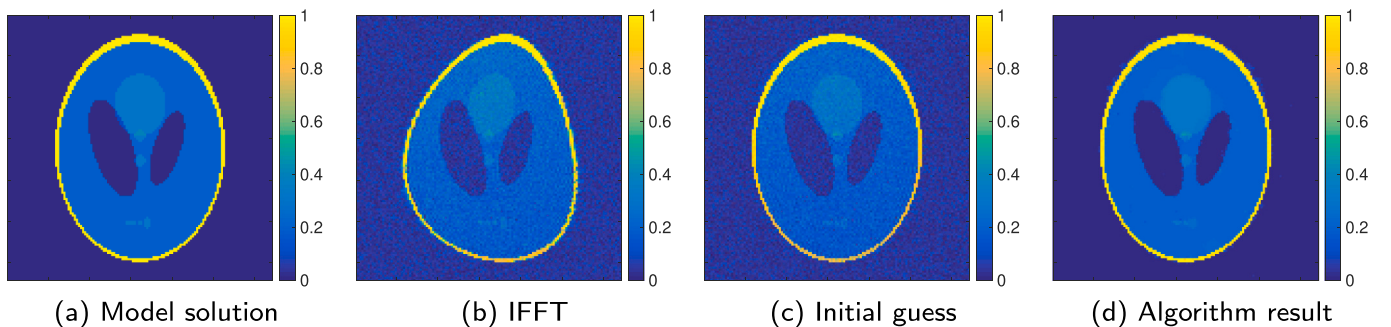


Fig. 3. Reconstruction results.




Fig. 4. Plots of the objective function value $F(x)$ (left), which is equal to the product of $F^{TV}(x)$ (middle) and $F^{data}(x)$ (right), for the two-dimensional Shepp-Logan image.

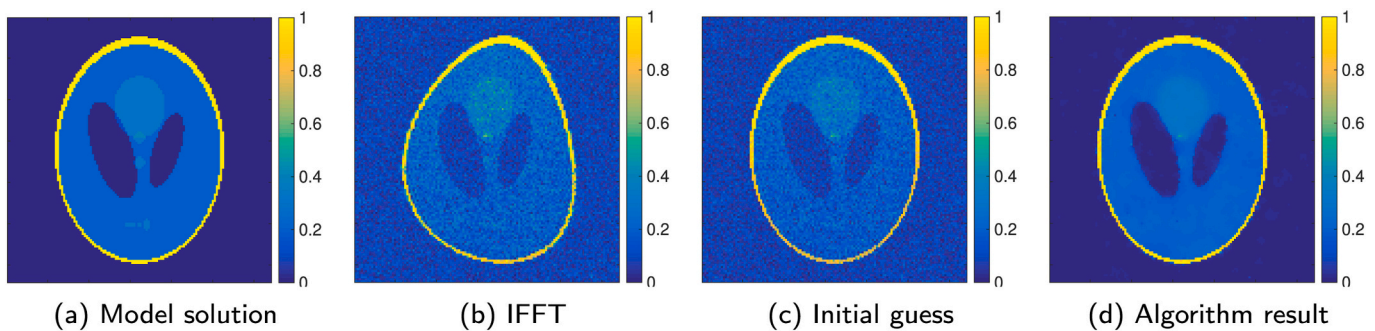


Fig. 5. Reconstruction results for an SNR of 5.




Fig. 6. Plots of the objective function value $F(x)$ (left), which is equal to the product of $F^{TV}(x)$ (middle) and $F^{data}(x)$ (right), for the two-dimensional Shepp-Logan image. In this case, the SNR was lowered to 5.

Table 1

Parameter settings for the apple imaging experiment.

Parameter	Value
Repetition time (TR)	3 s
Echo time (TE)	30 ms
Number of complex data points	$64 \times 64 \times 64$
Imaging domain/FoV	$128 \times 128 \times 128 \text{ mm}^3$
Pulse duration	100 μs
Acquisition bandwidth	10 kHz
Number of averages	1

3.2.1. Apple experiment

In our first experiment, an apple is imaged using a spin-echo sequence whose parameters are given in Table 1. As an initial guess, we

take a masked version of the three-dimensional inverse DFT of the signals obtained during the apple experiment. The 35th slice of this initial guess is shown in Fig. 7a. Clearly, the initial guess is contaminated by noise and the mask that is used is visible. We use this initial guess to determine $F_0^{TV}(x)$, after which we start the iterative process.

