

Accelerated MR Parameter Mapping with a union of local subspaces constraint

Sagar Mandava^{1,2}, Mahesh B. Keerthivasan^{1,2}, Zhitao Li^{1,2}, Diego R. Martin², Maria I. Altbach^{2,3} and Ali Bilgin^{1,2,3*}

¹ Department of Electrical and Computer Engineering, University of Arizona, Tucson, Arizona, USA

² Department of Medical Imaging, University of Arizona, Tucson, Arizona, USA

³ Department of Biomedical Engineering, University of Arizona, Tucson, Arizona, USA

Corresponding author:

Ali Bilgin, PhD

Department of Electrical and Computer Engineering,

University of Arizona, Tucson, AZ, USA

Email: bilgin@email.arizona.edu

Word count: 6590

Figure + Table count: 10

Institution:

University of Arizona,

Tucson, Arizona, USA

Running title: Parameter Mapping with a union of local subspaces constraint

Keywords: Multi-contrast; parameter mapping; clustering; sparsity constraint; union of subspaces constraint; image reconstruction

Abstract

Purpose: A new reconstruction method for multi-contrast imaging and parameter mapping based on a union of local subspaces constraint is presented.

Theory: Subspace constrained reconstructions use a pre-determined subspace to explicitly constrain the relaxation signals. The choice of subspace size (K) impacts the approximation error vs noise-amplification tradeoff associated with these methods. A different approach is used in the MOCCO framework to leverage the subspace model to enforce a softer penalty. Our proposed method, MOCCO-LS, augments the MOCCO model with a union of local subspaces (LS) approach. The union of local subspaces model is coupled with spatial support constraints and incorporated into the MOCCO framework to regularize the contrast signals in the scene.

Methods: The performance of the MOCCO-LS method was evaluated in vivo on T_1 and T_2 mapping of the human brain and with Monte-Carlo simulations and compared against MOCCO and the explicit subspace constrained models.

Results: The results demonstrate a clear improvement in the multi-contrast images and parameter maps. We sweep across the model order space (K) to compare the different reconstructions and demonstrate that the reconstructions have different preferential operating points. Experiments on T_2 mapping show that the proposed method yields substantial improvements in performance even when operating at very high acceleration rates.

Conclusions: The use of a union of local subspace constraints coupled with a sparsity promoting penalty leads to improved reconstruction quality of multi-contrast images and parameter maps.

Keywords: Multi-contrast; parameter mapping; clustering; sparsity constraint; union of subspaces constraint; image reconstruction

Accelerated MR Parameter Mapping with a union of local subspaces constraint

INTRODUCTION

Quantitative mapping of tissue parameters like T_1 and T_2 has emerged as a valuable tool for tissue characterization and pathological assessment (1-5). Classical parameter mapping experiments involve the acquisition of a series of contrast weighted images where each image is acquired under different acquisition conditions (like echo time (TE) for T_2 mapping or inversion time (TI) for T_1 mapping). These images are then fit pixel-wise to an underlying model to generate the parameter map. As the number of images needed for accurate parameter mapping increases so does the scan time making it challenging for practical clinical use. Accelerated imaging techniques using a combination of sparse sampling and constrained reconstructions have been proposed to lower the scan time in parameter mapping experiments. A variety of reconstruction schemes based on parallel imaging (6,7), sparsity constraints (8-10), low-rank constraints (11,12), manifold learning (13), and combinations of these methods (14-17) have been proposed. Furthermore, approaches that seek to estimate parameter maps directly from k-space data (bypassing the need to reconstruct contrast images for subsequent parameter fitting) also exist (18-21).

One variant of the low-rank constraint uses a pre-estimated subspace basis to constrain the relaxation signals in the contrast dimension and include methods like k-t PCA (11) and its regularized variants (14-17). These methods enforce an explicit low rank constraint by restricting the M dimensional contrast signals to a K dimensional subspace, typically $K \ll M$, with K governing the constraining power. Augmenting these methods with additional constraints, such as spatial sparsity (14,15) locally low-rank (16), or Hankel low-rank (17), improves the constraining power further at the cost of some spatial smoothing. The subspace basis is either estimated from experimental training data (15) or from synthetic training signals generated from a suitable signal model (14). The former variant has the advantage of yielding a basis tailored for the specific imaging scene but has the drawback of measurement noise corrupting the estimate of the subspace basis. In addition, acquisitions like the single shot T_1 mapping approach exist (21,22) where it is not feasible to estimate the basis from the experimental data motivating the use of signal model driven approaches for subspace basis design. Extensions of k-t PCA which create different sets of basis functions for different regions (compartments) in the anatomy have

been proposed under the name of compartment-based k-t PCA and have been applied for improved perfusion (23) and phase contrast imaging (24).

A different approach to leverage subspace driven models has been proposed with the model consistency condition (MOCCO) formulation for dynamic imaging applications (25) with preliminary results for parametric imaging applications reported in (26). MOCCO uses the subspace basis to create a penalty function that measures the “closeness” of the temporal signals to the subspace model and can be used to penalize signals that deviate from the model. It was also shown to be robust to model mismatches especially when used with an ℓ_1 penalty which enforces a form of spatio-temporal sparsity to regularize the reconstruction problem. However, it has been argued that the constraining power of explicit subspace models is higher than that of MOCCO especially in applications with relatively simple temporal signal evolutions (25). Velikina et al. present a compartment-driven extension of MOCCO for dynamic imaging that partitions the imaging scene and builds unique MOCCO penalty functions for each partition (27). This data-driven approach to spatially adapt the penalty functions results in improved performance over the regular MOCCO approach.

In this work, we present an extension of the MOCCO framework to work with a union of local subspaces constraint instead of a single subspace. Local subspaces (LS) are designed to cover smaller regions of the model signal space in contrast to a single ‘global’ subspace designed to cover the entire signal range of interest. We expect contrast signals with similar relaxation behavior to be better represented in an appropriate local subspace than with a ‘global’ subspace. The proposed model, MOCCO-LS, relies on clustering the model signal space to build a union of local subspace bases and a first pass reconstruction to fit the imaging scene to the local subspaces. We demonstrate the performance of MOCCO-LS and compare it to other relevant approaches in both T_2 and T_1 mapping applications.

THEORY

Data model: A parameter mapping experiment involves the acquisition of a series of images $\{X_j\}_{j=1}^M$ with appropriate contrast weighting. A general k-space data model for a parametric imaging experiment is given by:

$$Y = E(X), \quad [1]$$

where $Y \in \mathbb{C}^{MN_c N_k \times 1}$ represents the measured multi-coil (N_c coils), multi-contrast k-space (N_k k-space samples at each contrast frame) data. E denotes the Fourier and coil sensitivity encoding operator for the M frames and $X \in \mathbb{C}^{M \times N}$ denotes the multi-contrast images acquired at M different contrasts and sampled on N spatial locations.

