Using Spatial Prior Knowledge in the Spectral Fitting of Magnetic Resonance Spectroscopic Images

B. Michael Kelm[1,2], Frederik O. Kaster[1,3], Anke Henning[4], Marc-André Weber[5,6],
Peter Bachert[3], Peter Boesiger[4], Fred A. Hamprecht[1], and Bjoern H. Menze[1,7,8]

[1]Interdisciplinary Center for Scientific Computing (IWR), University of Heidelberg, Germany

[2]Corporate Technology, Siemens AG, Erlangen, Germany

[3]Physics in Radiology, German Cancer Research Center (DKFZ), Heidelberg, Germany

[4]Institute for Biomedical Engineering, University and ETH Zurich, Switzerland

[5]Radiology, German Cancer Research Center (DKFZ), Heidelberg, Germany

[6]Diagnostic and Interventional Radiology, University Hospital, Heidelberg, Germany

[7]Asclepios Project, INRIA Sophia-Antipolis, France

[8]Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology
(MIT), Cambridge, Massachusetts, USA

February 9, 2011

Spectral Fitting of MRSI using Spatial Prior Knowledge

Correspondence to:

B. Michael Kelm

Siemens AG, CT T DE TC4, San-Carlos-Str. 7, 91058 Erlangen, Germany

`michael.kelm@siemens.com`

Bjoern H. Menze

CSAIL MIT, 32 Vassar Street, Cambridge, MA 02139

`menze@csail.mit.edu`

Abstract

We propose a Bayesian smoothness prior in the spectral fitting of magnetic resonance spectroscopic images which can be used in addition to commonly employed prior knowledge. By combining a frequency-domain model for the free induction decay with a Gaussian Markov random field prior, a new optimization objective is derived that encourages smooth parameter maps. Using a particular parametrization of the prior, smooth damping, frequency and phase maps can be obtained while preserving sharp spatial features in the amplitude map. A Monte Carlo study based on two sets of simulated data demonstrates that the variance of the estimated parameter maps can be reduced considerably, even below the Cramér-Rao lower bound, when using spatial prior knowledge. Long echo time ^{1}H-MRSI at 1.5T of a patient with brain tumor shows that using the spatial prior resolves the overlapping peaks of choline and creatine also when a single voxel method fails to do so. Improved and detailed metabolic maps can be derived from high spatial resolution short echo ^{1}H-MRSI at 3T. Finally, the evaluation of four series of long echo time brain MRSI data with various signal-to-noise ratios shows the general benefit of the proposed approach. *(190 words)*

2

Nomenclature

AMARES	Advanced Method for Accurate Robust and Efficient Spectral fitting (5)
AQSES	Automated Quantitation of Short Echo time MRS Spectra (11, 12)
Cho	Choline
Cr	Creatine
CRLB	Cramér-Rao Lower Bound
DCE-MRI	Dynamic Contrast-Enhanced Magnetic Resonance Imaging
EEG	ElectroEncephaloGraphic (imaging)
FID	Free Induction Decay
FITT	quantification approach proposed in (8)
GAMMA	A General Approach To Magnetic Resonance Mathematical Analysis (29)
GMRF	Gaussian Markov Random Field
GSLIM	Generalized Spectral Localization by IMaging (14)
HSVD	Hankel Singular Value Decomposition
ICM	Iterated Conditional Modes (algorithm)
jMRUI	Java-based graphical user interface for processing MRS (39)
LCModel	quantification approach proposed in (9)
MAP	Maximum-A-Posteriori (estimate)
MEG	MagnetoEncephaloGraphic (imaging)
NAA	N-Acteylaspartate
PET	Positron Emission Tomography
PRESS	Point RESolved Spectroscopy
ProFit	PRior-knOwledge Fitting (13)
QUEST	quantitation based on semi-parametric QUantum ESTimation (10)

RMSE   Root Mean Square Error

SLIM   Spectral Localization by IMaging (14)

SNR   Signal-to-Noise Ratio

SPECT   Single Photon Emission Computed Tomography

1   Introduction

Accurate quantification is a crucial requirement for the study of metabolism *in vivo* by means of magnetic resonance spectroscopy (MRS). The metabolic profile of a magnetic resonance spectrum may be used in noninvasive diagnosis and the characterization of patho-physiological changes and, thus, may serve as an important tool in clinical research. In addition to this traditional field of application, recent approaches using hyperpolarized nuclei as tracers in molecular imaging also rely on MR spectroscopic imaging (MRSI) (1). The accurate quantification of such MR spectroscopic signals depends on stable approaches to spectral fitting, which is a challenging problem due to artifacts (most notably eddy currents or ghost), overlapping peaks and baseline distortions due to macromolecules and residual water or lipids present in MRS data (2, 3). Another, major problem is the notorious low signal-to-noise ratio (SNR) of the MRS signal due to the generally low abundance of the metabolites of interest (4).

A common approach to improve the quality of spectral model fitting is the use of prior knowledge. For example the AMARES fitting algorithm, already introduced by Vanhamme *et al.* in 1997 (5), allows to consider various forms of constraints on the model parameters in the non-linear optimization. This prior knowledge may be the specification of upper and lower limits, *i.e.* box constraints, for resonance line position, line width, or phase difference, but it may also impose limits on differences or ratios thereof in a signal model with multiple resonance lines. Both empirical evidence (*e.g.* (6)) and formal results on the decrease of the Cramér-Rao lower bound (CRLB) for box constraints show the benefit of imposing such prior knowledge in the fitting of the spectral model (7). More recent fitting algorithms also impose prior knowledge on the estimate of the spectral signal composition by the use of experimentally measured or simulated metabolite templates – as implemented, for example, in FITT (8), LCModel (9), QUEST (10), AQSES (11, 12) or ProFit (13).

Although current approaches in MR spectroscopy aim at recording *high resolution*[1] spectroscopic images rather than acquiring single voxel spectra, most quantification routines largely ignore the spatial context of the individual spectra in the MRSI and do little more than a sequential fitting of individual spectra. With reasonable *spatial smoothness assumptions* for certain parameters of the signal model at hand, however, one may want to introduce novel constraints in the spectral model fitting to further improve quality and reliability of the fitting procedure.

A number of methods use spatial information for an optimal reconstruction of the signal from k-space. The "spectral localization by imaging" (SLIM, GSLIM) by Hu *et al.* (14) and Liang *et al.* (15), for example, assumes homogeneous metabolite concentrations in anatomical regions identified in MRI, in order to reconstruct MR spectra from each region. Plevritis *et al.* (16) use the anatomical

---

[1]Note that by "high resolution" MRSI we refer to data with high *spatial* resolution as opposed to *spectral* resolution.

5

information to adapt the k-space sampling and to enhance resolution in spectroscopic images. Bao and Maudsley (17) propose an approach that estimates signal intensities for partial volume over MRI-defined anatomic regions as an extension to the SLIM approach. Haldard *et al.* (18) introduce a Bayesian prior in the reconstruction of MR signals to enable extended k-space sampling, also pointing out the applicability of this concept for MRSI model fitting. In Kornak *et al.* (19), an expectation maximization algorithm is proposed for the spatially regularized reconstruction of MRSI data using a segmented map of tissue regions.

Bayesian approaches are often used for integrating anatomical or spatial information in the analysis of functional data as well, for example, for positron emission tomography (PET), single photon emission computed tomography (SPECT) (20), or magneto-/electroencephalographic imaging (MEG/EEG) (21). In contrast to interpolation methods in k-space, these smoothing approaches work in the spatial domain. Kelm *et al.* (22, 23) and Schmidt *et al.* (24) proposed the use of such Bayesian spatial smoothness priors for the processing of dynamic contrast-enhanced MRI (DCE-MRI). Being a general regularization method for spatial parameter maps in the processing of nonlinear signals, this smoothness prior can be generalized to other types of vector-valued image data and also for a use with MR spectroscopic images as first proposed by Kelm in 2007 (25).

One of the early approaches to exploiting spatial prior knowledge was introduced with FITT (8) which proposes an iterative MRSI fitting methodology that includes a spatial smoothing step. After voxel-wise baseline removal and parameter estimation, selected parameters (line width, frequency shifts, and first and second order phase terms) are only accepted if consistent with a local neighborhood. Also LCmodel (9) provides a spatial fitting mode for MRSI. If activated, LCmodel first analyzes a central voxel and then proceeds outwards using the results from previously fitted voxels for initialization and as a soft constraint (26) for new fits. This additional feature may speed up and improve the analysis. An approach modeling the spatial variation of the phase parameter by a first order function, *i.e.*, a constraint which is linear in the three spatial directions, has been described by Bretthorst (27). Only recently, Croiter Sava *et al.* (28) successfully incorporated a spatial prior in the AQSES algorithm, following ideas from our initial study (25). A detailed comparison of spatial fitting approaches is deferred to the discussion in section 5.3.

1.1   Objective of this study

In the present work, we propose and examine *spatial smoothness assumptions* as spatial constraints on selected model parameters for improving the spectral fitting of MRSI data. We detail and extend on earlier work presented in (25).

