

Robust Water/Fat Separation in the Presence of Large Field Inhomogeneities Using a Graph Cut Algorithm

Diego Hernando,^{1,2*} P. Kellman,³ J. P. Haldar,^{1,2} and Z.-P. Liang^{1,2}

Water/fat separation is a classical problem for in vivo proton MRI. Although many methods have been proposed to address this problem, robust water/fat separation remains a challenge, especially in the presence of large amplitude of static field inhomogeneities. This problem is challenging because of the nonuniqueness of the solution for an isolated voxel. This paper tackles the problem using a statistically motivated formulation that jointly estimates the complete field map and the entire water/fat images. This formulation results in a difficult optimization problem that is solved effectively using a novel graph cut algorithm, based on an iterative process where all voxels are updated simultaneously. The proposed method has good theoretical properties, as well as an efficient implementation. Simulations and in vivo results are shown to highlight the properties of the proposed method and compare it to previous approaches. Twenty-five cardiac datasets acquired on a short, wide-bore scanner with different slice orientations were used to test the proposed method, which produced robust water/fat separation for these challenging datasets. This paper also shows example applications of the proposed method, such as the characterization of intramyocardial fat. Magn Reson Med 63:79–90, 2010. © 2009 Wiley-Liss, Inc.

Key words: water/fat; graph cut; field map; cardiac MRI; Dixon

In vivo proton MR images contain signals from water and fat protons. Separation of the water and fat signals is a problem of considerable practical importance. In some cases, the fat signal is of diagnostic interest (1–3), and in other circumstances it appears bright and obscures the water signal (4). A number of methods have been developed to address the water/fat separation problem. A straightforward approach is to suppress the fat signal during excitation, which can be done using fat saturation or spatial-spectral pulses (based on the difference in the resonance frequencies of water and fat protons) (5,6), or by signal nulling using a short-tau inversion recovery sequence (based on the short T_1 relaxation time of the fat signal) (7). However, fat suppression has well-known limitations, e.g., high sensitivity to amplitude of static field and amplitude of radio frequency field inhomogeneities, removal of fat signal information, or loss of signal-to-noise ratio (1,4,8).

An alternative is to separate the water and fat components by postprocessing chemical shift-encoded data, which is the crux of the celebrated Dixon method (9) and its many variants. In a chemical shift-based water/fat separation acquisition, a sequence of images is obtained with different echo time (TE) shifts, t_1, t_2, \dots, t_N (typically $N = 3$). The signal at an individual voxel q can be described by the simplified model:

$$s_q(t_n) = e^{i2\pi f_{B,q}t_n}(\rho_{W,q} + \rho_{F,q}e^{i2\pi f_F t_n}), \text{ for } n = 1, \dots, N, \quad [1]$$

where $f_{B,q}$ (in hertz) is the local frequency shift due to static field inhomogeneity, $\rho_{W,q}$ and $\rho_{F,q}$ are the amplitudes of the water and fat components, respectively, and f_F (in hertz) is the frequency shift of fat relative to the water, which is assumed to be known a priori (4,9,10). In this simplified model, T_2^* effects are ignored and the fat signal is considered to have a single spectral line (11,12). These simplifications can be removed if needed, as described in the Materials and Methods section.

The unknowns in the signal model of Eq. [1] are the nonlinear parameter $f_{B,q}$ and the linear parameters $\{\rho_{W,q}, \rho_{F,q}\}$, for $q = 1, \dots, Q$, where Q is the number of voxels. Clearly, estimation of $\{\rho_{W,q}, \rho_{F,q}\}$ is trivial if $f_{B,q}$ is known. However, estimation of $f_{B,q}$ is complicated by the nonlinearity of the signal model. Several practical factors make the problem even more challenging, including the large range of $f_{B,q}$, rapid spatial variation of $f_{B,q}$, presence of low-signal regions, “spectral aliasing” (especially for long TE spacing, or at high field), and ambiguities and inaccuracies in the signal model (for instance, the signal model in Eq. [1] is ambiguous in voxels containing only water or only fat) (4,8,11–14).

A number of methods have been proposed for water/fat separation, which differ essentially in how they address the effects of field inhomogeneities in the acquired signal. Dixon’s (9) original method assumes $f_{B,q} = 0$ and performs water/fat separation using only two images. Glover and Schneider (10) proposed a three-point method ($N = 3$) where the t_n are chosen such that $f_{B,q}$ can be estimated directly from the first and third images, avoiding the nonlinearity of the problem. Xiang and An (15) proposed a method (termed “direct phase encoding”) that allows analytical separation of water and fat for a broader choice of t_n than the original three-point method. An and Xiang (16) introduced a method for fitting multiple spectral components using nonlinear least squares. Ma (17) introduced an improved two-point method where phase errors due to field inhomogeneities are corrected using a region-growing algorithm. Reeder et al. (4) introduced a novel method for iterative decomposition of water and fat with echo asymmetry and least squares estimation (IDEAL) where $\{f_{B,q}, \rho_{W,q}, \rho_{F,q}\}$ are estimated independently at each voxel by an iterative nonlinear least squares fitting procedure.

¹Department of Electrical and Computer Engineering, University of Illinois at Urbana–Champaign, Urbana, Illinois, USA

²Beckman Institute for Advanced Science and Technology, University of Illinois at Urbana–Champaign, Urbana, Illinois, USA

³Laboratory of Cardiac Energetics, National Heart, Lung, and Blood Institute, National Institutes of Health, Department of Health and Human Services, Bethesda, Maryland, USA

*Correspondence to: Diego Hernando, Beckman Institute for Advanced Science and Technology, 405 N. Mathews Ave, Urbana, IL 61801. Phone: 1-217-244-4278, E-mail: dhernan2@illinois.edu

Received 1 October 2008; revised 9 July 2009; accepted 22 July 2009.

DOI 10.1002/mrm.22177

Published online 26 October 2009 in Wiley InterScience (www.interscience.wiley.com).

The IDEAL method has several desirable properties. For example, it works for arbitrary echo times and can result in the maximum-likelihood water/fat decomposition. However, the original IDEAL method has trouble dealing with large field inhomogeneities, due to the implicit assumption that the field inhomogeneity is moderate and the fact that only local convergence is guaranteed. Several extensions of IDEAL have been proposed in recent years to address the problem with large field inhomogeneities. Yu et al. (18) proposed a region-growing extension of IDEAL where field map smoothness is imposed by a region-growing process initialized with an automatically selected seed voxel. Tsao and Jiang (19) proposed a multiresolution method to help guide the selection of the correct decomposition at each voxel. Lu and Hargreaves (20) developed a method that combines region-growing and multiresolution by using region-growing at the coarsest resolution and propagating the resulting estimates to the finer resolutions.

This article reports a new method to estimate $\{f_{B,q}, \rho_{W,q}, \rho_{F,q}\}$, for $q = 1, \dots, Q$, jointly for all the voxels (in contrast to voxel-based estimation). Relative to a previous method presented in Hernando et al. (14) (and applied in Kellman et al. (3)), this article introduces: (i) a novel optimization method based on graph cuts, with improved theoretical properties and practical performance; (ii) a different weighting scheme for the cost function, designed to address problems with rapid field variations; and (iii) a novel and more detailed analysis of the spatial resolution properties of the estimated field map, as well as its effects on the resulting water/fat images. In the remainder of this paper, we will describe the proposed method and show some representative results from challenging cardiac imaging applications to demonstrate its performance.

THEORY

Joint Estimation of Water/Fat Images and Field Map

Under the usual assumption of white additive Gaussian noise, the maximum likelihood estimate of $\{\rho_{W,q}, \rho_{F,q}, f_{B,q}\}$ in Eq. [1] is obtained by minimizing the following cost function at each voxel q (as previously proposed (4,16)):

$$R_0(\rho_{W,q}, \rho_{F,q}, f_{B,q}; \mathbf{s}_q) = \sum_{n=1}^N |s_q(t_n) - e^{i2\pi f_{B,q} t_n} (\rho_{W,q} + \rho_{F,q} e^{i2\pi f_F t_n})|^2 \quad [2]$$

where $\mathbf{s}_q = [s_q(t_1) \dots s_q(t_N)]^T$.

