Magnetic Resonance Imaging 50 (2018) 78–83

Contents lists available at ScienceDirect

Magnetic Resonance Imaging

journal homepage: www.elsevier.com/locate/mri

Original contribution

Water-fat separation incorporating spatial smoothing is robust to noise

Jonathan Andersson[a,*,1], Håkan Ahlström[a,b,1,2], Joel Kullberg[a,b,1,2]

[a] Section of Radiology, Department of Surgical Sciences, Uppsala University, Uppsala, Sweden
[b] Antaros Medical, Mölndal, Sweden

ARTICLE INFO

Keywords:
Dixon
Water-fat separation
Graph cuts
Multi-scale
Chemical shift imaging
Quadratic pseudo-Boolean optimization

ABSTRACT

Purpose: To develop and evaluate a noise-robust method for reconstruction of water and fat images for spoiled gradient multi-echo sequences.

Methods: The proposed method performs water-fat separation by using a graph cut to minimize an energy function consisting of unary and binary terms. Spatial smoothing is incorporated to increase robustness to noise. The graph cut can fail to find a solution covering the entire image, in which case the relative weighting of the unary term is iteratively increased until a complete solution is found.

The proposed method was compared to two previously published methods. Reconstructions were performed on 16 cases taken from the 2012 ISMRM water-fat reconstruction challenge dataset, for which reference reconstructions were provided. Robustness towards noise was evaluated by reconstructing images with different levels of noise added. The percentage of water-fat swaps were calculated to measure performance.

Results: At low noise levels the proposed method produced similar results to one of the previously published methods, while outperforming the other. The proposed method significantly outperformed both of the previously published methods at moderate and high noise levels.

Conclusion: By incorporating spatial smoothing, an increased robustness towards noise is achieved when performing water-fat reconstruction of spoiled gradient multi-echo sequences.

1. Introduction

The idea of utilizing the property of chemical shift for separation of the water and fat signal in MRI was first introduced by Dixon [1]. The most difficult aspect of the signal separation is to take the amplitude of the static field (B_0) inhomogeneity into account. This inhomogeneity, known as the off-resonance, will cause a spatially dependent phase shift that varies linearly with time. The set of the off-resonance of all the voxels forms a field map. Typically, two reasonable off-resonance candidates can be calculated for each voxel, with one of them being correct. In the simple model employing only one fat peak and assuming equal effective transverse relaxation rates for water and fat, choosing the wrong off-resonance will result in the signal from the water being swapped with the signal from the fat in the calculated images. Therefore, this error is known as water-fat swaps. Choosing the correct off-resonance is necessary to calculate correct water and fat images.

Since Dixon's original paper was published, many similar methods have been developed to perform water and fat signal separation of images. In the original paper, two spin echoes were used, and the fat signal was modeled as a single peak. Today, spoiled gradient echo sequences are often used, and many methods have been developed that can make use of arbitrarily many echoes (multi-echo methods) [2–13]. Additionally, a multi-peak fat spectra [3–11,14–17] may be employed for improved signal separation.

Several different methods have been developed for finding the correct off-resonance, based on the assumption of a smooth field map. The enforcement of spatial smoothness has been performed using region growing [2–4,14,18–22], which operates locally. One problem with region growing approaches is that when swaps do occur, they typically propagate well beyond the voxels where they originate [5]. A potentially more robust alternative to region growing are graph cuts [6–11,15,16], which can optimize entire slices or whole volumes at once.

The graph cut method quadratic pseudo-Boolean optimization (QPBO) [23] has previously been used to find which off-resonance out of two candidates is correct for each voxel [7,10,15]. In these methods, an energy function is minimized using QPBO. The energy consists of a unary cost for each voxel based on how well the signal model fits to the

Abbreviations: GOOSE, Globally optimal surface estimation; ICM, Iterated conditional modes; MSGCA, Multi-scale graph cut algorithm; QPBO, Quadratic pseudo-Boolean optimization

* Corresponding author.
E-mail addresses: jonathan.andersson@surgsci.uu.se (J. Andersson), hakan.ahlstrom@radiol.uu.se (H. Ahlström), joel.kullberg@radiol.uu.se (J. Kullberg).
[1] Address: MRT, Entrance 24, Uppsala University Hospital, SE-751 85 Uppsala, Sweden.
[2] Address: BioVenture Hub, Pepparedsleden 1, SE-431 83 Mölndal, Sweden.

https://doi.org/10.1016/j.mri.2018.03.015
Received 24 October 2017; Received in revised form 19 March 2018; Accepted 26 March 2018

data, and a binary cost enforcing spatial smoothness of the field map. For the voxels where QPBO finds a solution, it is guaranteed to be part of an optimal global solution. It is, however, not guaranteed to find a solution for each voxel. In [7], this was resolved by choosing the voxel with the lowest unary cost. In [10], this method was improved by using a multi-scale approach. For voxels where no solution was found, the field map value found at a lower spatial resolution was used. This guaranteed a solution for all voxels and improved robustness towards noise.

Iterated conditional modes (ICM) [24], which converges locally, has been used on its own [12] or in conjunction with other methods [7,10] to optimize the water-fat separation.