In a formulation based on a Gaussian Markov random field (GMRF) prior, the smoothness of

6

selected parameter maps can be encouraged, and efficient methods for inference are presented. This novel spatial constraint can be applied in addition to any other, commonly employed single voxel prior knowledge. The spatial constraint can be highly selective: While encouraging a certain spatial smoothness for specific parameters such as phase or line width over the whole image, it allows to estimate parameters of diagnostic interest, such as the amplitude of the signal, in a totally unconstrained fashion. In contrast to approaches in k-space (14–18) the proposed constraint allows to use spatial neighborhood information in a very intuitive way; in contrast to (27) the proposed approach is not constrained to certain parametric spatial functions of the spectral model parameters.

In the following we will first detail on the signal model, the proposed spatial constraint, and on the inference algorithm used in the spatio-spectral parameter estimation (section 2). In section 3, the experimental data, parameters and methods for the evaluation of the proposed approach are presented. Results on resonance line fitting with and without spatial prior knowledge on real data, and in simulation experiments are presented in section 4. An interpretation and discussion of the results is provided in section 5 before the paper is concluded in section 6.

2   Theory

In the following, the signal model is described, a model for introducing spatial prior knowledge in the estimation of the parameters maps is presented and some details on the inference algorithm used for this task are provided.

2.1   The spectral signal model

Most quantification algorithms fit a nonlinear signal model $S_\theta(n)$ to the observed spectral data y_n, either in time-domain (5, 10, 11) or in frequency-domain (8, 9, 13), by minimizing the sum of squared residuals

$$l(\theta) \quad = \quad \sum_{n=1}^{N} (S_\theta(n) - y_n)^2 \tag{1}$$

which is the maximum likelihood solution under the assumption of additive white Gaussian noise. Fitting in the frequency-domain allows to exclude frequency ranges which are spoiled by unmodeled signals such as residual water or lipids. The corresponding residuals are then simply omitted from the sum in Eq. (1).

In QUEST (10) and AQSES (11), the free induction decay (FID) signal is modeled as a linear combination of M possibly damped, phase- and frequency-shifted metabolite templates $T_m(t_n)$. Here,

the same model in the frequency domain is used:

$$S_\theta(n) \quad = \quad \mathcal{FFT}\left\{ \mathrm{e}^{j\phi_0} \sum_{m=1}^{M} T_m(t_n)\, a_m \mathrm{e}^{(j2\pi\Delta f_m - \Delta d_m)t_n} \right\} \quad \text{with } t_n = n\Delta t,\, n = 1\ldots N \qquad (2)$$

with the imaginary unit $j = \sqrt{-1}$. The parameter vector θ contains a common zero-order phase correction ϕ_0 and for each of the M metabolites an amplitude a_m, frequency shift Δf_m and exponential damping Δd_m. All Δf_m and Δd_m are initialized to zero whereas initial guesses for a_m and ϕ_0 are obtained by linear least squares. Note that it is sufficient to define the model $S_\theta(n)$ at the N discrete time points t_n sampled at the same interval Δt as the recorded FID signal. Also the metabolite templates $T_m(t_n)$ only need to be known at these discrete time points. They may be defined by experimental measuring (*e.g.* *in vitro*), physically motivated modeling (*e.g.* Lorentzian, Voigt line shape for each resonance) or simulation using, for example, the GAMMA software (29).

2.2   The spatially constrained signal model

Spatial prior knowledge is incorporated using the Bayesian approach of specifying a *prior distribution* over the sought parameter maps. The maximum likelihood point estimate is then replaced by the mode of the posterior distribution, i.e. by the maximum a posteriori estimate (MAP). Here, a Gaussian Markov random field (GMRF) is used to model a prior distribution that favors smooth maps (23) for some model parameters ϕ_0, a_m, Δf_m and Δd_m (cf. Eq. (2)). In the following, these parameters are summarized in the vector θ_s with s indexing the voxel position. The optimization objective thus becomes

$$l(\Theta) \quad = \quad \sum_{s} \sum_{n=1}^{N} (S_{\theta_s}(n) - y_n^s)^2 + \sigma^2 \sum_{s \sim t} ||\theta_s - \theta_t||_W^2 \qquad (3)$$

where Θ now contains the nonlinear parameters θ_s from all MRSI voxels, σ is the standard deviation of the signal noise, $||\theta||_W^2 = \theta^T W \theta$ is a 2-norm weighted with the diagonal matrix W and $s \sim t$ denotes two voxels s and t that are neighbors according to some neighborhood system.

The first term in Eq. (3) just builds a sum over the squared residuals from all voxels s. Without the second term in Eq. (3), the parameter vectors θ_s at individual voxels would be independent, the optimization problem would decouple and would yield exactly the same solution as with the single voxel approach in Eq. (1). Adding the second term which builds a sum over the squared distance $||\cdot||_W$ between all neighboring voxel pairs $s \sim t$ encourages smooth solutions. Both σ and W determine the trade-off between fitting the individual signals against smoothing the parameter maps. Furthermore, W can be used to adjust the smoothness force of certain parameter maps individually, and even to turn off smoothing for parameters such as the amplitude a_m by using a zero weight.

Note that other signal models like those employed with FITT (8), LCModel (9) or ProFit (13) along with additional baseline models could be used together with the spatial regularization term in

8

Eq. (3). Similar improvements as those demonstrated with our model (Eq. (2)) should be obtained for these as well. Furthermore, spatial regularization terms other than the Gaussian could be used such as generalized Gaussian (22) or total variation (TV) terms that generate less penalty at step edges and might be considered more suitable spatial prior knowledge under some circumstances. While our proposed approach can be applied in the same way, it turns out that the resulting objective function cannot be optimized as efficiently as for a Gaussian prior. In our experience, the additional computational needs do not justify the hardly observable improvements of the fit results.

2.3   Inference

An exact joint optimization of the objective in Eq. (3) over all parameters Θ would be computationally intractable. Already an MRSI with, for example, 32×32 voxels and a signal model with just three metabolites (each of which has three parameters plus a common phase parameter, Equation (2)) would yield a parameter vector θ_s with $32*32*10 \approx 10{,}000$ entries, leading to an optimization problem in a 10,000-dimensional space.

Using a neighborhood system of 4 or 8 nearest neighbors as in our 2D implementation, however, yields a sparse optimization problem in the sense that most of the entries in Θ are not directly coupled. This sparseness can be exploited by special algorithms. One such algorithm is the *iterated conditional modes* (ICM) algorithm (30), exploiting the Markov properties of the spatial constraint. This algorithm optimizes single sites (here: spectra) in the image grid individually, conditioned on observations and fitted parameters θ_s in its neighborhood, and iterates over the whole grid (31, 32). For the optimization of Eq. (3), a generalized version of the ICM algorithm has been used which is considering collections of sites instead of single sites at each step of the iteration. Although, this *Block-ICM* algorithm has already been proposed by Besag (30), it has rarely been applied in the literature. In (33), Block-ICM has been applied for subvoxel tissue classification in NMR microscopic images, *i.e.* a discrete optimization problem. Spatially constrained fitting of dynamic contrast-enhanced MRI using Block-ICM, *i.e.* nonlinear least squares fitting, has been proposed in (23) and was shown to converge faster than the standard ICM approach in (22). The Block-ICM approach is closely related to domain decomposition methods (34) which are frequently used in spatial optimization problems. They are often used in combination with multigrid approaches (35) which may also provide directions for improving the efficiency of the parameter estimation, and to easily tackle inference in 3D settings with 6 or 26 neighborhood systems.

9

3 Methods

For a first qualitative comparison of standard single-voxel and spatially constrained GMRF approach, both methods were applied to exemplary clinical patient and experimental high resolution data (Sec. 3.1). For a more detailed understanding of the smoothing behavior and potential biases (cf. (22)) of the latter synthetic spectral data sets were designed (Sec. 3.2). Finally, the performance of both fitting approaches were evaluated using experimental MRSI series with increasing spatial resolution in a visual inspection of the fitting results (Sec. 3.3).

3.1 Fitting the models to real data

In a first test both fitting approaches were applied to two real ^{1}H-MRSI recordings of the brain, both with long and short echo time data. The first recording had been obtained from a clinical study with brain tumor patients (36). The clinical study was conducted after institutional review board agreement had been obtained. Patient enrollment was performed after informed consent had been obtained. Long echo time ^{1}H-MRSI data were acquired on a 1.5T clinical MR scanner (Siemens Magnetom Avanto) using a standard PRESS sequence (data points: 512, echo time: 135 ms, repetition time: 1000 ms, slice thickness: 15 mm, matrix size: 16×16, FOV: 160 mm×160 mm yielding a voxel size of 10 mm × 10 mm × 15 mm). The short echo time brain data were acquired from a healthy volunteer on a 3T clinical MR scanner (Phillips Achieva) at ultra-high spatial resolution (data points: 1024, echo time: 34 ms, repetition time: 1300 ms, averages: 1, slice thickness: 10 mm, matrix size: 32 × 32, FOV: 100 mm × 100 mm yielding a voxel size of 3.125 mm × 3.125 mm × 10 mm). The field-of-view reduction was achieved by inner-volume saturated PRESS based on a highly spatially selective T1- and B1- insensitive outer volume suppression (37).