To remove the noise from the initial image, we use the algorithm as a denoising algorithm and use a mixed differences approach with the Laplacian matrix of Eq. (29) in the discretized total variation functional. One iteration takes approximately 0.3 s. This is very fast compared to the Shepp-Logan problem, for which the problem is smaller (64×64 pixels instead of $64 \times 64 \times 64$ pixels), which can be explained by the fact that here, the FFT can be employed. We observe that the amount of noise in the image decreases as the iteration process continues, as can




Fig. 7. Reconstruction result of the 35th slice of the apple.




Fig. 8. Plots of the objective function value $F(x)$ (left), which is equal to the product of $F^{TV}(x)$ (middle) and $F^{data}(x)$ (right), for the three-dimensional apple image.

Table 2
Parameter settings for the melon imaging experiment.

Parameter	Value
Repetition time (TR)	2000 ms
Echo time (TE)	30 ms
Number of complex data points	$128 \times 128 \times 128$
Imaging domain/FoV	$200 \times 200 \times 200 \text{ mm}^3$
Pulse duration	$100 \mu\text{s}$
Acquisition bandwidth	20 kHz
Number of averages	1

be seen in Fig. 15 in the Appendix. After 30 iterations an image of good quality is obtained, with the apple's shape and seeds being clearly visible and the noise having been eliminated. Fig. 8 shows the convergence plots for our algorithm. Again, we see that F^{TV} converges to 1. Iterating further leads to a result that is somewhat oversmoothed, which is similar to the 2D results with a low SNR. Therefore, it is advisable to stop the iterative process early (semi-convergence). Finally, we note that the reconstructed images are obtained essentially in real-time and the reconstructions can be monitored as the iterative scheme progresses, which is of great importance in practice when a scanner is used for diagnostic purposes. Fig. 8 shows a three-dimensional visualization of the reconstructed apple.

3.2.2. Melon experiment

In our second experiment, a melon is imaged using a spin-echo sequence with the parameter settings given in Table 2. Here too, we

construct an initial guess by masking the image that is obtained by applying an inverse three-dimensional DFT on the data. Note that this data set has twice as many data points in each Cartesian direction as the data set in the apple experiment. The 64th slice of the initial guess is shown in Fig. 10a. The mask that is used is clearly visible and again a very noisy initial image is obtained. To remove this noise, we use mixed finite differences to implement the total variation functional and the resulting reconstruction of the melon is shown in Fig. 10b. As can be seen in Fig. 16 in the Appendix, the noise level in the images decreases with the iteration number. We note that some individual pixels appear overly bright or dark, which may be removed by iterating further but this leads to oversmoothing, since the SNR is very low. We therefore terminate the iterative process after 20 iterations. Finally, even though the complex data set of the melon is eight times larger than the apple data set, the method still produces images essentially in real time that can be monitored as the scheme progresses. A three-dimensional visualization of the melon is shown in Fig. 9b. Each iteration takes approximately 2 s. In Fig. 11, the convergence plots are shown, which follow the same pattern as before.

3.3. Additive and multiplicative regularization

Finally, in Fig. 12, our multiplicatively regularized reconstructions are shown next to additively regularized reconstructions. We tuned the regularization parameter λ in a heuristic manner such that the background noise disappears, while the different structures are still visible. For all experiments, we used 10 ADMM iterations and within each