In this paper, a method for reconstruction of spoiled gradient multi-echo scans with a constant echo time spacing, expanding upon the method in [7], is proposed and evaluated. We show that by consistently minimizing an energy function, similar to the one used in [7], at a decreased resolution, a greater robustness towards noise is achieved, without deteriorating the solution at low noise levels. Some insight is given as to why the method in [10] performs well at high noise levels in some cases, and fails in others. Noise is a problem at higher resolutions and one of the problems in highly accelerated scans.

2. Methods

2.1. Signal model

The signal model for an arbitrary voxel in a spoiled gradient echo sequence can be expressed as:

$$S_n = (W + a_n F)e^{(i\omega - R_2^*)t_n} \tag{1}$$

where S_n is the expected signal of echo n, excluding noise, in the voxel. The noise is expected to be additive complex Gaussian. W and F are the complex signals for water and fat in the voxel at the time of excitation. a_n is:

$$a_n = \frac{1}{\sum_{m=1}^{M}\alpha_m}\sum_{m=1}^{M}\alpha_m e^{i\,\gamma\,B_0\delta_m\,t_n} \tag{2}$$

where α_m are the relative magnitudes of the M different fat peaks and δ_m are their corresponding chemical shifts relative to water. These values are set to be the same as in the reference signal model of the 2012 ISMRM challenge [17]. That is: $\delta_m = -3.8,\ -3.4,\ -2.6,\ -1.94,\ -0.39$, and 0.6 ppm; with corresponding $\alpha_m = 87, 693, 128, 4, 39$, and 48. The remaining parameters are: t_n, the time of echo n, ω, the off-resonance shift of the voxel, and R_2^*, the effective transverse relaxation rate of the voxel. The gyromagnetic ratio of ^{1}H is denoted by γ.

Under the assumption of equally spaced echoes, Eq. (1) can be written as:

$$\mathbf{S} = \mathbf{BRAx} \tag{3}$$

With $\mathbf{S} = [S_1\ S_2\ \cdots\ S_N]^T$, $\mathbf{B} = \mathrm{diag}[1\ e^{i\omega\Delta t}\ \cdots\ e^{i\omega\Delta t(N-1)}]$, $\mathbf{R} = \mathrm{diag}[e^{-R_2^* t_1}\ e^{-R_2^* t_2}\cdots\cdots e^{-R_2^* t_N}]$, $\mathbf{A} = [a_{i,j}]_{N\times 2}$, with $a_{n,1} = 1$ and $a_{n,2} = a_n$, and $\mathbf{x} = [W\ F]^T$. Δt is the echo time spacing. Note that in this formulation $e^{i\omega t_1}$ has been merged into W and F. This influences the phase of these values, but not their amplitudes, and should therefore not be of concern.

2.2. Parameter estimation

Two complex, W and F, and two real, ω and R_2^*, parameters are to be estimated, meaning a minimum of three complex echoes are needed. The parameter estimation is similar to that of [7], but expands upon it in ways detailed in the Section 4.

For given values of R_2^* and ω, the least-squares estimates of W and F,$\widehat{W}$ and $\widehat{F}$, can be calculated as:

$$\begin{bmatrix}\widehat{W}\\ \widehat{F}\end{bmatrix} = \hat{\mathbf{x}} = \mathbf{A}^{+}\mathbf{R}^{-1}\mathbf{B}^{H}\mathbf{S} \tag{4}$$

where $\mathbf{A}^+$ is the pseudoinverse of $\mathbf{A}$, $\mathbf{R}^{-1}$ is the inverse of $\mathbf{R}$, and $\mathbf{B}^H$ is the conjugate transpose of $\mathbf{B}$. The squared error residual of $\hat{\mathbf{x}}$, J, can be calculated as:

$$J(R_2^*, \omega) = \|(\mathbf{I} - \mathbf{AA}^{+})\mathbf{R}^{-1}\mathbf{B}^{H}\mathbf{S}\| \tag{5}$$

The values of R_2^* and ω are unknown. They can be estimated jointly or decoupled [13]. In a joint estimation, both parameters would be estimated at the same time, while in a decoupled estimation, one parameter is estimated after the other. A joint estimation of the parameters could result in a marginally smaller residual, but for the method to be computationally efficient, they are estimated decoupled as follows:

At first, Eq. (5) is minimized with respect to ω, under the assumption of $R_2^* = 40\ \mathrm{s}^{-1}$. Thanks to the assumption of equally spaced echoes, the residual is periodic with respect to ω, with a period of $\Omega = 2\pi/\Delta t$. The residual is evaluated at 100 equally spaced values of ω over one period, and the two smallest local minima are considered as candidates. For each voxel, one candidate is chosen by minimizing a global energy (E). Let Q be the set of all voxels, N the set of all voxel pairs in the 4-, for 2D images, or 6-, for 3D images, -neighborhood of each other, and $\boldsymbol{\omega}$ the set of ω for all voxels, then define:

$$E(\boldsymbol{\omega}) = \sum_{q\in Q}\lambda_q J_q(R_2^* = 40\ \mathrm{s}^{-1}, \omega_q) + \sum_{(p,q)\in N} w_{p,q} V(\omega_p, \omega_q) \tag{6}$$