All spectra were subject to a water peak removal using the hankel singular value decomposition (HSVD) algorithm as implemented in the *CSI tools* software[2] (38). Resonance lines of choline (Cho), creatine (Cr) and N-acteylaspartate (NAA) singlets were fit using Lorentzian line models (cf. Tab. 1). The single-voxel fits were performed using the AMARES algorithm from jMRUI (39) as well as an analogous implementation from *CSI tools*. As both yielded the same results, only results obtained with one of them are provided in the following. The GMRF for the spatially constrained fitting was implemented as an extension to the *CSI tools* version; for details of the implementation of the inference algorithm via Block-ICM see (22).

An estimate of the noise variance, which is needed for the GMRF (cf. Eq. (3)) approach, was obtained from the residuals of manually verified single voxel fits. The task dependent regularization parameter in Eq. (3) is the diagonal weighting matrix W. It has been determined once from the single

---

[2]Matlab based toolbox; freely available from `http://hci.iwr.uni-heidelberg.de/download3/`

voxel fits of the simulated MRS images based on a robust variogram estimate at pixel distance (22) and then used throughout all experiments. Hence, the smoothing parameters for the dampings $W_d^{(m)} = .2$, the frequencies $W_f^{(m)} = 2$ and the phase map $W_{\phi_0} = 20/\pi$ were fixed. The amplitude smoothing parameters $W_a^{(m)}$ were set to zero which means that no smoothing constraint was imposed on the amplitude maps. For the experiments in section 3.3, GMRFs with rescaled $W' = \lambda \times W$ ($\lambda = 1, 2, 4$) were also tested.

3.2 Simulation study

As the introduction of constraints may bias the solutions significantly, the bias-variance behavior of a single-voxel nonlinear least-squares fit and the GMRF approach was compared. Since this requires knowledge of ground truth, synthetic test data had to be generated. To this end long echo brain MRSI with 32×32 voxels at 1.5T was simulated using three singlets representing choline (Cho), creatine (Cr) and N-acetylaspartate (NAA) with the parameters provided in Tab. 1. Based on three frequency, damping, and amplitude maps and a phase shift map noiseless MRSI data were generated according to the signal model in Eq. (2) (dwell time $\Delta t = 1\,\mathrm{ms}$, imaging frequency $64\,\mathrm{MHz}$). For a Monte Carlo study $R = 100$ versions of the generated MRSI data were simulated by adding isotropic white Gaussian noise (σ_n).

- - - - - - - - - - Tab. 1 about here - - - - - - - - - -

When introducing spatial constraints into the analysis of an image, over-smoothing may easily occur, obscuring or even removing relevant spatial features. Introducing spatial constraints to the processing of a *spectral* image, however, may also affect the spectral dimension of the data. To study both aspects, two data sets with different parameter maps were generated for the purpose of emphasizing these different effects of a spatial prior:

- **Sharp edges.** First, spatially correlated random maps for all parameters have been generated by smoothing maps of independent Gaussian random values. Sharp edges have been introduced by subtracting the original amplitude in regions of the generated maps from a constant, sufficiently large to ensure non-negative amplitudes (cf. Fig. 3). This data set has been simulated with $N = 256$ data points and noise standard deviation of $\sigma_n = .232$.

- **Overlapping peaks.** Two ramp-shaped amplitude maps running vertically from zero to one for Cho and horizontally for Cr were created in order to provide multiple combinations of Cho and Cr amplitudes (cf. Fig. 6). NAA was simulated with a spatially correlated random amplitude map and a mean amplitude of 1.5. All frequency and phase shifts were set to zero, whereas all

11

dampings were increased by 10 Hz above the models in Tab. 1 as to cause overlapping spectral peaks. This data set was simulated with $N = 512$ data points and noise standard deviation of $\sigma_n = .155$.

Both synthetic data sets were fit using the known noise variance $\sigma^2 = \sigma_n^2$ and W as detailed above (Sec. 3.1). From the ground truth parameter θ and the R fit results $\hat{\theta}_i$ ($i = 1 \ldots R$) for the repeatedly simulated MRSI data the root-mean-squared error (RMSE) can be calculated for each voxel as

$$\text{rmse} = \left(\frac{1}{R} \sum_{i=1}^{R} (\hat{\theta}_i - \theta)^2 \right)^{1/2}. \tag{4}$$

Note that this is an estimate for the root of the expected squared error of the estimator $\hat{\theta}$ under the data distribution, $\sqrt{\mathrm{E}[(\hat{\theta} - \theta)^2]}$. More information can be gained from the decomposition of the RMSE into a bias and a standard deviation term which is provided by a *bias-variance decomposition* (40):

$$\text{rmse}^2 = \text{bias}^2 + \text{stdev}^2 \tag{5}$$

where

$$\text{bias} = \bar{\theta} - \theta \tag{6}$$

$$\text{stdev} = \left(\frac{1}{R} \sum_{i=1}^{N} (\bar{\theta} - \hat{\theta}_i)^2 \right)^{1/2} \tag{7}$$

with $\bar{\theta} = \frac{1}{R} \sum_{i=1}^{R} \hat{\theta}_i$. Root-mean-squared-error, bias and standard deviation were calculated for both synthetic MRSI data sets.

3.3 Expert rating of fit quality

To compare the performance of the standard single voxel and the spatially constrained fitting four series of MRSI image data were acquired with increasing spatial resolution, and decreasing signal-to-noise ratio. Spectral images of the brain were acquired from healthy volunteers on a 3T Siemens Trio Scanner (data points: 512, echo time: 135 ms, repetition time: 1000 ms). The four series varied in slice thickness and number of repeats. For each of these series three MRSI images were acquired with increasing spatial resolution overall within the range from 2 to 0.2 cm^3 nominal volume. This resulted in a total of 1433 spectra. Table 2 provides experimental details for the different data sets. The mean signal-to-noise ratios (SNRs) have been computed in frequency domain as $\text{SNR} = \text{I}/\sigma$ where I denotes the height of the NAA peak in the phase-corrected real spectrum (the peak intensity) and σ the standard deviation of the noise (2). All spectra were fit by both approaches as described above (section 3.1).

12

- - - - - - - - - - Tab. 2 about here - - - - - - - - - -

Fitting results were presented to two physicists with 2 years (FOK) and 7 years (BHM) of experience in MRS data analysis in a randomized and blinded fashion. In a first round spectra were inspected for their quality, i.e. the presence of artifacts or additional resonances, for example from lipids. Spectra which were deemed non-evaluable by one of the raters were assigned to an 'artifact' class and discarded for the evaluation. Then, in a second round, spectra and fits were inspected for quality. Criteria for a good fit were a match between spectrum and model fit in linewidth, position of the maximum, and amplitude.

In a first test the unconstrained single-voxel fitting results (i.e., $\lambda = 0$) were compared with the spatially constrained fitting results ($\lambda = 1$). After the inspection of the spectra the fraction of good (*i.e.* successful) fits for every MRSI data set (Tab. 2) was calculated, individually for each rater. The values for the single voxel fitting and the constrained fitting were then compared in a paired Wilcoxon test at 5% significance level. Also the *relative* decrease in poor fits resulting from employing the spatial constraint was calculated.

In a second test data from spatially constrained fits with $\lambda = 2$ and $\lambda = 4$ (section 3.1) were presented to the second rater twice, also in a randomized and blinded fashion. Again, spectra were evaluated for the goodness of fit. The analysis for $\lambda = 0$ (unconstrained, single voxel) and $\lambda = 1$ (the standard value for the GMRF fit) was also repeated another time. As a result every spectrum was labeled twice from the same rater for $\lambda = 0, 1, 2, 4$. (We decided to use the same rater twice, as the inter-operator variation showed to be significantly larger than variation from different signal processing). This second test aimed at evaluating the sensitivity of the fitting result with respect to the choice of the regularization parameter λ.

4 Results

Large differences were observed between the sequential single voxel fitting and the coupled GMRF approach on the real data sets (Sec. 4.1). To understand these differences the smoothing behavior (Sec. 4.2) and peak resolution (Sec. 4.3) on the synthetic data was studied, and the fit quality as a function of the employed GMRF model was evaluated on real data in section 4.4.

4.1 Real data

Figure 1 shows nine adjacent voxels from a clinical 1.5T long echo time brain data set (36) with clear Cho, Cr and NAA peaks. A tumor is present in the upper part of the array as visible from the T1-

13

weighted MRI and the drastically elevated Cho levels, and low values for NAA in the MRSI. In four of the nine voxels the single voxel fit (left) explains Cho and Cr with only one resonance component. Here, the GMRF prior can resolve the two peaks in all voxels (right), and irrespectively of the changes of the spectral pattern at the interface ("edge") of normal and tumorous tissue, by enforcing that the damping (line width) of the resonance lines should not change rapidly in neighboring voxels.

- - - - - - - - - - Fig. 1 about here - - - - - - - - - -

Figure 2 provides another comparison between sequential single voxel and spatial MRSI fitting using a larger, high-resolution brain MRSI data set acquired with short echo time at 3T. While the phase, damping and frequency maps are significantly smoothed from the GMRF prior on these quantities, the amplitude map, indicating the metabolite concentrations, maintains sharp spatial features. Most speckles in the amplitude map of the single voxel approach (top) are removed in the spatial fitting. Overall, much clearer metabolite maps are obtained for the GMRF approach, with clear-cut edges and in good agreement with the physiological structures visible from MRI.

