

Robust Reconstruction of Accelerated Perfusion MRI Using Local and Nonlocal Constraints

Cagdas Ulas^{1,3}, Pedro A Gómez^{1,3}, Felix Krahmer², Jonathan I Sperl³,
Marion I Menzel³, Bjoern H Menze^{1,4}

¹Computer Science, Technische Universität München, Munich, Germany

²Applied Numerical Analysis, Technische Universität München, Munich, Germany

³GE Global Research, Munich, Germany

⁴Institute for Advanced Study, Technische Universität München, Munich, Germany

cagdas.ulas@tum.de

Abstract. Dynamic perfusion magnetic resonance (MR) imaging is a commonly used imaging technique that allows to measure the tissue perfusion in an organ of interest via assessment of various hemodynamic parameters such as blood flow, blood volume, and mean transit time. In this paper, we tackle the problem of recovering perfusion MR images from undersampled k-space data. We propose a novel reconstruction model that jointly penalizes spatial (local) incoherence on temporal differences obtained based on a reference image and the patch-wise (non-local) dissimilarities between spatio-temporal neighborhoods of MR sequence. Furthermore, we introduce an efficient iterative algorithm based on a proximal-splitting scheme that solves the joint minimization problem with fast convergence. We evaluate our method on dynamic susceptibility contrast (DSC)-MRI brain perfusion datasets as well as on a publicly available dataset of in-vivo breath-hold cardiac perfusion. Our proposed method demonstrates superior reconstruction performance over the state-of-the-art methods and yields highly accurate estimation of perfusion time profiles, which is very essential for the precise quantification of clinically relevant perfusion parameters.

1 Introduction

Medical diagnosis and research heavily employ perfusion-weighted magnetic resonance imaging (MRI) techniques to estimate the blood flow and volume through examination of the spatio-temporal changes of the signal intensities following the injection of a blood bolus via exogenous paramagnetic tracers. In neuroimaging, these techniques have become widespread clinical tools in the diagnosis of stroke – for the assessment of the tissue at risk –, and the treatment of patients with cerebrovascular disease. One of the major obstacles in the clinical use of perfusion imaging is the need to track the rapid kinetics of contrast agent (tracer) uptake for accurate perfusion quantification [6]. Moreover, the short scanning time available for each frame often results in limited spatial and temporal resolution, or poor signal-to-noise ratio (SNR) images. In order to improve the

spatial or temporal resolution, one widely used approach is to accelerate the acquisition of each frame through the undersampling of k-space by acquiring only a subset of k-space lines [3,15]. However, the undersampling often results in aliasing artifacts in image space and in the context of perfusion MRI, accurate reconstruction of the complete temporal perfusion signal with its peak and high dynamics becomes an even more challenging task.

In recent years, various approaches have been proposed to solve the reconstruction problem in related dynamic imaging tasks, considering, such as piecewise smoothness in the spatial domain [17], high correlation and sparsity in the temporal domain [10,4,3], sparse representations of local image regions via learned dictionaries [3] and low-rank property of MR sequences in the full spatio-temporal space [10,14,17]. Although these methods allow highly accurate reconstructions from fewer k-space data, the main drawback is that their performance is very sensitive to motion and rapid intensity changes occurring over the duration of image acquisition as encountered in perfusion MRI. In addition, these methods often result in over smooth and blurry image regions that are lacking finer details when the acquired data are contaminated with high noise.

In this paper, we integrate two fundamentally different approaches that both increase the robustness of the reconstruction for perfusion MRI: (i) we enforce pixel-wise local sparsity constraint on the temporal differences that limits the overall dynamic of the perfusion time series, (ii) we enforce patch-wise similarity constraint on the spatio-temporal neighborhoods of whole MR sequence, which provides smooth spatial image regions with less temporal blurring especially when there is significant inter-frame motion and noise. We present the main optimization problem in a joint formal framework and introduce a new proximal splitting strategy that benefits from the weighted-average of proximals – thus, overcome a key limitation of the widely used Fast Composite Splitting Algorithm (FCSA) [9] –, and efficiently solves the joint minimization problem with fast convergence. The proposed method is validated on different types of MR perfusion datasets in comparison with the state-of-the-art methods and extensive experiments demonstrate the superior performance of our method in terms of reconstruction accuracy and accurate estimation of perfusion time profiles from undersampled k-space data even when being presented with high noise levels.