The unary residual term, $J_q(R_2^* = 40\ \mathrm{s}^{-1}, \omega_q)$, will assert data fidelity, while the binary term, $V(\omega_p, \omega_q)$, imposes spatial smoothness of the field map. It is defined as:

$$V(\omega_p, \omega_q) = \min(|\omega_p - \omega_q|^2, (\Omega - |\omega_p - \omega_q|)^2) \tag{7}$$

The weights $w_{p,q}$ are defined as:

$$w_{p,q} = \frac{1}{d_{p,q}}\min(J_p''(R_2^* = 40\ \mathrm{s}^{-1}, \omega_p^{\min}), J_q''(R_2^* = 40\ \mathrm{s}^{-1}, \omega_q^{\min})) \tag{8}$$

where $d_{p,q}$ is the Euclidian distance between the voxels p and q in millimeters. This measure is included in the weight to decrease the influence of neighboring voxels by their distance. The double prime symbol ($''$) represents the second derivative with respect to ω. These derivatives are included to have both the unary and the binary terms scale with the signal strength. The superscript of $\omega^{\min}$ indicates that it is the value of ω resulting in the smallest minimum of Eq. (5).

The term λ_q is a regularization parameter that can be spatially varying. It is initially set to 10 for all voxels, which has empirically been found to be a good value [10].

To minimize the energy function in Eq. (6), QPBO [23] is used. The energy function is non-submodular [23]. Unlike certain graph cut optimization methods [25], QPBO can minimize non-submodular energies. QPBO is guaranteed to find the solution, or label, that minimizes the energy for all voxels where it succeeds, but may fail to produce labels for some voxels. A novel approach is used in this work this work to resolve this problem [15]. If there are any unlabeled voxels after running QPBO, the regularization parameter, λ_q, for these voxels is doubled, which modifies the unary cost locally, and QPBO is run again. This process is repeated iteratively until all voxels have become labeled.

To be more robust to noise, a spatial Gaussian smoothing is applied on the residuals before running QPBO. The standard deviation of the Gaussian kernel, σ, is defined per millimeter, and used to decide the degree of smoothing. The amount of smoothing to apply was determined using the first four included cases. Reconstructions with σ set to 0, i.e. no smoothing, and 25 values logarithmically spaced from 0.125 to 8 mm^{-1}, were performed. The value resulting in the lowest average percentage of swaps, calculated as described in Section 2.3., over all tested noise levels, also described in Section 2.3., was used for all further reconstructions.

After this, the solution is fine-tuned using 10 iterations of ICM [24], for the same values of ω and the same assumption of R_2^* as above. Candidates of ω are checked in order of proximity to the previous solution. In each iteration, ω is restricted to be within Ω/10 from its value in the previous iteration. ICM is performed using the original residuals, which lessens the impact of the Gaussian smoothing of the residuals on the resulting images. Additionally, the regularization parameter, λ_q, is set to its original value, since there is no risk of not finding a solution.

Once a value of ω has been chosen for each voxel, Eq. (5) is minimized with respect to R_2^*. The equation is extremely likely to only have one minimum for R_2^* outside of background [6,7].

The residual is minimized under the assumption of a single minimum. Beginning at $R_2^* = 0$, the value is increased by $1\,s^{-1}$ until a local minimum is found, or the maximal considered value, $R_2^* = 144\,s^{-1}$, is reached. If needed the maximal considered value can be increased, resulting in a small increase of the total runtime.

2.3. Evaluation

The proposed method was evaluated against the publically available dataset provided for the 2012 ISMRM challenge on water-fat reconstruction [26]. Performance was compared to two previously published methods [9,10] at varying noise levels. A total of 17 cases are provided in the challenge set from varying vendors, protocols, field strengths, and anatomies. Multi-echo images, foreground masks, and reference fat fraction images are provided. The reference fat fraction images were calculated using an expanded data set with additional echos. Case #3 had to be excluded since the proposed method assumes equally spaced echoes, which was not fulfilled for that case. Case #9 had varying center frequencies between its slices, and was therefore reconstructed slicewise, all other cases were reconstructed in 3D. Since no voxel sizes were provided, they were assumed to be $1.5 \times 1.5 \times 5\,mm^3$ for all cases.

To evaluate for robustness towards noise, complex Gaussian noise was added to the cases. The noise was generated using a pseudorandom number generator. Noise was added at different noise levels, defined as the standard deviation of the added noise divided by the median of the sum of the magnitude of all the echoes of a case within the provided foreground mask. Noise levels from 0 to 0.2 were used, in increments of 0.025.

To measure the quality of the reconstructions, the percentage of swapped voxels in the foregrounds of the images were calculated. Fat fraction (*FF*) images were calculated as $FF = |\widehat{F}|/(|\widehat{F}| + |\widehat{W}|)$. Voxels were considered to be swapped if both the calculated fat fraction differed > 10 percentage points from the reference fat fraction, and if the dominant species differed.