- - - - - - - - - - Fig. 2 about here - - - - - - - - - -

4.2   Simulation – Sharp Edges

The simulation study provides further quantitative insights in the smoothing behavior and the resolution of single peaks. Figure 3 presents results on the first simulated data set ("sharp edges"). The first row shows the employed ground truth amplitude maps, *i.e.* the amplitude parameters a_m for Cho, Cr and NAA, and the following rows show difference images to the solutions obtained for the same simulated data with different fitting approaches. The second row shows results when using the spatial prior (GMRF) whereas the third row shows results when using a single voxel method (AMARES). The results in the last row have been obtained by smoothing the amplitude maps from AMARES with a Gaussian filter where the kernel width has been chosen for each metabolite as to minimize the squared error to the ground truth image. Note that such an optimization is certainly not possible given real data and the obtained results are overoptimistic in favor of the smoothed single voxel approach.

- - - - - - - - - - Fig. 3 about here - - - - - - - - - -

Compared with the single voxel, the spatial approach visibly improves the estimate of Cho and Cr in many voxels. The smoothed single voxel solution seems to be considerably better but also exhibits

14

spatially correlated errors predominantly at the image border and at the edges. For NAA hardly any difference can be seen between the three amplitude maps in Figure 3.

The Monte Carlo study using all $R = 100$ repetitions of the "sharp edges" data reveals that the root-mean-squared error of the NAA amplitude estimate actually improves when using the spatial prior (Fig. 4a). Even more improvement is obtained for the remaining parameters. As apparent from the presented bias-variance decomposition of the RMSE in Figure 4 the gain is mainly due to a reduction in standard deviation. Furthermore, the GMRF prior does not seem to introduce significantly more bias than AMARES for most voxels in the "sharp edges" data.

- - - - - - - - - - Fig. 4 about here   - - - - - - - - - -

A comparison of the standard deviations of the parameter estimates with the corresponding Cramér-Rao lower bound (CRLB) for all metabolites in the "sharp edges" data is presented in Figure 5. As expected from theory, the unbiased single voxel estimates do not beat the CRLB since (nearly) all points fall above the diagonal in the scatter plots of Figure 5a. Occasional points below the diagonal are due to the fact that the standard deviations obtained from the Monte Carlo study are only estimates based on $R = 100$ repetitions and therefore exhibit some randomness themselves. Systematic improvements are obtained with the spatial prior as apparent from Figure 5b. In line with the observed clear improvement of the Cho and Cr difference images in Figure 3, most gain in terms of standard deviation is obtained for these two metabolites (red and blue dots).

- - - - - - - - - - Fig. 5 about here   - - - - - - - - - -

4.3   Simulation – Overlapping Peaks

Concentrating on the overlapping peaks of Cho and Cr, the bias-variance analysis on the second simulated data set ("overlapping peaks") confirms that the spatial prior heavily affects the standard deviation of the amplitude estimates which make the main contribution to the RMSE (cf. Fig. 6). Surprisingly, however, the single voxel approach also shows more bias than the GMRF approach for small amplitudes at the left and upper border of the images. This effect is even more pronounced when, at the same time, one metabolite has a high amplitude and the other a low amplitude. Then also the standard deviation raises significantly. The GMRF approach does not exhibit such behavior.

- - - - - - - - - - Fig. 6 about here   - - - - - - - - - -

15

An analysis of these border voxels reveals that the observed effect is due to the fact that the single voxel approach sometimes explains both the Cho and Cr resonances with only one component and sets the amplitude estimate for the smaller metabolite peak to zero as for the example shown in Figure 7. In contrast, the GMRF approach uses information from neighboring voxels to infer that an additional small peak is more likely which is in accordance with the known ground truth of this simulated example.

- - - - - - - - - - Fig. 7 about here - - - - - - - - - -

4.4 Differences in the experts' rating

In Figure 8 three example spectra of adjacent voxels from the long echo time brain MRSI data set (Tab. 2, series 4.3) are shown together with the single-voxel and GMRF fits. In the middle voxel both the Cho and the Cr peaks are fit with a shifted frequency using the single-voxel approach while the GMRF prior prevents the sudden frequency jump by comparison with the neighboring voxels. Also, the phase estimates are much more consistent with the GMRF prior.

- - - - - - - - - - Fig. 8 about here - - - - - - - - - -

Having observed a large benefit from the GMRF prior in individual cases as well as in the simulation study, it is interesting to evaluate its effect for different data sets and different experimental conditions. Instead of looking at fit results for individual voxels, the effect of the GMRF prior is evaluated on different MRSI data sets (cf. Tab. 2) in the following.

The fraction of successful fits correlates well with the SNR and the volume of the voxels for all series (Fig. 9, left and middle). Differences can also be seen between series 3 and 4 which differ by the number of scan averages and, hence, with respect to their SNR (cf. Tab. 2). Large differences are observed in particular for the four images with the highest spatial resolution – all of which with similar nominal voxel volume (Fig. 9). This may be attributed to the different partial volume effects. One might notice that rater 1 discarded approx. 20% more spectra than rater 2, implicitly using much stricter criteria to assign a spectrum to the "evaluable" class. Despite this difference the decisions made by both raters follow the same general trends.

- - - - - - - - - - Fig. 9 about here - - - - - - - - - -

The spatially constrained fit performs significantly better than the unconstrained (p-value on the "evaluable" data: rater 1 $\leq$.005, rater 2 $\leq$.05; on the complete data set: rater 1 $\leq$.01, rater 2 $\leq$.05).

16

The spatial constraint improves the quality of the fits in ten out of the twelve MRSI images (Fig. 9, right). In the worst case the constrained fit results in a 3% increase of low quality fits when compared to the single voxel approach; in the best case poor fits are removed completely when using the spatial prior (-100%). On average, the percentage of poor fits is reduced by approx. 10% when using a spatial constraint as indicated by the horizontal lines in Figure 9, right (9.5% rater 1, 13.7% rater 2). This value is relatively constant over all image series, and both raters.

- - - - - - - - - - Fig. 10 about here - - - - - - - - - -

A high spatial regularization (Fig. 10) leads to fitting results which, for large values of λ, are worse than unconstrained fits. Unconstrained fits, with $\lambda = 0$, are not optimal either; an intermediate value of $\lambda = 1$ is found to be a good choice for improving fitting results on a wide range of experimental settings, resolution, and rating criteria.

Finally, Tab. 2 also lists computation times for both approaches along with their ratios. For series 1 and 2 the computation times increase with falling SNR. In contrast, computation times for series 3 and 4 are decreased for the intermediate SNR level and increased for very high and very low SNRs. Overall, the GMRF approach follows the same trends and takes between 3.5 and 6.9 times longer than the SV approach. However, the ratio increases with SNR for series 1 and 2 while it decreases for series 3 and 4. This effect might be due to the fact that series 3 and 4 exhibit more point spread which makes the smoothness assumption of the GMRF more valid.

5   Discussion

5.1   Benefit of a biased estimator

The simulation study shows that by using spatial prior knowledge the estimation variance can be decreased significantly, even below the Cramér-Rao lower bound (CRLB) of the single voxel approach. The CRLB, however, provides a lower bound for *unbiased* estimators only. It can be beaten when using a *biased* estimate. This is also well known for an AMARES-type single voxel fitting and can be observed with (slightly misspecified) parameter constraints. For example, if all frequency shifts in the simulated "sharp edges" data set had been constrained to vary by only $\pm 1\,\mathrm{Hz}$, the standard deviation of the frequency estimates could not have exceeded $2\,\mathrm{Hz}$. This would be better than the CRLB for some of the voxels (cf. Fig. 5a). However, if the true frequency lay outside the interval of the constraints the estimate would be biased. In general, prior knowledge and specific constraints on the fitting model may reduce estimation variance, but may also largely increase bias if not appropriate.

17

This emphasizes that specifying a correct spectral model is as important for a spatially constrained fitting, as in a single voxel approach. A learning-based approach for testing whether spectra of an MR spectroscopic image are in accordance with the predefined spectral model, and potentially to be used here, has been proposed by Menze *et al.* (3). It may be used to minimize bias arising from imperfect spectral models. Still, the CRLB assumes statistical independence between the FID parameters and a spatial model such as the proposed GMRF prior certainly violates this assumption.

So, the question arises, whether *spatial* prior knowledge in form of the proposed GMRF is appropriate or not, *i.e.* whether it evokes bias or not. This question cannot be answered satisfactorily by a Monte Carlo study since it would require the full and correct specification of the data generating process. However, a number of arguments can be put forward for justifying that parameter maps are truly smooth and thus for justifying the use of a spatial prior: First, the employed MR sequences always cause a certain point spread inevitably leading to FID signals that are spatially correlated. Second, parameters such as frequency shift and damping very much depend on the homogeneity of the applied B0-field which is optimized in preceding shimming procedures; ideally these parameters may therefore be constant across the whole region of interest. Third, it has been observed that the zero-phase parameter only varies smoothly within homogeneous tissue regions (in accordance with (27)). Finally, exemplary fits seem to be much more plausible when using the proposed spatial prior to encourage similar resonance line shift or widths in neighboring spectra (Fig. 1c), than without it (Fig. 1b).