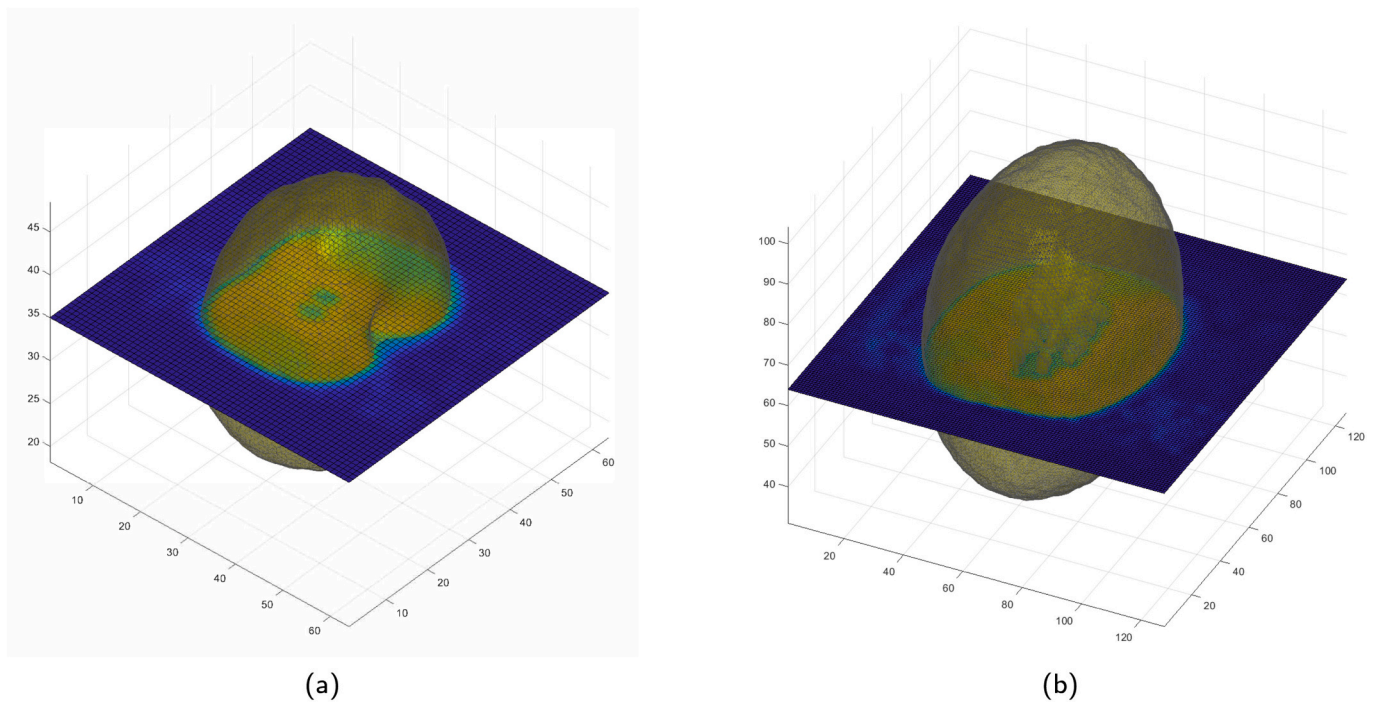


Fig. 9. Three-dimensional visualization of the reconstructed apple (a) and melon (b) obtained after 20 iterations using mixed finite differences.




Fig. 10. Reconstruction result of the 64th slice of the melon.




Fig. 11. Plots of the objective function value $F(x)$ (left), which is equal to the product of $F^{TV}(x)$ (middle) and $F^{data}(x)$ (right), for the three-dimensional melon image.




Fig. 12. Images of the Shepp-Logan phantom using simulated data with an SNR of 20 (top row) and an SNR of 5 (second row), an apple based on measured data (third row), and a melon based on measured data (bottom row), using additive regularization (left column) and multiplicative regularization (right column).

ADMM iteration, we used 10 iterations of the Conjugate Gradient (CG) algorithm to solve the first minimization problem. We observe that, in terms of image quality, additive and multiplicative regularization yield comparable results, but no parameter tuning is required in the multiplicative case. This is reflected in the peak signal-to-noise ratios (PSNR),

shown in Table 3, which we were only able to calculate for the simulated case because no ground truth was available for the measured data. We see that multiplicative regularization achieves a somewhat higher PSNR for low SNR and a marginally lower PSNR for high SNR. Furthermore, in Table 4, the computing times for the 2 different algorithms (our algorithm and the ADMM algorithm solving the additively regularized problem) are shown. We note that it might be possible to lower the computation time for the additively regularized problem by lowering the amount of either ADMM or CG iterations. This table shows that, for our test problems, the computation times for multiplicative regularization are quite competitive compared to additive regularization. Additionally, the computation time shown for additive regularization does not include the time spent searching for an appropriate regularization parameter (requiring repeated execution of the algorithm in general), which is a problem that is completely eliminated when using a multiplicatively regularized approach.