The noise sensitivity of the proposed method was compared to two previously published graph cut methods. These methods will be referred to as multi-scale graph cut algorithm (MSGCA) [10], and globally optimal surface estimation (GOOSE) [9]. These were the two most accurate methods for the 2012 ISMRM challenge on water-fat reconstruction [26] when compared with several others [10]. For both methods, the implementations provided by the authors, and the accompanying default parameters, were used.

Statistical comparisons between the different methods were performed using Wilcoxon signed-rank test. *P*-values < 0.05 were considered significant.

2.4. Implementation

The proposed algorithm was implemented in MATLAB R2016b (MathWorks, Natick, MA), except for QPBO which was called using a C++ MEX file [23]. The code is available online at https://github.com/Snubben-B/FW-Recon-Spatial-Smoothing. All reconstructions were performed on a computer using an Intel Core i7-3770 CPU with 3.40 GHz.

Reconstruction times were measured, excluding the time taken to read and write the data.

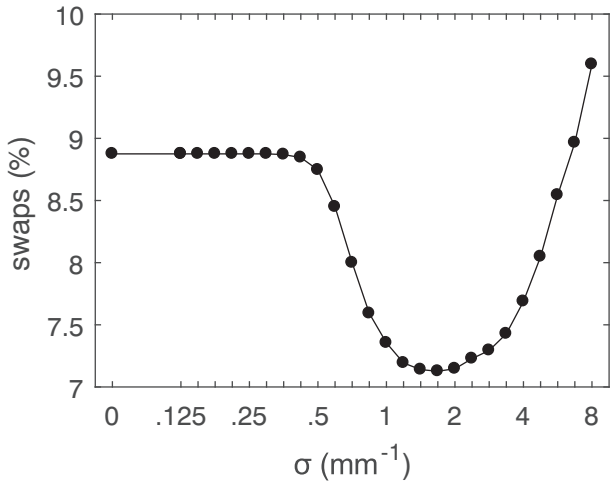

Fig. 1. The average percentage of swaps of the proposed method for the first four included cases, over all the tested noise levels, for different values of σ. Evaluation was performed for the values indicated in the figure, see Section 2.3. for the exact values.

3. Results

3.1. Optimization of the parameter σ

By calculating the average percentage of swaps for the first four included cases over all tested noise levels, the value of σ producing the least swaps was found to be $1.68\,mm^{-1}$. Therefore, this value was used for all further reconstructions. In Fig. 1, a plot of the percentage of swapped voxels can be seen for the different amounts of smoothing.

3.2. Noise sensitivity

In Fig. 2a, a box plot is shown, showing the percentage of swaps for all the included cases for GOOSE, MSGCA, and the proposed method at the different noise levels. The box plot in Fig. 2b shows percentage of swaps for GOOSE respectively MSGCA, minus the percentage of swaps for the proposed method, calculated for each dataset. The percentage of swapped voxels when using the proposed method compared to MSGCA was reduced by a statistically significant amount for noise levels 0.1 and higher ($P < 0.003$ for all). When compared to GOOSE, the proposed method produced fewer swaps for all noise levels ($P < 0.02$ for all).

In Figs. 3 and 4, examples are shown with resulting fat fraction maps using MSGCA and the proposed method. Fig. 3 shows the case where the proposed method performed the worst compared to MSGCA at noise level 0.2, with 0.3 percentage points more swaps, which is difficult, if not impossible, to notice by visual inspection. Fig. 4 shows the case where the proposed method performed the second best compared to MSGCA at noise level 0.2, with 6.1 percentage points less swaps, which is clearly visible. Also shown are the reference fat fraction images and *level maps*, which shows at which resolution MSGCA found its solution. The concept of the different resolutions of MSGCA is