5.2   Optimally adjusting the spatial prior

In the presented experimental study no direct smoothness prior was imposed on the amplitude maps, since this is the main parameter of interest in quantification. Variation is expected for the amplitude parameter which should not be smoothed away. To this end, the "sharp edges" data set (cf. Fig. 3) demonstrates that edges can be reconstructed in amplitude maps without blurring. By contrast the simple alternative of a post-hoc smoothing of single voxel estimates leads to blurring at the edges.

If no sharp edges are expected one could decide to perform some spatial smoothing of the amplitude maps as well, which might further improve results in homogeneous image regions. This is easily achieved by choosing a small positive weight W_a. Further work is necessary to devise a procedure for fine-tuning the weighting matrix W that is better than the proposed variogram-based approach.

Some MRS image data, for example, from the acquisition of short echo time [1]H-MRSI of the brain, require appropriate baseline handling for obtaining optimal results, and most single voxel fitting algorithms allow to consider baselines stemming from macromolecules (11, 12). The proposed spatial-fitting framework using GMRFs would also allow to consider similar nonparametric model extensions. An interesting direction could be the spatial regularization of baseline parameters. This will be left to

be a subject of future work as well.

5.3   Related approaches

Related approaches for the incorporation of spatial prior knowledge in the fitting of MRSI have been proposed before. One such approach, which is followed in FITT (8), for example, adjusts the initialization of the nonlinear fitting problem based on estimates obtained in neighboring voxels from a previous run or sweep. While this approach might reduce the risk of running into non-global optima it does not spatially regularize the fit and thus does not help reduce the statistical variance of the parameter estimates as the GMRF approach (cf. Fig. 4, Fig. 5b). The Block-ICM algorithm as described in (22) avoids non-global optima by an iterative optimization scheme that prevents premature convergence in early sweeps.

A second approach, followed in LCModel (9, 26), for example, performs single voxel fitting using box or soft constraints with parameters from previously fitted neighboring voxels according to a certain visiting schedule (*e.g.* sequential, spiral or checkerboard). Compared to the spatial MRSI mode of LCModel (26), our approach mainly differs in three ways. Firstly, the voxels are grouped into blocks and the optimization problem is solved exactly for each block (given its surroundings). Secondly, instead of processing the voxels in a radial order (from the interior to the outward voxels), a checkerboard visiting schedule is used. Thirdly, multiple passes over the entire volume are performed: hence every voxel is influenced by the fitting results of all others, predominantly its close neighbors. In contrast to the GMRF approach, LCModel very much depends on the choice of order and the fitting results may deteriorate if spectra with poor signal quality are processed first.

In contrast to the linear model proposed by Bretthorst (27) the GMRF approach allows to capture higher order, nonlinear spatial variations in phase, damping or frequency and to impose intermediate smoothness priors as illustrated in Figure 2.

Only recently, an interesting approach has been proposed in (28) which combines three of the above methods. Using a sequential visiting schedule several sweeps are performed. In the first round of sweeps, a parameter initialization approach similar to (8) is employed using median values from a local neighborhood. In the second round of sweeps, spatially dependent box constraints are applied. In the third and last round of sweeps spatial soft constraints are emplyoed, *i.e.* the GMRF objective is solved using the (single-site) ICM algorithm. While the overall optimization scheme finally attempts to solve the same GMRF objective that we propose it may end up in different optima and exhibit different convergence behavior. As compared to our proposed Block-ICM algorithm, however, the approach in (28) has many more hyper-parameters that may need careful tuning. Unlike our work, (28) and (41) provide results on short echo time brain MRSI quantifying more than three, *i.e.* eleven,

19

metabolites and show the beneficial effect of spatial prior knowledge in this setting.

5.4 Improvements and further directions

Our experiments were performed with a basic, AMARES-type fitting routine for the sole purpose of proving the benefit of using spatial context in fitting the spectral model. One should note that the spatial constraint proposed here is *additive* to any other spectral fitting constraint. It can straightforwardly be used in conjunction with any other single voxel signal model including approaches based on experimentally measured template spectra. As a consequence the benefits of spatial fitting can either be traded with other constraints – relaxing, for example, an overly narrow range on expected phase, damping, or frequencies – or can be used to improve the overall quality of a fit in a high resolution MRSI with low SNR.

The use of spatial context from morphological MRI has proven to be beneficial in the evaluation of the metabolic information from MRSI. In the grading of tumors, for example, this joint analysis proved to be more reliable than the analysis of the individual modalities alone (42–44); it also allows the image-guided estimation of pure tissue spectra (45). Here, a probabilistic spatial fitting approach as the one proposed in this study opens new directions for considering spatial variation and, hence, considering spatial context in MRSI signal processing: Multi-modal spatial information from morphologic MR images may easily be used to adjust the coupling of neighboring voxels. In the present study W (and λ) was set to one single, constant value. One could easily think of an image-guided local adjustment of $\lambda(x)$ by using a segmented tissue map obtained from MRI, for example. This is left to future extensions of our proposed approach.

6 Conclusions

The use of *spatial* prior knowledge in addition to commonly employed parameter constraints has been proposed for an improved spectral fitting of magnetic resonance spectroscopic images. Using a Gaussian Markov random field that favors smooth maps for selected parameters such as phase, line width and frequency shift, the standard deviation of parameter estimates can be reduced below the Cramér-Rao lower bound that is obtained for the single voxel approach. An important advantage of using spatial prior knowledge is that overlapping peaks can often be resolved reliably, also in cases where the single voxel approach fails to do so (Fig. 1).

The proposed method is found particularly useful for high-resolution MRSI and it allows to derive detailed metabolic images. A relatively constant improvement of 10% is found in the fit quality when using a small spatial constraint, with considerable differences in particular for images with low signal-to-noise ratio. One may recommend to use a spatial prior as default on such data.

A multi-modal analysis of image content may provide further directions for an integrated use of metabolic information from MRSI together with structural information from morphological MRI. A principled probabilistic approach for the spatial fitting of MRSI – such as the one based on Gaussian Markov random fields proposed here – provides the technical means for stepping further into this direction.

7   Acknowledgments

References

1. Schröder L, Lowery TJ, Hilty C, Wemmer DE, Pines A. Molecular Imaging using a Targeted Magnetic Resonance Hyperpolarized Biosensor. *Science* 2006; **314**: 446–449. doi:10.1126/science.1131847.

2. Kreis R. Issues of spectral quality in clinical ^{1}H-magnetic resonance spectroscopy and a gallery of artifacts. *NMR Biomed.* 2004; **17**: 361–381. doi:10.1002/nbm.891.

3. Menze BH, Kelm BM, Weber MA, Bachert P, Hamprecht FA. Mimicking the Human Expert: Pattern Recognition for an Automated Assessment of Data Quality in MRSI. *Magn. Reson. Med.* 2008; **59**: 1457–1466. doi:10.1002/mrm.21519.

4. Soher BJ, Maudsley AA. Evaluation of variable line-shape models and prior information in automated ^{1}H spectroscopic imaging analysis. *Magn. Reson. Med.* 2004; **52**: 1246–54. doi:10.1002/mrm.20295.

5. Vanhamme L, van den Boogaart A, Van Huffel S. Improved Method for Accurate and Efficient Quantification of MRS Data with Use of Prior Knowledge. *J. Magn. Reson.* 1997; **129**: 35–43. doi:10.1006/jmre.1997.1244.

6. Mierisová Š, van den Boogaart A, Tkáč I, Van Hecke P, Vanhamme L, Liptaj T. New approach for quantitation of short echo time *in vivo* ^{1}H MR spectra of brain using AMARES. *NMR Biomed.* 1998; **11**: 32–39.

7. Cavassila S, Deval S, Huegen C, van Ormondt D, Graveron-Demilly D. Cramér-Rao bounds: an evaluation tool for quantitation. *NMR Biomed.* 2001; **14**: 278–283. doi:10.1002/nbm.701.

8. Soher BJ, Young K, Govindaraju V, Maudsley AA. Automated spectral analysis III: application to in vivo proton MR spectroscopy and spectroscopic imaging. *Magn. Reson. Med.* 1998; **40**: 822–831. doi:10.1002/mrm.1910400607.

9. Provencher S. Automatic quantitation of localized *in vivo* ^{1}H spectra with LCModel. *NMR Biomed.* 2001; **14**: 260–264. doi:10.1002/nbm.698.

10. Ratiney H, Sdika M, Coenradie Y, Cavassila S, van Ormondt D, Graveron-Demilly D. Time-domain semi-parametric estimation based on a metabolite basis set. *NMR Biomed.* 2005; **18**: 1–13. doi:10.1002/nbm.895.

11. Poullet JB, Sima DM, Simonetti AW, De Neuter B, Vanhamme L, Lemmerling P, Van Huffel S. An automated quantitation of short echo time MRS spectra in an open source software environment: AQSES. *NMR Biomed.* 2006; **20**: 493–504. doi:10.1002/nbm.1112.

12. Sima DM, Van Huffel S. Regularized semiparametric model identification with application to nuclear magnetic resonance signal quantification with unknown macromolecular base-line. *J. Roy. Stat. Soc. B Met.* 2006; **68**: 383–409. doi:10.1111/j.1467-9868.2006.00550.x.