At moderate acceleration rates, it is possible to solve the problem in Eqn. 1 with a direct inversion while removing most of the aliasing artifacts and with minimal noise amplification using parallel imaging techniques (6,7). However, most practical applications of parametric imaging require much higher acceleration rates which makes direct inversion impractical. In these situations, image reconstruction is typically formulated as:

$$\hat{X} = \arg \min_X \|Y - E(X)\|_2^2 + R(X), \quad [2]$$

where the first term promotes data fidelity and the second term enforces constraints on X (like sparsity or smoothness or low rank) and is called the regularizer. After reconstruction of the contrast weighted images the underlying parameter maps can be extracted by a pixel-wise non-linear least squares fit to a suitable signal model.

Model based subspace constraint: A subspace constraint for the multi-contrast images takes the following form:

$$X \approx \phi_K \phi_K' X = \Phi_K X, \quad [3]$$

where ϕ_K is a K -dimensional subspace basis. The operator $\phi_K' \in \mathbb{C}^{K \times M}$ maps the M different contrast weighted images to K subspace images and $\phi_K \in \mathbb{C}^{M \times K}$ maps these subspace images back to the M -dimensional contrast signal space. The projection and re-projection operations can be combined to yield a single operator denoted by $\Phi_K \in \mathbb{C}^{M \times M}$ which restricts the contrast signals in the scene to a K -dimensional subspace. A subspace constraint can be added to the reconstruction problem in Eqn. 2 as follows:

$$\hat{X} = \arg \min_X \|Y - E(\Phi_K X)\|_2^2 + R(X), \quad [4]$$

When no additional regularization is used, Eqn. 4 is a pure subspace-constrained reconstruction that yields low-rank solutions when $K < M$ and is referred to as k-t PCA reconstruction (11). This constraint can be used to augment standard parallel imaging strategies like SENSE and can be used for parameter mapping at moderate acceleration factors but can lead to significant noise amplification and residual aliasing artifacts at high acceleration rates. Zhao et

al. proposed to augment the k-t PCA reconstruction with a joint sparsity constraint enforced with the mixed ℓ_2/ℓ_1 norm, $\|\cdot\|_{2,1}$, which promotes correlated edges across the contrast images (15). The mixed ℓ_2/ℓ_1 norm is defined as follows: $\|A\|_{2,1} = \sum_{n=1}^N \|A^{(n)}\|_2$ where $A^{(n)}$ represents the n^{th} row of A . They demonstrate that the noise amplification seen in pure k-t PCA reconstructions can be significantly mitigated by adding the joint sparse penalty to the k-t PCA reconstruction. The reconstruction problem, referred to as L12 in subsequent usage, is formulated as follows

$$\hat{X} = \arg \min_X \|Y - E(\Phi_K X)\|_2^2 + \lambda \|D_x x\|_{2,1} + \lambda \|D_y x\|_{2,1}, \quad [5]$$

where D_x, D_y are block horizontal and vertical finite differences operators that extract the spatial gradients on all the contrast images (vectorized notation for X denoted as $x \in \mathbb{C}^{MN \times 1}$). This problem can be solved in the alternating direction method of multipliers (ADMM) framework with variable splitting (28). Zhao et al. have a comprehensive discussion on implementing a closely related algorithm (which works with ϕ_K instead of Φ_K) which can be adapted to solve Eqn. 5 (15).

The subspace basis ϕ_K can either be estimated from acquired calibration data (15) or from a suitable signal model (14) using singular value decomposition (SVD). Estimating the basis from acquired calibration data would in theory be ideal as the basis is tailored for the underlying scene. But in reality, the calibration data is low-resolution (and reused to reconstruct the high-resolution data) and corrupted with noise. There are also acquisition schemes like single shot T_1 scans for instance where it is not practical to acquire any calibration data (21,22). This motivates the use of model driven subspaces where training curves from a suitable signal model can be used to generate the subspace basis. Simple fitting strategies based on the mono-exponential model have been used for this purpose (14,26). More sophisticated models that can account for deviations from ideal signal behavior have also been proposed (16,22,29).

For a chosen signal model, we can generate an ensemble of relaxation signals by evaluating the model with a suitable discretization of the model parameter space. For example, in 2D T_2 mapping with fast spin echo sequences (FSE), the slice resolved extended phase graph (SEPG) model is known to model the relaxation curves accurately (30). This model uses acquisition parameters like echo spacing, echo train length, excite and refocus slice profiles, along with model parameters like T_2 and B_1 (T_1 is also a parameter but it is usually fixed to a large value

as the signal is relatively insensitive to variations in T_1). In the absence of any other information, one usually samples the T_2 (M_1 samples) and B_1 (M_2 samples) model parameters uniformly within the range of expected values to generate an ensemble of relaxation curves and collect them into a dictionary $D \in \mathbb{C}^{M \times M_1 M_2}$. When we have information on the relative density of any of the model parameters, we can draw samples weighted by this density information to fill the dictionary. For a given dictionary, the full subspace basis ($\phi \in \mathbb{C}^{M \times M}$) can be computed by the SVD of either D or the correlation matrix DD^T (14). The columns of ϕ are usually ordered to follow decaying singular values and the dominant K columns are collected into the matrix $\phi_K \in \mathbb{C}^{M \times K}$. An illustration of this process is shown in Figure 1(a).

Model consistency (MOCCO) constraint: Velikina et al. propose a novel variant of the subspace constraint where instead of hard constraining the contrast signals to a subspace, they propose penalizing signals that deviate from the subspace (25). They introduce a new operator $\Psi_K = \Phi_K - I^{(M)}$ where $\Psi_K \in \mathbb{C}^{M \times M}$ and $I^{(M)} \in \mathbb{C}^{M \times M}$ is the identity matrix. This operator measures the similarity between the actual contrast signals and the model constrained signals and allows the formulation of a penalty term. With this operator, the reconstruction problem is posed as:

$$\hat{X} = \arg \min_X \|Y - E(X)\|_2^2 + \lambda \|\Psi_K X\|_1, \quad [6]$$

where the penalty is enforced in the form of the ℓ_1 norm and allows some contrast signals to deviate from the model. Velikina et al. solve this problem using an iteratively reweighted least squares algorithm (25). We solve it using ADMM with variable splitting where the ℓ_1 norm is enforced via soft-thresholding.