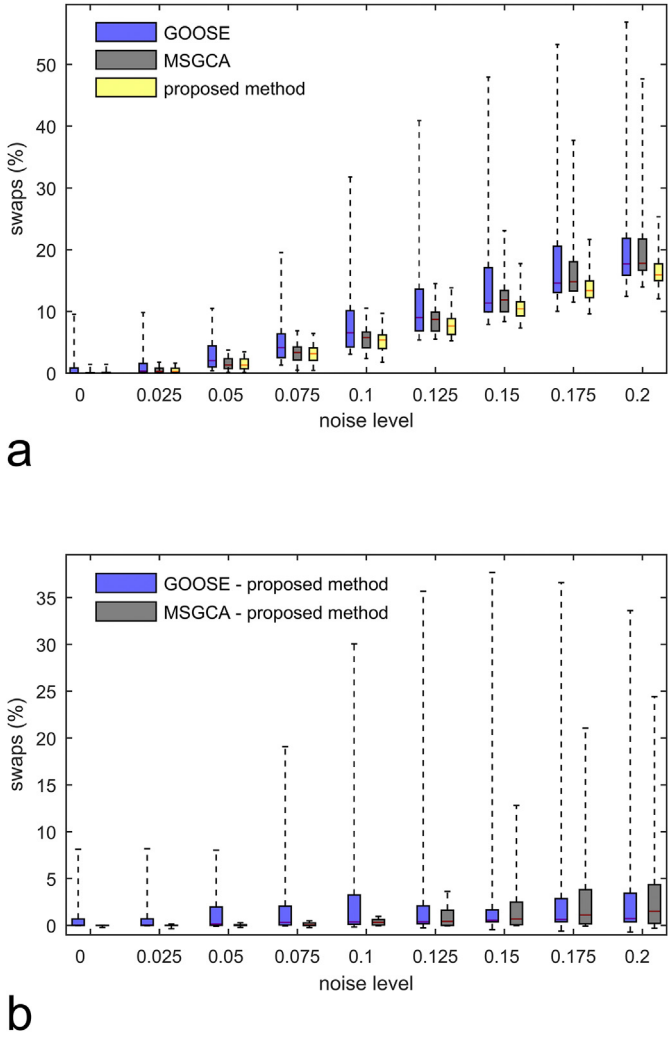

Fig. 2. Box plots of the percentages of swaps for all the included cases at the different noise levels. a: Percentages of swaps for GOOSE, MSGCA, and the proposed method separately. b: percentage of swaps for GOOSE respectively MSGCA, minus the percentage of swaps for the proposed method, calculated for each dataset.

discussed in Section 4. In the level maps, black represents the finest resolution, i.e. at the voxel level, and increasingly bright colors represent increasingly coarse resolutions. Details about the dimensions of the different resolutions have previously been described [10].

In Fig. 5, reconstructions of the fat fraction using MSGCA and the proposed method are shown for noise levels 0.1, 0.15, and 0.2, together with the reference fat fraction image. It can be seen that there is an increased number of swapped patches with increased noise level for the images reconstructed using MSGCA, but no swapped patches are seen for the proposed method.

In Fig. 6, reconstructions of the fat fraction using MSGCA, GOOSE, and the proposed method are shown for noise level 0.2, together with the reference fat fraction image. It can be noticed that the reconstruction performed using MSGCA has several smaller patches that have been swapped, GOOSE has one large swapped patch, and the proposed method has no major swapped patches.

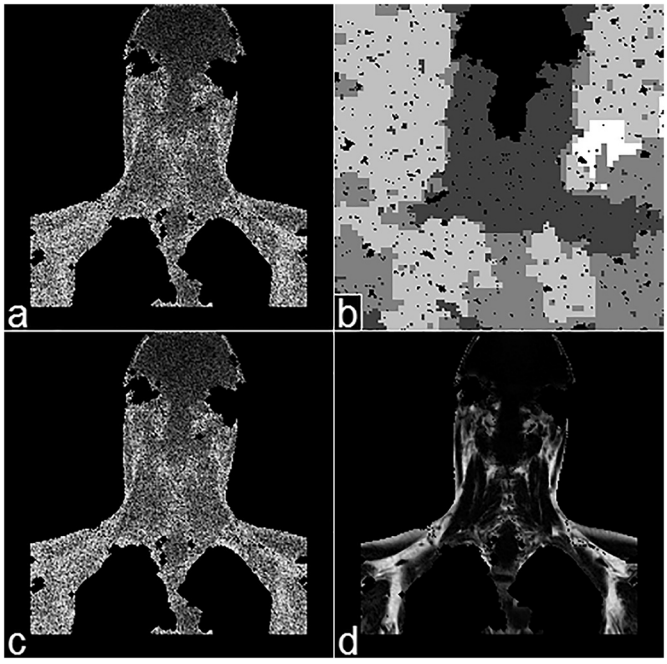

Fig. 3. Fat fraction example of case #2, reference reconstruction, and reconstructed at noise level 0.2. Background has been removed for the fat fraction images. a: Reconstructed with MSGCA, percentage swaps: 25.0; b: level map, notice that for the larger part of the image there was no solution found at the voxel level; c: reconstructed with the proposed method, percentage swaps: 25.3; d: provided reference reconstruction.