4. Conclusions

In this paper we applied a multiplicative regularization approach to low-field MR imaging. By multiplying a least-squares data fidelity function by a regularizing total variation function that is differentiable, we avoid the problem of having to determine and compute a regularization parameter as is required for additive regularization. The resulting multiplicative regularization problem is nonlinear and solved by a nonlinear conjugate gradient scheme with Polak-Ribière update directions. Furthermore, we showed that the algorithm can be used as an image reconstruction and denoising algorithm by applying the method to two-dimensional noise-corrupted MR data obtained with a nonlinear gradient field and to three-dimensional measured data obtained with a low-field Halbach scanner. We demonstrated that multiplicative regularization yields very promising results, converging within a few iterations, whether we are dealing with two-dimensional noisy data for which a Fourier signal representation is no longer valid, or with three-dimensional measured data for which a Fourier transform relationship between signal and image can be assumed. Moreover, accurate reconstructions are obtained essentially in real time in case a Fourier signal model is applicable. We observed that in case of a low SNR, it is better to stop the iterative process somewhat early (semi-convergence), to maintain the edges in the image. For a high SNR, edges are preserved while noise is eliminated, even for a large number of iterations.

In this work, we focused on a low-field MRI setting with one single receiver coil. However, multiplicative regularization may also be applied to multicoil high-field MRI, of course, and we intend to test its performance against more standard high-field reconstruction techniques used today.

Moreover, in future work we will also focus on incorporating measured background and gradient fields of practical low-field MR scanners into the data model that is used in our multiplicative regularization scheme. Obviously, this is particularly important in case the background and gradient fields are such that a standard Fourier transform signal representation is no longer valid, since otherwise significant distortions in the image are obtained as we demonstrated for two-dimensional simulated MR data. Standard FFTs can no longer be used in this case and accurate images may only be obtained by solving the image reconstruction problem as an optimization problem. In our multiplicative scheme, forming matrix-vector products with the model matrix \mathbf{A} is then the main computational bottleneck, since computing the action of the Laplacian of the total variation functional on a vector involves sparse finite difference operators only. Consequently, to reduce the reconstruction time of the method, efficient routines to compute matrix-vector products with the data matrix \mathbf{A} have to be developed, possibly involving nonuniform FFTs. Finally, we intend to include compressed sensing techniques into our multiplicative regularization framework as well, since this may lead to reduced scan times and a reduction of motion artifacts, for example.