However, minimizing $R_0(\rho_{W,q}, \rho_{F,q}, f_{B,q}; \mathbf{s}_q)$ voxel by voxel (as is done in conventional voxel-based water/fat separation methods) is undesirable because (a) $R_0(\rho_{W,q}, \rho_{F,q}, f_{B,q}; \mathbf{s}_q)$ has multiple local and global minimizers (18,20), and (b) the maximum likelihood estimates from Eq. [2] are sensitive to noise and often require postestimation smoothing of the field map (4). To address both of these issues, we minimize $R_0(\rho_{W,q}, \rho_{F,q}, f_{B,q}; \mathbf{s}_q)$ for $q = 1, \dots, Q$ jointly, which allows us to impose spatial smoothness on the field map. Invoking the penalized maximum likelihood framework, we can formulate the estimation of the complete field map $\mathbf{f}_B = \{f_{B,q}\}_{q=1}^Q$ and water/fat images $\rho_W = \{\rho_{W,q}\}_{q=1}^Q$ and $\rho_F = \{\rho_{F,q}\}_{q=1}^Q$ as:

$$\{\hat{\rho}_W, \hat{\rho}_F, \hat{\mathbf{f}}_B\} = \arg \min_{\substack{\mathbf{f}_B \in \mathbb{R}^Q \\ \rho_W, \rho_F \in \mathbb{C}^Q}} \sum_{q=1}^Q R_0(\rho_{W,q}, \rho_{F,q}, f_{B,q}; \mathbf{s}_q) + \mu \sum_{q=1}^Q \sum_{j \in \delta_q} w_{q,j} V(f_{B,q}, f_{B,j}) \quad [3]$$

where δ_q is the local neighborhood of voxel q , μ is a regularization parameter balancing data consistency and smoothness of the solution, $w_{q,j}$ are spatially dependent weights, and $V(f_{B,q}, f_{B,j})$ penalizes the roughness of the field map. In this work, δ_q is the second-order neighborhood (which, in two dimensions, includes the eight voxels surrounding q) (21), and a quadratic penalty, $V(f_{B,q}, f_{B,j}) = (f_{B,q} - f_{B,j})^2$, is chosen to promote field map smoothness (14,22,23). The selection of μ and $w_{q,j}$ is discussed in the next section.

Optimization Algorithm

Joint estimation of $\{\rho_W, \rho_F, \mathbf{f}_B\}$ using the penalized maximum likelihood formulation in Eq. [3] has several significant computational challenges:

- *High dimension.* The space of all possible solutions has $5Q$ dimensions because each voxel contains two complex-valued parameters ($\rho_{W,q}, \rho_{F,q}$) and one real-valued parameter ($f_{B,q}$). In practice, the solution space has on the order of 10^5 dimensions for the datasets considered in this paper.
- *Nonconvexity.* The cost function is nonconvex and presents the usual difficulties of nonconvex optimization (e.g., gradient-based methods only guarantee local convergence and depend heavily on the initialization) (24).
- *Multiple local minima.* The cost function has a very large number of local (and often global) minima due to the complex exponential form of the signal model (Eq. [1]). Convergence to suboptimal local minima typically results in inaccurate water/fat separation (18).

We have developed a novel method to address these challenges. The proposed method is based on the following key components: (a) use of variable projection (VARPRO) for dimensionality reduction, (b) conversion of Eq. [3] to a discrete optimization problem, and (c) use of a novel graph cut-based algorithm to efficiently solve the discretized problem. These components are discussed next.

Dimensionality Reduction Using VARPRO

$R_0(\rho_{W,q}, \rho_{F,q}, f_{B,q}; \mathbf{s}_q)$ has a particular mathematical structure that lends itself straightforwardly to the VARPRO formulation. Specifically, the nonlinear parameter $f_{B,q}$ can be estimated by minimizing (14):

$$R(f_{B,q}; \mathbf{s}_q) = \|\mathbf{I} - \Psi(f_{B,q})\Psi^\dagger(f_{B,q})\mathbf{s}_q\|_2^2 \quad [4]$$

where $\Psi(f_{B,q})$ is a $N \times 2$ matrix with entries $[\Psi(f_{B,q})]_{(n,1)} = e^{i2\pi f_{B,q} t_n}$ and $[\Psi(f_{B,q})]_{(n,2)} = e^{i2\pi (f_F + f_{B,q}) t_n}$, for $n = 1, \dots, N$, and \dagger denotes pseudoinverse.

Note that VARPRO effectively isolates the key component of water/fat separation: field map estimation. Thus,

the field map estimate for the regularized problem in Eq. [3] can be equivalently expressed as:

$$\hat{\mathbf{f}}_{\mathbf{B}} = \arg \min_{\mathbf{f}_{\mathbf{B}} \in \mathbb{R}^Q} \sum_{q=1}^Q R(f_{\mathbf{B},q}; \mathbf{s}_q) + \mu \sum_{q=1}^Q \sum_{j \in \delta_q} w_{q,j} V(f_{\mathbf{B},q}, f_{\mathbf{B},j}), \quad [5]$$

where the dimension of the problem is now reduced to Q . Estimation of $\{\rho_{\mathbf{W}}, \rho_{\mathbf{F}}\}$ is performed subsequently by solving the corresponding linear problem at each voxel (Eq. [1]), which can be done very efficiently (14).

Problem Discretization

$R(f_{\mathbf{B},q}; \mathbf{s}_q)$ contains many local minima at each voxel, so gradient-based methods often converge to a suboptimal solution. This limitation can be overcome by discretizing the problem (14). The proposed method constrains $f_{\mathbf{B},q}$ to a discrete set of possible values $\Omega = \{\psi_l\}_{l=1}^L$, where the ψ_l are uniformly spaced with spacing 2–4 Hz over a range ± 1500 Hz. This spacing was found to introduce only negligible errors in water/fat separation. The wide range of Ω accounts for the potentially very large field inhomogeneities that often appear near the edges of the field of view (FOV), particularly in short, wide-bore scanners. Note that, for the usual acquisitions with uniformly spaced TEs ($t_n = t_0 + n\Delta t$) (4,13), $R(f_{\mathbf{B},q}; \mathbf{s}_q)$ is periodic with period $1/\Delta t$. In this case, even though Ω spans ± 1500 Hz, it suffices to evaluate $R(f_{\mathbf{B},q}; \mathbf{s}_q)$ on the interval $[0, 1/\Delta t]$ (20). Limiting $\mathbf{f}_{\mathbf{B}} \in \Omega^Q$ yields the following discrete optimization problem:

$$\hat{\mathbf{f}}_{\mathbf{B}} = \arg \min_{\mathbf{f}_{\mathbf{B}} \in \Omega^Q} \sum_{q=1}^Q R(f_{\mathbf{B},q}; \mathbf{s}_q) + \mu \sum_{q=1}^Q \sum_{j \in \delta_q} w_{q,j} V(f_{\mathbf{B},q}, f_{\mathbf{B},j}). \quad [6]$$

Next, we describe a graph cut-based algorithm to solve this optimization problem.

Solution Using Graph Cuts

The size of the set Ω^Q in Eq. [6] is L^Q (i.e., the total number of possible field maps in the formulation). For typical image sizes (e.g., 256×256 , so $Q = 65,536$), and discretization levels (e.g., $L = 1000$), $L^Q = 1000^{65,536}$. Therefore, the set Ω^Q is too large for any exhaustive search. This paper presents an algorithm that subdivides the problem in Eq. [6] into a sequence of binary decision problems and solves each of them efficiently using a graph cut algorithm at each iteration. Specifically, let Γ be a subset of Ω^Q , defined as:

$$\Gamma = \hat{\Omega}_1 \times \hat{\Omega}_2 \times \cdots \times \hat{\Omega}_Q \quad [7]$$

where $\hat{\Omega}_q = \{\hat{f}_{\mathbf{B},q}, \hat{f}'_{\mathbf{B},q}\}$, $q = 1, \dots, Q$ are binary sets. We further assume that $\hat{f}_{\mathbf{B},q}$ is the current field map estimate at voxel q , and $\hat{f}'_{\mathbf{B},q}$ is a potential update of $\hat{f}_{\mathbf{B},q}$ for the next iteration. Limiting $\mathbf{f}_{\mathbf{B}} \in \Gamma$ yields the following discrete optimization problem at each iteration:

$$\hat{\mathbf{f}}_{\mathbf{B}} = \arg \min_{\mathbf{f}_{\mathbf{B}} \in \Gamma} \sum_{q=1}^Q R(f_{\mathbf{B},q}; \mathbf{s}_q) + \mu \sum_{q=1}^Q \sum_{j \in \delta_q} w_{q,j} V(f_{\mathbf{B},q}, f_{\mathbf{B},j}). \quad [8]$$

Even though Γ is still too large (with size 2^Q) for exhaustive search, Eq. [8] can be solved very efficiently by mapping it to an equivalent graph cut problem (25–30); details on how to perform the mapping are provided in the Appendix.