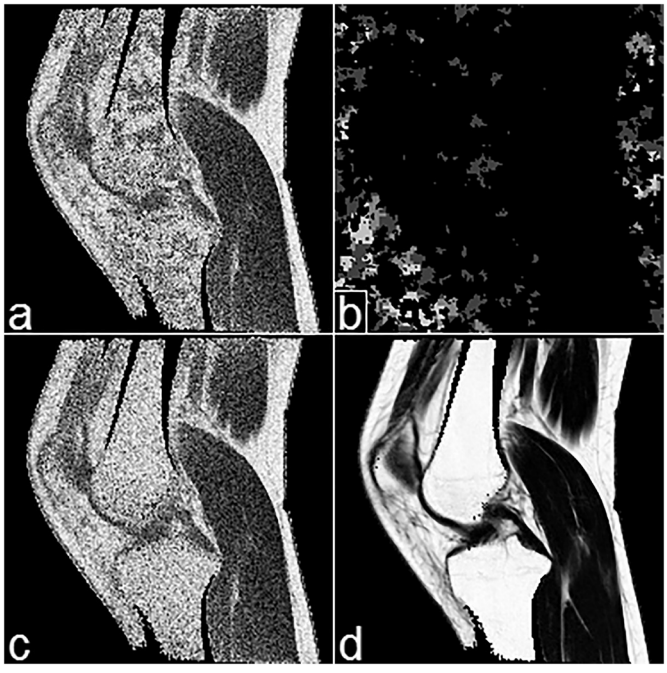

Fig. 4. Fat fraction example of case #4, reference reconstruction, and reconstructed at noise level 0.2. Background has been removed for the fat fraction images. a: Reconstructed with MSGCA, percentage swaps: 18.1, the swaps are especially visible in the cortical bone, and to a lesser degree in the anterior adipose tissue; b: level map, notice that for the larger part of the image there was a solution found at the voxel level; c: reconstructed with the proposed method, percentage swaps: 12.0; d: provided reference reconstruction.

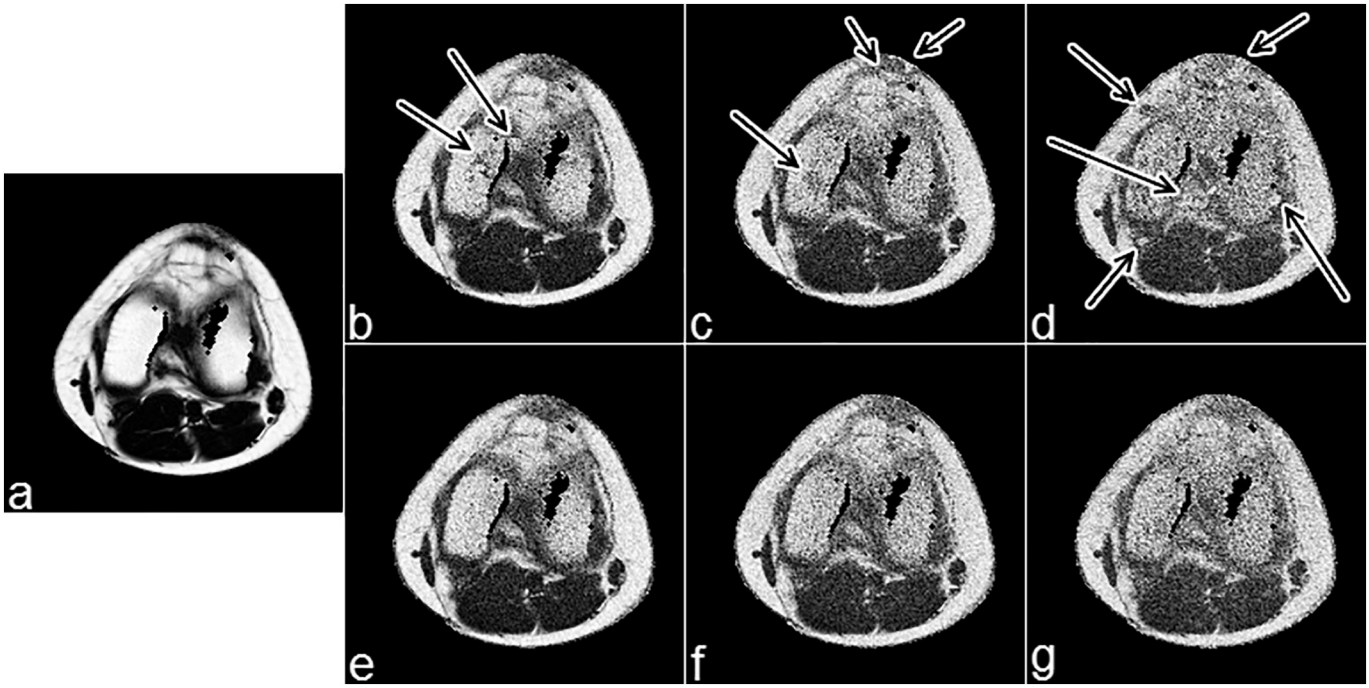

Fig. 5. Fat fraction example of case #1, reference reconstruction and reconstructed at noise levels 0.1, 0.15, and 0.2. Arrows point to patches that have been swapped. Background has been removed for all images. a: Provided reference reconstruction; b–d: reconstructed with MSGCA at noise levels 0.1, 0.15, and 0.2; e–g: reconstructed with the proposed method at noise levels 0.1, 0.15, and 0.2.

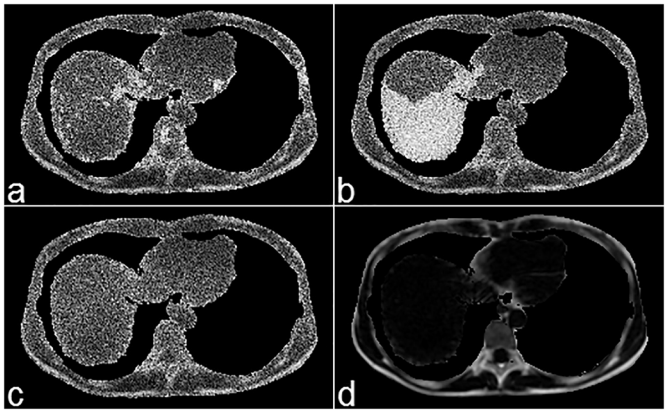

Fig. 6. Fat fraction example of case #4, reference reconstruction and reconstructed at noise level 0.2. Background has been removed for all images. a: Reconstructed with MSGCA, percentage swaps: 22.9, several smaller patches that have been swapped can be seen; b: Reconstructed with GOOSE, percentage swaps: 22.2, one large swapped patch can be seen; c: reconstructed with the proposed method, percentage swaps: 17.4; d: provided reference reconstruction.