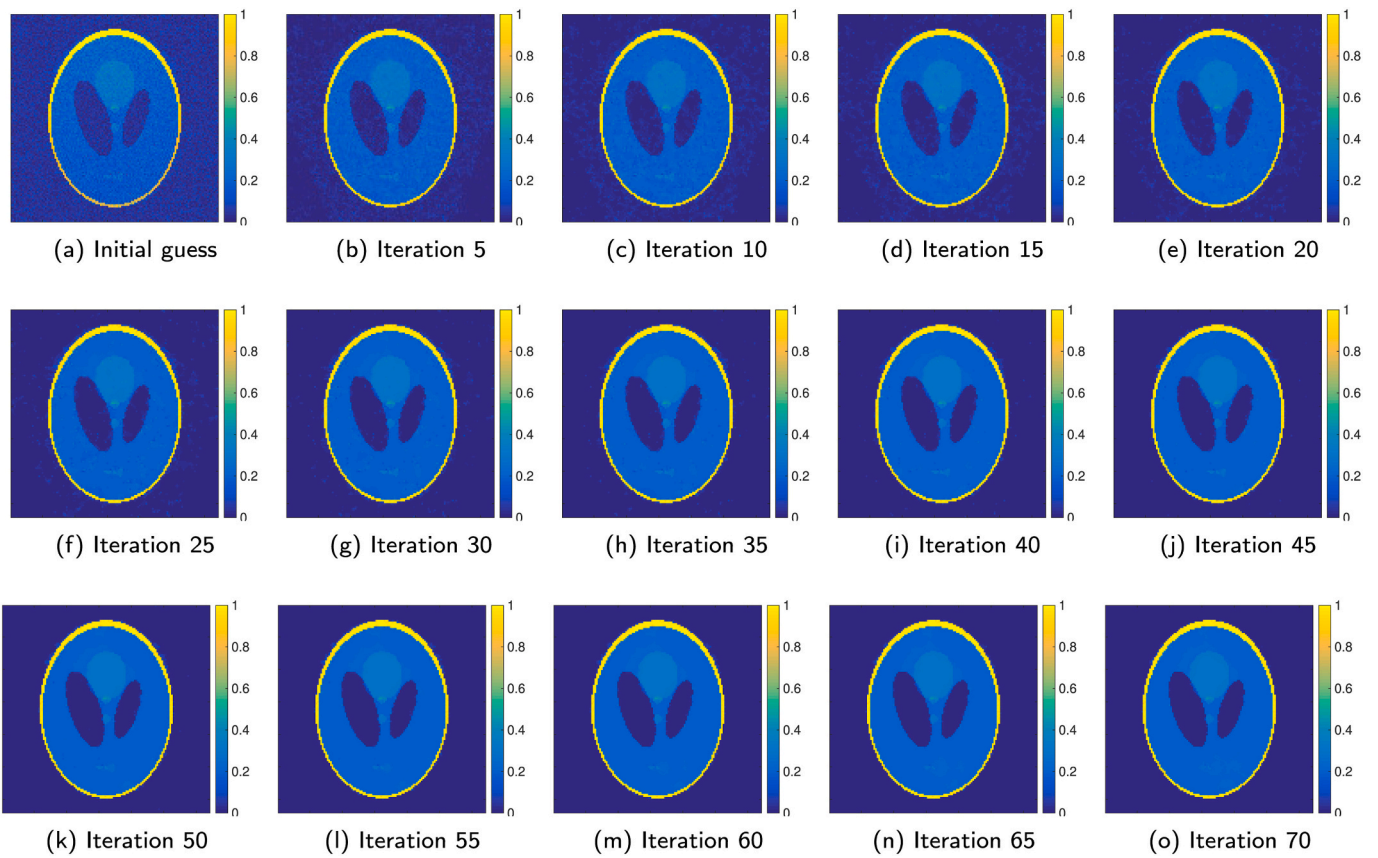


Fig. 13. Reconstruction of the Shepp-Logan phantom, with an SNR of 20, for different iteration numbers.