13. Schulte R, Boesiger P. ProFit: two-dimensional prior-knowledge fitting of J-resolved spectra. *NMR Biomed.* 2006; **19**: 255–263. doi:10.1002/nbm.1026.

14. Hu X, Levin DN, Lauterbur PC, Spraggins T. SLIM: Spectral localization by imaging. *Magn. Reson. Med.* 1988; **8**: 314–322. doi:10.1002/mrm.1910080308.

15. Liang ZP, Lauterbur PC. A generalized series approach to MR spectroscopic imaging. *IEEE T. Med. Imaging* 1991; **10**: 132–137. doi:10.1109/42.79470.

16. Plevritis SK, Macovski A. MRS imaging using anatomically based k-space sampling and extrapolation. *Magn. Reson. Med.* 1995; **34**: 686–693. doi:10.1002/mrm.1910340506.

17. Bao Y, Maudsley AA. Improved Reconstruction for MR Spectroscopic Imaging. *IEEE T. Med. Imaging* 2007; **26**: 686–695. doi:10.1109/TMI.2007.895482.

18. Haldar JP, Hernando D, Song SK, Liang ZP. Anatomically constrained reconstruction from noisy data. *Magn. Reson. Med.* 2008; **59**: 810–818. doi:10.1002/mrm.21536.

19. Kornak J, Young K, Soher BJ, Maudsley AA. Bayesian k-space-time reconstruction of MR spectroscopic imaging for enhanced resolution. *IEEE T. Med. Imaging* 2010; **29**: 1333–1350. doi:10.1109/TMI.2009.2037956.

20. Gindi G, Lee M, Rangarajan A, Zubal IG. Bayesian reconstruction of functional images using anatomical information as priors. *IEEE T. Med. Imaging* 1993; **12**: 670–680. doi:10.1109/42.251117.

21. Baillet S, Garnero L. A Bayesian approach to introducing anatomo-functional priors in the EEG/MEG inverse problem. *IEEE T. Bio-Med. Eng.* 1997; **44**: 374–385. doi:10.1109/10.568913.

22. Kelm BM, Menze BH, Nix O, Zechmann C, Hamprecht FA. Estimating Kinetic Parameter Maps From Dynamic Contrast-Enhanced MRI Using Spatial Prior Knowledge. *IEEE T. Med. Imaging* 2009; **28**: 1534–1547. doi:10.1109/TMI.2009.2019957.

23. Kelm BM, Müller N, Menze BH, Hamprecht FA. Bayesian Estimation of Smooth Parameter Maps for Dynamic Contrast-Enhanced MR Images with Block-ICM. *Proceedings of the 7th IEEE Computer Society Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA).* New York, USA, 2006; 96–103. doi:10.1109/CVPRW.2006.41.

24. Schmid V, Whitcher B, Padhani A, Taylor N, Yang G. Bayesian Methods for Pharmacokinetic Models in Dynamic Contrast-Enhanced Magnetic Resonance Imaging. *IEEE T. Med. Imaging* 2006; **25**: 1627–1636. doi:10.1109/TMI.2006.884210.

25. Kelm BM. *Evaluation of Vector-Valued Clinical Image Data Using Probabilistic Graphical Models: Quantification and Pattern Recognition.* Ph.D. thesis, University of Heidelberg, 2007.

26. Provencher S. *LCModel & LCMgui User's Manual*, 6.2 edn., 2010. URL `http://s-provencher.com/pub/LCModel/manual/manual.pdf`. Accessed on November 15, 2010.

27. Bretthorst GL. Automatic phasing of MR images. Part I: Linearly varying phase. *J. Magn. Reson.* 2008; **191**: 184–192. doi:10.1016/j.jmr.2007.12.010.

28. Sava AC, Sima DM, Poullet JB, Wright AJ, Heerschap A, Huffel SV. Exploiting spatial information to estimate metabolite levels in 2D MRSI of heterogeneous brain lesions. Internal Report 09-182, ESAT-SISTA, K.U.Leuven (Leuven, Belgium), 2009.

29. Smith SA, Levante TO, Meier BH, Ernst RR. Computer Simulations in Magnetic Resonance. An Object-Oriented Programming Approach. *J. Magn. Reson. Ser. A* 1994; **106**: 75–105. doi:10.1006/jmra.1994.1008.

30. Besag JE. On the Statistical Analysis of Dirty Pictures. *J. Roy. Stat. Soc. B Met.* 1986; **48**: 259–302.

31. Bouman CA, Sauer K. A unified approach to statistical tomography using coordinate descent optimization. *IEEE T. Image Process.* 1996; **5**: 480–492. doi:10.1109/83.491321.

32. Winkler G. *Image Analysis, Random Fields and Dynamic Monte Carlo Methods.* Springer: Berlin, 2003.

33. Wu Z, Chung HW, Wehrli FW. A Bayesian approach to subvoxel tissue classification in NMR microscopic images of trabecular bone. *Magn. Reson. Med.* 1994; **3**: 302–308. doi:10.1002/mrm.1910310309.

34. Toselli A, Widlund O. *Domain Decomposition Methods – Algorithms and Theory.* Springer: Berlin, 2005.

35. Trottenberg U, Oosterlee CW, Schuller A. *Multigrid.* Academic Press: San Diego, 2000.

36. Weber MA, Zoubaa S, Schlieter M, Juttler E, Huttner HB, Geletneky K, Ittrich C, Lichy MP, Kroll A, Debus J, Giesel FL, Hartmann M, Essig M. Diagnostic performance of spectroscopic and perfusion MRI for distinction of brain tumors. *Neurology* 2006; **66**: 1899–1906. doi:10.1212/01.wnl.0000219767.49705.9c.

37. Henning A, Schar M, Schulte RF, Wilm B, Pruessmann KP, Boesiger P. SELOVS: Brain MRSI localization based on highly selective T_1- and B_1-insensitive outer-volume suppression at 3T. *Magn. Reson. Med.* 2008; **59**: 40–51. doi:10.1002/mrm.21374.

38. Kelm BM, Menze BH, Neff T, Zechmann C, Hamprecht FA. CLARET: a tool for fully automated evaluation of MRSI with pattern recognition methods. *Proceedings of Bildverarbeitung für die Medizin (BVM).* Hamburg, Germany, 2006; 51–55.

39. Naressi A, Couturier C, Castang I, de Beer R, Graveron-Demilly D. Java-based graphical user interface for MRUI, a software package for quantitation of in vivo/medical magnetic resonance spectroscopy signals. *Comput. Biol. Med.* 2001; **31**: 269–286.

40. Hastie T, Tibshirani R, Friedman J. *The Elements of Statistical Learning.* Springer: New York, 2001.

41. Sima DM, Sava AC, Huffel SV. Adaptive alternating minimization for fitting magnetic resonance spectroscopic imaging signals. Internal Report 09-232, ESAT-SISTA, K.U.Leuven (Leuven, Belgium), 2009.

42. Nelson SJ. Analysis of volume MRI and MR spectroscopic imaging data for the evaluation of patients with brain tumors. *Magn. Reson. Med.* 2001; **46**: 228–239. doi:10.1002/mrm.1183.

43. Simonetti AW, Melssen WJ, van der Graaf M, Postma GJ, Heerschap A, Buydens LMC. A Chemometric Approach for Brain Tumor Classification Using Magnetic Resonance Imaging and Spectroscopy. *Anal. Chem.* 2003; **75**: 5352–5361. doi:10.1021/ac034541t.

44. Luts J, Laudadio T, Idema AJ, Simonetti AW, Heerschap A, Vandermeulen D, Suykens JA, Van Huffel S. Nosologic imaging of the brain: segmentation and classification using MRI and MRSI. *NMR Biomed.* 2009; **22**: 374–390. doi:10.1002/nbm.1347.

45. Gruber S, Mlynárik V, Moser E. High-resolution 3D proton spectroscopic imaging of the human brain at 3T: SNR issues and application for anatomy-matched voxel sizes. *Magn. Reson. Med.* 2003; **49**: 299–306. doi:10.1002/mrm.10377.

Tables

metabolite	$f_m(\mathrm{Hz})$	$d_m(\mathrm{Hz})$	$a_m(\text{a.u.})$	$\phi_m(\mathrm{rad})$
Choline (Cho)	$-94.2 \ (-188.4)$	12 (8)	1	0
Creatine (Cr)	$-107.5 \ (-215.0)$	12 (8)	1	0
N-Acetyl Aspartate (NAA)	$-171.2 \ (-342.4)$	12 (8)	1	0

Table 1: Lorentzian metabolite models for long echo brain spectra at 1.5T and short echo brain spectra at 3T in parentheses.