3.3. Computation times

All reconstructions with the proposed method were performed in < 75 s. The average reconstruction took 31 s, with a standard deviation of 16 s.

4. Discussion

We have developed a reasonably fast, noise-robust method for reconstruction of water and fat images for spoiled gradient multi-echo sequences. The method has been evaluated against a fat-water challenge dataset provided by ISMRM, and has been compared to two previously published methods, MSGCA and GOOSE, at different noise levels. The proposed method produced better results than GOOSE at all tested noise levels, and better results than MSGCA at moderate to high noise levels, while maintaining MSGCA's good performance at lower noise levels. The proposed method improves upon the same method that MSGCA expands upon [7] in two ways.

The main improvement of the proposed method is the spatial smoothing of the residuals, which decreases the influence of noise. The method has also been improved in that it locally updates the unary cost of the energy function in Eq. (6) iteratively when QPBO fails to label all voxels, until all voxels are labeled. This increases the enforcement of spatial smoothness, as compared to simply choosing the label corresponding to the smallest residual for the unlabeled voxels, as was done in [7].

As a comparison, MSGCA expands upon the method in [7] by the way it acts when QPBO fails to find a solution. To solve the problem with unlabeled voxels, a multi-scale approach was employed. Whenever a voxel would not be labeled, the residuals would be scaled down in a Cartesian fashion, and the energy would be minimized using QPBO at the new resolution. This would be iterated until all voxels would have been assigned an off-resonance frequency. For voxels where off-resonance frequencies were assigned at several resolutions, the value at the finest of these was used. In essence, this multi-scale approach is a coarse form of smoothing.

The good performance at lower noise levels of MSGCA is maintained with the proposed method. The proposed method is shown to be less sensitive to noise than MSGCA, especially the worst cases were greatly improved. MSGCA was in turn shown to be superior at higher noise levels compared to the method it expands upon [7], that uses a single-scale graph cut algorithm. In the cases where MSGCA resulted in fewer swaps than the single-scale graph cut algorithm, this was due to QPBO not finding a solution at the voxel level, since the solution would

otherwise be identical. In these cases, a solution found at a lower resolution was used.

In Fig. 3, an example is shown where MSGCA did not find a solution at the voxel level for most of the image, as can be seen in the level map. Instead, solutions were found at lower resolutions. This resulted in a correct reconstruction that is very similar to the one obtained with the proposed method. In Fig. 4, an example is shown where MSGCA did find a solution at the voxel level for most of the image. This resulted in several swaps, as can be seen in Fig. 4a. However, the proposed algorithm produced a proper reconstruction, as seen in Fig. 4c.

This tells us that MSGCA can provide good results at high noise levels if no solution is found at the finer resolutions. However, since a global minimization of the energy in Eq. (6) is not necessarily optimal for correct water-fat separation, especially at higher noise levels, QPBO risks finding solutions that result in swaps. This means that, at higher noise levels, if solutions are found at the finer resolutions, swaps are likely to follow. This effect is mitigated in the proposed method by always applying Gaussian smoothing to the residuals, which decreases the influence of noise on the energy.

In the proposed method, the ICM step is performed on residuals that have not been smoothed, i.e. they are at full resolution. As mentioned in Section 2.2., this lessen the impact of the Gaussian smoothing of the residuals on the resulting images. It might seem counterintuitive to do this considering that a global minimization of the energy might result in swaps. However, the ICM is merely a fine-tuning step, converging to a local minimum, and should therefore not cause any swaps. In the proposed method, the QPBO and ICM optimization algorithms are performed at different resolutions, the proposed method can therefore be considered to be a multi-scale method.

The proposed method required at least three echoes, although many reconstruction methods, including two-echo methods, could likely be combined with some kind of spatial smoothing, and any method utilizing QPBO could use the same method as presented to iteratively update the energy if not all voxels are labeled.

The proposed method has only been evaluated on one challenge set. While the set contains images from varying vendors, protocols, field strengths, and anatomies, it does not necessarily cover all potential challenges found in clinical practice. Furthermore, the value chosen for the degree of spatial smoothing was chosen using only four of the cases, and might not be optimal in practice.

Gaussian noise was added to evaluate the method's performance in more challenging settings. It can be noted that the more extreme noise levels that the method was evaluated at are not expected to be encountered in a clinical setting, although an improvement is seen already at the clinically more relevant lower noise levels. Moreover, the noise is likely more spatially invariant in reality than the added noise, for example in parallel imaging where it varies with the g-factor.