Union of local subspaces constraint: Model based subspace design involves the generation of a set of training curves that covers a certain range in the model parameter space which are then used to create a single subspace basis. In contrast, we propose to generate several subspace bases from ‘local’ regions in the model parameter space. The motivation for this is that signal curves representing different kinds of tissues (such as gray matter, white matter, and CSF) are quite different and it is possible to represent them more efficiently using subspaces derived from a smaller subset of the training curves that are targeted to represent them. The proposed method, MOCCO-LS, involves two operations to be performed prior to reconstruction: 1) Clustering the training curves to create local subspace basis for each cluster, 2) Fit the current

estimate of the scene pixel-wise to these local subspaces and choose the best cluster/subspace (in the ℓ_2 sense). This results in a collection of local subspaces, (ϕ_K^j) where $j = 1 \dots L$, and a spatial support map for the L clusters. This process is presented as a flow chart in Figure 1(b).

The clustering of training curves can be augmented by using additional information on the relative distribution of the model parameters. In practice, we do not have access to this distribution information exactly, but an approximation can be used. Tamir et al. (16) have used an existing scan to approximate the density of T_2 s of interest in the scene and use this information during the training curve design; we follow the same approach in this work. We demonstrate this entire process on SEPG training curves in Figure 2 which are parameterized by the T_2 and B_1 indices and displayed in the form of an image for visualization (each pixel represents a contrast training curve). The T_2 histogram shows the relative discretization of the T_2 model space. In contrast, the B_1 space is discretized uniformly due to the lack of any prior information. These curves are clustered using k-means with four clusters and color coded for visualization. Local subspace bases for these four clusters are generated using the process in Figure 1(a). An initial estimation of the imaging scene from MOCCO is generated and the contrast signal at each spatial location is constrained using all the local subspaces and the one yielding the smallest ℓ_2 error is assigned to it. In this four-cluster example, the ℓ_2 distance based assignment is shown with the same color coding scheme. Binary spatial support maps can then be generated for each cluster.

Given a set of local subspace bases, we can create MOCCO operators for each subspace by subtracting the identity matrix from them and create the corresponding set of operators (Ψ_K^j) where $j = 1 \dots L$. The MOCCO-LS operator is now defined as:

$$\Omega(X) = \sum_{j=1}^L \Psi_K^j X S_j, \quad [7]$$

where $S_j \in \mathbb{C}^{N \times N}$ is a diagonal matrix with ones on pixel locations that are assigned to the j^{th} cluster and zeros elsewhere. $X S_j$ extracts the contrast signals that need to be constrained using the j^{th} subspace (Ψ_K^j) . Note that MOCCO is a special case of MOCCO-LS where $L = 1$ (and S_1 is the identity matrix). The reconstruction problem is now given as:

$$\hat{X} = \arg \min_X \|Y - E(X)\|_2^2 + \lambda \|\Omega(X)\|_1, \quad [8]$$

and can be solved with the ADMM framework.

METHODS

In Vivo Data: Data for in vivo experiments were acquired on a 3T MRI scanner (Skyra, Siemens Healthcare, Malvern, PA, USA). All in vivo experiments were performed in compliance with the human research subject guidelines and the institutional review board. The proposed methods were tested on in vivo brain datasets for T_2 mapping and for T_1 mapping. For T_2 mapping, a radial FSE pulse sequence (31) with echo-spacing (ESP) = 8.78ms, echo-train-length (ETL) = 16, TR=4000ms, slice thickness=5mm, field-of-view (FOV) = 220mm was used. Data were acquired with a 16-channel receiver head coil, 256 views/TE, and 256 readout points per view. This dataset is used to form the reference images and parameter maps used for the study. The same dataset was retrospectively undersampled to 16 views/TE (R=16) and 8 views/TE (R=32) to simulate acceleration.

For T_1 mapping, data were acquired with an inversion-recovery (IR) steady-state free-precession (SSFP) pulse sequence (22) with TR=4.38ms, FOV=220mm, TI to the first excitation=70ms, slice thickness=5mm, flip angle=20°, using a 16-channel receiver head coil. Following the inversion pulse 512 radial views are acquired with the tiny golden angle ordering (32). These views are grouped into 32 TIs (16 views/TI) leading to a temporal footprint of about 60ms which is consistent with existing work based on single shot IR-FLASH (21). To generate a reference, this sequence was repeated 28 times with a delay time of 15 sec between the IR modules and with the tiny golden angle view ordering continuing across the inversion modules to generate fully sampled images (created from 448 radial views) at each TI.

Image reconstruction and evaluation:

The SEPG model driven subspace design for the T_2 mapping experiment used a prior distribution on the T_2 space (between 10 and 350ms) while the B_1 space was discretized (min:step:max) as 0.5:0.01:1.5 and a fixed T_1 value set to infinity in the model. The T_1 mapping experiment used a Bloch simulation driven model that is sensitive to T_1 , T_2 and B_0 (off-resonance). The T_1 space used a prior distribution to discretize it non-uniformly between 100 and 5000ms and the training curves were generated with the following T_2 (10:10:300) and B_0 (-400:20:400)

values. The subspace and MOCCO operators were created from the subspace bases generated from these training curves for the T_1 and T_2 mapping experiments. For MOCCO-LS, the training curves were clustered using the k-means algorithm and the first pass reconstruction necessary to generate the spatial support maps for the clusters was done with the MOCCO reconstruction.

The coil sensitivity maps were estimated from the central region of the k-space datasets using the ESPIRiT technique (33). ADMM with variable splitting was used to solve the regularized reconstructions while k-t PCA was implemented using the CG-SENSE (34) framework. The non-uniform fast Fourier transforms in the Gadgetron library (35) were used for the forward and inverse Fourier transforms for the radial data. Consistent with existing work on stopping criterion, we let the relative image residual fall below a threshold $\{5e^{-4}\}$ or a preset number of iterations $\{50\}$ are exceeded (15,25). The actual number of iterations used varies depending on the acceleration rates as well as the reconstruction model but all the models converge well before the preset maximum iterations. We note that the k-t PCA and L12 reconstructions converge faster (15-20 iterations) than the MOCCO and MOCCO-LS reconstructions (30-35 iterations) in general due to the stronger de-noising that is characteristic of subspace constrained reconstructions. Image reconstruction was performed on a workstation with a 3.6 GHz Intel Xeon processor E5-1620, 64GB RAM and an NVIDIA GeForce GTX 780 GPU. The reconstruction time was within 20 minutes for all the reconstructions except MOCCO-LS. MOCCO-LS requires an initial estimate for the generation of the spatial support constraints. In this work, we used MOCCO to obtain these initial estimates. There is also additional overhead associated with the use of multiple subspaces. Therefore, MOCCO-LS required roughly twice the reconstruction time as the other techniques in these experiments.