With the graph cut algorithm guaranteeing the global minimum of Eq. [8], the key to solving Eq. [6] is the design of Γ at each iteration (i.e., choosing $\hat{f}'_{\mathbf{B},q}$). In this work, we use three different constructions for Γ , corresponding to different choices of $\hat{f}'_{\mathbf{B},q}$:

$$\begin{aligned} \Gamma_{\beta}: \hat{f}'_{\mathbf{B},q} &= \hat{f}_{\mathbf{B},q} + \beta \\ \Gamma_{+}: \hat{f}'_{\mathbf{B},q} &= \min_m \{f_{\mathbf{B},q}^{\min,m}\} \text{ s.t. } f_{\mathbf{B},q}^{\min,m} > f_{\mathbf{B},q} \\ \Gamma_{-}: \hat{f}'_{\mathbf{B},q} &= \max_m \{f_{\mathbf{B},q}^{\min,m}\} \text{ s.t. } f_{\mathbf{B},q}^{\min,m} < f_{\mathbf{B},q} \end{aligned} \quad [9]$$

where β is a constant, and $\{f_{\mathbf{B},q}^{\min,m}\}$ is the set of local minimizers of $R(f_{\mathbf{B},q}; \mathbf{s}_q)$ at voxel q . In noise-only voxels (identified using a threshold on the signal amplitude), the locations of local minima are meaningless, and thus the “jumps” corresponding to the separation between local minimizers in a voxel with a single component are used in Γ_{+} and Γ_{-} . Note that Γ_{β} corresponds to a uniform “jump” with step size β applied to all the voxels (31,32), whereas Γ_{+} and Γ_{-} correspond to voxel-dependent jumps (see Fig. 1). In practice, iterations based on Γ_{+} and Γ_{-} provide rapid convergence to the correct “valley” of $R(f_{\mathbf{B},q}; \mathbf{s}_q)$ at each voxel (their role is similar to the search for the correct local minima performed in Lu and Hargreaves (20); in practice, we set the first few, e.g., 15, iterations to be of these kinds), and iterations based on Γ_{β} perform fine-tuning. A simple proof of the equivalence of the proposed iterations to a graph cut problem (28) is given in the Appendix.

In this work, we employ a randomized scheduling of the proposed iterations, where, at each iteration, Γ_{β} (with random step size β in the range ± 20 Hz), Γ_{+} or Γ_{-} is used (33). Upon convergence, the solution $\hat{\mathbf{f}}_{\mathbf{B}}$ is optimal with respect to an exponentially large set (33). An example of the evolution of $\hat{\mathbf{f}}_{\mathbf{B}}$ in the proposed algorithm is shown in Fig. 2. A key advantage over previous methods (4,14,18,20) is the ability to simultaneously update $\hat{\mathbf{f}}_{\mathbf{B}}$ for arbitrary sets of voxels, thus enabling the proposed algorithm to escape suboptimal solutions, where methods that consider one voxel at a time may be trapped.

Selection of the Regularization Parameters μ and $w_{q,j}$

Selection of the regularization parameter μ and the spatial weights $w_{q,j}$ is based on the resolution properties of the estimated field map. In this work, the weights are set to:

$$w_{q,j} = \min \left(\left. \frac{\partial^2 R(f_{\mathbf{B},q}; \mathbf{s}_q)}{\partial f_{\mathbf{B},q}^2} \right|_{f_{\mathbf{B},q} = \hat{f}_{\mathbf{B},q}^{\min}}, \left. \frac{\partial^2 R(f_{\mathbf{B},j}; \mathbf{s}_j)}{\partial f_{\mathbf{B},j}^2} \right|_{f_{\mathbf{B},j} = \hat{f}_{\mathbf{B},j}^{\min}} \right) \quad [10]$$

where $\hat{f}_{\mathbf{B},q}^{\min}$ is the minimizer of $R(f_{\mathbf{B},q}; \mathbf{s}_q)$, for $q = 1, \dots, Q$. This choice is obtained by approximating $R(f_{\mathbf{B},q}; \mathbf{s}_q)$ by a quadratic function near its minimizer and results in approximately uniform spatial smoothing of the field map (34). The second derivatives in Eq. [10] are easily approximated after the discretization. The degree of smoothing is




FIG. 1. Example of $R(f_B; \mathbf{s}_q)$ at an individual voxel. Note the nonconvexity of $R(f_B; \mathbf{s}_q)$, which contains multiple local minimizers, and no unique global minimizer. Given $\hat{f}_{B,q}$ as the current field inhomogeneity estimate at voxel q , $\hat{\Omega}_q = \{\hat{f}_{B,q}, \hat{f}'_{B,q}\}$ is the binary set for Γ in the proposed algorithm. There are three choices for $\hat{f}'_{B,q}$, corresponding to Γ_+ , Γ_- , and Γ_β (with $\hat{f}'_{B,q} = \hat{f}_{B,q} + \beta$).

then determined by μ , which is empirically set to 0.02 for all the datasets processed in this work. The effect of varying μ is analyzed in the Discussion section.

Advanced Signal Models

In addition to the standard signal model (Eq. [1]), the proposed method can easily be extended to handle more advanced signal models:


- T_2^* decay. The presence of significant T_2^* decay can severely bias the estimates of the water/fat images, if not included in the signal model. Generally, the water

and fat components in a given voxel experience different T_2^* decay rates. However, estimating two separate decay rates significantly increases noise sensitivity. Even though separate rates can be estimated with more images (35), it is common to model the decay by a single decay rate $R_{2,q}^* = 1/T_2^*$ at each voxel q . The corresponding signal model becomes (11,36–39):

$$s_q(t_n) = e^{-R_{2,q}^* t_n} e^{i2\pi f_{B,q} t_n} (\rho_{W,q} + \rho_{F,q} e^{i2\pi f_F t_n}),$$

for $n = 1, \dots, N$. [11]

FIG. 2. Results to illustrate convergence of the proposed method. (Top) Estimated field map at several iterations. (Center) Corresponding water images. (Bottom) Corresponding fat images. The ability of the graph cut algorithm to update a large set of voxels at any iteration results in rapid convergence of the proposed method, even in the presence of large field inhomogeneities. Additionally, no complicated initialization heuristics are necessary (the field inhomogeneity map can simply be initialized to zero).



The above signal model can easily be included in the proposed method by redefining $R(f_{B,q}; \mathbf{s}_q)$ as:

$$R(f_{B,q}; \mathbf{s}_q) = \min_{\substack{R_{2,q}^* \in \mathbb{R} \\ \rho_{W,q}, \rho_{F,q} \in \mathbb{C}}} \sum_{n=1}^N |s_q(t_n) - e^{-R_{2,q}^* t_n} e^{i2\pi f_{B,q} t_n} (\rho_{W,q} + \rho_{F,q} e^{i2\pi f_F t_n})|^2, \quad [12]$$

where $\{\rho_{W,q}, \rho_{F,q}\}$ can be removed using VARPRO, and the minimization with respect to $R_{2,q}^*$ is performed by discretizing the $R_{2,q}^*$ values (note that estimation of $R_{2,q}^*$ does not present the same difficulty as f_B , associated with the multitude of local minima). Therefore, the field map estimation algorithm, which depends only on $R(f_{B,q}; \mathbf{s}_q)$, remains unchanged (36).

