



Published in final edited form as:

Magn Reson Med. 2018 January ; 79(1): 83–96. doi:10.1002/mrm.26639.

Low Rank Alternating Direction Method of Multipliers Reconstruction for MR Fingerprinting

Jakob Assländer^{1,2,3}, Martijn A Cloos^{1,2}, Florian Knoll^{1,2}, Daniel K Sodickson^{1,2}, Jürgen Hennig³, and Riccardo Lattanzi^{1,2}

¹Bernard and Irene Schwartz Center for Biomedical Imaging, Department of Radiology, New York University School of Medicine, New York, NY, USA

²Center for Advanced Imaging Innovation and Research (CAI²R), Department of Radiology, New York University School of Medicine, New York, NY, USA

³Dept. of Radiology, Medical Physics, Medical Center – University of Freiburg, Faculty of Medicine, University of Freiburg, Freiburg, Germany

Abstract

Purpose—The proposed reconstruction framework addresses the reconstruction accuracy, noise propagation and computation time for Magnetic Resonance Fingerprinting (MRF).

Methods—Based on a singular value decomposition (SVD) of the signal evolution, MRF is formulated as a low rank inverse problem in which one image is reconstructed for each singular value under consideration. This low rank approximation of the signal evolution reduces the computational burden by reducing the number of Fourier transformations. Also, the low rank approximation improves the conditioning of the problem, which is further improved by extending the low rank inverse problem to an augmented Lagrangian that is solved by the alternating direction method of multipliers (ADMM). The root mean square error and the noise propagation are analyzed in simulations. For verification, in vivo examples are provided.

Results—The proposed low rank ADMM approach shows a reduced root mean square error compared to the original fingerprinting reconstruction, to a low rank approximation alone and to an ADMM approach without a low rank approximation. Incorporating sensitivity encoding allows for further artifact reduction.

Conclusion—The proposed reconstruction provides robust convergence, reduced computational burden and improved image quality compared to other MRF reconstruction approaches evaluated in this study.

Keywords

quantitative MRI; parameter mapping; MRF; SVD; parallel imaging; augmented Lagrangian

Introduction

The dynamics of large spin-1/2 ensembles in a magnetic field are generally described by the Bloch Equation, which captures spin-lattice and spin-spin interactions by relaxation terms with the characteristic time constants T_1 and T_2 , respectively. Quantifying these parameters accurately, reproducibly and within a clinically feasible measurement time is desirable, e.g., for long term studies and for creating synthetic contrasts.

Following departure from thermal equilibrium, otherwise undisturbed magnetization relaxes exponentially, which makes an inversion- or saturation-recovery experiment the natural choice for T_1 -mapping. In order to keep the measurement time within limits, one usually acquires multiple data points during a single relaxation process (1). The same goes for T_2 -mapping, where a multi-spin echo sequence is usually employed. It has also been proposed to map T_1 and T_2 by acquiring multiple images with different flip angles, which proved to be very time efficient (2, 3, 4, 5). However, these fast mapping techniques necessitate multiple interactions with the spin ensemble, which hampers the free relaxation and introduces systematic errors due to magnetization transfer and stimulated echoes amongst other effects.

In recent years, model based parameter mapping has become increasingly popular for multiple reasons. Doneva et al. (6) used a pre-learned orthogonal dictionary to accelerate inversion-recovery and multi-spin echo based T_1 and T_2 mapping experiments. They proposed to reconstruct images from undersampled data by enforcing sparsity in the temporal domain, i.e., they assumed that few entries of the dictionary suffice for describing the signal evolution. Alternatively, model based parameter mapping can be used to remove systematic errors in fast imaging techniques. For example, model based T_2 mapping with multi-spin echo experiments has been shown to increase precision when incorporating signal contributions from stimulated echoes (7, 8).

The model based approach Magnetic Resonance Fingerprinting (MRF) (9) aims at combining acceleration and accuracy. MRF departs from traditional sequence design concepts by deliberately avoiding a steady-state magnetization and relies on the Bloch equations to interpret the measured signal. Without the constraints imposed by a steady-state (or exponential relaxation), MRF opens up a rich new sequence design space that could fundamentally transform the way MR experiments are performed. For example, multiple transmit channels can be incorporated directly into the sequence to enable quantitative imaging with heterogeneous RF fields in high-field MR systems (10) and gradient waveforms can be modulated freely to transform the sound of a traditional MR systems into music (11).

In MRF, the excitation state is, in general, distinct at each TR . Thus, different repetitions cannot easily be combined, making Nyquist-sampled gradient encoding impossible at a clinically feasible spatial resolution and within a reasonable measurement time. The original MRF paper addresses this issue with a two step procedure: First, the undersampled k-space data are projected back into image space, resulting in a series of aliased images. Second, the time series of each voxel is matched to a precomputed dictionary. Under the premises of a

vanishing correlation between the aliasing artifacts in the time series of each voxel and the dictionary, the atom (dictionary entry) with the highest correlation is determined by a brute force search within the dictionary of pre-simulated spin evolutions. The relaxation times are then extracted via a lookup table and the relative proton density is given by the correlation coefficient multiplied by a normalizing factor.

Recent developments suggest that significant improvements in image quality can be gained by feeding the data model to the image reconstruction (12, 13, 14). Due to the non-linear nature of the Bloch equation, the MRF reconstruction problem is generally non-convex. Therefore, global convergence of MRF reconstruction algorithms - including the proposed one - is usually not guaranteed and a good initial guess is essential. The present paper addresses this issue by a low rank (LR) approximation (15, 16, 17, 18, 19, 20, 21, 22) of the voxels' signal evolution, to which the alternating direction method of multipliers (ADMM) (13, 14, 23) is applied. The proposed algorithm splits the MRF reconstruction problem into two sub-problems that are solved alternately: A linear inverse problem with data consistency at its core and a dictionary matching step. Describing the signal evolution with full rank, the inverse problem is under-determined in practical MRF cases. It has been suggested to address this by adding a low rank promoting regularization term (20). Here, we explicitly limit the search to a low rank space, which allows us to formulate an over-determined inverse problem. In the limit of a well-conditioned inverse problem, the first ADMM iteration solves the MRF reconstruction problem. In practice, however, the inverse problem is typically ill-conditioned even in the low rank approximation. In this case, the matched signal provides regularization to the data consistency term in consecutive ADMM iterations, compared to the mere data consistency term. We demonstrate the capabilities and convergence behavior of the proposed algorithm with simulations and experiments under practical scenarios.

Theory

In the following we derive a formalism that describes the entire MRF experiment in a single equation. This might seem a little cumbersome at first, but it facilitates the derivation of the proposed method.

Dynamic MRI Signal Equation

This section describes the signal formation assuming an arbitrary time series of images, which is described by a single vector $\mathbf{x} \in \mathbb{C}^{NT}$, where N denotes the number of voxels and T is the number of time-frames. The time series of images has the following structure:

$$\mathbf{x} = (x_{1,1}, \dots, x_{N,1}, \dots, x_{1,T}, \dots, x_{N,T})' \quad [1]$$

where $'$ denotes a vector transpose. We further allow for the k-space trajectory to vary between different TR s. For simplicity we assume the trajectories have the same length K in each TR , which does not imply a loss of generality. The measured signal is denoted by the vector $\mathbf{S} \in \mathbb{C}^{KT}$. With the formalism of the non-uniform fast Fourier transformation (nuFFT), we describe the measurement process by the simple linear equation

$$\mathbf{S} = \mathbf{G} \cdot \mathbf{F} \cdot \mathbf{x}, \quad [2]$$

where $\mathbf{F} \in \mathbb{C}^{NT \times NT}$ is a block-diagonal matrix containing T blocks on the diagonal, each representing the same Fourier transformation along all considered spatial dimensions (see Eq. [23] in the Appendix). Note that \mathbf{F} can be implemented as an FFT of each time-frame in each spatial dimension. The gridding operator $\mathbf{G} \in \mathbb{C}^{KT \times NT}$ grids the Cartesian k-space data onto the non-Cartesian trajectory. Since gridding is a frame-by-frame operation, \mathbf{G} is also a block-diagonal matrix.

MRF Reconstruction Problem

A key element of MRF is modeling the signal with a precomputed dictionary, where each atom describes the signal evolution over time given a particular set of relaxation parameters (and potentially other effects). Thus, we can denote the dictionary by $\boldsymbol{\delta} \in \mathbb{C}^{T \times A}$, where A is the number of atoms in the dictionary. After normalizing all atoms to have a unit ℓ_2 -norm, we can identify T_1 and T_2 of each voxel by finding the atom that maximizes the inner product with the estimated time series (9). The best fitting atoms of all voxels can be combined into a single matrix $\mathbf{D} \in \mathbb{C}^{NT \times N}$ which has the following structure:

$$\mathbf{D} = \begin{pmatrix} d_{1,1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & d_{N,1} \\ \vdots & \vdots & \vdots \\ d_{1,T} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & d_{N,T} \end{pmatrix}. \quad [3]$$

The non-zeros elements of each column in \mathbf{D} correspond to a dictionary atom. The product $\mathbf{D}^H \mathbf{x}$ is a vector of scaling factors that reflect the proton density in each voxel, where H denotes the Hermitian conjugate (or conjugate transpose). Consequently, $\mathbf{D} \mathbf{D}^H \mathbf{x}$ is a time series of images composed by dictionary entries that are scaled to match \mathbf{x} . MRF generally assumes the signal evolution of all voxels to be described by the dictionary and the complete forward model is given by

$$\mathbf{S} = \mathbf{G} \cdot \mathbf{F} \cdot \mathbf{D} \cdot \mathbf{D}^H \cdot \mathbf{x}. \quad [4]$$

The general MRF reconstruction problem can then be formulated as

$$\min_{\mathbf{D}, \mathbf{x}} \|\mathbf{G} \cdot \mathbf{F} \cdot \mathbf{D} \cdot \mathbf{D}^H \cdot \mathbf{x} - \mathbf{S}\|_2^2. \quad [5]$$

The two unknown \mathbf{D} and \mathbf{x} are multiplicative such that Eq. [5] is a non-convex optimization problem.