The performances of the proposed methods were evaluated quantitatively using the following metrics: (a) Normalized error: $|x_n - \hat{x}_n|/x_n$, where x_n is the reference parameter and \hat{x}_n is the estimated parameter at the n^{th} voxel, (b) Overall error: $\|x - \hat{x}\|_2/\|x\|_2$, where x and \hat{x} contain the true and estimated parameter maps for all the image voxels, and (c) HFEN(x, \hat{x}) = $\left(\frac{\|\text{LoG}(x) - \text{LoG}(\hat{x})\|_F^2}{\|\text{LoG}(x)\|_F^2}\right)^{\frac{1}{2}}$ where LoG(.) is a rotationally symmetric Laplacian of Gaussian filter with a filter kernel size of 15 pixels and standard deviation of 1.5 pixels (36). Spatial maps of normalized error allow visualization of the error at each pixel. Overall error is a form of normalized mean

squared error while the high frequency error norm (HFEN) gives a measure of the preservation of high frequency content. As these error metrics capture global performance which can be dominated by errors in regions of limited clinical interest, we also present the % error between the true and estimated mean parameter values in ROIs for both T_1 and T_2 mapping.

The reconstructions were performed at different model orders and cluster sizes (for MOCCO-LS) and optimized for the overall error metric to compare the performance across these dimensions. MOCCO-LS experiments were also conducted where individual clusters were allowed to have different model orders to study the impact of model order on reconstruction performance. Overall error evaluated on a per-cluster level was used to first identify the K value for each cluster resulting in the lowest error and a new MOCCO-LS reconstruction was executed with these optimized model orders for each cluster. Note that this is not a practical strategy in general and is used only for evaluation. For all the models, after reconstruction of the multi-contrast images, the variable projection (VARPRO) algorithm was used to perform voxel wise non-linear fitting to the signal model to yield the parameter maps (15,37).

Monte-Carlo study:

In order to study the accuracy and precision attained by different reconstruction techniques, Monte-Carlo (MC) simulations were carried out for a T_2 mapping experiment. The simulations were setup according to the framework described in Velikina et al. (10). A 2D radial FSE acquisition with the same acquisition parameters used in the in-vivo study was simulated on a single slice from the Brainweb database (38) with radial sampling at $R = 16$. MC experiments were carried out using $J = 100$ i.i.d. Gaussian noise realizations at SNR = 20 for the fully sampled case where SNR is defined as the ratio of the mean signal intensity of the contrast weighted images to the noise standard deviation. Noise was added to the k-space data and reconstructions were obtained for each MC realization. An initial reconstruction at different values of the regularization parameter was carried out for all methods (except for k-t PCA). For each method, the regularization parameter that minimizes the overall error in T_2 was identified and used in all MC reconstructions. For each method, the pixel-wise mean of the T_2 maps were generated using $\langle T_2 \rangle = \frac{1}{J} \sum_{j=1}^J T_2^{(j)}$ where $T_2^{(j)}, j = 1, \dots, J$ denotes the T_2 values obtained in each MC realization.

A second MC experiment was also set up where i.i.d. Gaussian noise at the same SNR level was added to fully sampled data and parameter maps (denoted by $T_{2,Full}^{(j)}, j = 1, \dots, J$) were computed for each noise realization. These datasets were mapped back to the image space using the SENSE approach (34) to yield coil combined data followed by a pixel wise fit to yield the maps. A pixel-wise mean $\langle T_{2,Full} \rangle = \frac{1}{J} \sum_{j=1}^J T_{2,Full}^{(j)}$ was then calculated. Pixel-wise noise variance was computed as follows: $\sigma_R^2 = \frac{1}{(J-1)} \sum_{j=1}^J (T_{2,Full}^{(j)} - \langle T_{2,Full} \rangle)^2$ and σ_F^2 is defined for the fully sampled data in a similar manner. Noise amplification factor, referred to as g-factor, was computed using: $g_{T_2} = \frac{1}{\sqrt{R}} \frac{\sigma_R}{\sigma_F}$. Normalized error between $\langle T_2 \rangle$ of each reconstruction method and a T_2 reference (fully sampled noiseless data) were calculated to quantify reconstruction accuracy.

RESULTS:

Figure 3 shows the results of the MC study. We report results for k-t PCA, L12, and MOCCO methods at the operating point (in K) where they yield the lowest overall error. The k-t PCA reconstruction, which does not require a regularization parameter, can be considered as a baseline against which the other models can be compared. The L12 reconstruction yields very low noise amplification (see g-factor map) but at the cost of reduced accuracy which is captured in both the error map and the overall error. The g-factor map for MOCCO shows lower noise amplification and error than k-t PCA demonstrating its advantages over subspace constrained reconstructions. MOCCO-LS results at two different cluster sizes and at different model orders are presented. Fixing $K = 2$, we see that increasing L from 4 to 8 clusters allows a reduction in both noise amplification as well as normalized error. Note that despite not being optimal in the overall error, MOCCO-LS ($L = 8, K = 2$) has good performance in the gray and white matter regions (see Table 1) and low g-factor. MOCCO-LS ($L = 8, K = 2$ or 3) yields the lowest overall error but the noise amplification is higher than the MOCCO-LS ($L = 8, K = 2$) case. The g-factor and error maps for the K and L values that are not shown here but reported in Table 1 are visually similar to the MOCCO-LS ($L = 8, K = 2$ or 3) case.

The overall error as well as error evaluated separately for gray matter, white matter, and CSF for all the methods studied in the MC experiments are reported in Table 1. These results

demonstrate the benefits of using the union of local subspaces model. The advantages of using different model orders for different clusters is also evident from these results.

Figure 4 shows the reconstructed T_2 maps of the brain from data at $R=16$ for the k-t PCA and MOCCO reconstructions at $K=2,3$, and 4. The normalized error maps with the overall error (on the bottom left) are also presented. The error performance of k-t PCA at $K=2$ is dominated by the model error, i.e. 2 coefficients are not sufficient to represent the contrast signal ensemble. The performance drastically improves with the use of $K=3$ and improves a little further with the use of $K=4$. The use of larger values of K will improve the model approximation but comes at the cost of increased noise variance. MOCCO reconstructions exhibit their best performance at $K=2$ and the reconstructions deteriorate in quality both visually and as measured by the overall error metric with increasing K . Qualitatively, the MOCCO reconstructions appear sharper than the k-t PCA reconstructions while quantitatively k-t PCA yields higher accuracy at this acceleration rate.

Figure 5 shows the reconstructions and the corresponding normalized error maps for the T_2 mapping data at $R=16$ and $R=32$. At $R=16$, the performances of all methods are comparable. The use of L12 constraints improves the performance over k-t PCA slightly but comes at the cost of some spatial smoothing (which is visually evident but not well-captured by the HFEN metric). The MOCCO reconstruction yields slightly noisier maps but retains the sharpness qualitatively. MOCCO-LS yields a map with the lowest overall error and retains the image sharpness. At $R=32$, k-t PCA yields a result with a significant amount of residual aliasing. The use of L12 constraints improves performance slightly but the resulting map is heavily smoothed out. MOCCO yields a result with an overall error comparable to the L12 reconstruction but with better preserved high frequency content. MOCCO-LS yields a reconstruction with significantly improved quality that is clearly reflected in the normalized error map, overall error, as well as the visual quality. The spatial support maps used in the MOCCO-LS reconstruction are also shown in this figure. ROI level improvements are provided in Supporting Information Table S1.