- **Multipeak fat model.** This model allows the fat signal to have M distinct peaks (often, $M = 3$). As a result, the signal at voxel q can be expressed as (39,40):

$$s_q(t_n) = e^{i2\pi f_{B,q} t_n} \left(\rho_{W,q} + \rho_{F,q} \sum_{m=1}^M \alpha_{m,q} e^{i2\pi f_{F,m} t_n} \right), \quad \text{for } n = 1, \dots, N, \quad [13]$$

where $|\alpha_{1,q}| + |\alpha_{2,q}| + \dots + |\alpha_{M,q}| = 1$ for $q = 1, \dots, Q$, and $\{f_{F,m}\}_{m=1}^M$ are the (known) frequency shifts of the M individual fat peaks.

This model requires $N \geq M + 2$ acquisitions to estimate $f_{B,q}$, $\rho_{W,q}$ and $\{\rho_{F,q} \alpha_{m,q}\}_{m=1}^M$, which may not be practical. Alternatively, we may assume that $\{\alpha_{m,q}\}_{m=1}^M$ are known (or calibrated from the data themselves) and $\alpha_{m,q} = \alpha_m$ for $q = 1, \dots, Q$. Under this assumption, the multipeak model (Eq. [13]) contains the same number of unknown parameters as the original signal model in Eq. [1], so $N = 3$ acquisitions are sufficient to perform the separation (12,37,41). Since the only non-linear parameter in the multipeak model is the field map, the proposed method applies naturally to this model. The only modification necessary is substituting the second column in $\Psi(f_{B,q})$ by the corresponding linear combination of fat peaks from Eq. [13].

MATERIALS AND METHODS

Data for quantitative evaluation were acquired on a Siemens Magnetom Espree (Siemens AG Medical Solutions, Erlangen, Germany) 1.5-T scanner using a phased-array coil, in accordance with the local institutional review board. Twenty-five cardiac datasets were acquired (from 21 subjects), of which 15 were short-axis slice orientation and 10 were long-axis orientation. Imaging was performed with an ECG-triggered GRE sequence, using an echo-train with monopolar readout. Typical parameters included FOV = 36 cm \times 27 cm; bandwidth = 977 Hz/pixel; pulse repetition time = 11.2 ms; flip angle = 20° to 25°; matrix size = 256 \times 126, TE spacing between 1.9 ms and 3.07 ms (3). Usually four echoes were collected (often selected to provide nearly optimal noise properties (13)), but only three are used in this work, to conform more closely to the common conditions used in water/fat separation (4,10). One

additional dataset, not included in the quantitative results, was acquired on a Siemens Avanto 1.5-T scanner with TEs {3.6, 5.8, 7.9}ms.

The proposed algorithm was run on each of the acquired two-dimensional slices, for 50 iterations in all cases, at which point the changes in the estimated field map were negligible. Multicoil data were processed jointly, as described in Hernando et al. (14). In order to evaluate the reliability of the proposed method, water/fat separation was performed on 25 cardiac datasets acquired with various slice orientations. Three echoes ($N = 3$) were used for each dataset. For comparison, the same datasets were also processed using a previously proposed method, based on VARPRO and the iterated conditional modes (ICM) algorithm, where the voxels are updated one at a time (14). Both methods included T_2^* decay in the signal model. By visual inspection of the resulting decompositions, we counted the number of images containing errors (e.g., localized water/fat swaps). These swaps are defined as estimation errors where the main signal component in a voxel is assigned to the wrong chemical species (e.g., identifying as mostly fat a voxel that contains mostly water). Additionally, some of the datasets were processed using our own implementation of the voxel-independent IDEAL algorithm (without region growing or any other advanced features that may have been added to the current commercial implementation) (4). Note that the data used in the comparison have a different set of TEs from those suggested in Reeder et al. (4) and Pineda et al. (13). The TEs employed in this work are not signal-to-noise-ratio optimal (which would require a TE spacing of nearly 1.6 ms at 1.5 T) due to the monopolar readout with gradient flyback.

RESULTS

For short, wide-bore scanners, the field variation at the edge of the FOV was found to be on the order of ± 1000 Hz. The central FOV excluding the border was more well behaved. Nevertheless, the frequency variation across the heart was in the range of 100–150 Hz. This variation may be due to tissue-air interfaces (42) or the presence of deoxygenated blood in large epicardial veins (43). The central FOV was identified on a per-slice basis as the region having field inhomogeneities in the range ± 300 Hz. Fat and water swaps using the ICM method were observed in 18 of the 25 cases, with five occurring in the central FOV. Fat and water swaps using the graph cut method were observed in two slices of the 25 cases, with a single fat/water swap in the central FOV, in a region of low signal.

Figure 3 shows representative results from a sagittal view of the heart, comparing the proposed method, ICM, and voxel-independent IDEAL. Images were acquired with TEs {4.2, 6.7, 9.2}ms. Since the original IDEAL method did not include T_2^* decay, the modified algorithm T_2^* -IDEAL was used (11). It must be noted that the water/fat images obtained with voxel-independent T_2^* -IDEAL were similar to a voxel-independent VARPRO, where ambiguities are resolved by forcing the field map to be in the range $(-f_F/2, f_F/2)$. For this reason, these results are denoted simply “voxel-independent” in the figures. The proposed method provided significantly improved results, particularly in regions with rapid field variations, where previous




FIG. 3. Comparison of the proposed method with two previously proposed methods (ICM and voxel independent). This dataset contains large field inhomogeneities near the edges of the FOV. (Top) Estimated field maps (in Hz). (Center) Water images. (Bottom) Fat images. The proposed method produced accurate field map estimation throughout the FOV, providing uniformly good water/fat separation. Previous methods were not able to track the field variations in the regions of high inhomogeneity, resulting in incorrect water/fat separation (indicated by arrows). Note that the field inhomogeneity reached ± 1000 Hz, but the color scale was kept in the range ± 600 Hz to show better contrast throughout most of the image.

methods produced water/fat swaps. Even though the heart was the region of interest in this application, artifacts in other areas of the FOV are undesirable since they may erroneously lead to incidental findings.


Figure 4 shows another case, acquired on a Siemens Avanto 1.5-T scanner with TEs {3.6, 5.8, 7.9}ms. Note that both ICM and the voxel-independent method produced water/fat swaps in the central FOV (see arrows), whereas the proposed method produced much better water/fat decomposition throughout the FOV.

Figure 5 shows multipeak fat modeling results from a 13-point acquisition with echo spacing 1.9 ms (with the first TE at 1.4 ms), using $M = 3$ fat peaks with known frequency shifts $\{-210, -159, 47\}$ Hz at 1.5 T. An independent decomposition of the three fat peaks and the water peak was performed using all 13 TEs. Water/fat decompositions with multipeak and single peak fat modeling, respectively, were obtained from the first five TEs. All cases were processed accounting for T_2^* decay. For the multipeak decomposition, the relative amplitudes α_m were

estimated from the data themselves (as proposed in Yu et al. (37) under “self-calibration for 6-point T_2^* -IDEAL acquisitions”), and were found to be $\alpha_1 \approx 0.77$, $\alpha_2 \approx 0.13e^{-i0.08\pi}$, and $\alpha_3 \approx 0.10e^{-i0.04\pi}$. Additionally, the $\alpha_{m,q}$ obtained with the independent peak model were averaged over the fat region, and the results were in good agreement with the multipeak estimates (the averages of the independent peak model produced $\alpha_1 \approx 0.77$, $\alpha_2 \approx 0.13e^{i0.06\pi}$, and $\alpha_3 \approx 0.10e^{-i0.12\pi}$). Multipeak modeling has two main advantages over single peak modeling (see the arrows in Fig. 5): (a) improved water/fat separation, which is clearly noticeable in fat-only regions (e.g., the subcutaneous fat layer), and (b) reduced ambiguity in the estimation (37).