Contributions. Our main contributions are four-fold: (1) We present a new reconstruction scheme which cannot only produce high-quality spatial images for dynamic MRI but also enable to reconstruct the complete temporal signal dynamics for perfusion MRI from undersampled k-space data (Sect. 2). (2) We introduce an efficient proximal-splitting algorithm based on a generalized forward-backward splitting scheme [13]. This algorithm provides fast convergence and can be easily applied to various medical image applications that consider optimization problems where the objective function is the sum of several convex regularization terms (Sect. 3). (3) We demonstrate the efficiency and effectiveness of our method by comparing with state-of-the-art techniques on clinical datasets (Sect. 4). (4) Our proposed reconstruction model can enable ac-

curate quantification of clinically useful perfusion parameters while accelerating the acquisition through the use of fewer k-space samples.

2 Formulation

For the sake of simplicity, throughout the paper, we consider the reconstruction problem only on 2D+t data (i.e., on a single slice followed over time), however the idea presented here can also easily be extended to 3D+t data. We assume that $X^{3D} \in \mathbb{C}^{M \times N \times T}$ is a 2D perfusion image series represented as a spatio-temporal 3D volume. Let $x_t \in \mathbb{C}^{M \times N}$ denote one perfusion MR frame at time t with $M \times N$ pixels, y_t is the corresponding undersampled k-space measurements of x_t , and $\mathbb{T} = \{1, 2, \dots, T\}$ is the set of frame number indices in the sequence. The main goal is to recover all x_t 's from the collected k-space measurements y_t 's. The observation model between x_t and y_t can be mathematically formulated as

$$y_t = R_t(\mathcal{F}_{2D}x_t + \eta) \quad (1)$$

where R_t denotes the undersampling mask to acquire only a subset of k-space, \mathcal{F}_{2D} is the 2D Fourier Transform operator and η is additive Gaussian noise in k-space. We also denote the partial 2D Fourier operator for frame t as $\mathcal{F}_t = R_t\mathcal{F}_{2D}$, and stack the \mathcal{F}_t for all frames of the sequence as $\mathcal{F}_u = \text{diag}\{\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_T\}$.

We propose solving the following optimization problem for the reconstruction of perfusion MR sequences:

$$\hat{X} = \arg \min_X \left\{ \frac{1}{2} \|\mathcal{F}_u X - Y\|_2^2 + \lambda_1 \mathcal{G}_1(X) + \lambda_2 \mathcal{G}_2(X) \right\} \quad (2)$$

where $X \in \mathbb{C}^{MN \times T}$ denotes the whole perfusion MRI sequence and $Y \in \mathbb{C}^{MN \times T}$ represents the collection of all the k-space measurements. λ_1 and λ_2 are the tuning parameters for two regularization terms.

Local (\mathcal{G}_1) regularizer: The first regularization term in (2) penalizes the sum of spatial finite differences on the difference images calculated based on a reference for every image frame $x_t \in \mathbb{C}^{M \times N}$, and this term is named as dynamic total variation (TV) [4] and for the whole MR sequence, it can be defined as

$$\mathcal{G}_1(X) = \sum_{t \in \mathbb{T}} \sum_{n=1}^{M \times N} \sqrt{(\nabla_x(x_t - \bar{x}))_n^2 + (\nabla_y(x_t - \bar{x}))_n^2} \quad (3)$$

where \bar{x} is the reference image computed by averaging all the frames in MR sequence, ∇_x and ∇_y represent the finite-difference operators along the x and y dimensions, respectively. The intuition behind using dynamic TV over standard TV is that it is better adjusted to the variation in time, and this regularizer serves as a penalty on the overall dynamic of the temporal perfusion signal.

Nonlocal (\mathcal{G}_2) regularizer: The second regularization term in (2) penalizes the weighted sum of ℓ_2 norm distances between spatio-temporal neighborhoods