Figure 6 shows the reconstructions and the corresponding normalized error maps for T_1 mapping. The preferred operating point for both MOCCO and MOCCO-LS is at $K=2$. Again, the L12 reconstruction improves the reconstruction quality over k-t PCA with no noticeable loss in spatial resolution. The MOCCO reconstruction shows slightly higher error as seen in the

normalized error map. MOCCO-LS ($K=2$) yields a reconstruction that agrees quite well with the reference map and with an overall error comparable to the L12 reconstruction. The overall error metric plotted on the bottom left corner of the normalized error maps shows performance improvement trends similar to what was seen with the T_2 mapping case at $R=16$. MOCCO-LS ($K=3$) and MOCCO-LS ($K=2$ or 3) reconstructions are also presented.

The % T_1 errors for the ROIs indicated on the reference image are tabulated in Table 2 with MOCCO-LS ($K=2$ or 3) yielding results that are better than those obtained from the L12 reconstruction in two of the ROIs and slightly worse in the third ROI.

Figure 7(a) presents the overall error and the HFEN metrics for the T_2 mapping data at $R=16$ and $R=32$ as a function of the model order (K). At $R=16$, both metrics demonstrate the preferred operating points for the two distinct reconstruction types: $K = 4$ for k-t PCA and L12 and $K = 2$ for MOCCO and MOCCO-LS. L12 exhibits a clear improvement over k-t PCA as expected. MOCCO reconstruction comes close to reaching the performance of the L12 model as captured by the proposed metrics but never exceeds it. MOCCO-LS performs better than all other algorithms at $K = 2$ at $R=16$ but its performance is slightly worse than L12 at $K=3$ and 4 . At $R=32$, MOCCO-LS outperforms L12 across all model orders but its optimal performance is still at $K=2$.

Figure 7(b) shows the performance of MOCCO-LS as a function of the number of clusters (L) on the T_2 mapping data. We include the case of $L = 1$ i.e. the MOCCO reconstruction as well to illustrate the performance change. At $R=16$, the performance increases with the number of clusters with the dominant gains in performance at $L = 8$ after which the performance saturates. At $R=32$, the overall error actually starts to worsen when the number of clusters increases beyond 8 . This suggests that there is a preferred operating point in the cluster dimension. It can also be observed that increasing the number of clusters beyond 8 does not provide a substantial improvement in the HFEN metric.

Figure 7(c) presents the overall error and the HFEN metrics for the T_1 mapping data as a function of the model order (K). The performance trends of k-t PCA and L12 are similar to what was seen in the T_2 mapping case with the preferred operating point being at $K = 4$ although the performance of L12 starts to saturate at $K = 3$. MOCCO and MOCCO-LS perform best when $K =$

2 and their performances degrade with increasing K . Overall, the performance of MOCCO-LS is comparable to the performance of the L12 reconstruction on this dataset.

Figure 8(a) presents a qualitative comparison of representative TE images for the reconstructions at $R=16$. The k-t PCA and L12 reconstructions both exhibit spatial blurring across the TEs while MOCCO and MOCCO-LS yield visually sharper images. Note that the artifact level in the k-t PCA reconstructions is not uniform across the TEs. The TE2 image exhibits a substantially larger amount of the aliasing artifact than the TE8 image. This is a consequence of the decaying nature of the subspace basis functions which distribute the noise and artifact non-uniformly across the TEs with the early TEs suffering the most (see (16) for details). The use of L12 constraints can effectively suppress this noise behavior, as seen in the improved quality on the TE2 image, but comes at the cost of some loss in resolution. MOCCO and MOCCO-LS yield images that retain a greater amount of high frequency information. Figure 8(b) shows some representative TI images for the T_1 mapping data. Consistent with the T_2 mapping case, the early TI images of k-t PCA suffer from noise amplification that reduces as one moves to longer TIs. The L12 reconstruction clearly reduces this noise effect as seen in the significant amount of de-noising it creates in all the TI images. Although there is some loss in spatial resolution due to the L12 constraints on the TI frames, we do not observe this effect on the T_1 maps (likely due to the non-linear fitting algorithm). While the MOCCO reconstruction is sharper than the L12 reconstruction it retains a significant amount of noise. The MOCCO-LS reconstructions agree very well with the reference images at all TIs and the improved performance of MOCCO-LS is best seen on the TI1 image.

DISCUSSION In this paper, we present a novel method to augment subspace driven reconstructions using a union of local subspaces model. Local subspaces cover narrower regions of the model signal space and can be used to represent the relaxation signals more effectively than a single ‘global’ subspace basis. Intuitively, the use of a single subspace will lead to some regions in the signal space being represented well at the cost of others. Local subspace bases end up representing the signals in a ‘local’ region of the signal space well at the cost of performing worse in other regions. By coupling local subspace bases with spatial support constraints, we can

leverage higher performance than what a single subspace basis can achieve. The effectiveness of incorporating local subspace constraints on both T_1 and T_2 mapping applications has been demonstrated.

MOCCO-LS adds to the existing body of model-driven reconstructions (8,10,13,14,16-22,26,29) for multi-contrast imaging and parameter mapping. Model-driven reconstructions seek to enforce useful prior information drawn from a signal model to improve image reconstruction and allow parameter mapping from highly accelerated scans. Direct parameter mapping approaches (18-21) directly estimate the parameter maps from highly undersampled k-space measurements. Two-stage reconstruction models, such as the ones discussed in this work, first reconstruct the contrast weighted images and then fit them to a suitable signal model. Model-driven approaches in the two-stage setting use relaxation models to build different forms of constraints. Doneva et al. (8) use training curves generated from the signal model to learn a dictionary and the reconstructed signals are constrained to have a sparse representation in this dictionary. Other approaches use training curves to learn a low-dimensional subspace which can accurately represent the relaxation signals (14,22,26,29). During reconstruction, the reconstructed signals are constrained to this subspace which acts as a low-rank constraint. Similarly, Zhou et al. (13) use the relaxation model to generate training images to learn a low-dimensional manifold instead of a low-dimensional subspace. The reconstructed signals are then constrained to be close to this manifold. In (17), Peng et al. express T_2 decay as a linear combination of exponential functions. Using the fact that a linear combination of L exponentials is linearly predictable to the L -th order, the low-rankness of a Hankel matrix representing the signal decay is set as a constraint during reconstruction. Following the model-driven theme, the proposed approach, MOCCO-LS, uses local subspaces which can represent the signal more accurately in a local region in comparison to a global subspace approaches which aim to represent the signal over the entire range of interest.