Water/fat separation is useful for tissue characterization in cardiac MRI, where it has been shown to allow robust detection of fibrofatty infiltration of the myocardium (1,3), as well as characterization of tumors and masses, including lipomas. Figure 6 shows results from a patient with intramyocardial fat, which was processed previously using ICM (3), reconstructed here using the proposed graph cut

FIG. 4. Comparison of the proposed method with two previously proposed methods (ICM and voxel independent). In this dataset, field inhomogeneities near the edges of the FOV are relatively moderate because it was not acquired on a wide-bore scanner. (Top) Field maps. (Center) Water images. (Bottom) Fat images. ICM and voxel-independent methods resulted in water/fat swaps (indicated by arrows) in the liver under the dome of the diaphragm, as well as in the subcutaneous fat, but the proposed method produced good water/fat separation.



method. Images were acquired with TEs {1.5, 3.6, 5.7}ms. The separation was performed including multiplex modeling of the fat signal, as well as T_2^* decay. The intramyocardial fat is clearly visible in the fat-only image (Fig. 6b), which has positive contrast (i.e., fat against dark background) but is difficult to detect in the conventional fat-saturated turbo spin echo image (Fig. 6d), which has negative contrast (3). Figure 7 shows another example application. Images were acquired from a three-chamber view using TEs {2.5, 4.7, 7.0, 9.2}ms. The water/fat separated images (Fig. 7a and 7b, respectively) clearly show a large lipoma.

One desirable feature of the proposed penalized maximum likelihood formulation is our ability to characterize the spatial resolution properties of the estimates. This is important in order to improve confidence in our results, as well as to provide a criterion for choosing the regularization parameter. In the case of field map estimation, it is desirable to know the amount of smoothing introduced by the spatial regularization. Even though a complete characterization is challenging due to the nonlinearity of the signal model, one can analyze the local properties of the estimation by calculating the local impulse response (LIR), as derived in Fessler and Rogers. (34). The $LIR^q(\mathbf{f}_B)$ is defined as the change in the mean estimated field map caused by a perturbation in the true field map \mathbf{f}_B at voxel q . The expression for the $LIR^q(\mathbf{f}_B)$ is given in Eq. [16] of Fessler and Rogers. (34). Evaluation of this expression for Eq. [5] can be done efficiently by using a quadratic approximation of $R(\mathbf{f}_B; \mathbf{s}_q)$ at each voxel (since the regularization is also quadratic, this leads to a closed-form solution). The quadratic approximation is shown in Fig. 8a. Figure 8b shows an example of the $LIR^q(\mathbf{f}_B)$ at two different voxels for the dataset shown in Fig. 5.

To study the field map smoothing introduced by the spatial regularization, we simulated a dataset where the field

map contained an abrupt transition. Subsequently, field map estimation was performed using several values of the regularization parameter μ . Results demonstrating different levels of smoothing as a function of μ are shown in Fig. 8c-d. As can be seen from the figure, spatial regularization with $\mu = 0.02$ (the value used in this work) results in only moderate smoothing, which is important in regions of rapid field variation. Errors in the field map result in inaccurate water/fat separation, which is shown quantitatively by simulation in Fig. 8d. Note that the abrupt field map transition used in the simulation is more severe than the field maps observed in practice across the heart (43). It is expected that even with a worst-case gradient of 15 Hz/pixel, based on experimental measurements within the heart, a frequency error < 3 Hz would result using regularization with parameter $\mu = 0.02$, corresponding to an erroneous fat signal with amplitude equal to 2% of the water signal (Fig. 8e). In this case, with a water signal-to-noise ratio in the range of 20, the artifactual fat signal would be well below the noise level.

In the proposed algorithm, the bulk of the computation time is spent solving Eq. [8] (via the equivalent graph cut problem) at each iteration. On an Intel Xeon-based desktop personal computer with 48 GB of random-access memory and a 3.16 GHz central processing unit, solving this problem at each iteration requires 0.3 sec for images of size 192×144 and 0.9 sec for images of size 192×256 (image sizes from the results shown in this paper). A moderate number of iterations suffices to produce good results: the field map estimate converges rapidly and the improvements are negligible after 50 iterations for all the datasets processed in this work. The total processing time for the proposed algorithm is typically around 60 sec (90 sec if the model includes T_2^*). Additionally, the proposed method can be parallelized to improve speed (33).

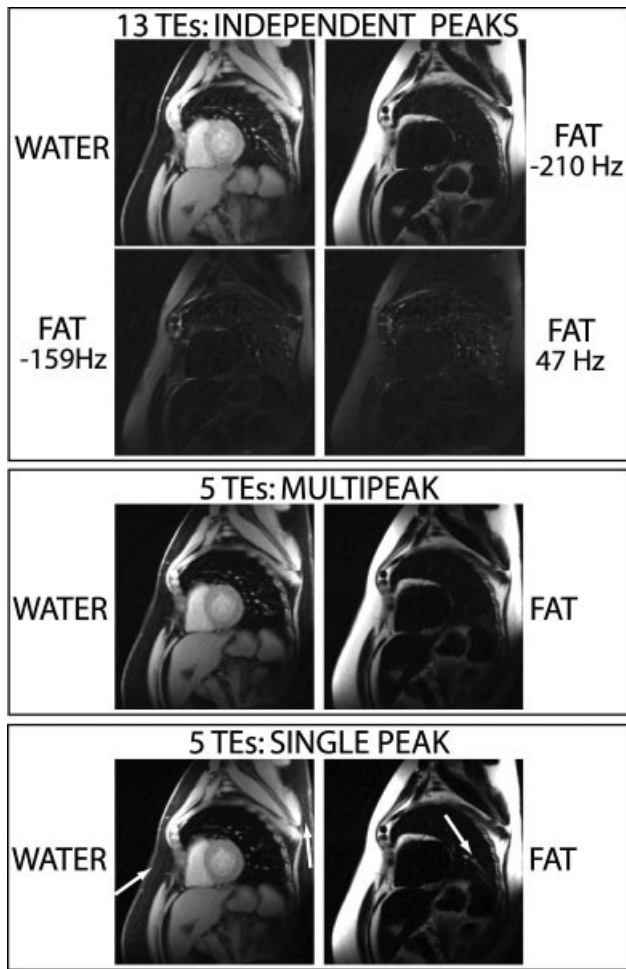


FIG. 5. Multipeak fat modeling. Data were acquired at 13 TEs, with uniform spacing 1.9 ms. The presence of several fat peaks in the signal is shown by performing an independent fit (without fixing $\alpha_{m,q}$) using all 13 TEs. The data corresponding to the first five TEs are then processed using a multipeak model and the standard single peak model (both including T_2^* decay). Multipeak modeling results in better water/fat separation, particularly in the regions with high fat signal such as the subcutaneous layer (see arrows in single peak water image). Additionally, the multipeak model helps resolve ambiguities in isolated signal regions (see arrow in single peak fat image).

DISCUSSION

Field map estimation is a critical step for accurate water/fat separation. However, the problem is severely ill posed when voxels are considered individually, which makes spatial regularization necessary. This has led to a variety of methods that impose field map smoothness, e.g., using multiresolution or region-growing algorithms (17–20). The proposed method has two main desirable properties: (i) the use of a penalized maximum likelihood formulation that allows a local characterization of the spatial resolution (smoothing) properties of the resulting field map; and (ii) the introduction of a novel optimization algorithm based on graph cuts, which allows the update of field map estimates at all the voxels simultaneously. This is quite different from algorithms where voxels are visited one at a time, even if information from previously visited voxels is used to constrain/initialize the estimate at the current voxel.

For the datasets considered in this work, it is important that the regularization term be spatially varying (there is an effective weighting based on local signal intensity, imposed through the $w_{q,j}$ s in Eq. [3]) because of the widely varying signal intensities observed in different regions of the image. If the $w_{q,j}$ s were constant, then it would not be possible to achieve regularization in the high signal regions without oversmoothing the field map in the low signal regions. The effect of spatially varying regularization can be well characterized using the LIR.

The performance of the proposed method depends on the accuracy of the signal model. For instance, the presence of T_2^* decay or multiple fat peaks can, if not accounted for, result in not only small perturbations on the water/fat estimates but also water/fat swaps. Additionally, multipeak fat modeling reduces the ambiguity in water/fat separation because the water and fat signals become different in this improved model (instead of being the same signal model with different frequency shift) (37). These effects are observed clearly in Fig. 5.