Original MRF Reconstruction

Ma et al. (9) used a two-step procedure to solve Eq. [5]. First, they used a filtered back-projection (BP) algorithm for reconstructing the time series of images. In the notation of Eq. [2], such a reconstruction algorithm is given by

$$\mathbf{x}_{\text{BP}} = \mathbf{F}^H \cdot \mathbf{G}_{dc}^H \cdot \mathbf{S}, \quad [6]$$

where \mathbf{G}_{dc}^H is the density compensated regridding operator.

Given the time series of images \mathbf{x}_{BP} , the dictionary is used to determine the relaxation times for each voxel. While the original MRF reconstruction finds the best dictionary match using the maximum correlation, the following ℓ_2 -norm can equivalently be minimized:

$$\mathbf{D}_{\text{BP}} = \arg \min_{\mathbf{D}} \|\mathbf{x}_{\text{BP}} - \mathbf{D} \cdot \mathbf{D}^H \cdot \mathbf{x}_{\text{BP}}\|_2^2. \quad [7]$$

Here, each voxel (column of \mathbf{D}) can be treated separately. By an exhaustive search in the dictionary, the atom is chosen that best matches the time series of the corresponding voxel in the least square sense.

Low Rank MRF Reconstruction Problem

As shown in (24), the dictionary can be compressed by a singular value decomposition:

$$\delta = \mathbf{u} \sum \mathbf{v}^H. \quad [8]$$

A low rank approximation of the dictionary matrix is given by

$$\tilde{\delta} = \mathbf{u}_R^H \cdot \delta, \quad [9]$$

where $\mathbf{u}_R \in \mathbb{C}^{T \times R}$ contains the first R columns of the unitary matrix \mathbf{u} and R is the rank of the approximation. From here on, a tilde will denote variables associated with the low rank approximation.

Assuming that the time series of each voxel in \mathbf{x} can be described by one atom of the dictionary, \mathbf{x} can be approximated by

$$\tilde{\mathbf{x}} = \mathbf{U}_R^H \cdot \mathbf{x}, \quad [10]$$

where $\mathbf{U}_R \in \mathbb{C}^{NT \times NR}$ is a block matrix composed of weighted identities:

$$\mathbf{U}_R = \begin{pmatrix} u_{1,1} & \cdots & u_{1,R} \\ \vdots & \ddots & \vdots \\ u_{T,1} & \cdots & u_{T,R} \end{pmatrix}. \quad [11]$$

Here, $u_{t,r} \in \mathbb{C}$ are the entries of the first R columns of the matrix \mathbf{u} and $\mathbb{1}$ is the identity of size $N \times N$. The compressed dictionary matrix $\tilde{\mathbf{D}} \in \mathbb{C}^{NR \times N}$, can be calculated in the same way ($\tilde{\mathbf{D}} = \mathbf{U}_R^H \cdot \mathbf{D}$) and has the same structure as \mathbf{D} . Together with the low rank approximation of the image series $\tilde{\mathbf{x}} \in \mathbb{C}^{NR}$, the MRF reconstruction problem can be formulated as

$$\min_{\tilde{\mathbf{D}}, \tilde{\mathbf{x}}} \|\mathbf{G} \cdot \mathbf{F} \cdot \mathbf{U}_R \cdot \tilde{\mathbf{D}} \cdot \tilde{\mathbf{D}}^H \cdot \tilde{\mathbf{x}} - \mathbf{S}\|_2^2. \quad [12]$$

A block diagonal matrix composed out of R identical Fourier matrices is denoted by $\tilde{\mathbf{F}} \in \mathbb{C}^{NR \times NR}$ similar to $\mathbf{F} \in \mathbb{C}^{NT \times NT}$. The block diagonal structure of $\tilde{\mathbf{F}}$ and \mathbf{F} in combination with \mathbf{U}_R being composed of weighted identity matrices can be used to show that $\mathbf{F} \cdot \mathbf{U}_R = \mathbf{U}_R \cdot \tilde{\mathbf{F}}$ (see Appendix). As mentioned earlier, one usually implements $\tilde{\mathbf{F}}$ as an FFT along all spatial dimensions. By changing the order of the Fourier transform and the SVD-compression, the computation time for the FFTs is reduced by a factor R/T . Further, \mathbf{G} and \mathbf{U}_R can be combined into the sparse matrix

$$\tilde{\mathbf{G}} \equiv \mathbf{G} \cdot \mathbf{U}_R \quad [13]$$

prior to the actual reconstruction. This reduces the low rank MRF reconstruction problem to

$$\min_{\tilde{\mathbf{D}}, \tilde{\mathbf{x}}} \|\tilde{\mathbf{G}} \cdot \tilde{\mathbf{F}} \cdot \tilde{\mathbf{D}} \cdot \tilde{\mathbf{D}}^H \cdot \tilde{\mathbf{x}} - \mathbf{S}\|_2^2. \quad [14]$$

The evaluation of this objective function requires R FFT operations (one for each considered singular value) and one sparse matrix-vector multiplication. The matrix $\tilde{\mathbf{G}}$ grids each of the resulting R Cartesian k-space data sets onto the non-Cartesian k-space trajectory of the entire experiment. Therefore, the computational burden of the gridding operation is increased by a factor of R compared to \mathbf{G} , which grids the k-space data frame by frame.

Low Rank Back-Projection

McGivney et al. (24) introduced the SVD compression for MRF in order to reduce the computation burden by matching $\tilde{\mathbf{x}}$ with the compressed dictionary. In the spirit of the original MRF reconstruction (9), they proposed to first reconstruct \mathbf{x}_{BP} with Eq. [6], then compress the time series with Eq. [10] and match the resulting $\tilde{\mathbf{x}}_{BP}$. This procedure is equivalent to the low rank back-projection

$$\tilde{\mathbf{x}}_{BP} = \tilde{\mathbf{F}}^H \cdot \tilde{\mathbf{G}}_{dc}^H \cdot \mathbf{S}. \quad [15]$$

The matching procedure of Eq. [7] is analogously given in the low rank case by

$$\tilde{\mathbf{D}}_{BP} = \arg \min_{\tilde{\mathbf{D}}} \|\tilde{\mathbf{x}}_{BP} - \tilde{\mathbf{D}} \cdot \tilde{\mathbf{D}}^H \cdot \tilde{\mathbf{x}}_{BP}\|_2^2. \quad [16]$$

Low Rank Inversion

In general, a filtered back-projection, as used in Eq. [6] and [15], does not solve the given inverse problem. An inversion of the forward operation in [6] is not possible without additional constraints, since it is highly under-determined in practical MRF implementations ($KT \ll NT$)¹. On the other hand, using the low rank approximation [15], an overdetermined system ($KT \gg NR$) is obtained provided that $R \ll T$. In the derived formalism it is straightforward to formulate the inverse problem:

$$\tilde{\mathbf{x}}_{inv} = \arg \min_{\tilde{\mathbf{x}} \in \mathbb{C}^{NR}} \|\tilde{\mathbf{G}} \cdot \tilde{\mathbf{F}} \cdot \tilde{\mathbf{x}} - \mathbf{S}\|_2^2. \quad [17]$$

This is a linear system and can be solved with a conjugate gradient (CG) algorithm. This procedure searches for the low rank approximation $\tilde{\mathbf{x}}$ that best describes the measured signal in the least square sense. After solving Eq [17], Eq. [16] is used for the dictionary matching.

Low Rank Alternating Directions Method of Multipliers

Depending on the particular implementation of the MRF experiment and the reconstruction, the condition $KT \gg NR$ might not be fulfilled. And even if it is, the condition number of the forward operation $\tilde{\mathbf{G}}\tilde{\mathbf{F}}$ with respect to inversion is likely high, making the above approach unfeasible.

The recently proposed application of ADMM (23) to the original MRF (13, 14) addresses the reconstruction problem by solving Eq. [14] via variable splitting. This approach can be denoted by an augmented Lagrangian:

¹Reminder: K is the length of the trajectory of each time-frame, N is the number of voxels of the image of each time-frame, T is the number of time-frames and R is the rank of the approximation.

$$\{\tilde{\mathbf{x}}_{\text{ADMM}}, \tilde{\mathbf{D}}_{\text{ADMM}}, \tilde{\mathbf{y}}_{\text{ADMM}}\} = \arg \min_{\tilde{\mathbf{x}}, \tilde{\mathbf{D}}, \tilde{\mathbf{y}}} \|\tilde{\mathbf{G}} \cdot \tilde{\mathbf{F}} \cdot \tilde{\mathbf{x}} - \mathbf{S}\|_2^2 + \mu \|\tilde{\mathbf{x}} - \tilde{\mathbf{D}}\tilde{\mathbf{D}}^H \tilde{\mathbf{x}} + \tilde{\mathbf{y}}\|_2^2. \quad [18]$$

The first summand is the data consistency term, which is equivalent to Eq. [17]. The second summand compares the splitted variable $\tilde{\mathbf{x}}$ to its projection onto the dictionary $\tilde{\mathbf{D}}\tilde{\mathbf{D}}^H\tilde{\mathbf{x}}$. The Lagrangian multiplier $\tilde{\mathbf{y}} \in \mathbb{C}^{NR}$ is denoted here in the scaled dual form (23) and μ is the ADMM penalty parameter.