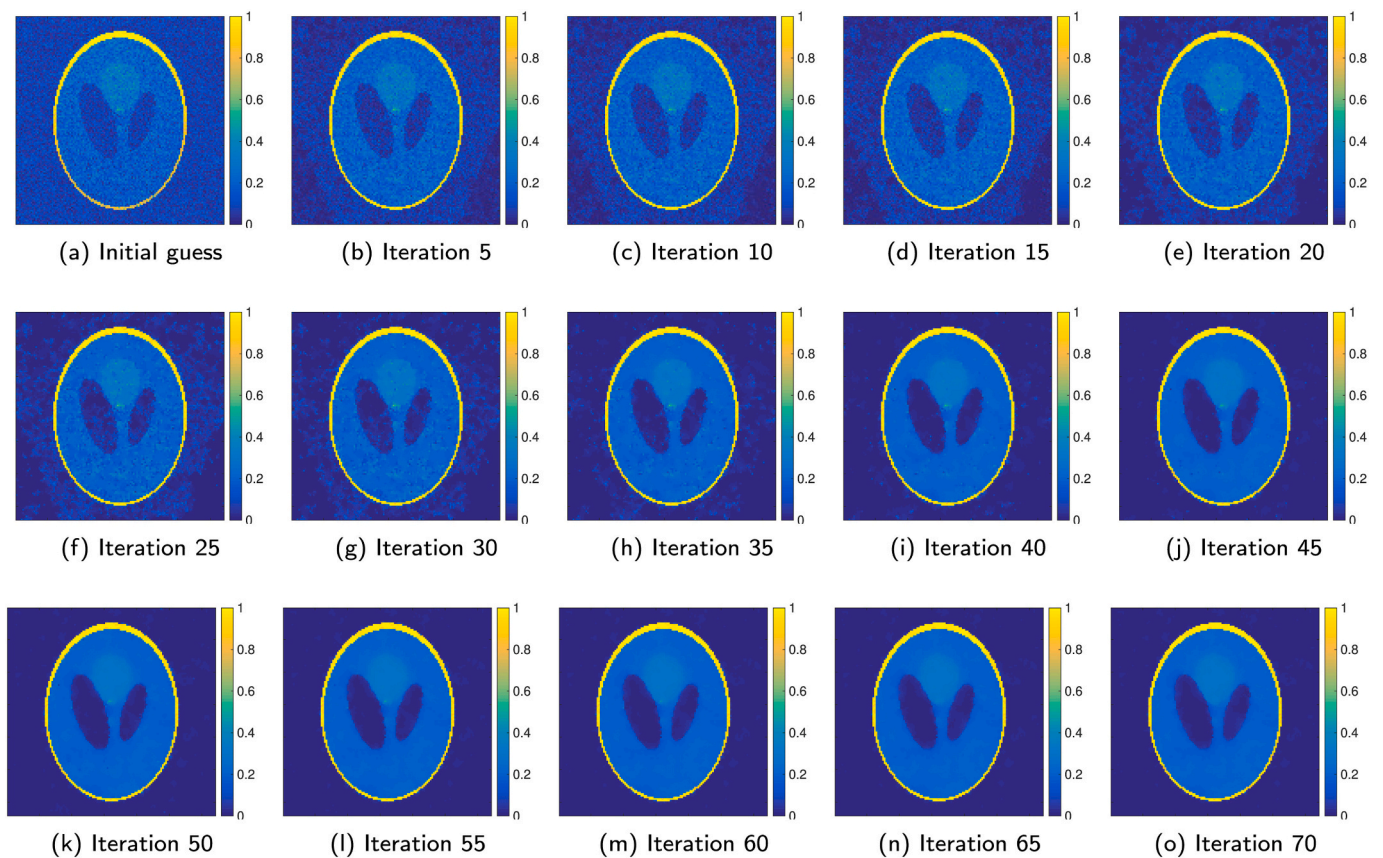


Fig. 14. Reconstruction of the Shepp-Logan phantom, with an SNR of 5, for different iteration numbers.




Fig. 15. Reconstruction of the 35th slice of the apple for different iteration numbers.




Fig. 16. Reconstruction of the 64th slice of the melon for different iteration numbers.

Table 3
PSNR for different reconstructions of the Shepp-Logan phantom.

	Multiplicative regularization	Additive regularization
SNR 20	37.30	37.34
SNR 5	33.83	30.41

Declaration of competing interest

There are no conflicts of interest to disclose.

Acknowledgments

This research is supported by NWO-WOTRO (Netherlands Organization for Scientific Research) under grant W07.303.101 and by

Table 4
Computing time (and total number of iterations) for the different reconstructions in seconds.

	Multiplicative regularization	Additive regularization
Shepp-Logan (SNR 20)	94.19 (50 iterations)	194.81 (100 iterations)
Shepp-Logan (SNR 5)	89.87 (50 iterations)	193.29 (100 iterations)
Apple	5.42 (30 iterations)	6.3 (100 iterations)
Melon	36.50 (20 iterations)	81.20 (100 iterations)

the TU Delft | Global Initiative, a program of the Delft University of Technology to boost Science and Technology for Global Development. The authors would like to thank Tom O'Reilly for carrying out the experiments and all low-field MRI team members of the Leiden University Medical Center, the Electronic and Mechanical Support Division

(DEMO) in Delft, Pennsylvania State University and Mbarara University of Science and Technology for their insight and expertise.

Appendix A. Increasing the number of iterations

In Figs. 13 and 14, the Shepp-Logan reconstructions (for an SNR of 20 and an SNR of 5) are shown for a larger number of iterations. We see that in the low SNR case, the result becomes very blurry and we lose all the small structures in the image if we iterate too long. In the case of a higher SNR of 20, some oversmoothing is taking place when the iteration number exceeds 50 but much less severely than in the low SNR case. We see something similar happening in Figs. 15 and 16, where the apple and melon reconstructions are shown for a larger number of iterations. We see that, for the apple experiment, increasing the number of iterations to 45 yields an image that is slightly blurrier than the image obtained after 30 iterations, which is the result that is shown in Fig. 7. However, for the melon, we see that when the number of iterations exceeds approximately 20, the amount of blur in the image becomes excessive.

These observations suggest that stopping the algorithm before full convergence is attained yields a solution of better quality. This is especially important in case of a low SNR.