It must be noted that, in addition to the voxel-independent and ICM-based methods shown in this article, there are several recent methods that impose spatial constraints on the field map to improve water/fat separation (18–20,23). A comparative study with these recent methods is beyond the scope of this article (such a study should be performed involving the different research groups, so that the comparison is fair and accurate). This article is focused on comparing the proposed method to voxel-independent





FIG. 6. Results showing the application of the proposed method for the detection of fatty infiltration in the myocardium. **a,b**: Water/fat images obtained with the proposed method. The fatty infiltration is clearly visible in the fat image. **c**: Standard turbo spin echo acquisition, without fat saturation. **d**: turbo spin echo including conventional fat saturation. The fatty infiltration appears as a decrease in intensity in the fat-saturated image but is difficult to discern due to the negative contrast.

FIG. 7. Example of lipoma in the anteroseptal region of the myocardium, seen clearly in cardiac three-chamber view. **a:** Water image. **b:** Fat image.



separation and ICM to highlight two important points: (i) the need for spatial regularization of the field map, which is well addressed using a joint estimation approach; and (ii) the advantage of considering all voxels simultaneously when solving the joint estimation problem.

The proposed method has several limitations. First, it provides a locally (as opposed to globally) optimal solution to Eq. [6] (it is optimal with respect to an exponentially large set). Perhaps surprisingly, the globally optimal solution to Eq. [6] can also be found in polynomial time using graph cut methods. As shown in Ishikawa (44), the solution can be achieved by solving a graph cut problem on a different graph (larger than the ones used in the proposed method). This result holds as long as the regularization functional $V(f_{B,q}, f_{B,j})$ is convex. Direct application of the method proposed in Ishikawa (44) to Eq. [6] requires the manipulation of a very large graph (containing on the order of QL^2 edges if V is quadratic), making it less practical for realistic problem sizes. However, we have implemented it for a penalty $V(f_{B,q}, f_{B,j}) = |f_{B,q} - f_{B,j}|$, which requires a smaller graph (on the order of QL edges), resulting in good water/fat separation (results not shown). Even though a quadratic penalty is more appropriate for field map estimation, it is remarkable that this type of high-dimensional, nonconvex cost function can be globally optimized (in its discretized version) with an efficient algorithm.

Second, the discretization required for applying the proposed graph cut algorithm imposes a limit on the accuracy of the estimated field map. Even though the discretization is fine enough that we have not found it to be significant, it can be overcome by running a descent algorithm (such as the one proposed in Huh et al. (23)), initialized with the outcome of the proposed method.

Third, the proposed method uses a penalized maximum likelihood formulation (Eq. [3]) to regularize the field map estimate by penalizing nonsmooth solutions. This regularization is useful and has desirable properties in terms of characterizing its resolution properties, as discussed above. However, it is only a crude model if viewed as imposing a

prior distribution on the field map. For instance, in regions of extremely rapid field variation (such as when imaging near metal implants), this smoothness assumption would not be adequate, and the corresponding image distortions would make the current signal model inaccurate.

Extension of the proposed graph cut method to handle three-dimensional datasets is conceptually straightforward. Field map smoothness can be imposed along all three dimensions by using a three-dimensional neighborhood δ_q at each voxel q in the dataset (see Eq. [3]). However, computational requirements may make a multiresolution approach more practical in the three-dimensional case (19,20). This extension is currently under investigation.

In MRI, there is a variety of applications requiring the regularized estimation of nonlinear parameters (e.g., relaxation rates, amplitude of radio frequency field) (45,46). The method presented in this paper, based on VARPRO followed by a graph cut optimization algorithm, may also prove useful in these scenarios. The most important restriction on the proposed method is that the data term of the cost function is defined voxel by voxel (or, at least, that the interactions between different voxels are very limited (30)). In this case, the proposed optimization method provides a powerful tool for overcoming the nonlinearity of the model and the nonconvexity of the cost function (29).

CONCLUSION

This paper has presented a novel method for robust water/fat separation in the presence of large field inhomogeneities. The proposed method uses a statistically motivated formulation and solves the underlying optimization problem by subdividing it into a sequence of binary decision problems, which are solved efficiently using a graph cut algorithm. The method has good theoretical properties and has been shown to perform robustly in a number of challenging cardiac imaging studies. With these capabilities, it should prove particularly useful for the clinical applications where large field inhomogeneities currently prevent reliable water/fat separation.