The simultaneous search for the optimal $\tilde{\mathbf{x}}$, $\tilde{\mathbf{D}}$ and $\tilde{\mathbf{y}}$ in their product space is a challenging nonlinear problem. ADMM addresses this by alternately solving

$$\tilde{\mathbf{x}}_{j+1} = \arg \min_{\tilde{\mathbf{x}}} \|\tilde{\mathbf{G}} \cdot \tilde{\mathbf{F}} \cdot \tilde{\mathbf{x}} - \mathbf{S}\|_2^2 + \mu \|\tilde{\mathbf{x}} - \tilde{\mathbf{D}}_j \tilde{\mathbf{D}}_j^H \tilde{\mathbf{x}} + \tilde{\mathbf{y}}_j\|_2^2 \quad [19]$$

$$\tilde{\mathbf{D}}_{j+1} = \arg \min_{\tilde{\mathbf{D}}} \|\tilde{\mathbf{x}}_{j+1} - \tilde{\mathbf{D}}\tilde{\mathbf{D}}^H \tilde{\mathbf{x}}_{j+1} + \tilde{\mathbf{y}}_j\|_2^2 \quad [20]$$

$$\tilde{\mathbf{y}}_{j+1} = \tilde{\mathbf{y}}_j + \tilde{\mathbf{x}}_{j+1} - \tilde{\mathbf{D}}_{j+1} \tilde{\mathbf{D}}_{j+1}^H \tilde{\mathbf{x}}_{j+1}. \quad [21]$$

Eq. [19] is a linear optimization problem and can be solved with a CG algorithm. The minimization in Eq. [20] is solved by exhaustive search within the dictionary for each voxel. Note that Eq. [20] differs from the maximum correlation approach due to the Lagrangian multiplier, which increases the computational burden by a factor of two, whereas the time complexity remains the same. The update of the Lagrangian multiplier (Eq. [21]) has the standard form used for ADMM algorithms (23) and increasingly addresses those errors in $\tilde{\mathbf{x}}_j - \tilde{\mathbf{D}}_j \tilde{\mathbf{D}}_j^H \tilde{\mathbf{x}}_j$ that remain unchanged over multiple iterations. In the light of a regularized inverse problem, the Lagrangian multiplier in Eq. [18] might be considered nonessential. However, in order to solve Eq. [14], it avoids the necessity to continuously increase μ towards $\mu \rightarrow +\infty$ as j increases (23). Ultimately, it speeds up the convergence of the ADMM algorithm.

The optimization problem in Eq. [18] is non-convex, such that the convergence of the algorithm depends on the initial guess (13). Here, we initialize the algorithm with $\tilde{\mathbf{x}}_0 = 0$, $(\tilde{\mathbf{D}}\tilde{\mathbf{D}}^H)_0 = \mathbb{1}$ and $\tilde{\mathbf{y}}_0 = 0$. Note that $(\tilde{\mathbf{D}}\tilde{\mathbf{D}}^H)_0 = \mathbb{1}$ is not described by the dictionary and was chosen to transform Eq. [18] into Eq. [17] for the first iteration. In the limit of a well conditioned $\tilde{\mathbf{G}}\tilde{\mathbf{F}}$ with respect to inversion, the first ADMM iteration solves the MRF reconstruction problem. With a slightly ill-conditioned Eq. [17], we expect the first ADMM iteration to approximate the solution well enough to result in good convergence. This is not

the case, where the data consistency term is highly ill-conditioned, including cases, where it is under-determined as used in (13, 14).

Sensitivity Encoding

Parallel imaging (25) can help to further improve the conditioning of the reconstruction. Here, CG sensitivity encoding (SENSE) (26) is employed by reformulating the data consistency term:

$$\{\tilde{\mathbf{x}}_{\text{ADMM}}, \tilde{\mathbf{D}}_{\text{ADMM}}, \tilde{\mathbf{y}}_{\text{ADMM}}\} = \arg \min_{\tilde{\mathbf{x}}, \tilde{\mathbf{D}}, \tilde{\mathbf{y}}} \sum_c \|\tilde{\mathbf{G}} \cdot \tilde{\mathbf{F}} \cdot \tilde{\mathbf{E}}_c \cdot \tilde{\mathbf{x}} - \mathbf{S}_c\|_2^2 + \mu \|\tilde{\mathbf{x}} - \tilde{\mathbf{D}}\tilde{\mathbf{D}}^H \tilde{\mathbf{x}} + \tilde{\mathbf{y}}\|_2^2. \quad [22]$$

Before the Fourier transformation, the search vector $\tilde{\mathbf{x}}$ is multiplied by the diagonal matrix $\tilde{\mathbf{E}}_c \in \mathbb{C}^{NR \times NR}$ which contains R repetitions of the sensitivity profile of coil c and is then compared to the signal \mathbf{S}_c that was measured by the coil under consideration. The sum over all coil elements provides the total cost of the data-consistency term. Note that SENSE can also be incorporated in the low rank inversion (Eq. [17]) in the same way.

Methods

In the spirit of reproducible research, the source code of the proposed algorithm is available at https://bitbucket.org/asslaender/nyu_mrf_recon.

Simulations

All simulations were performed with the pseudo steady-state free precession (pSSFP) pattern described in (14). The flip angle scheme is displayed in Fig. 1a and consists of an inversion pulse followed by RF-pulses with varying flip angles and a phase increment of π between consecutive pulses. The repetition and echo time of each repetition was calculated based on the pSSFP approach with $TR_0 = 5$ ms. A dictionary was computed with Bloch simulations with $T_1(s) = 0.3 \cdot 1.02^j \forall j \in \{0, 1, \dots, 152\}$, thus covering the range between 300 ms and 6 s in steps of 2%. The dictionary covers T_2 between 50 ms and 3 s in steps of 2%, i.e. $T_2(s) = 0.05 \cdot 1.02^j \forall j \in \{0, 1, \dots, 207\}$. For simplicity, off-resonance effects were neglected, as well as B_1 inhomogeneities and other interfering factors such as slice profile imperfections and magnetization transfer. The resulting dictionary $\mathcal{D} \in \mathbb{R}^{841 \times 24921}$ has a rank of 62 at a tolerance of $2.8 \cdot 10^{-9}$. The dictionary is real valued due to the nature of the sequences and the assumption of on-resonant spins. An SVD was performed according to Eq. [8]. The first five left-singular vectors (columns of \mathbf{u}) are depicted in Fig. 1b and the corresponding singular values in (c).

A numerical brain phantom from the Brainweb database (<http://brainweb.bic.mni.mcgill.ca/bra>) (27) was used to create synthetic data sets. Segments of fat were associated with the “original” values $PD^{\text{org}} = 0.9$, $T_1^{\text{org}} = 0.37$ s, $T_2^{\text{org}} = 0.13$ s, white matter (WM) with $PD^{\text{org}} = 0.65$, $T_1^{\text{org}} = 1.08$ s, $T_2^{\text{org}} = 0.07$ s, gray matter (GM) with $PD^{\text{org}} = 0.8$, $T_1^{\text{org}} = 1.82$ s, $T_2^{\text{org}} = 0.10$ s and cerebrospinal fluid (CSF) with $PD^{\text{org}} = 1.0$, $T_1^{\text{org}} = 4.5$ s, $T_2^{\text{org}} = 2.2$ s. Bloch simulations were used to create the original time series \mathbf{x}_{org} .

Unless stated otherwise, a matrix size of 128×128 was used, which corresponds to $N=16384$. If not stated otherwise, one radial k-space spoke with a golden angle increment (28) was acquired for each TR . To demonstrate the impact of sampling density on the reconstruction quality, a subset of the simulations was repeated with multiple spokes per TR . These additional spokes were distributed using the tiny golden angle of approximately 32.0397° (29). All radial spokes were associated with an extend of $k_{\max} = \sqrt{2}\pi/\text{voxel}$ contrary to the usual convention of radial MRI ($k_{\max} = \pi/\text{voxel}$). This avoids artifacts due to unsampled corners of k-space in the simulation. The signal \mathbf{S} was simulated according to Eq. [2], where the forward operation \mathbf{GF} was implemented as a nuFFT operator with Kaiser-Bessel interpolation and optimal scaling in the min-max sense (30). Three nearest neighbors were taken into account in each spatial direction and an oversampling factor of 2 was used.

Image Reconstruction—Four different reconstructions were performed on the same noise-free synthetic data: A filtered back-projection (Eq. [6]) as proposed in (9), a filtered back-projection in the low rank domain (Eq. [15]) as proposed in (24), a low rank approximation of the signal (Eq. [17]) and the proposed low rank ADMM reconstruction (Eq. [18]). The back-projection reconstructions were performed by first gridding the data onto an enlarged k-space (factor 2) with linear density compensation. Thereafter, k-space was cut to cover $\pm\pi/\text{voxel}$ followed by an FFT. This procedure ensures correct density compensation with the lengthened radial spokes. All CG-based methods do not rely on density compensation such that this procedure can be avoided. The low rank inversion was performed for $R \in \{2, 3, \dots, 6, 8, \dots, 20\}$. Eq. [17] is solved by a CG algorithm with 100 iterations after initializing it with zero. The low rank ADMM reconstructions were performed with 20 CG steps for solving Eq. [19] in each ADMM iteration. A simulation was performed with $R \in \{2, 3, \dots, 6, 8, \dots, 20\}$ at a fixed $\mu = 1.26 \cdot 10^{-3}$ and 10 ADMM iterations. Further, at a fixed $R = 5$, reconstructions were performed for the product space of $\mu = 10^j \forall j \in \{-4, -3.8, \dots, -2\}$ and the number of ADMM iterations $\in \{1, 2, \dots, 31\}$.

Root Mean Square Error—The root mean square error (RMSE) was calculated from noise-free reconstructions. A region of interest was drawn over the entire white matter and the reconstructed quantitative maps were compared to the original maps of PD^{org} , T_1^{org} and T_2^{org} . The RMSE was normalized by the original value and is then referred to as NRMSE.

Noise Analysis—The reconstructions with $R = 5$, $\mu = 1.26 \cdot 10^{-3}$, 10 ADMM and 20 CG iterations were repeated 100 times each with pseudo random white noise added to the signal \mathbf{S} (31). The input noise level was $SNR_{\text{in}} = 100$ for each case, where $SNR_{\text{in}} = 1$ corresponds an equal mean and standard deviation in a voxel with $PD^{\text{org}} = 1$ in a hypothetical experiment with $\alpha = 90^\circ$, $TR \rightarrow +\infty$ and $TE \rightarrow +0$.

Two different analyses were performed on the 100 pseudo replicas: First, the $PDNR$, T_1NR and T_2NR were calculated by dividing the mean PD , T_1 and T_2 of all white matter voxels in all pseudo replicas by the corresponding standard deviation. This analysis was performed for the low rank inversion and the low rank ADMM reconstruction and the result was normalized by SNR_{in} .