(patches) of MR sequence, and this penalty term can be specified by [16]

$$\mathcal{G}_2(X) = \sum_{(p_x, p_y, p_t) \in \Omega} \sum_{(q_x, q_y, q_t) \in \mathcal{N}_p} w(p, q) \|P_p(X^{3D}) - P_q(X^{3D})\|_2^2 \quad (4)$$

where $p = (p_x, p_y, p_t)$ and $q = (q_x, q_y, q_t)$ are two voxels, and the voxel of interest is $p \in \Omega$, where $\Omega = [0, M] \times [0, N] \times [0, T]$. The term $P_p(X^{3D})$ denotes a spatio-temporal 3D patch of the MR sequence centered at voxel p . We represent \mathcal{N}_p as a 3D search window around voxel p , and the size of the patch should be smaller than the size of the search window. We simply denote N_p and N_w as the size of a patch and search window, respectively. The weights $w(p, q)$ are determined based on ℓ_2 norm distance between the patches and calculated as

$$w(p, q) = e^{-\frac{\|P_p(X^{3D}) - P_q(X^{3D})\|_2^2}{h^2}} \quad (5)$$

where h is a smoothing parameter controlling the decay of the exponential function. The use of exponential weighting ensures that a voxel which is more similar to the voxel of interest in terms of Euclidean distance receives higher weight.

This regularizer is capable of exploiting the similarities between patch pairs in adjacent frames and it can enforce smooth solutions in the spatio-temporal neighbourhoods of MR sequence even when there is significant inter-frame motion and high noise introduced during acquisition.

3 Algorithm

To efficiently solve the primal problem (2), we propose to apply a proximal-splitting framework to this problem. Before describing our proximal-splitting based algorithm, we should first give the definition of a proximal map.

Proximal map: Given a continuous convex function $g(x)$ and a scalar $\rho > 0$, the proximal operator associated to convex function g can be defined as [9]

$$prox_\rho(g)(z) := \arg \min_{x \in \mathcal{H}} \left\{ \frac{1}{2\rho} \|x - z\|_2^2 + g(x) \right\} \quad (6)$$

Now we can reformulate the problem (2) as the following denoising problem

$$\hat{X} = \arg \min_X \left\{ \frac{1}{2} \|X - X_g\|_2^2 + \rho\lambda_1 \mathcal{G}_1(X) + \rho\lambda_2 \mathcal{G}_2(X) \right\} \quad (7)$$

Since each of the regularization term in the cost function (2) is convex, the problem (7) can be represented as the proximal map of the sum of two regularization terms as described in Fast Composite Splitting Algorithm (FCSA) [9]

$$\hat{X} = prox_\rho(\lambda_1 \mathcal{G}_1 + \lambda_2 \mathcal{G}_2)(X_g) \quad (8)$$

The problem (7) admits to a unique solution as given in (8). However, the proximity operator of the sum of two functions is usually intractable. To compute it

iteratively, one can adopt an efficient proximal-splitting method to this problem. Proximal-splitting methods are first-order iterative algorithms that solve relatively large-scale optimization problems with several nonsmooth penalties. They operate by splitting the convex objective function to minimize and generating individual subproblems which are evaluated easily via proximal operators.

To solve our main problem in (7), we split the objective function into two individual subproblems that we term \mathcal{G}_1 -subproblem and \mathcal{G}_2 -subproblem.

\mathcal{G}_1 -subproblem: The proximal map for this subproblem can be defined as

$$X_{\mathcal{G}_1} = \text{prox}_\rho(\lambda_1 \mathcal{G}_1)(X_g) = \arg \min_X \left\{ \frac{1}{2\rho} \|X - X_g\|_2^2 + \lambda_1 \mathcal{G}_1(X) \right\} \quad (9)$$

In order to solve the subproblem (9), we first reformulate it by introducing new variables $d_t = x_t - \bar{x}$ and $d_g^t = X_g^t - \bar{x}$, in this way the problem turns into

$$\hat{d} = \arg \min_d \left\{ \sum_{t \in \mathbb{T}} \left(\frac{1}{2\rho} \|d_t - d_g^t\|_2^2 + \lambda_1 \|d_t\|_{TV} \right) \right\} \quad (10)$$

where $d = \{d_1, \dots, d_T\}$ and $\|d_t\|_{TV} = \|[Q_1 d_t, Q_2 d_t]\|_{2,1}$, where Q_1 and Q_2 are two $MN \times MN$ first order finite difference matrices in vertical and horizontal directions, and $\ell_{2,1}$ norm is the sum of the ℓ_2 norm of each row of given matrix.