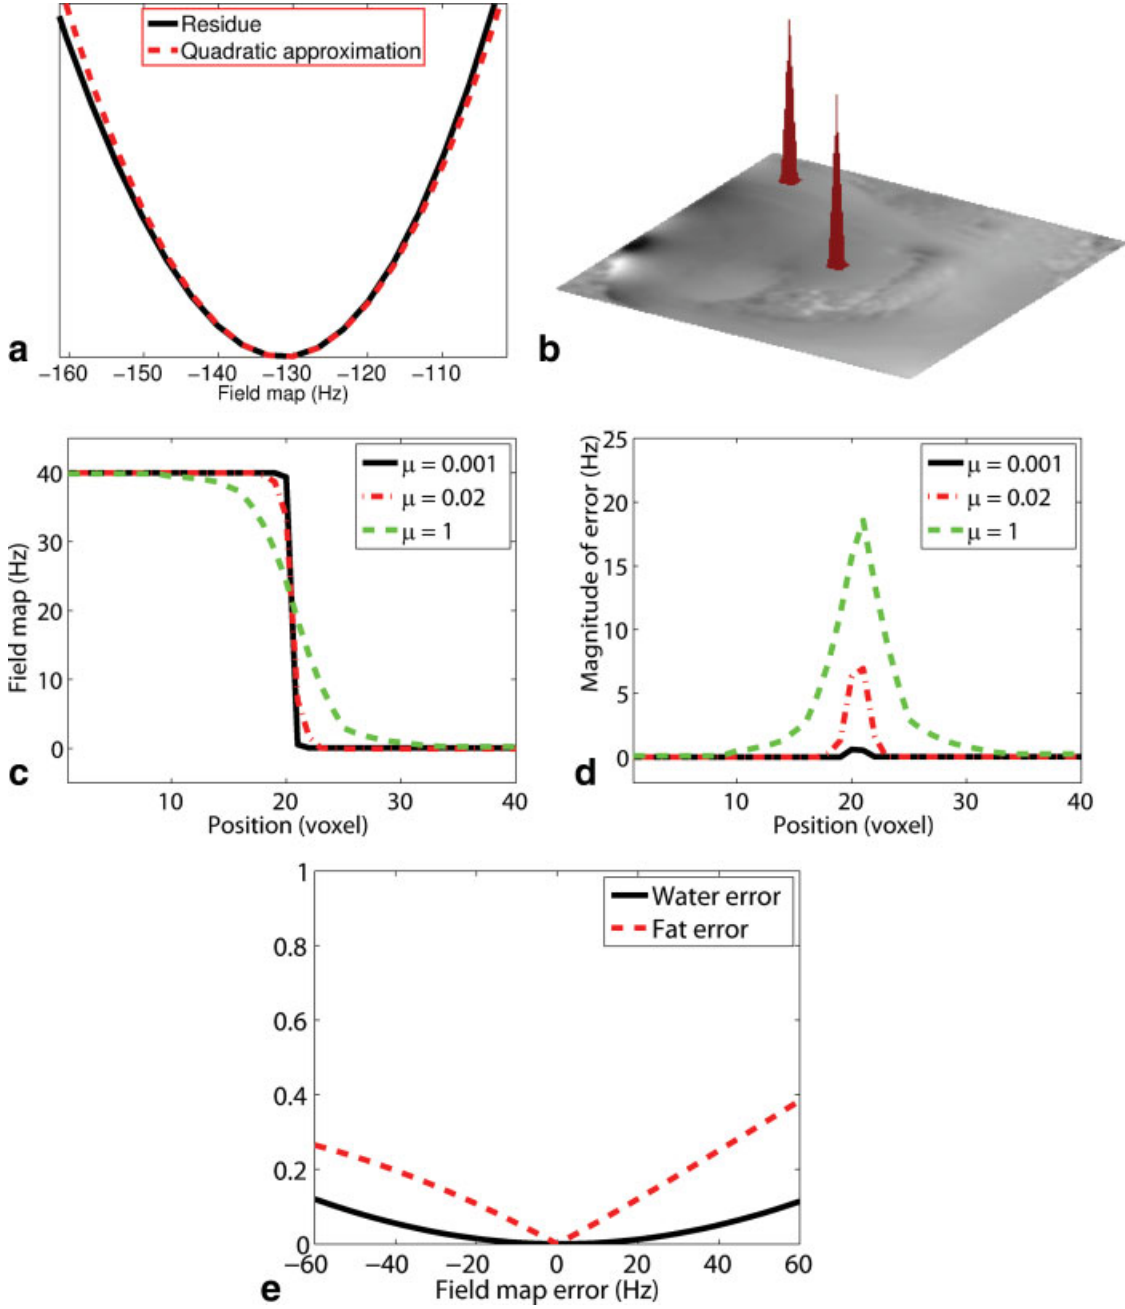


FIG. 8. Analysis of the spatial resolution properties of the proposed regularized field map estimate, and the associated errors in the water/fat decomposition. **a**: Residue $R(f_{B,q}; s_q)$ at an individual voxel and local approximation using a quadratic function. **b**: LIR for field map perturbations at different locations (shown in logarithmic scale over the true field map). **c**: Simulation demonstrating the field map smoothing that results from different values of μ . The true field map contains a sharp jump in the center of the image. In this work, we use $\mu = 0.02$. **d**: Absolute field map errors corresponding to varying values of μ in the previous example. **e**: Errors in the estimation of the water/fat magnitudes at a single voxel, as a function of the error in the field map (in the absence of noise). The simulated TEs are $\{6.76, 8.36, 9.96\}$ ms. The true water/fat amplitudes are $\rho_W = 1, \rho_F = 0$.

ACKNOWLEDGEMENTS

The work presented in this paper was supported in part by the following research grants: NIH-P41-RR023953-01, NIH-P41-EB001977-21 and NSF-CBET-07-30623. We acknowledge the use of the Matlab Boost Graph Library (MatlabBGL) package, written by David Gleich.

APPENDIX: CONVERSION OF EQ. [8] TO A GRAPH CUT PROBLEM

Recall the definition of the subset $\Gamma \subset \Omega^Q$ for the optimization problem in Eq. [8]: $\Gamma = \hat{\Omega}_1 \times \hat{\Omega}_2 \times \dots \times \hat{\Omega}_Q$, where $\hat{\Omega}_q = \{\hat{f}_{B,q}, \hat{f}'_{B,q}\}$, $q = 1, \dots, Q$. The necessary and sufficient condition for graph-representability (i.e., the existence of an equivalent graph cut problem that can be efficiently

Graph Cut Problem




FIG. A1. Example of graph used in this work, including a graph cut. The number of vertices in the graph is $Q+2$, i.e., one vertex per voxel in the corresponding image, plus two additional vertices s (source) and t (sink). The edge weights d_{qj} are determined by the cost function and the set Γ of possible field maps.

solved) of Eq. [8] derived in Kolomogorov and Zibah (28) and can be stated as:

$$V(\hat{f}_{B,q}, \hat{f}_{B,j}) + V(\hat{f}'_{B,q}, \hat{f}'_{B,j}) \leq V(\hat{f}_{B,q}, \hat{f}_{B,j}) + V(\hat{f}_{B,q}, \hat{f}'_{B,j}) \quad [14]$$

for $q = 1, \dots, Q, j \in \delta_q$. Clearly, in our problem graph representability depends only on the regularization penalty function $V(\hat{f}_{B,q}, \hat{f}_{B,j})$ and the choice of $\hat{f}'_{B,q}$.

For a quadratic penalty $V(\hat{f}_{B,q}, \hat{f}_{B,j}) = (\hat{f}_{B,q} - \hat{f}_{B,j})^2$, we can easily show that any choice of Γ where $\hat{f}'_{B,q} - \hat{f}_{B,q}$ has the same sign for all voxels $q = 1, \dots, Q$ (as is the case for the proposed iterations Γ_β, Γ_+ and Γ_-), is graph representable. Denoting $\Delta_0 = \hat{f}_{B,q} - \hat{f}_{B,j}$, $\Delta_q = \hat{f}'_{B,q} - \hat{f}_{B,q}$, and $\Delta_j = \hat{f}'_{B,j} - \hat{f}_{B,j}$, then Eq. [14] will be satisfied if

$$(\Delta_0 + \Delta_q)^2 + (\Delta_0 - \Delta_j)^2 - \Delta_0^2 - (\Delta_0 + \Delta_q - \Delta_j)^2 > 0 \quad [15]$$

i.e., if $\Delta_q \Delta_j > 0$ (or equivalently, if Δ_q and Δ_j have the same sign). Therefore, the iterations employed in this paper are graph representable.

To construct the equivalent graph for Eq. [8] (see Fig. A1), we assign to each voxel one vertex v_q , plus one source (s) and one sink (t) vertex. Thus, the total number of vertices is $Q + 2$. The edges of the graph, accounting for $R(\hat{f}_{B,q}; \mathbf{s}_q)$ (called data edges) and the regularization term (called regularization edges) in Eq. [8], are defined as follows (28):

- **Data edges.** $R(\hat{f}_{B,q}; \mathbf{s}_q)$ generates one edge for each voxel q . This edge is (s, v_q) with weight $d_{sq} = R(\hat{f}_{B,q}; \mathbf{s}_q) - R(\hat{f}'_{B,q}; \mathbf{s}_q)$ if $R(\hat{f}_{B,q}; \mathbf{s}_q) - R(\hat{f}'_{B,q}; \mathbf{s}_q) > 0$, and (v_q, t) with weight $d_{qt} = R(\hat{f}'_{B,q}; \mathbf{s}_q) - R(\hat{f}_{B,q}; \mathbf{s}_q)$ otherwise.
- **Regularization edges.** Each term $V(\hat{f}_{B,q}, \hat{f}_{B,j})$ generates three edges (if a data edge already exists, the new weight is added to the existing weight). Defining $A_q = V(\hat{f}_{B,q}, \hat{f}_{B,j}) - V(\hat{f}_{B,q}, \hat{f}'_{B,j})$ and $A_j = V(\hat{f}_{B,q}, \hat{f}'_{B,j}) - V(\hat{f}_{B,q}, \hat{f}_{B,j})$, the following edges are added:
 - Edge (s, v_q) with weight $d_{sq} = A_q$ if $A_q > 0$, or edge (v_q, t) with weight $d_{qt} = -A_q$ otherwise.
 - Edge (s, v_j) with weight $d_{sj} = A_j$ if $A_j > 0$, or edge (v_j, t) with weight $d_{jt} = -A_j$ otherwise.

- Edge (v_q, v_j) with weight

$$d_{qj} = V(\hat{f}'_{B,q}, \hat{f}_{B,j}) + V(\hat{f}_{B,q}, \hat{f}'_{B,j}) - V(\hat{f}_{B,q}, \hat{f}_{B,j}) - V(\hat{f}'_{B,q}, \hat{f}'_{B,j}).$$

As formulated, the solution to Eq. [8] is given by the minimum cut problem (28). Note that cut \mathcal{C} of a graph is a partition of the vertices of the graph into two disjoint subsets, \mathcal{S} and \mathcal{T} , such that $s \in \mathcal{S}$ and $t \in \mathcal{T}$. Every remaining vertex is either in \mathcal{S} or in \mathcal{T} (see Fig. A1). The cost $|\mathcal{C}|$ of \mathcal{C} is defined as:

$$|\mathcal{C}| = \sum_{q \in \mathcal{S}, j \in \mathcal{T}} d_{qj}. \quad [16]$$

The minimum cut problem is defined as solving:

$$\mathcal{C}_{\min} = \arg \min_{\mathcal{C}} |\mathcal{C}|, \quad [17]$$

which can be solved with worst-case complexity $O(Q^3)$ for the graph defined above (26).