A second analysis was performed on the pseudo replicas reconstructed with a rank $R = 5$. For each tissue, the mean and the standard deviation of all voxels and all pseudo replicas were calculated. They were both normalized by the original parameters used in the simulation.

Sensitivity Encoding—The low rank ADMM simulation (Eq. [18]) was repeated with an increased matrix size of 384×384 at the same number of spokes in order to increase undersampling artifacts. It is compared to a reconstruction that incorporates SENSE (Eq. [22]). In vivo coil sensitivity profiles of 12 virtual coils of a 44 channel head coil (Siemens, Erlangen, Germany) were employed. The numerical phantom was multiplied by the sensitivity profile of each coil element. Thereafter, the signal of each coil was simulated with Eq. [2] followed by a reconstruction based on Eq. [22] with the same parameters as used before.

In Vivo Experiments

Single slice MR fingerprinting data of a healthy volunteer's brain were acquired with a 3T Skyra scanner (Siemens, Erlangen, Germany). Written informed consent was obtained and the protocol was approved by the Institutional Review Board. For signal reception, the manufacturer's 16-channel head coil was used, which were reduced to 8 virtual coils via SVD compression. The previously described pSSFP acquisition scheme was employed. The inversion pulse was implemented as a 10.24 ms secant pulse followed by a train of sinc-pulses with a time bandwidth product of 2 and a duration of 800 μ s. The in plane spatial resolution was 1 mm \times 1 mm at a $FOV = 192$ mm \times 192 mm with a slice thickness of 3 mm. The dwell time of the readout was set to 2.4 μ s. With these parameters, the timing of the i^{th} readout can be realized only when $\min\{TE_i, TR_i - TE_i\} > 1.3$ ms. The current implementation of the sequence skips the readout whenever this requirement is not fulfilled. In total, 841 of the designated 850 spokes were acquired within 4.1 s. After acquiring the data along one RF-train, a 10 s pause was inserted to ensure full relaxation before the next repetition of the RF-train is started. Relative to one another, subsequent spokes within the RF-train are rotated by the golden angle (111.25°) (28). Each time the RF-train is repeated, all spokes are rotated by the tiny golden angle of approximately 32.0397° (29). A total of 32 repetitions were measured. The data set was reconstructed in two different ways, using all samples and using only the first RF-train. For more details about the acquisition, please refer to (14), where a similar experimental setup was used.

For image reconstruction, the nuFFT was implemented with Kaiser-Bessel interpolation and optimal scaling in the min-max sense (30). Three nearest neighbors were considered in each direction and an oversampling of 2 was employed. The low rank inversion was performed with 20 CG steps and the low rank ADMM reconstruction was performed with 5 ADMM iterations and each with 10 CG iterations. The ADMM penalty parameter was set to $\mu = 0.1$ based on the visual impression of the results.

Results

Simulations

As shown in Fig. 1c, the singular values of the dictionary matrix δ decrease rapidly. Reconstructions were performed in low rank approximations with $R \in \{2, 3, \dots, 6, 8, \dots, 20\}$. The corresponding NRMSE based on noise free reconstructions are shown in Fig. 2a–c. The signal to noise ratio based on pseudo replicas is displayed in (d–f). Using the low rank back-projection reconstruction (Eq. [15] and [16]), the NRMSE and the SNR are hardly influenced by the rank of the approximation (Fig. 2a,d). However, solving the inverse problem in the low rank approximation (Eq. [17]) reduces the overall NRMSE (b). The minimum NRMSE is reached at $R = 3$, where also a local maximum of the SNR can be found, which is greater than the SNR of the back-projection reconstruction. At higher rank approximations, a non-monotonic behavior can be observed: The NRMSE of all parameters increases for $3 < R < 11$ and for $R > 11$ the NRMSE of T_2 and PD decreases again. Also, the SNR shows a non-monotonic behavior, which favors proton density over the relaxation times at $R > 11$. Applying the ADMM algorithm to the low rank approximation, a further decreased NRMSE can be observed compared to the low rank inverse problem. Also, the minimum of the NRMSE is shifted towards $R = 5$ and the SNR shows a more stable behavior with a maximum at $R = 5$. Note that large R values approach the behavior of an ADMM reconstruction in the temporal domain without a low rank approximation. In particular, at $R = 841$ both reconstructions would be equivalent.

Supporting Fig. 1a repeats the results of Fig. 2c with a different scaling and compares them to the same simulation with 4 radial spokes per TR instead of one. Since all other parameters were kept equivalent, this effectively reduces the undersampling. As a consequence, the NRMSE is reduced over all and the minimum is shifted to $R = 8$. Note that the NRMSE of T_1 is bound by the discretization of the dictionary in this particular example.

Reconstructions with $R = 5$ —Based on the NRMSE minimum, $R = 5$ was selected for a more detailed analysis of the data with one radial spoke per TR . The NRMSE and the residuum of the objective function are displayed in Fig. 3 as a function of the ADMM penalty parameter μ and the number of ADMM iterations. After 1 iteration (dark blue), μ has no influence on the NRMSE and the residuum, which is expected due to the initialization of the algorithm. With an increasing number of iterations, a monotonous decrease of the NRMSE and the residuum can be observed until the point of convergence, after which the NRMSE increases slightly in some cases. In general, the reduction speed of the NRMSE depends on μ , as does the minimum NRMSE reached after convergence. The minimal NRMSE varies slightly as a function of μ . For bigger μ values, the NRMSE abruptly increases. The optimal μ also depends on the number of iterations as well as the considered parameter (PD , T_1 , T_2). For practical reasons, 10 iterations (black) were chosen here heuristically, for which $\mu = 1.26 \cdot 10^{-3}$ (arrows in Fig. 3) minimizes the NRMSE of T_1 , while PD and T_2 are close to optimal, justifying its choice for reconstructions throughout the paper. The residuum (Eq. [14]) converges very fast and shows only little variations within a certain range of μ and the minimum at 10 iterations is shifted with respect to the NRMSE (d).

Fig. 4 shows the five elements of \mathbf{x} reconstructed from noise free data with $\mu = 1.26 \cdot 10^{-3}$. The rows correspond to the singular values in decreasing order (the color coding corresponds to Fig. 1b,c). The right hand side of all images are displayed using the same scale and the decreasing brightness reflects the magnitude of the singular values. The second column depicts the back-projection reconstruction (Eq. [15]). Contrary to the other reconstructions, the individual images do not depend on the rank of the approximation. The images are increasingly dominated by noise-like artifacts for decreasing singular values. The third column displays \mathbf{x}_{inv} , which was reconstructed with Eq. [17] and exhibits considerably reduced artifacts. However, residual undersampling artifacts can be observed when increasing the brightness (left half of the images), especially for smaller singular values (last two rows). The artifacts are further reduced, when employing the proposed ADMM reconstruction (Eq. [18]).

The corresponding quantitative maps are shown in Fig. 5. They were reconstructed with Eq. [7] (second column), [16] (column 3 and 4) and [18] (column 5). The back-projection reconstruction (column 2) exhibits significant undersampling artifacts in all three parameter maps. Noise-like artifacts can be seen in the proton density, as well as in the T_1 map and most prominently in the T_2 map. The quantitative maps reconstructed from the low rank back-projection (third column of Fig. 5) and the time domain back-projection (second column) are visually the same. In fact, the low rank back-projection does not promise improved image quality and is purely motivated by computation time (24). On the contrary, solving the inverse problem in the low rank approximation (Eq. [17]) results in significantly improved quantitative maps: The noise-like artifacts are reduced, which is even more the case when employing the ADMM algorithm in the low rank domain.

When analyzing the reconstruction of noisy data within all four simulated tissue types, systematic errors can be observed in the two cases of the back-projection reconstruction, most prominently for T_2 (Fig. 6). On the contrary, both the low rank inversion and the low rank ADMM reconstruction accurately reproduce the original parameters, except in the case of CSF, where systemic errors in T_1 and T_2 can still be noticed. The employed excitation pattern in combination with the short acquisition time (4.1 s) results in a reduced sensitivity to long relaxation times. As a consequence, the relaxation times of CSF are not measured precisely. Furthermore, the standard deviation is decreased in all tissue types when employing the low rank ADMM reconstruction. Due to the non-linearity of the Bloch equation, the noise is distributed in a non-Gaussian manner (not shown here). Therefore, mean and standard deviation have a limited information value. For the cases with strong deviations from a Gaussian distribution, the standard deviation is not displayed.

Reconstructions with $R = 20$ —Fig. 7 compares the low rank inversion and the ADMM algorithm for $R = 20$ in order to further analyze the influence of the rank. In the case of the low rank inverse problem (Eq. [17]), spatial blurring can be observed, most prominently in the T_1 map. This is also reflected in Fig. 2b, where the NRMSE increase at high R values is strongest for T_1 . When employing the ADMM algorithm, some noise-like artifacts can be observed, but the blurring is reduced compared to the low rank inversion.

Sensitivity Encoding—The benefits of SENSE in the low rank ADMM reconstruction is demonstrated in Supporting Fig. 2. Without parallel imaging, the ADMM reconstruction of higher resolution data with same number of radial spokes results in increased undersampling artifacts. These artifacts are reduced by incorporating CG SENSE (26) in the proposed low rank ADMM reconstruction.

In Vivo Experiment

The in vivo maps reconstructed from one radial spoke per TR (Fig. 8) appear noisy when reconstructed with the standard back-projection and the low rank back-projection. However, the artifact level is reduced compared to the simulations. Since the data were acquired with a multi-coil array, the back-projection includes a multiplication with the complex conjugate of the coil sensitivity profiles. Even though this does not represent a parallel imaging algorithm, it still mitigates the undersampling artifacts to some extent. An improvement of the reconstruction quality can be observed when employing a low rank inversion and particularly when using the proposed low rank ADMM reconstruction. However, the noise level is still high, especially in the T_2 map. This can be attributed to the short acquisition time (4.1 s) in combination with a suboptimal acquisition pattern. A similar effect can also be observed in the simulation data shown in Figs. 2f and 6, which also show a relatively high variability in the T_2 maps. Measuring the mean and standard deviation over a region of interest in the frontal white matter shows a decrease in the standard deviation when employing the proposed LR ADMM algorithm (Tab. 1). Note that the mean values are affected by slice profile imperfections and magnetization transfer effects.