Series	$\triangle x$ [mm]	$\triangle y$ [mm]	$\triangle z$ [mm]	volume [cm^3]	voxels	avg	SNR	SV time [ms/voxel]	GMRF time [ms/voxel]	ratio
1	10.00	10.00	20.00	2.00	10^2	3	17.62	15.06	52.37	3.5
	7.00	7.00	20.00	0.98	13^2	3	9.80	15.44	70.99	4.6
	3.44	3.44	20.00	0.24	13^2	3	4.31	20.75	143.15	6.9
2	10.00	10.00	10.00	1.00	8^2	3	11.72	11.07	60.70	5.5
	6.88	6.88	10.00	0.47	11^2	3	6.23	17.39	88.10	5.1
	5.00	5.00	10.00	0.25	12^2	3	4.46	24.37	158.44	6.5
3	10.00	10.00	10.00	1.00	8^2	6	35.69	9.68	57.87	6.0
	8.00	8.00	8.00	0.51	10^2	6	21.03	8.75	50.83	5.8
	6.00	6.00	6.00	0.22	13^2	6	9.75	11.52	55.91	4.9
4	10.00	10.00	10.00	1.00	8^2	3	28.33	9.45	64.52	6.8
	8.00	8.00	8.00	0.51	10^2	3	16.72	8.80	54.85	6.2
	6.00	6.00	6.00	0.22	13^2	3	7.75	13.59	60.67	4.5

Table 2: Resolution series used to compare the spatially constrained and the unconstrained fitting. Side lengths, volume and number of MRSI voxels as well as number of averages and estimated mean SNR are given. All spectra were acquired on a rectangular grid. The last three columns list computation times with and without spatial prior and their ratio. Three example spectra from series 4.3 along with their single-voxel and GMRF fits are shown in Fig. 8.

27

Figure legends

Figure 1: Nine adjacent voxels from a patient with brain tumor and the corresponding AMARES and GMRF fits. The overlapping Cho and Cr peaks lead to erroneous single voxel fits in four of the nine voxels which are well captured when using spatial prior knowledge.

Figure 2: Metabolite maps from the brain of a healthy volunteer at 3T (32×32 voxels). The upper row shows the parameter maps obtained with single voxel fitting and the bottom row shows the results from the GMRF approach. While the phase, damping and frequency maps are significantly smoothed from the spatial prior, the amplitude map that indicates the metabolite concentrations maintains sharp spatial features.

Figure 3: Simulated MRSI data with sharp amplitude edges. From top to bottom: ground truth; difference to GMRF estimate; difference to AMARES; difference to optimally smoothed AMARES. Since the GMRF approach does not smooth amplitudes, sharp edges are not smeared. With post-hoc smoothing (last row) the estimates of Cho and Cr can be improved, however, at the cost of border effects and the oversmoothing of edges. As expected for high signal-to-noise-ratios, all methods yield similar results for NAA.

Figure 4: Scatter plots of root-mean-squared error (RMSE), bias and standard deviation (stdev) of NAA amplitudes as estimated from $R = 100$ repetitions of the simulated "sharp edges" MRSI data. Each dot corresponds to one of the 32×32 voxels. The GMRF approach yields lower RMSE than AMARES in all parameters. Since the bias is comparable, the gain must be entirely ascribed to a reduction in standard deviation.

Figure 5: Scatter plots of Cramér-Rao lower bounds (CRLB) against standard deviations (stdev) of the parameter estimates obtained with AMARES and the proposed GMRF approach. The results are based on 100 repetitions of simulated MRSI images containing the three metabolites Cho (red), Cr (blue) and NAA (green). Each dot corresponds to one of the 32×32 voxels. While the single voxel approach does not beat the CRLB, using spatial prior knowledge can reduce the estimation variance below the theoretical CRLB.

Figure 6: Bias, standard deviation (stdev) and root-mean-squared error (rmse) for amplitude estimates using the single voxel (sv, AMARES) and spatial (sp, GMRF) approaches when fitting two overlapping peaks: Choline (Cho) and Creatine (Cr). Example spectra are shown in Figure 7.

Figure 7: Four example spectra from the simulated data analyzed in Fig. 6. The examples are taken from the corners of the simulated spectral image and contrast the behavior of single voxel and spatially regularized fit for overlapping peaks with varying amplitudes. The examples show that with sufficient amplitude both, the Cho and the Cr-component can be fitted reliably. If one of the peaks gets small, however, the single voxel approach runs into problems while the spatial prior produces stable fits using information from a local neighborhood.

Figure 8: Spectra of three consecutive voxels from an MRSI brain scan at 3T with long echo time (Series 4.3, Tab. 2) and spectral line fits for Cho (3.2 ppm), Cr (3.0 ppm) and NAA (2.0 ppm). The top row shows results from single-voxel fitting whereas the bottom row shows the corresponding fits obtained with using spatial prior knowledge. The single voxel approach fails to fit the Cho and Cr peaks in the middle voxel whereas the GMRF uses information from the neighboring spectra to stabilize the fit. Furthermore, the GMRF fit shows a much more stable phase estimate than the single-voxel fit.

Figure 9: Evaluation of the spatially constrained and the unconstrained fitting routines on real data. The left and the central figure show the fraction of successful fits as a function of voxel volume for both raters, and on all resolution series from Tab. 2. The GMRF fitting (with $\lambda = 1$) performs better in nearly all tests. The right figure presents this benefit in values relative to the performance of the unconstrained SV results. On average the GMRF constraint reduced the fraction of failed fits by approx. 10%.

Figure 10: Effect of the regularization parameter λ on the fraction of successful fits. The six lines show the results for all data of resolution series 1 and 2 (Tab. 2) after double evaluation by rater 2. The notches indicate the extrema (*i.e.* the two obtained percentages) while the lines follow the average. A value of $\lambda = 1$ performs well for nearly all data sets, over all resolutions and corresponding signal-to-noise ratios.

Figures

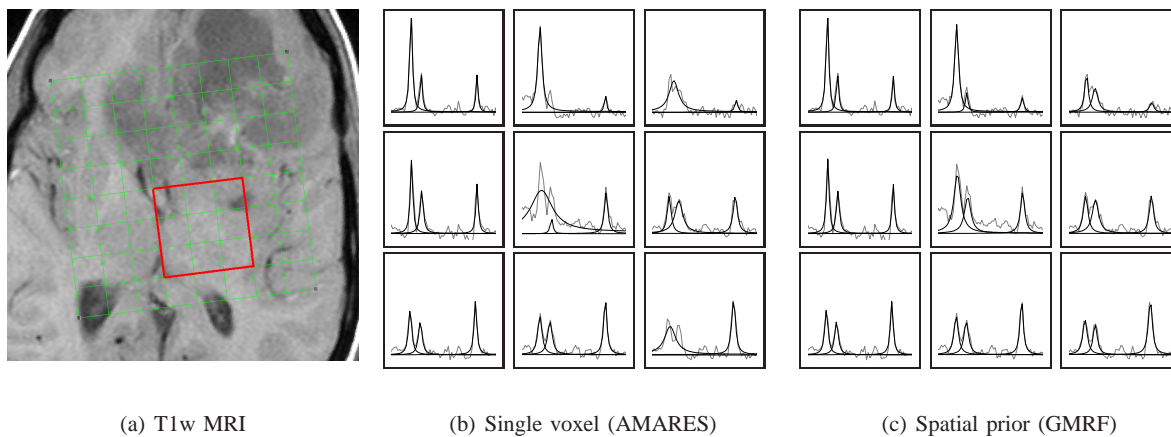(a) T1w MRI   (b) Single voxel (AMARES)   (c) Spatial prior (GMRF)

Figure 1: Nine adjacent voxels from a patient with brain tumor and the corresponding AMARES and GMRF fits. The overlapping Cho and Cr peaks lead to erroneous single voxel fits in four of the nine voxels which are well captured when using spatial prior knowledge.
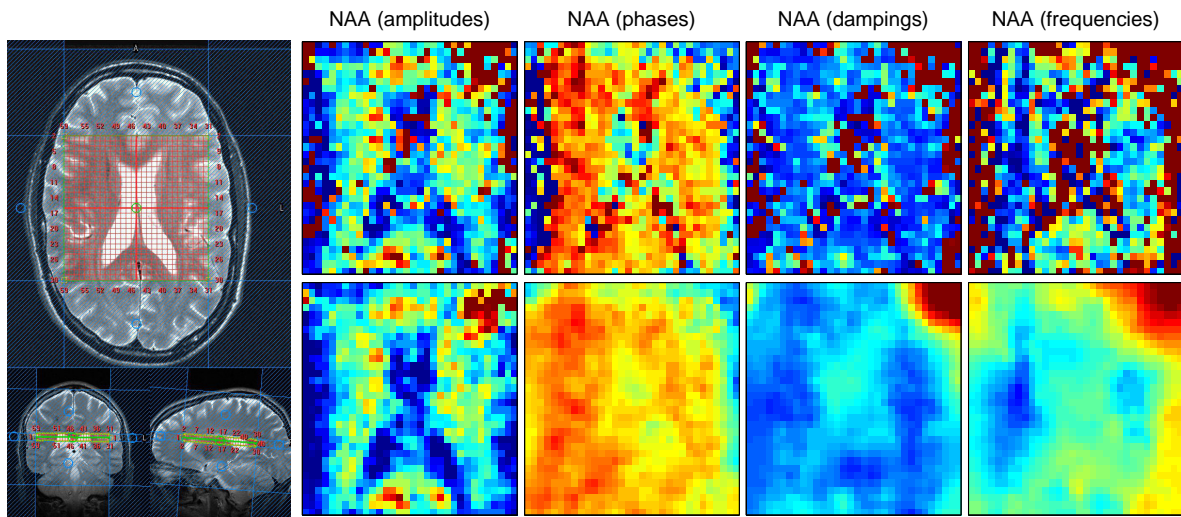
Figure 2: Metabolite maps from the brain of a healthy volunteer at 3T (32×32 voxels). The upper row shows the parameter maps obtained with single voxel fitting and the bottom row shows the results from the GMRF approach. While the phase, damping and frequency maps are significantly smoothed from the spatial prior, the amplitude map that indicates the metabolite concentrations maintains sharp spatial features.