5. Conclusions

The proposed method is robust to noise and runs reasonably fast. It outperforms GOOSE on the ISMRM challenge dataset at all tested noise levels. It also outperforms MSGCA at moderate and high noise levels, while having the same good performance at lower noise levels.

Funding

This work was supported by the Swedish Research Council [grant number 2016-01040].

References

[1] Dixon WT. Simple proton spectroscopic imaging. Radiology 1984;153:189–94.
[2] Lu W, Hargreaves BA. Multiresolution field map estimation using golden section search for water-fat separation. Magn Reson Med 2008;60:236–44.
[3] Yu H, Reeder SB, Shimakawa A, McKenzie CA, Brittain JH. Robust multipoint water-fat separation using fat likelihood analysis. Magn Reson Med 2012;67:1065–76.
[4] Cheng C, Zou C, Liang C, Liu X, Zheng H. Fat-water separation using a region-growing algorithm with self-feeding phasor estimation. Magn Reson Med 2017;77:2390–401.
[5] Eggers H, Börnert P. Chemical shift encoding-based water–fat separation methods. J Magn Reson Imaging 2014;40:251–68.
[6] Hernando D, Kellman P, Haldar JP, Liang Z-P. Robust water/fat separation in the presence of large field inhomogeneities using a graph cut algorithm. Magn Reson Med 2010;63:79–90.
[7] Berglund J, Kullberg J. Three-dimensional water/fat separation and T_2^* image estimation based on whole-image optimization—application in breathhold liver imaging at 1.5 T. Magn Reson Med 2012;67:1684–93.
[8] Soliman AS, Yuan J, Vigen KK, White JA, Peters TM, CA McKenzie. Max-IDEAL: a max-flow based approach for IDEAL water/fat separation. Magn Reson Med 2014;72:510–21.
[9] Cui C, Wu X, Newell JD, Jacob M. Fat water decomposition using globally optimal surface estimation (GOOSE) algorithm. Magn Reson Med 2015;73:1289–99.
[10] Berglund J, Skorpil M. Multi-scale graph-cut algorithm for efficient water-fat separation. Magn Reson Med 2017;78:941–9.
[11] Cui C, Shah A, Wu X, Jacob M. A rapid 3D fat–water decomposition method using globally optimal surface estimation (R-GOOSE). Magn Reson Med 2018;79:2401–7.
[12] Hernando D, Haldar JP, Sutton BP, Ma J, Kellman P, Liang Z-P. Joint estimation of water/fat images and field inhomogeneity map. Magn Reson Med 2008;59:571–80.
[13] Hernando D, Kellman P, Haldar J, Liang Z-P. Estimation of water/fat images, B_0 field map and T_2^* map using VARPRO. Proceedings of the 16th Annual Meeting of ISMRM, Toronto, Canada. 2008. p. 1517.
[14] Berglund J, Johansson L, Ahlström H, Kullberg J. Three-point Dixon method enables whole-body water and fat imaging of obese subjects. Magn Reson Med 2010;63:1659–68.
[15] Andersson J, Malmberg F, Ahlström H, Kullberg J. Analytical three-point Dixon method using a global graph cut. Proceedings of the 24th Annual Meeting of ISMRM, Singapore, Singapore. 2016. [Abstract nr 3278].
[16] Stinson EG, Trzasko JD, Fletcher JG, Riederer SJ. Dual echo Dixon imaging with a constrained phase signal model and graph cuts reconstruction. Magn Reson Med 2017;78:2203–15.
[17] The 2012 ISMRM Challenge on water-fat reconstruction, Judging. http://challenge.ismrm.org/node/17; 2012 [accessed 17 May 2017].
[18] Xiang Q-S, An L. Water-fat imaging with direct phase encoding. J Magn Reson Imaging 1997;7:1002–15.
[19] An L, Xiang Q-S. Chemical shift imaging with spectrum modeling. Magn Reson Med 2001;46:126–30.
[20] Ma J. Breath-hold water and fat imaging using a dual-echo two-point Dixon technique with an efficient and robust phase-correction algorithm. Magn Reson Med 2004;52:415–9.
[21] Yu H, Reeder SB, Shimakawa A, Brittain JH, Pelc NJ. Field map estimation with a region growing scheme for iterative 3-point water-fat decomposition. Magn Reson Med 2005;54:1032–9.
[22] Ma J. A single-point Dixon technique for fat-suppressed fast 3D gradient-echo imaging with a flexible echo time. J Magn Reson Imaging 2008;27:881–90.
[23] Kolmogorov V, Rother C. Minimizing nonsubmodular functions with graph cuts — a review. IEEE Trans Pattern Anal Mach Intell 2007;29:1274–9.
[24] Besag J. On the statistical analysis of dirty pictures. J Roy Stat Soc B Met 1986;48:259–302.
[25] Kolmogorov V, Zabih R. What energy functions can be minimized via graph cuts? IEEE Trans Pattern Anal Mach Intell 2004;26:147–59.
[26] The 2012 ISMRM Challenge on water-fat reconstruction, Overview. http://challenge.ismrm.org/node/8; 2012 [accessed 1 August 2017].