When using 32 spokes per TR , all reconstruction techniques converge to very similar results (Fig. 9). This is also reflected by the comparably constant relaxation times measured within a region of interest in the frontal white matter (Tab. 1). Note that B_0 drifted significantly during the time needed to obtain all 32 spokes. Consequently, the relaxation times are biased and deviate from those obtained when reconstructed from the first repetition only (Fig. 8). This effect originates from the sensitivity of balanced SSFP sequences to B_0 variations and a detailed study of these effects is beyond the scope of this paper.

Discussion

In this paper we propose a reconstruction framework for MRF based on a low rank approximation of the signal evolution in combination with the alternating direction method of multipliers. The original MRF reconstruction (9) and SVD-based dictionary matching (24) are described within a formalism that identifies SVD-based matching as a filtered back-projection into the space of few singular vectors. By exploiting the commutative property of the spatial Fourier transformation and the SVD-compression in the temporal domain (see Appendix), we show that the computation time of the FFTs can be reduced by a factor R/T , which is 5/841 in the particular case under investigation. The gridding operator and the SVD-compression are combined into a single sparse matrix, which allows for gridding the Cartesian k-space data in the domain of the singular vectors directly onto the non-Cartesian trajectory in the temporal domain. The overall computational cost of this sparse matrix-vector multiplication is a factor of R larger than the frame-by-frame gridding in the standard

reconstruction, since $\tilde{\mathbf{G}}$ grids each singular value image onto the entire trajectory. However, an overall speedup is expected when setting $R \ll T$. For the in vivo data, a speedup of a factor of 12 was observed for executing $\tilde{\mathbf{F}}^H \tilde{\mathbf{G}}^H \tilde{\mathbf{G}} \tilde{\mathbf{F}} \mathbf{x}$ compared to $\mathbf{F}^H \mathbf{G}^H \mathbf{G} \mathbf{F} \mathbf{x}$. However, the speedup depends on R , T and the implementation of the nuFFT. For large T values as used, e.g., for 3D fingerprinting, even larger speedups are expected. Solving Eq. [20] has roughly double the computational cost compared to maximizing the correlation, since $\tilde{\mathbf{D}}^H \mathbf{x}$ and $\tilde{\mathbf{D}}^H \mathbf{y}$ need to be calculated. However, the low rank approximation reduces the complexity of these matrix-vector multiplications and speeds up the matching process by the same factor as in the low rank back-projection proposed in (24). The reconstruction of the in vivo data (Fig. 8) took approximately 43 s to reconstruct on a desktop computer (3.3 GHz Intel i5-6600 processor). The reconstruction is implemented in MATLAB (The MathWorks, Natick, MA, USA). Further speedups could be achieved with a multi core or GPU sparse matrix multiplication and a more efficient dictionary search (18, 32).

The formalism introduced here recasts the reconstruction of the image series into the form of a straightforward inverse problem that is over-determined depending on the choice of parameters. When the condition $KT \gg NR$ is fulfilled, artifact free images can theoretically be reconstructed with Eq. [17]. Our low rank inversion embodies some concepts akin to the methods described in (15) in the context of parameter mapping and more specifically in (17, 19) in the context of MRF. However, the method described in (19) does not provide the here described speed-up related the number FFT operations, while the method proposed in (17) lacks a straightforward incorporation of sensitivity encoding. Fig. 2 shows that very small R values result in approximation errors. When R is increased, the problem first becomes increasingly ill-conditioned, which results in strong noise amplification (Fig. 2). If the rank is increased beyond $R = 11$, the problem becomes under-determined ($KT < NR$). Employing a CG algorithm to this under-determined problem leads to blurring, particularly in the T_1 map (Fig. 7). On the other hand, the SNR of T_1 decreases for $R > 11$ (Fig. 2). This identifies the blurring not to be a simple smoothing effect, where an increased SNR for T_1 would be expected. This behavior might be closely linked to employed k-space trajectory. However, a detailed discussion is beyond the scope of this paper.

The described issues can be mitigated by applying the proposed ADMM algorithm to the low rank approximation of the time series. ADMM, like other iterative reconstructions proposed for MRF, constrains the reconstruction directly to atoms of the dictionary (12, 13, 14, 20). However, these approaches are non-convex problems which are prone to convergence issues. Examining Fig. 2c and f, an increased NRMSE and a decreased SNR can be observed for an increasing R . This shows that the ADMM algorithm without a low rank approximation converges to a solution that neither minimizes NRMSE nor maximizes SNR. Note that a reconstruction with $R = 841$ would be equivalent to an ADMM reconstruction in the temporal domain without a low rank approximation similar to (13, 14). The application of the ADMM algorithm to $\tilde{\mathbf{x}}$ in the space of the first R singular vectors reduces the NRMSE and increases the SNR compared to a low rank approximation alone (Fig. 2b vs. c and e vs. f) and also compared to the ADMM algorithm without a low rank approximation (Fig. 2c and f: low R vs. high R). In general, choosing the rank of the approximation is a trade-off between not describing the signal evolution accurately enough

(small R) and solving an ill-posed problem (large R). The first part of the trade-off depends on the dictionary and ultimately on the RF-, TR - and TE -train, as well as the types of tissue under consideration, while the latter part depends on the spatial encoding (Supporting Fig. 1). Consequently, the best choice of R depends on the pulse sequence. However, Fig. 2 indicates a reduced sensitivity of the ADMM approach to large R values compared to the low rank inversion alone, which is expected due to extra penalty term that is added in Eq. [18] compared to Eq. [17].

In the case of convex optimization, the ADMM penalty parameter μ influences the speed of convergence, but not the solution (23). This is not the case for non-convex optimization problems as Eq. [18], where μ can also influence the solution after convergence (Fig. 3). The extreme values give some intuition about the effect of μ . The ADMM reconstruction problem (Eq. [18]) approaches the low rank inversion (Eq. [17]) for small μ , while large μ emphasize matching \mathbf{x} to $\tilde{\mathbf{D}}_j \tilde{\mathbf{D}}_j^H \tilde{\mathbf{x}}$ after the first iteration. As a consequence, data consistency is neglected and the resulting parameter maps exhibit little change after the initial guess. This is likely the cause of the abrupt increase of the NRMSE when increasing μ beyond a certain point. This abrupt increase suggests to choose smaller μ in practical cases. If the algorithm is terminated before convergence, the optimal μ is slightly shifted (black dots and arrows in Fig. 3). The minimal residuum does in general not coincide with the smallest NRMSE, which limits the feasibility of using it to tune μ . For the in vivo data, μ was chosen heuristically. The step size of the relaxation times in the dictionary does generally influence convergence as well. A finer grid allows $\tilde{\mathbf{D}}$ to be updated even after small changes in \mathbf{x} , potentially improving the convergence. However, the present implementation penalizes the distance to the dictionary, rather than strictly tying \mathbf{x} to the dictionary, as done in the BLIP algorithm (12). Therefore, the step size in the dictionary is not expected to be crucial for convergence and a detailed analysis is omitted here.

Due to the non-linearity of the Bloch equation and the reconstruction, the noise in the computed quantitative maps is not expected to be Gaussian distributed. A general reduction of the noise can be observed for the proposed low rank ADMM approach compared to the other reconstructions under investigation (Fig. 6). In the case of CSF, T_1 and T_2 exhibit noise-like variations that deviate strongly from a Gaussian distribution (not shown here). However, the SNR and the distribution of the noise depends strongly on the acquisition pattern and a careful consideration of these effects is beyond the scope of this paper.

Formulating the reconstruction as a minimization problem (Eq. [18]) provides the flexibility to add different constraints. Here, iterative SENSE was incorporated (Eq. [22], Supporting Fig. 2, 8 and 9). Additional priors in the spirit of compressed sensing (33) can also be added as additional summands to Eq. [18]. The provided source code implements an $\ell_{1,2}$ -penalty of \mathbf{x} (16) (not shown here). Future work could also explore the potential advantages of incorporating the nuclear norm which has been shown to provide superior properties for multi-contrast image reconstruction (34).

Conclusion

A reconstruction framework for MRF is presented that incorporates the data model in two different ways. A low rank approximation of the signal evolution is computed based on a singular value decomposition of the dictionary matrix. Additionally, the best matching atom of the dictionary is used to constrain the signal evolution in each voxel. An ADMM algorithm is then used to alternately optimize the low rank approximation of the signal evolution and its projection onto the dictionary. The resulting quantitative maps show a reduced error and superior SNR compared to the original MRF reconstruction, a low rank approximation alone, as well as the application of the ADMM algorithm in the temporal domain rather than in a low rank space. The proposed algorithm is a general reconstruction framework for MRF and can be readily applied to plug and play parallel transmission (10) and chemical exchange MRF-X (35) amongst others.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported in part by the research grants NIH R21 EB020096 and R01 AR070297, and was performed under the rubric of the Center for Advanced Imaging Innovation and Research (CAI2R, www.cai2r.net), a NIBIB Biomedical Technology Resource Center (NIH P41 EB017183).