Figure 3: Simulated MRSI data with sharp amplitude edges. From top to bottom: ground truth; difference to GMRF estimate; difference to AMARES; difference to optimally smoothed AMARES. Since the GMRF approach does not smooth amplitudes, sharp edges are not smeared. With post-hoc smoothing (last row) the estimates of Cho and Cr can be improved, however, at the cost of border effects and the oversmoothing of edges. As expected for high signal-to-noise-ratios, all methods yield similar results for NAA.

(a) root-mean-squared error   (b) bias   (c) standard deviation

Figure 4: Scatter plots of root-mean-squared error (RMSE), bias and standard deviation (stdev) of NAA amplitudes as estimated from $R = 100$ repetitions of the simulated "sharp edges" MRSI data. Each dot corresponds to one of the 32×32 voxels. The GMRF approach yields lower RMSE than AMARES in all parameters. Since the bias is comparable, the gain must be entirely ascribed to a reduction in standard deviation.

(a) Single voxel (AMARES) (b) Spatial prior (GMRF)

Figure 5: Scatter plots of Cramér-Rao lower bounds (CRLB) against standard deviations (stdev) of the parameter estimates obtained with AMARES and the proposed GMRF approach. The results are based on 100 repetitions of simulated MRSI images containing the three metabolites Cho (red), Cr (blue) and NAA (green). Each dot corresponds to one of the 32×32 voxels. While the single voxel approach does not beat the CRLB, using spatial prior knowledge can reduce the estimation variance below the theoretical CRLB.

Figure 6: Bias, standard deviation (stdev) and root-mean-squared error (rmse) for amplitude estimates using the single voxel (sv, AMARES) and spatial (sp, GMRF) approaches when fitting two overlapping peaks: Choline (Cho) and Creatine (Cr). Example spectra are shown in Figure 7.

(a) Single voxel (AMARES)

(b) Spatial prior (GMRF)

Figure 7: Four example spectra from the simulated data analyzed in Fig. 6. The examples are taken from the corners of the simulated spectral image and contrast the behavior of single voxel and spatially regularized fit for overlapping peaks with varying amplitudes. The examples show that with sufficient amplitude both, the Cho and the Cr-component can be fitted reliably. If one of the peaks gets small, however, the single voxel approach runs into problems while the spatial prior produces stable fits using information from a local neighborhood.
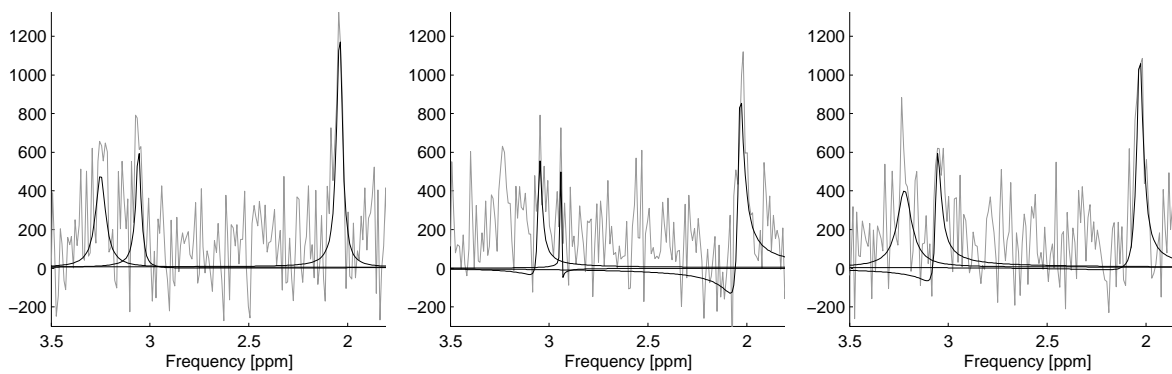
(a) Single voxel
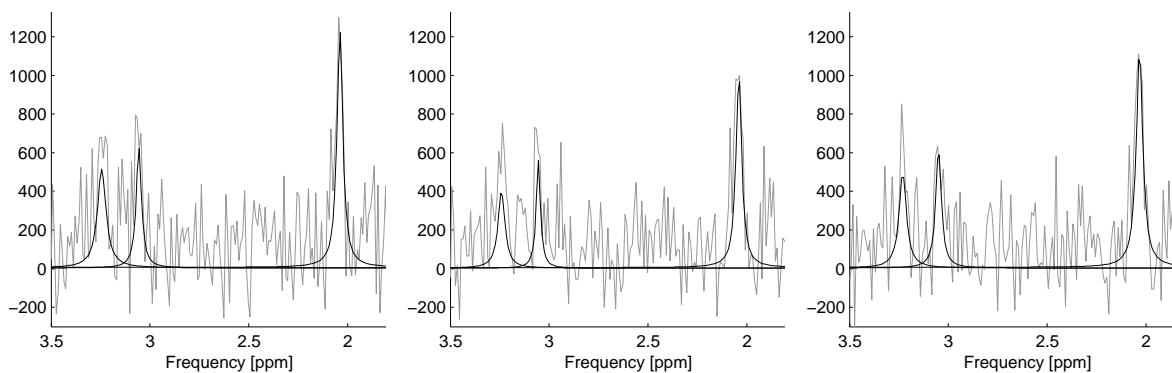
(b) Spatial prior

Figure 8: Spectra of three consecutive voxels from an MRSI brain scan at 3T with long echo time (Series 4.3, Tab. 2) and spectral line fits for Cho (3.2 ppm), Cr (3.0 ppm) and NAA (2.0 ppm). The top row shows results from single-voxel fitting whereas the bottom row shows the corresponding fits obtained with using spatial prior knowledge. The single voxel approach fails to fit the Cho and Cr peaks in the middle voxel whereas the GMRF uses information from the neighboring spectra to stabilize the fit. Furthermore, the GMRF fit shows a much more stable phase estimate than the single-voxel fit.
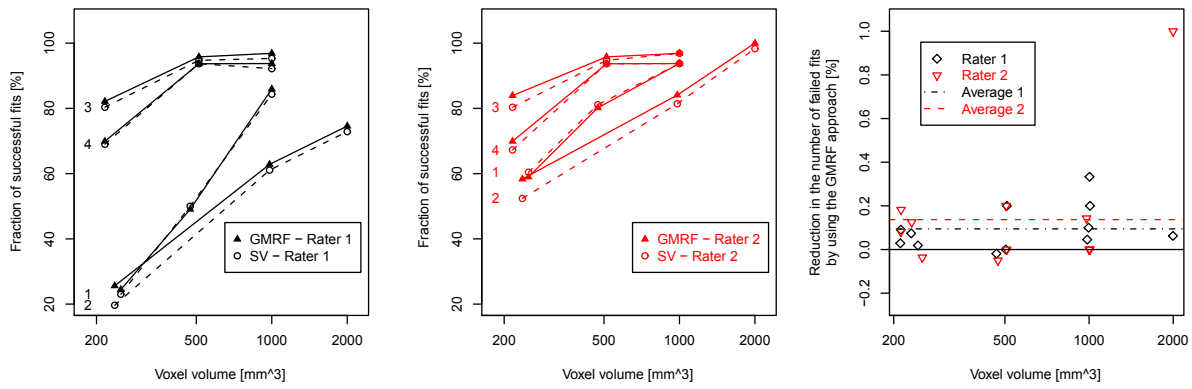
Figure 9: Evaluation of the spatially constrained and the unconstrained fitting routines on real data. The left and the central figure show the fraction of successful fits as a function of voxel volume for both raters, and on all resolution series from Tab. 2. The GMRF fitting (with $\lambda = 1$) performs better in nearly all tests. The right figure presents this benefit in values relative to the performance of the unconstrained SV results. On average the GMRF constraint reduced the fraction of failed fits by approx. 10%.
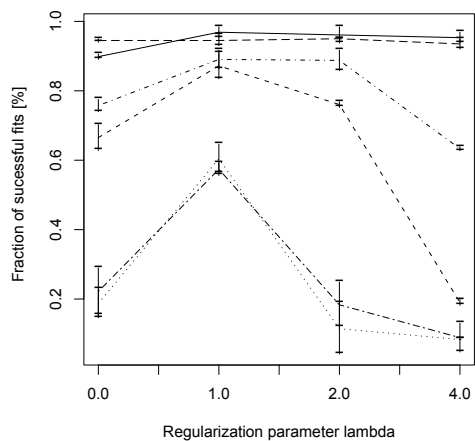
Figure 10: Effect of the regularization parameter λ on the fraction of successful fits. The six lines show the results for all data of resolution series 1 and 2 (Tab. 2) after double evaluation by rater 2. The notches indicate the extrema (*i.e.* the two obtained percentages) while the lines follow the average. A value of $\lambda = 1$ performs well for nearly all data sets, over all resolutions and corresponding signal-to-noise ratios.