The proposed method requires a first pass reconstruction of the imaging scene and in this work, we used the MOCCO reconstructions for this purpose. This reconstruction is necessary to generate spatial support constraints for the local subspaces. An interesting alternative could be to use view-sharing driven reconstructions (39) as the first pass reconstruction. These

reconstructions have the advantage of being fast but can suffer from artifacts due to sharing of k-space data across different contrasts. The use of these fast reconstruction techniques during initialization would reduce the reconstruction time of MOCCO-LS by roughly a half, bringing it closer to the reconstruction time of MOCCO.

The clustering of training curves was done using the k-means algorithm. This algorithm requires the number of clusters as an input and can be sensitive to initialization but this sensitivity can be mitigated using restarts. As our results suggest, this clustering is sufficient to yield improved results. We have observed that cluster assignment errors increase as the number of clusters increases however the impact of such misclassification on the overall error is small. It is likely that more sophisticated algorithms based on subspace clustering (40) can yield even better performance but their use is currently limited by their significantly higher computational complexity.

MOCCO-LS shares similarities with the spatially adaptive version of MOCCO which was proposed for dynamic imaging applications (27). The method proposed in (27) relies on direct clustering of the first-pass reconstruction to obtain the cluster assignments as well as the subspaces for each cluster. Therefore, it can be stated that MOCCO-LS is a model-driven approach whereas the method proposed in (27) is data-driven. This fundamental difference points to advantages and disadvantages of each approach: Data-driven bases can be contaminated by noise but are otherwise tailored to the signal in the scene. MOCCO-LS uses model-driven subspaces which are not contaminated by noise but may be more susceptible to model mismatch.

All the reconstructions discussed in this work with the exception of k-t PCA require the selection of the regularization parameter. The optimal selection of this parameter is still an open problem and an area of active study by the research community (41). One heuristic that can be used in practice is to refine this parameter on some reference dataset and then reuse it for acquisitions that follow the same imaging protocol. In this work, we evaluated the reconstructions over a wide range of the regularization parameter and selected the parameters that yield the least overall error. This is the reason why the L12 reconstructions, particularly in the T_2 mapping case, appear smoothed out. The MOCCO and MOCCO-LS reconstructions rely on

spatio-temporal sparsity controlled by a regularization parameter. The in-vivo results presented in this work were obtained using regularization parameters with the least overall error in parameter maps. However, this choice of the regularization parameter does not necessarily yield optimum contrast-weighted images. It is noted that the non-linear fitting step used to obtain parameter maps from contrast-weighted images is robust to errors that are not consistent across contrast-weighted images and, thus, artifacts that may be present in contrast-weighted images are often not visible in the resulting parameter maps.

Although the goal of parameter mapping is to enable quantitative evaluation, we have asked two board-certified neuroradiologists to qualitatively evaluate the T_1 and T_2 -weighted images obtained from the proposed method using different cluster dimensions (shown in Supporting Information Figure S1). The consensus opinion of the neuroradiologists was that the images provided by the proposed method were a little noisier than the fully sampled reference but that the differences were subtle. Interestingly, when asked to pick one out of the three images corresponding to different cluster dimensions, the neuroradiologists picked the image with the lowest overall error for both the T_1 and T_2 -weighted cases, although they did not have access to the error values.

We have also evaluated the reconstructions at different model orders to demonstrate that different reconstruction schemes have different preferential operating points. For the MOCCO-LS reconstruction, we also studied the performance as a function of the number of clusters and demonstrate that there exists a preferential operating point even on the cluster size. In general, the optimal number of clusters depends not only on the imaging scene but also on the noise level of the imaging experiment as well as the acceleration rate. We observe that the dominant gains in performance are achieved at small cluster sizes ($L = 4$) and performance tends to saturate by $L = 8$. In some cases, performance even worsens when L increases beyond 8 suggesting that simply increasing the number of clusters may not be desirable. One source of error which increases with increasing number of clusters is the cluster assignment error. An increase in cluster assignment error can lead to an increase in overall error. It should be noted that changes in the noise level and acceleration rate can lead to increased errors in the first pass reconstruction which in turn lead to increased cluster assignment errors since the first pass

reconstruction is used to cluster the imaging scene. However, the use of the ℓ_1 norm in the reconstruction formulation provides a level of robustness to these assignment errors and small changes in the number of clusters do not lead to drastic changes in the overall error. While determination of the optimal number of clusters for each image seems non-trivial, we note that our results demonstrate that a small number of clusters can be sufficient to elicit improved performance

We also compared the impact of allowing different clusters to use different model orders in MC experiments as well as the in-vivo T_1 dataset and demonstrate that additional performance gains that can be leveraged. However, selecting the optimal size of the subspace is a non-trivial problem and is further complicated in a union of subspaces model. Velikina et al. have recently proposed a model where the Bayesian Information Criterion (BIC) is used to select the model order for the MOCCO operator at each spatial pixel (42). The MOCCO reconstruction is coupled with a separate update step which seeks to select the model order to use at each spatial pixel location using the BIC criterion and solve an alternating updated reconstruction problem. Their work uses a fixed set of basis functions and the goal is to adaptively select the model order at each spatial location. In contrast, our work uses a collection of local subspaces and a decision needs to be made on which local subspace is most suitable for the specific spatial pixel. We have started exploring the combination of such alternately updated reconstructions to further improve MOCCO-LS reconstructions.

CONCLUSION In this work, we presented a new reconstruction model for accelerated MR parameter mapping. The proposed scheme augments model driven subspaces with a union of local subspaces and spatial support constraints. We incorporated this into the form of a penalty function which was motivated by MOCCO and solved the reconstruction problem with the ADMM algorithm. The performance of the proposed scheme was demonstrated on both T_1 and T_2 mapping applications and improvements over several related methods have been demonstrated both qualitatively and quantitatively.

ACKNOWLEDGEMENT This work was supported in part by Technology and Research Initiative Fund (TRIF) – Improving Health and Arizona Biomedical Research Commission (ABRC) grant ADHS14-082996.