References

1. Look DC, Locker DR. Time saving in measurement of NMR and EPR relaxation times. *Review of Scientific Instruments*. 1970; 41(2):250–251.
2. Crawley AP, Henkelman RM. A comparison of one-shot and recovery methods in T1 imaging. *Magnetic Resonance in Medicine*. 1988; 7(1):23–34. [PubMed: 3386519]
3. Cheng HLM, Wright GA. Rapid high-resolution T1 mapping by variable flip angles: Accurate and precise measurements in the presence of radiofrequency field inhomogeneity. *Magnetic Resonance in Medicine*. 2006; 55(3):566–574. [PubMed: 16450365]
4. Deoni SCL, Rutt BK, Peters TM. Rapid combined T1 and T2 mapping using gradient recalled acquisition in the steady state. *Magnetic Resonance in Medicine*. 2003; 49(3):515–526. [PubMed: 12594755]
5. Deoni SCL, Peters TM, Rutt BK. Determination of Optimal Angles for Variable Nutation Proton Magnetic Spin-Lattice, T1, and Spin-Spin, T2, Relaxation Times Measurement. *Magnetic Resonance in Medicine*. 2004; 51(1):194–199. [PubMed: 14705061]
6. Doneva M, Börner P, Eggers H, Stehning C, Sénégas J, Mertins A. Compressed sensing reconstruction for magnetic resonance parameter mapping. *Magnetic Resonance in Medicine*. 2010; 64(4):1114–1120. [PubMed: 20564599]
7. Sumpf TJ, Petrovic A, Uecker M, Knoll F, Frahm J. Fast T2 mapping with improved accuracy using undersampled spin-echo MRI and model-based reconstructions with a generating function. *IEEE Transactions on Medical Imaging*. 2014; 33(12):2213–2222. [PubMed: 24988592]
8. Ben-Eliezer N, Sodickson DK, Shepherd T, Wiggins GC, Block KT. Accelerated and motion-robust in vivo T2 mapping from radially undersampled data using bloch-simulation-based iterative reconstruction. *Magnetic Resonance in Medicine*. 2016; 75:1346–1354. [PubMed: 25891292]
9. Ma D, Gulani V, Seiberlich N, Liu K, Sunshine JL, Duerk JL, Griswold MA. Magnetic resonance fingerprinting. *Nature*. 2013; 495(7440):187–192. [PubMed: 23486058]

10. Cloos MA, Knoll F, Zhao T, Block K, Bruno M, Wiggins C, Sodickson DKD. Multi-parametric imaging with heterogenous radiofrequency fields. *Nature Communication*. 2016; 7doi: 10.1038/ncomms12445
11. Ma D, Pierre EY, Jiang Y, Schluchter MD, Setsompop K, Gulani V, Griswold MA. Music-based magnetic resonance fingerprinting to improve patient comfort during MRI examinations. *Magnetic Resonance in Medicine*. 2016; 75:2303–2314. [PubMed: 26178439]
12. Davies M, Puy G, Vanderheyne P, Wiaux Y. A Compressed Sensing Framework for Magnetic Resonance Fingerprinting. *SIAM Journal on Imaging Sciences*. 2014; 7(4):27.
13. Zhao B, Setsompop K, Ye H, Cauley S, Wald L. Maximum Likelihood Reconstruction for Magnetic Resonance Fingerprinting. *IEEE Transactions on Medical Imaging*. 2016 epub ahead of print.
14. Assländer J, Glaser SJ, Hennig J. Pseudo Steady-State Free Precession for MR-Fingerprinting. *Magnetic Resonance in Medicine*. 2016 epub ahead of print.
15. Petzschner FH, Ponce IP, Blaimer M, Jakob PM, Breuer FA. Fast MR parameter mapping using k-t principal component analysis. *Magnetic Resonance in Medicine*. 2011; 66(3):706–716. [PubMed: 21394772]
16. Zhao B, Lu W, Hitchens TK, Lam F, Ho C, Liang ZP. Accelerated MR parameter mapping with low-rank and sparsity constraints. *Magnetic Resonance in Medicine*. 2015; 74:489–498. [PubMed: 25163720]
17. Doneva M, Amthor T, Koken P, Sommer K, Börnert P. Low Rank Matrix Completion-based Reconstruction for Undersampled Magnetic Resonance Fingerprinting Data. *Proc Intl Soc Mag Reson Med*. 2016; 24:0432.
18. Cline CC, Chen X, Mailhe B, Wang Q, Nadar M. AIR-MRF: Accelerated iterative reconstruction for magnetic resonance fingerprinting. *Proc Intl Soc Mag Reson Med*. 2016; 24:0434.
19. Zhao B, Setsompop K, Gagoski B, Ye H, Adalsteinsson E, Grant P, Wald L. A Model-Based Approach to Accelerated Magnetic Resonance Fingerprinting Time Series Reconstruction. *Proc Intl Soc Mag Reson Med*. 2016; 24:871.
20. Zhao B. Model-Based Iterative Reconstruction for Magnetic Resonance Fingerprinting. *IEEE International Conference on Image Processing (ICIP)*. 2015:3392–3396.
21. Mazor, G., Weizman, L., Tal, A., Eldar, YC. Low Rank Magnetic Resonance Fingerprinting. *Engineering in Medicine and Biology Society (EMBC), IEEE 38th Annual International Conference of the*, 646804; 2016. p. 2-5.
22. Liao C, Cao X, Ye H, Chen Y, He H, Chen S, Ding Q, Liu H, Zhong J. In this study, a low rank and sparsity based MRF reconstruction scheme (L&S MRF) is proposed for reducing the artifacts off each time point with a fraction off acquisition times. *Proc Intl Soc Mag Reson Med*. 2016; 24:4227.
23. Boyd S, Parikh N, Chu EBP, Eckstein J. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends in Machine Learning*. 2011; 3(1):1–122.
24. McGivney D, Ma D, Saybasili H, Jiang Y, Griswold M. Singular Value Decomposition for Magnetic Resonance Fingerprinting in the Time Domain. *IEEE transactions on medical imaging*. 2014; 33(12):2311–2322. [PubMed: 25029380]
25. Sodickson DK, Manning WJ. Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*. 1997; 38(4):591–603. [PubMed: 9324327]
26. Pruessmann KP, Weiger M, Börnert P, Boesiger P. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magnetic Resonance in Medicine*. 2001; 46(4):638–651. [PubMed: 11590639]
27. Collins DL, Zijdenbos AP, Kollokian V, Sled JG, Kabani NJ, Holmes CJ, Evans AC. Design and construction of a realistic digital brain phantom. *IEEE transactions on medical imaging*. 1998; 17(3):463–468. [PubMed: 9735909]
28. Winkelmann S, Schaeffter T, Koehler T, Eggers H, Doessel O. An Optimal Radial Profile Order Based on the Golden Ratio for Time-Resolved MRI. *IEEE Transactions on Medical Imaging*. 2007; 26(1):68–76. [PubMed: 17243585]

29. Wundrak S, Paul J, Ulrici J, Hell E, Rasche V. A small surrogate for the golden angle in time-resolved radial MRI based on generalized fibonacci sequences. *IEEE Transactions on Medical Imaging*. 2015; 34(6):1262–1269. [PubMed: 25532172]
30. Fessler J, Sutton B. Nonuniform fast fourier transforms using min-max interpolation. *IEEE Transactions on Signal Processing*. 2003; 51(2):560–574.
31. Robson PM, Grant AK, Madhuranthakam AJ, Lattanzi R, Sodickson DK, McKenzie CA. Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. *Magnetic Resonance in Medicine*. 2008; 60(4):895–907. [PubMed: 18816810]
32. Cauley SF, Setsompop K, Ma D, Jiang Y, Ye H, Adalsteinsson E, Griswold MA, Wald LL. Fast group matching for MR fingerprinting reconstruction. *Magnetic Resonance in Medicine*. 2014; 528:523–528.
33. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*. 2007; 58(6):1182–1195. [PubMed: 17969013]
34. Knoll F, Holler M, Koesters T, Otazo R, Bredies K, Sodickson DK. Joint MR-PET reconstruction using a multi-channel image regularizer. *IEEE Transactions on Medical Imaging*. 2016 epub ahead of print.
35. Hamilton JI, Deshmane A, Hougen S, Griswold M, Seiberlich N. Magnetic Resonance Fingerprinting with Chemical Exchange (MRF-X) for Quantification of Sub-voxel T1, T2, Volume Fraction, and Exchange Rate. *Proc Intl Soc Mag Reson Med*. 2015; 23:329.
36. Jiang Y, Ma D, Seiberlich N, Gulani V, Griswold MA. MR fingerprinting using fast imaging with steady state precession (FISP) with spiral readout. *Magnetic Resonance in Medicine*. 2015; 74:1621–1631. [PubMed: 25491018]

Appendix

Commutation of the Fourier transformation and the compression matrix

In order to prove $\tilde{\mathbf{F}} \cdot \mathbf{U}_R = \mathbf{U}_R \cdot \mathbf{F}$, we denote \mathbf{F} as a $T \times T$ block matrix containing the Fourier matrix $\mathbf{f} \in \mathbb{C}^{N \times N}$ on the diagonal and zero-matrices elsewhere:

$$\mathbf{F} = \begin{pmatrix} \mathbf{f} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{f} & \ddots & \mathbf{0} \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{f} \end{pmatrix}. \quad [23]$$

In combination with the notation in Eq. [11], the product of \mathbf{U}_R and \mathbf{F} can be written as

$$\mathbf{U}_R \cdot \mathbf{F} = \begin{pmatrix} u_{1,1} \cdot \mathbf{f} & \cdots & u_{T,1} \cdot \mathbf{f} \\ \vdots & \ddots & \vdots \\ u_{1,R} \cdot \mathbf{f} & \cdots & u_{T,R} \cdot \mathbf{f} \end{pmatrix}. \quad [24]$$

If we denote $\tilde{\mathbf{F}}$ as an $R \times R$ block-diagonal matrix with the same matrix \mathbf{f} on the diagonal, we can equivalently denote

$$\tilde{\mathbf{F}} \cdot \mathbf{U}_R = \begin{pmatrix} \mathbf{f} \cdot u_{1,1} \cdot & \cdots & \mathbf{f} \cdot u_{T,1} \cdot \\ \vdots & \ddots & \vdots \\ \mathbf{f} \cdot u_{1,R} \cdot & \cdots & \mathbf{f} \cdot u_{T,R} \cdot \end{pmatrix}. \quad [25]$$

Since the scalars $u_{t,r}$ and the unit matrix $\mathbb{1}$ always commute, Eq. [24] and [25] can be combined to yield

$$\tilde{\mathbf{F}} \cdot \mathbf{U}_R = \mathbf{U}_R \cdot \mathbf{F}. \quad [26]$$