REFERENCES

1. Reeder SB, Markl M, Yu H, Hellinger JC, Herfkens RJ, Pelc NJ. Cardiac CINE imaging with IDEAL water-fat separation and steady-state free precession. *J Magn Reson Imaging* 2005;22:44–52.
2. Goldfarb JW. Fat-water separated delayed hyperenhanced myocardial infarct imaging. *Magn Reson Med* 2008;60:503–509.
3. Kellman P, Hernando D, Shah S, Zuehlsdorff S, Jerecic R, Mancini C, Liang ZP, Arai A. Multi-echo dixon fat and water separation method for detecting fibro-fatty infiltration in the myocardium. *Magn Reson Med* 2009;61:215–221.
4. Reeder SB, Wen Z, Yu H, Pineda AR, Gold GE, Markl M, Pelc NJ. Multicoil Dixon chemical species separation with an iterative least squares estimation method. *Magn Reson Med* 2004;51:35–45.
5. Haase A, Frahm J, Hänicke W, Matthaei D. ^1H NMR chemical shift selective (CHESS) imaging. *Phys Med Biol* 1985;30:341–344.
6. Meyer CH, Pauly JM, Macovskiand A, Nishimura DG. Simultaneous spatial and spectral selective excitation. *Magn Reson Med* 1990;15:287–304.
7. Bydder GM, Steiner RE, Blumgart LH. MR imaging of the liver using short TI inversion recovery. *J Comput Assist Tomogr* 1985;9:1084–1089.
8. Ma J. Dixon techniques for water and fat imaging. *J Magn Reson Imaging* 2008;28:543–558.
9. Dixon WT. Simple proton spectroscopic imaging. *Radiology* 1984;153:189–194.
10. Glover GH, Schneider E. Three-point Dixon technique for true water/fat decomposition with B_0 inhomogeneity correction. *Magn Reson Med* 1991;18:371–383.
11. Yu H, McKenzie CA, Shimakawa A, Vu AT, Brau ACS, Beatty PJ, Pineda AR, Brittain JH, Reeder SB. Multiecho reconstruction for simultaneous water-fat decomposition and $T2^*$ estimation. *J Magn Reson Imaging* 2007;26:1153–1161.
12. Brodsky EK, Holmes JH, Yu H, Reeder SB. Generalized k-space decomposition with chemical shift correction for non-cartesian water-fat imaging. *Magn Reson Med* 2008;59:1151–1164.
13. Pineda A, Reeder S, Wen Z, Pelc NJ. Cramér-Rao bounds for three-point decomposition of water and fat. *Magn Reson Med* 2005;54:625–635.
14. Hernando D, Haldar JP, Sutton BP, Ma J, Kellman P, Liang ZP. Joint estimation of water/fat images and field inhomogeneity map. *Magn Reson Med* 2008;59:571–580.
15. Xiang QS, An L. Water-fat imaging with direct phase encoding. *J Magn Reson Imaging* 1997;7:1002–1015.
16. An L, Xiang QS. Chemical shift imaging with spectrum modeling. *Magn Reson Med* 2001;46:126–130.
17. 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;52:415–419.

18. 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;54:1032–1039.
19. Tsao J, Jiang Y. Hierarchical IDEAL: robust water-fat separation at high field by multiresolution field map estimation. In: Proceedings of the 16th Annual Meeting of ISMRM, Toronto, Canada, 2008. p. 653.
20. Lu W, Hargreaves BA. Multiresolution field map estimation using golden section search for water-fat separation. *Magn Reson Med* 2008;60:236–244.
21. Geman S, Geman D. Stochastic relaxation, Gibbs distributions and the bayesian restoration of images. *IEEE Trans Pattern Anal Mach Intell* 1984;6:712–741.
22. Fessler JA, Yao D, Noll DC. Regularized fieldmap estimation in MRI. In: Proceedings of the 3rd IEEE International Symposium on Biomedical Imaging: Macro to Nano, Arlington, VA, 2006. p. 706–709.
23. Huh W, Fessler JA, Samsonov AA. Water-fat decomposition with regularized field map. In: Proceedings of the 16th Annual Meeting of ISMRM, Toronto, Canada, 2008. p. 1382.
24. Boyd S, Vandenberghe L. Convex optimization. Cambridge University Press; 2004.
25. Boykov Y, Veksler O, Zabih R. Fast approximate energy minimization via graph cuts. *IEEE Trans Pattern Anal Machine Intell* 2001;23:1222–1239.
26. Cormen TH, Leiserson CE, Rivest RL, Stein C. Introduction to algorithms. 2nd ed. MIT Press; 2001.
27. Ahuja RK, Ergun Ö, Orlin JB, Punnen AP. A survey of very large-scale neighborhood search techniques. *Discrete Appl Math* 2002;123:75–102.
28. Kolmogorov V, Zabih R. What energy functions can be minimized via graph cuts? *IEEE Trans Pattern Anal Machine Intell* 2004;26:147–159.
29. Darbon J, Sigelle M. Image restoration with discrete constrained total variation. Part I: fast and exact optimization. *J Math Imaging Vis* 2006;26:261–276.
30. Raj A, Singh G, Zabih R, Kressler B, Wang Y, Schuff N, Weiner M. Bayesian parallel imaging with edge-preserving priors. *Magn Reson Med* 2007;57:8–21.
31. Veksler O. Efficient graph-based energy minimization methods in computer vision. PhD thesis, Cornell University, 1999.
32. Bioucas-Dias JM, Valadao G. Phase unwrapping via graph cuts. *IEEE Trans Image Proc* 2007;16:698–709.
33. Boykov YY, Veksler O. Graph cuts in vision and graphics: theories and applications. In: Paragios N, Chen Y, Faugeras, editors. Handbook of mathematical methods in computer vision. New York: Springer-Verlag; 2005. p 79–96.
34. Fessler JA, Rogers WL. Spatial resolution properties of penalized-likelihood image reconstruction: space-invariant tomographs. *IEEE Trans Image Proc* 1996;5:1346–1358.
35. O'Regan DP, Callaghan MF, Wylezinska-Arridge M, Fitzpatrick J, Naoumova RP, Hajnal JV, Schmitz SA. Liver fat content and T2*: simultaneous measurement by using breath-hold multiecho MR imaging at 3.0 T—feasibility. *Radiology* 2008;247:550–557.
36. Hernando D, Kellman P, Haldar JP, Liang ZP. Estimation of water/fat images, B0 field map and T2* map using VARPRO. In Proceedings of the 16th Annual Meeting of ISMRM, Toronto, Canada, 2008. p. 1517.
37. Yu H, Shimakawa A, McKenzie CA, Brodsky EK, Brittain JH, Reeder SB. Multiecho water-fat separation and simultaneous R2* estimation with multifrequency fat spectrum modeling. *Magn Reson Med* 2008;60:1122–1134.
38. Xiang QS. Simultaneous water-fat and T2* mapping with 3-point acquisitions. In: Proceedings of the 16th Annual Meeting of ISMRM, Toronto, Canada, 2008. p. 1383.
39. Bydder M, Yokoo T, Hamilton G, Middleton MS, Chavez AD, Schwimmer JB, Lavine JE, Sirlin CB. Relaxation effects in the quantification of fat using gradient echoimaging. *Magn Reson Imaging* 2008;26:347–359.
40. Reeder SB, Brittain JH, Grist TM, Yen YF. Least-squares chemical shift separation for 13C metabolic imaging. *J Magn Reson Imaging* 2007;26:1145–1152.
41. Reeder SB, Robson P, Yu H, Shimakawa A, Gard CD, McKenzie CA, Brittain JH. Quantification of hepatic steatosis with MRI: the effects of accurate fat spectral modeling. In: Proceedings of the 16th Annual Meeting of ISMRM, Toronto, Canada, 2008. p. 710.
42. Atalay MK, Poncelet BP, Kantor HL, Brady TJ, Weisskoff RM. Cardiac susceptibility artifacts arising from the heart-lung interface. *Magn Reson Med* 1999;45:341–345.
43. Reeder SB, Faranesh AZ, Boxerman JL, MvVeigh ER. In vivo measurement of T2* and field inhomogeneity maps in the human heart at 1.5 T. *Magn Reson Med* 1998;39:988–998.
44. Ishikawa H. Exact optimization for Markov random fields with convex priors. *IEEE Trans Pattern Anal Machine Intell* 2003;25:1333–1336.
45. Bernstein MA, King KF, Zhou XJ. Handbook of MRI pulse sequences. Academic Press; 2004.
46. Fessler JA, Yao D, Noll DC. Regularized B1+ map estimation in MRI. In: Proceedings of the 4th IEEE International Symposium on Biomedical Imaging: Macro to Nano, Metro Washington, DC, USA, 2007. p. 616–619.