REFERENCES

1. Baudrexel S, Nürnberger L, Rüb U, Seifried C, Klein JC, Deller T, Steinmetz H, Deichmann R, Hilker R. Quantitative mapping of T_1 and T_2^* discloses nigral and brainstem pathology in early Parkinson's disease. *NeuroImage* 2010;51:512–520.
2. Vymazal J, Righini A, Brooks RA, Canesi M, Mariani C, Leonardi M, Pezzoli G. T_1 and T_2 in the brain of healthy subjects, patients with Parkinson disease, and patients with multiple system atrophy: relation to iron content. *Radiology* 1999;211:489–495.
3. Cieszanowski A, Szeszkowski W, Golebiowski M, Bielecki DK, Grodzicki M, Pruszynski B. Discrimination of benign from malignant hepatic lesions based on their T_2 -relaxation times calculated from moderately T_2 -weighted turbo SE sequence. *Eur Radiol* 2002;12 (9):2273–2279.
4. Shah NJ, Neeb H, Zaitsev M, Steinhoff S, Kircheis G, Amunts K, Häussinger D, Zilles K. Quantitative T_1 mapping of hepatic encephalopathy using magnetic resonance imaging. *Hepatology* 2003;38:1219–1226.
5. Iles L, Pfluger H, Phrommintikul A, Cherayath J, Aksit P, Gupta SN, Kaye DM, Taylor AJ. Evaluation of diffuse myocardial fibrosis in heart failure with cardiac magnetic resonance contrast-enhanced T_1 mapping. *J Am Coll Cardiol* 2008;52:1574–1580.
6. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med* 2002;47:1202–1210
7. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med* 1999;42:952–962.
8. Doneva M, Börnert P, Eggers H, Stehning C, Senegas J, Mertins A. Compressed sensing reconstruction for magnetic resonance parameter mapping. *Magn Reson Med* 2010;64:1114–1120.
9. Bilgic B, Goyal VK, Adalsteinsson E. Multi-contrast reconstruction with Bayesian compressed sensing. *Magn Reson Med* 2011;66:1601–1615.
10. Velikina JV, Alexander AL, Samsonov A. Accelerating MR parameter mapping using sparsity-promoting regularization in parametric dimension. *Magn Reson Med* 2013;70:1263–1273.
11. Petzschner FH, Ponce IP, Blaimer M, Jakob PM, Breuer FA. Fast MR parameter mapping using k-t principal component analysis. *Magn Reson Med* 2011;66:706–716.
12. Zhang T, Pauly JM, Levesque IR. Accelerating parameter mapping with a locally low rank constraint. *Magn Reson Med* 2014;73:655-661.
13. Zhou Y, Shi C, Ren F, Lyu J, Liang D, Ying L. Accelerated MR parameter mapping using nonlinear manifold learning and

supervised-preimaging. In Proceedings of IEEE International Symposium on Biomedical Imaging, New York, NY, USA, 2015, pp. 897-900

14. Huang C, Graff CG, Clarkson EW, Bilgin A, Altbach MI. T2 mapping from highly undersampled data by reconstruction of principal component coefficient maps using compressed sensing. *Magn Reson Med* 2012;67:1355–1366.
15. Zhao B, Lu W, Hitchens TK, Lam F, Ho C, Liang ZP. Accelerated MR parameter mapping with low-rank and sparsity constraints. *Magn Reson Med* 2015;74: 489-498.
16. Tamir JI, Uecker M, Chen W, Lai P, Alley MT, Vasanawala SS, Lustig M. T2 shuffling: Sharp, multicontrast, volumetric fast spin-echo imaging *Magn Reson Med* 2016;77: 180-195.
17. Peng X, Ying L, Liu Y, Yuan J, Liu X, Liang D. Accelerated exponential parameterization of T2 relaxation with model-driven low rank and sparsity priors (MORASA). *Magn Reson Med* 2016;76:1865–1878.
18. Roeloffs V, Wang X, Sumpf TJ, Untenberger M, Voit D, Frahm J. Model-based reconstruction for T1 mapping using single-shot inversion-recovery radial FLASH. *International Journal of Imaging Systems and Technology* 2016;26:254-263.
19. Peng X, Liu X, Zheng H, Liang D. Exploiting parameter sparsity in model-based reconstruction to accelerate proton density and T2 mapping. *Medical Engineering and Physics* 2014;36:1428-1435.
20. Zhao B, Lam F, Liang ZP. Model-based MR parameter mapping with sparsity constraints: parameter estimation and performance bounds. *IEEE Trans Med Imaging* 2014;33:1832–1844.
21. Wang X, Roeloffs V, Klosowski J, Tan Z, Voit D, Uecker M, Frahm J. Model based T1 mapping with sparsity constraints using single-shot inversion-recovery radial FLASH. *Magn Reson Med* 2017; doi:10.1002/mrm.26726
22. Li Z, Berman BP, Galons JP, Bilgin A, Altbach MI, Martin DR. Rapid high-resolution T1 mapping using highly accelerated radial steady-state free-precession acquisition. In Proceedings of the 24th Annual Meeting of ISMRM, Singapore, 2016. p. 4196.
23. Vitanis V, Manka R, Giese D, Pedersen H, Plein S, Boesiger P, Kozerke S. High resolution three-dimensional cardiac perfusion imaging using compartment based k-t principal component analysis. *Magn Reson Med* 2011;65:575–587.
24. Giese D, Schaeffter T, Kozerke S. Highly undersampled phase-contrast flow measurements using compartment based k-t principal component analysis. *Magn Reson Med* 2013;69:434–443.
25. Velikina JV, Samsonov AA. Reconstruction of dynamic image series from undersampled MRI data using data-driven model consistency condition (MOCCO). *Magn Reson Med* 2015;74:1279-90.
26. Samsonov A. A Novel Reconstruction Approach Using Model Consistency Condition for Accelerated Quantitative MRI (MOCCA). In Proceedings of the 20th Annual Meeting of ISMRM, Melbourne, 2012. p. 0358.
27. Velikina J, Samsonov A. A novel spatio-temporal adaptive technique for reconstruction of dynamic MRI series. In Proceedings of the 25th Annual Meeting of ISMRM, Honolulu, 2014. p. 4376.

28. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found Trends Mach Learn* 2011;3:1–122.
29. Huang C, Bilgin A, Barr T, Altbach MI. T2 relaxometry with indirect echo compensation from highly undersampled data. *Magn Reson Med* 2013;70:1026–1037.
30. Lebel RM, Wilman AH. Transverse relaxometry with stimulated echo compensation. *Magn Reson Med* 2010;64:1005–1014.
31. Theilmann RJ, Gmitro AF, Altbach MI, Trouard TP. View-ordering in radial fast spin-echo imaging. *Magn Reson Med* 2004;51:768–774.
32. Wundrak S, Paul J, Ulrici J, Hell E, Rasche V. A small surrogate for the golden angle in time resolved radial MRI based on generalized Fibonacci sequences. *IEEE Trans Med Imaging* 2014;34:1262–1269.
33. Uecker M, Lai P, Murphy MJ, Virtue P, Elad M, Pauly JM, Vasanawala SS, Lustig M. ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA. *Magn Reson Med* 2014;71:990–1001.
34. Pruessmann KP, Weiger M, Börnert P, Boesiger P. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magn Reson Med* 2001;46:638–651.
35. Hansen MS, Sørensen TS. Gadgetron: An Open Source Framework for Medical Image Reconstruction. *Magn Reson Med* 2013;69:1768-76.
36. Ravishankar S, Bresler Y. MR image reconstruction from highly undersampled k-space data by dictionary learning. *IEEE Trans Med Imaging* 2011;30:1028–1041.
37. Golub G, Pereyra V. Separable nonlinear least squares: the variable projection method and its applications. *Inverse Probl* 2003;19:R1–R26.
38. Collins DL, Zijdenbos AP, Kollokian V, Sled JG, Kabani NJ, Holmes CJ, Evans AC. Design and construction of a realistic digital brain phantom. *IEEE Trans Med Imaging* 1998;17:463–486.
39. Altbach MI, Bilgin A, Li Z, Clarkson EW, Trouard TP, Gmitro AF. Processing of radial fast spin-echo data for obtaining T₂ estimates from a single k-space data set. *Magn Reson Med* 2005;54:549–559.
40. Vidal R. Subspace clustering. *IEEE Signal Process Mag* 2011;28:52–68.
41. Ramani S, Liu Z, Rosen J, Nielsen JF, Fessler J. Regularization parameter selection for nonlinear iterative image restoration and MRI reconstruction using GCV and SURE-based methods. *IEEE Trans Image Process* 2012;21: 3659–3672.
42. Velikina J, Samsonov A. Accurate, rank-adaptive reconstruction of undersampled dynamic MRI data using Bayesian information criterion. In *Proceedings of the 25th Annual Meeting of ISMRM, Honolulu, 2017*. p. 1203.

Figures

Figure 1: (a) Subspace basis generation from an ensemble of training curves using SVD, (b) Sequence of operations that need to be performed before MOCCO-LS reconstruction.

Figure 2: In the top row: The T_2 histogram shows the relative density of the discretization of the T_2 model space used in the training curve design and the SEPG training curves clustered with four clusters and displayed as an image where each pixel represents a training curve. In the bottom row: first pass estimate of the contrast images from a MOCCO reconstruction is shown along with the ℓ_2 distance based assignment of each spatial pixel to a specific local subspace. Binary spatial support maps extracted from the fit are also shown.

Figure 3: The results from the Monte-Carlo study at $R=16$ with radial sampling. g-factor (noise amplification) maps are shown in the first row and normalized error maps are shown in the second row. The overall error is indicated on the bottom left of the error maps. These results complement the results reported in Table 1.

Figure 4: The reconstructed T_2 maps and the normalized errors maps at $R=16$ generated from the k-t PCA and MOCCO methods at different model orders (K). The overall errors for the reconstruction is indicated at the bottom left of the error maps.

Figure 5: The reconstructed T_2 maps and the normalized errors maps at $R=16$ and $R=32$ generated from k-t PCA, L12, MOCCO and MOCCO-LS. The model order K used for each reconstruction is indicated at the top of the figure. The overall errors for the reconstruction are indicated at the bottom left of the error maps. Representative spatial support maps ($R=16, L=8$)

used in the MOCCO-LS reconstructions are also shown. The reference T_2 map shows the ROIs highlighted in black and the ROI results are provided in Supporting Information Table S1.

Figure 6: The reconstructed T_1 maps and the normalized errors maps generated from k-t PCA, L12, MOCCO and MOCCO-LS at $R=28$ relative to the reference. MOCCO-LS reconstructions at different model orders are shown. The model order used for the reconstructions is also indicated. The overall errors for the reconstruction is indicated at the bottom left of the error maps. The reference T_1 map shows the ROIs highlighted in black. The spatial support maps used in the MOCCO-LS reconstructions are also shown. These results complement the results in Table 2.

Figure 7: (a) Overall error and HFEN plots for all the T_2 mapping reconstructions at $R=16$ and $R=32$ as a function of the model order and (b) Overall error and HFEN plots of T_2 mapping with MOCCO-LS reconstructions at $R=16,32$ as a function of the number of clusters at $K = 2$. (c) Overall error and HFEN plots for all the T_1 mapping reconstructions as a function of the model order.

Figure 8: Representative TE ($R=16$) and TI images ($R=28$) from k-t PCA, L12, MOCCO, and MOCCO-LS reconstructions.

Overall Error									
Tissue	k-t PCA	L12	MOCCO	MOCCO-LS (L=4)			MOCCO-LS (L=8)		
	K=4	K=4	K=3	K=2	K=2 or 3	K=3	K=2	K=2 or 3	K=3
Gray	5.3	8.4	4.9	5.2	5.0	5.1	5.2	4.6	4.9
White	6.8	7.3	5.4	4.1	5.1	4.8	4.2	4.9	4.4
CSF	14.2	16.9	13.1	15.3	11.3	11.5	14.3	11.6	11.4
All tissues	13.1	15.6	12.0	13.9	10.5	10.6	13.1	10.7	10.9

Table 1: Overall error in T_2 evaluated over the entire imaging scene and over specific tissues for the MC experiments (R=16). Results for MOCCO-LS are reported at different K and L values.

		% T_1 error					
ROI#	Reference T_1 (ms)	k-t PCA	L12	MOCCO	MOCCO-LS (K=2)	MOCCO-LS (K=2 or 3)	MOCCO-LS (K=3)
1	826.1	2.81	2.35	5.41	2.38	1.56	2.15
2	1081.1	2.89	2.06	2.74	2.14	2.63	3.31
3	1493.4	2.73	2.88	2.64	2.15	1.04	1.35

Table 2: Percentage error between the T_1 estimates and the reference values for the brain ROIs indicated on the reference T_1 map in Figure 6.

Accelerated MR Parameter Mapping with a union of local subspaces constraint: Supporting Information

Supporting Figure S1: Results from the in-vivo T_1 and T_2 mapping experiments at different cluster sizes are shown along with the overall error in the parameter maps indicated on the bottom left corner.

			% T ₂ error			
R	ROI#	Reference T ₂ (ms)	k-t PCA	L12	MOCCO	MOCCO-LS
16	1	62.64	1.23	-0.88	-3.69	-1.19
	2	67.75	-1.98	-2.69	-3.62	-0.13
	3	75.7	3.34	-3.31	-2.38	-1.67
32	1	62.64	-16.35	-11.20	-6.94	-4.63
	2	67.75	-13.15	-3.42	-4.55	3.44
	3	75.7	-4.62	-6.38	-3.12	-2.55

Supporting Table S1: Differences between the T₂ estimates and the reference values for the brain ROIs indicated on the reference T₂ map in Figure 5 of the paper.