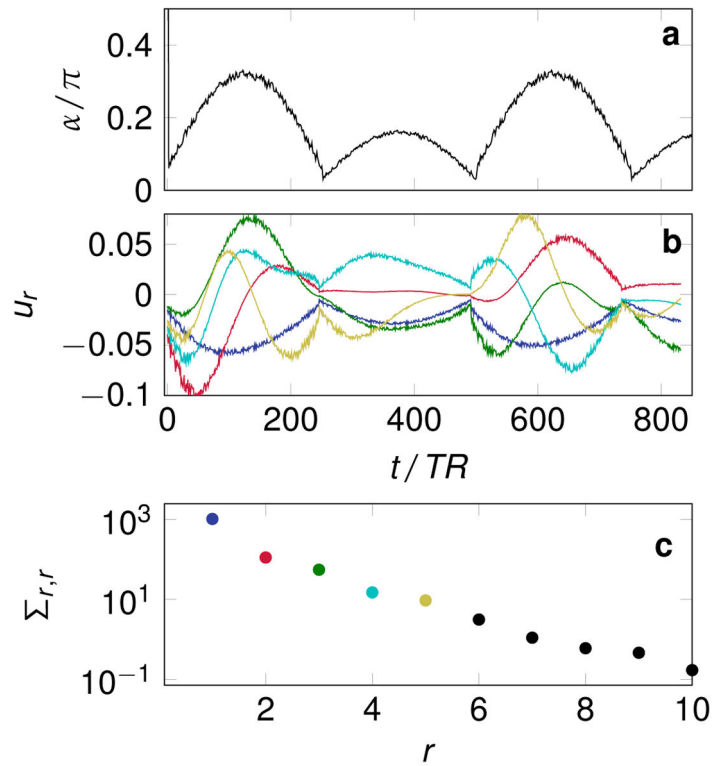


Figure 1. The flip angle pattern is depicted in (a), while (b) shows the first five left singular vectors (columns of u), which result from the singular value decomposition of the dictionary and were used for the low rank approximation unless stated otherwise. Their color-coding corresponds to (c), where the first ten singular values of the dictionary are shown.

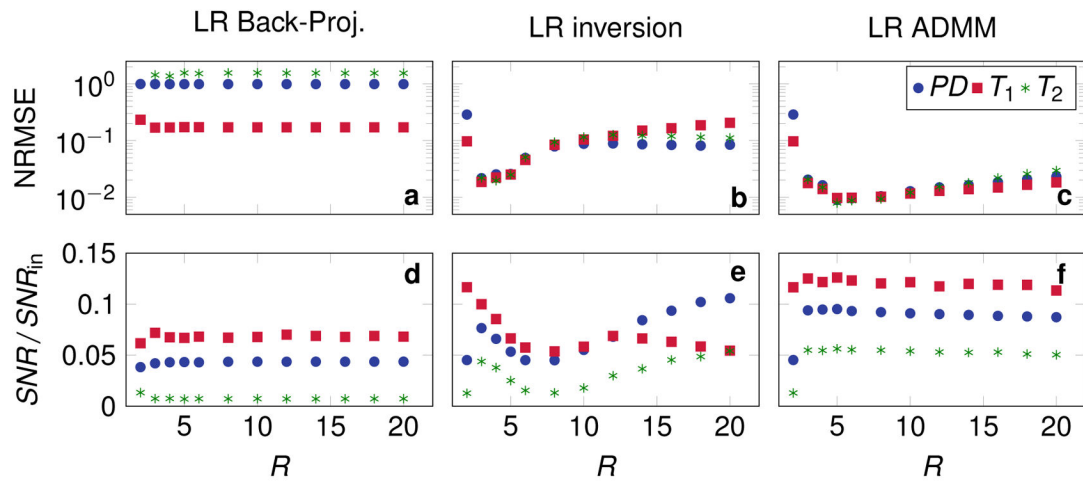


Figure 2. The normalized root mean square error (NRMSE) in white matter is shown (a, b, c) as a function of the rank used to approximate the signal evolution. The NRMSE for a rank 2 approximation is not displayed for T_2 since the values are too large for the displayed area. The PD - T_1 - and T_2 -to noise ratio in white matter are shown in (d, e, f) normalized by the input SNR_{in} .

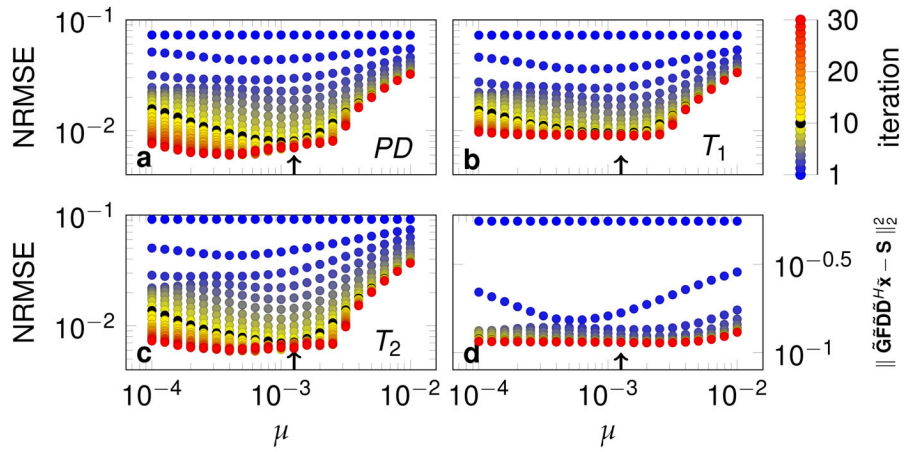


Figure 3. The convergence of the low rank ADMM algorithm is visualized in terms of the normalized root mean square error (NRMSE) of the quantitative maps (**a,b,c**) and the residuum (Eq. [14]) (**d**). The reconstructions were performed with $R = 5$. The color scale visualizes the number of ADMM iterations, where simulations with 10 iterations were highlighted by the color black and have been used for all other simulations in the present paper. The arrows indicate $\mu = 1.26 \cdot 10^{-3}$, which was used for all simulations in this paper, if not stated otherwise.

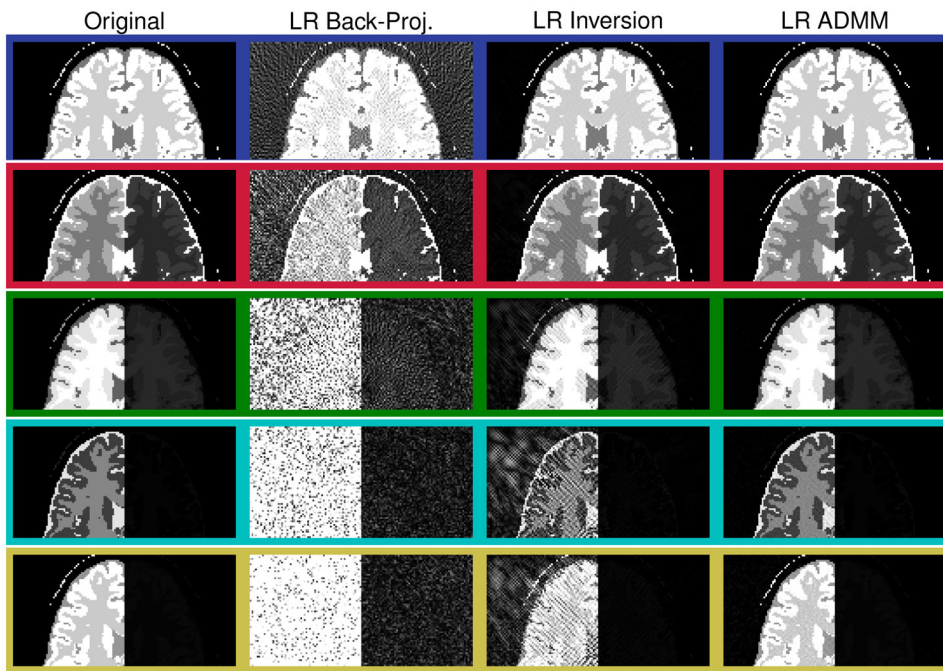


Figure 4. The series of images $\tilde{\mathbf{x}}$ reconstructed from noise free synthetic data are shown. The rows correspond to the singular values in decreasing order and their color coding corresponds to Fig. 1(b, c). The columns show the different low rank (LR) reconstruction techniques discussed in this paper. Only the top half of the images are shown for better depiction. The right hand side of all images are scaled equally, while the left hand side is scaled equally only within the same row.

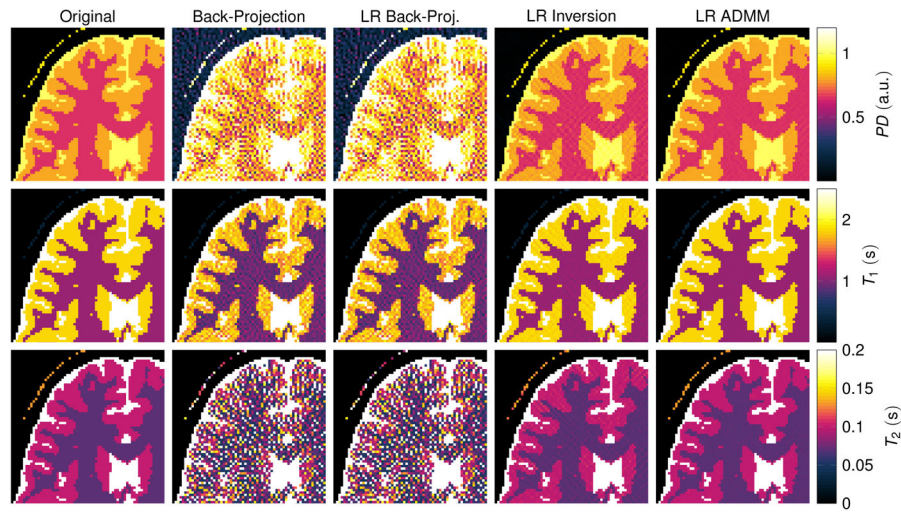


Figure 5.

The displayed excerpts of quantitative maps compare the discussed reconstruction methods. The simulations are based on a numerical phantom and the maps correspond to the singular value images shown in Fig. 4. The maps were reconstructed from noise free data using $R = 5$ singular values.