References

- [1] Cooley CZ, Stockmann JP, Armstrong BD, Sarraçanie M, Lev MH, Rosen MS, et al. Two-dimensional imaging in a lightweight portable MR scanner without gradient coils. *Magn Reson Med* 2015;73:872–83.
- [2] O'Reilly T, Teeuwisse WM, Webb AG. Three-dimensional MRI in a homogenous 27 cm diameter bore Halbach array magnet. *J Magn Reson* 2019;307:106578.
- [3] Obungoloch J, Harper JR, Consevage S, Savukov IM, Neuberger T, Tadigadapa S, et al. Design of a sustainable prepolarizing magnetic resonance imaging system for infant hydrocephalus. *MAGMA* 2018;31:665–76.
- [4] Marques JP, Simonis FF, Webb AG. Low-field MRI: an MR physics perspective. *J Magn Reson Imaging* 2019;49:1528–42.
- [5] O'Reilly T, Teeuwisse W, de Gans D, Koolstra K, Webb A. In vivo three-dimensional brain and extremity MRI at 50 mT using a permanent magnet Halbach array arXiv preprint arXiv:2005.02834; 2020.
- [6] Fessler JA. Optimization methods for magnetic resonance image reconstruction: key models and optimization algorithms. *IEEE Signal Process Mag* 2020;37:33–40.
- [7] Fessler JA. Model-based image reconstruction for MRI. *IEEE Signal Process Mag* 2010;27:81–9.
- [8] de Leeuw den Bouter ML, van Gijzen MB, Remis RF. Conjugate gradient variants for ℓ_p -regularized image reconstruction in low-field MRI. *SN Appl Sci* 2019;1:1736.
- [9] Janke A, Zhao H, Cowin GJ, Galloway GJ, Doddrell DM. Use of spherical harmonic deconvolution methods to compensate for nonlinear gradient effects on MRI images. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 2004;52:115–22.
- [10] Block KT, Uecker M, Frahm J. Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint. *Magn. Reson. Med. Offi. J.Int. Soc. Magn. Reson. Med.* 2007;57:1086–98.
- [11] Koolstra K, Van Gemert J, Börnert P, Webb A, Remis R. Accelerating compressed sensing in parallel imaging reconstructions using an efficient circulant preconditioner for cartesian trajectories. *Magn Reson Med* 2019;81:670–85.
- [12] Liu J, Sun Y, Xu X, Kamilov US. Image restoration using total variation regularized deep image prior. *ICASSP 2019–2019 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE; 2019. p. 7715–9.
- [13] Ma S, Yin W, Zhang Y, Chakraborty A. An efficient algorithm for compressed MR imaging using total variation and wavelets. *2008 IEEE conference on computer vision and pattern recognition*. IEEE; 2008. p. 1–8.
- [14] Hansen PC, Nagy JG, O'Leary DP. *Deblurring images: matrices, spectra, and filtering*. Philadelphia: Society for Industrial and Applied Mathematics; 2006.
- [15] van den Berg PM, Abubakar A, Fokkema JT. Multiplicative regularization for contrast profile inversion. *Radio Sci* 2003;38:23.1–23.10.
- [16] van den Berg PM, Abubakar A. Contrast source inversion method: state of art. *Pr Electromagn Res* 2001;34:189–218.
- [17] van den Berg PM, Abubakar A. Optical microscopy imaging using the contrast source inversion method. *J Mod Optic* 2010;57:756–64.
- [18] Balidemaj E, van den Berg CA, Trinks J, van Lier AL, Nederveen AJ, Stalpers LJ, et al. CSI-EPT: a contrast source inversion approach for improved MRI-based electric properties tomography. *IEEE T Med Imaging* 2015;34:1788–96.
- [19] Nocedal J, Wright SJ. *Numerical optimization*. Springer; 2006.
- [20] Abubakar A, Van Den Berg PM, Habashy TM, Braunsch H. A multiplicative regularization approach for deblurring problems. *IEEE Trans Image Process* 2004;13:1524–32.
- [21] de Vos B, Fuchs P, O'Reilly T, Webb A, Remis R. Gradient coil design and realization for a halbach-based MRI system. *IEEE T Magn* 2020;56:1–8.
- [22] Liang Z-P, Lauterbur PC. *Principles of magnetic resonance imaging. A signal processing perspective*, IEEE Press series in biomedical engineering IEEE Press; 2000.
- [23] Hansen PC. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Rev* 1992;34:561–80.
- [24] van den Berg PM, van Broekhoven AL, Abubakar A. Extended contrast source inversion. *Inverse Probl* 1999;15:1325–44.
- [25] Gorgin E. An analysis of multiplicative regularization, dissertation. Michigan Technological University; 2015.
- [26] Haffinger PR. Seismic broadband full waveform inversion by shot/receiver refocusing, Ph.D. thesis. Delft University of Technology; 2013.
- [27] Bouman CA. *Model based image processing*. Purdue University; 2013.