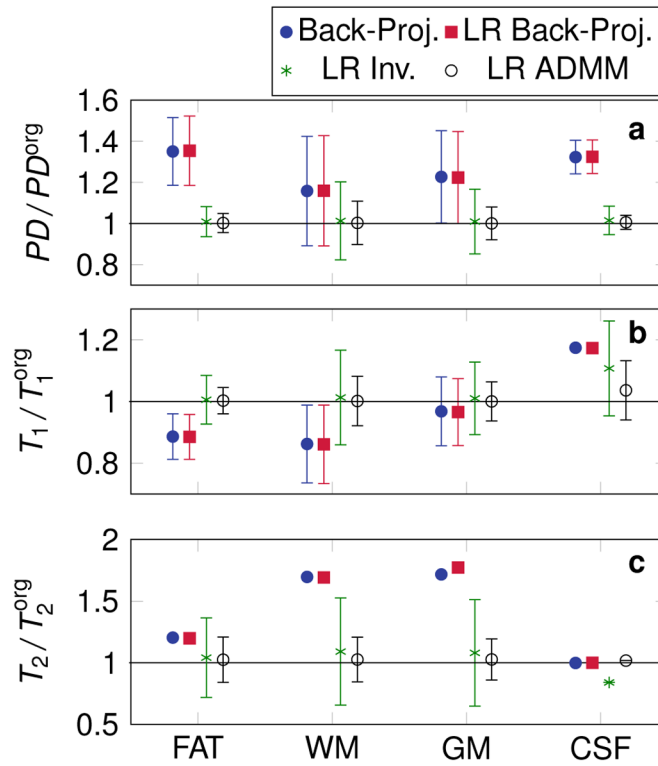


Figure 6.

The mean value and the standard deviation (error bars) within fat, white matter (WM), gray matter (GM) and cerebrospinal fluid (CSF) are shown. The values are normalized by the original parameters used in the simulation. All low rank (LR) reconstructions were performed with $R = 5$. Error bars are not displayed for those cases, where strong deviations from a Gaussian distribution make the calculation of the standard deviation questionable.

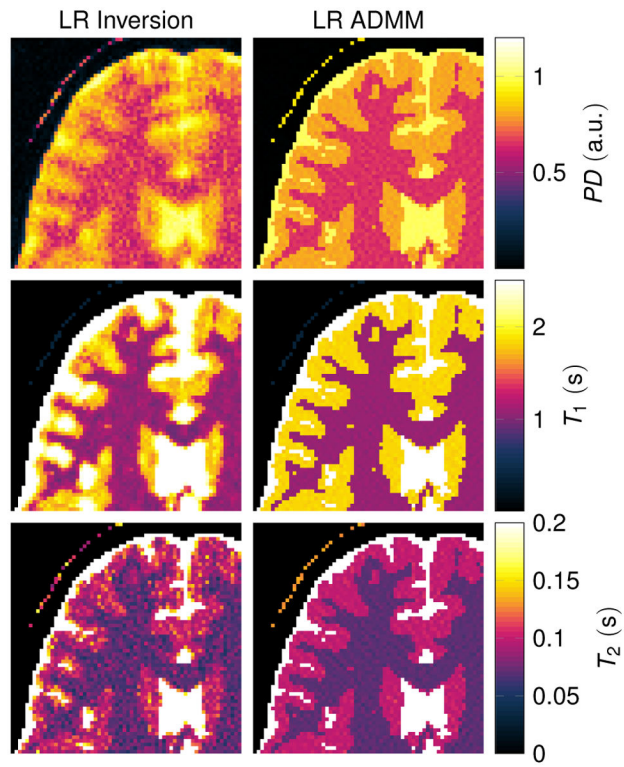


Figure 7. The displayed excerpts of quantitative maps compare low rank inverse reconstructions with $R = 20$. The simulations are based on a numerical phantom and were reconstructed from noise free data.

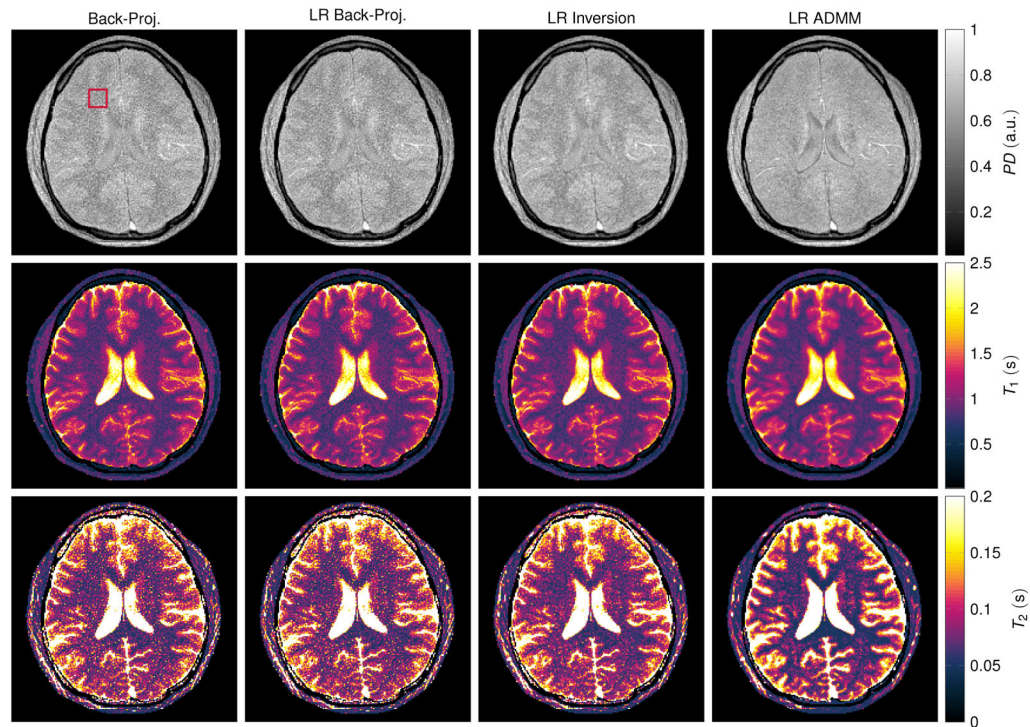


Figure 8.

The depicted parameter maps with a matrix size of 192×192 were reconstructed from 841 radial spokes, which were acquired in vivo within 4.1 s. Coil sensitivity profiles of a 16 channel head coil (compressed to 8 virtual coil elements) were incorporated. The red rectangle at the top left indicates the region of interest in which mean and standard deviation of the relaxation times were calculated (Tab. 1). Note that the relaxation times are biased by slice profile imperfections and by magnetization transfer effects.

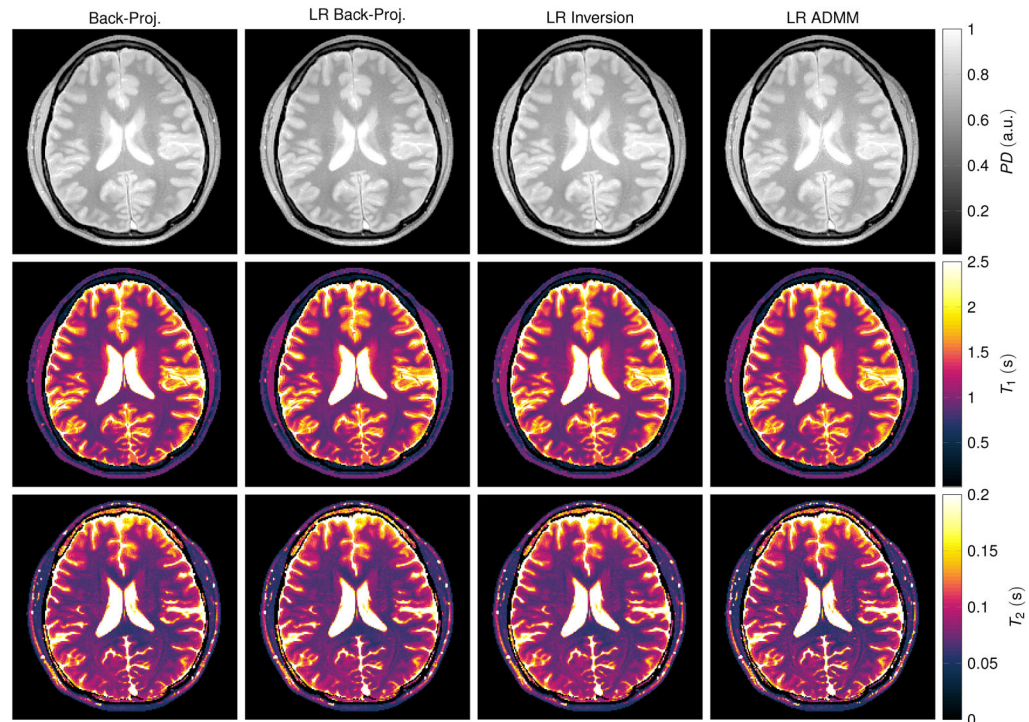


Figure 9.

The depicted parameter maps with a matrix size of 192×192 were reconstructed from 32×841 radial spokes, which were acquired in vivo within 7.35 minutes. Coil sensitivity profiles of a 16 channel head coil (compressed to 8 virtual coil elements) were incorporated. Note that the relaxation times are biased by B_0 -drifts during the scan, by slice profile imperfections and by magnetization transfer effects.

Table 1

Mean value and standard deviation of the relaxation times were calculated over a region of interest in the frontal white matter (Fig. 8). Note that values calculated from 32 repetitions are strongly biased by B_0 -drifts during the scan, while all values are affected by slice profile imperfections and magnetization transfer effects.

Repetitions	1	32	1	32
	T_1 (ms)		T_2 (ms)	
Back-Proj.	835 \pm 99	909 \pm 48	72 \pm 20	70 \pm 5
LR Back-Proj.	835 \pm 100	909 \pm 48	73 \pm 20	70 \pm 5
LR Inversion	838 \pm 81	917 \pm 45	71 \pm 13	71 \pm 4
LR ADMM	826 \pm 62	917 \pm 48	64 \pm 11	70 \pm 6
Literature (36)	781 \pm 61		65 \pm 6	