Given a reference image \bar{x} , the cost function in (10) can be minimized individually for every frame x_t of MR sequence. This guarantees that the sum of the costs in (10) is also minimized. The cost function can be efficiently minimized by using the fast iteratively reweighted least squares (FIRLS) algorithm [5] based on preconditioned conjugate gradient method. This algorithm enables fast convergence and low computational cost by adopting a new preconditioner which can be accurately approximated using the diagonally dominant structure of the matrix $\mathcal{F}_t^H \mathcal{F}_t$, where H is the conjugate transpose. Once the problem (10) is solved, the estimated solution for problem (9) can be calculated as

$$\hat{X}_{\mathcal{G}_1} = \left[\hat{d}_1 + \bar{x}, \hat{d}_2 + \bar{x}, \dots, \hat{d}_T + \bar{x} \right] \quad (11)$$

\mathcal{G}_2 -subproblem: The proximal map for \mathcal{G}_2 subproblem can be specified by

$$X_{\mathcal{G}_2} = \text{prox}_\rho(\lambda_2 \mathcal{G}_2)(X_g) = \arg \min_X \left\{ \frac{1}{2\rho} \|X - X_g\|_2^2 + \lambda_2 \mathcal{G}_2(X) \right\} \quad (12)$$

The problem (12) can be solved using a two-step alternating minimization scheme in an iterative projections onto convex sets (POCS) framework [11]. In each iteration, the first step involves the projection of image estimate onto the data fidelity term via a steepest descend update and the second step performs the minimization of the neighborhood penalty term on the projected data. The minimization of the penalty function in (12) is equivalent to applying non-local means (NLM) filter [2] to the projected images. This is mathematically derived in [12] with a single assumption that only one sub-iteration of the penalty term is performed

Algorithm 1: Proposed algorithm

Input: Undersampled k-space data $Y, \mathcal{F}_u, \lambda_1, \lambda_2$
Initialize: $z_1^0 = z_2^0 = \mathcal{F}_u^H Y, w_1, w_2, X^0 = \sum_{i=1}^2 w_i z_i^0, \alpha_0 = 0.5, \gamma = 1, k = 0$
while *stopping criteria not met* **do**
 $X_g = X^k - \gamma \mathcal{F}_u^H (\mathcal{F}_u X^k - Y)$;
 $z_1^{k+1} = z_1^k + \alpha_k (\text{prox}_{\frac{\gamma}{w_1}} (2\lambda_1 \mathcal{G}_1)(X^k + X_g - z_1^k) - X^k)$;
 $z_2^{k+1} = z_2^k + \alpha_k (\text{prox}_{\frac{\gamma}{w_2}} (2\lambda_2 \mathcal{G}_2)(X^k + X_g - z_2^k) - X^k)$;
 $X^{k+1} = w_1 z_1^{k+1} + w_2 z_2^{k+1}$;
 $\alpha_{k+1} = 1 + 2(\alpha_k - 1)/(1 + \sqrt{1 + 4(\alpha_k)^2})$;
 $k \leftarrow k + 1$;
end
Output: Reconstructed image data X

with constant and predetermined weights. The mathematical formulation of a NLM filter is given as [12]

$$\hat{X}(p_x, p_y, p_t) = \frac{\sum_{(q_x, q_y, q_t) \in \mathcal{N}_p} w(p, q) X(q_x, q_y, q_t)}{\sum_{(q_x, q_y, q_t) \in \mathcal{N}_p} w(p, q)} \quad (13)$$

We have now iterative solvers for each subproblem \mathcal{G}_1 and \mathcal{G}_2 . In this work, we have developed an efficient algorithm by adopting a generalized forward-backward splitting (GFBS) framework [13] that minimizes the sum of multiple convex functions. Basically, FCSA and GFBS are operator-splitting algorithms and they both use forward-backward schemes. The main difference between GFBS and FCSA is that GFBS enables the use of weighted-average of the outputs of individual proximal mappings for finitely many convex functions, whereas FCSA only applies simple averaging. The weighted-average of the outputs of proximals may practically yield better results depending on the effectiveness of each penalty (regularization) term employed in various applications.

We further accelerate the convergence of the algorithm with an additional acceleration step similar to the Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) [1]. This step adaptively increases the value of step size parameter (α_k) through iterations and make it sufficiently close to 1. Our proposed reconstruction algorithm is outlined in Algorithm 1. The most computationally expensive step of our algorithm is solving each proximal map. Fortunately, the computation of proximal maps can be done in parallel since there is no dependency between the inputs of proximity operators. All the other steps involve adding and multiplying vectors or scalars, and are thus very cheap in terms of computational complexity. The GFBS method has been shown to converge when $\gamma < 2/L$ if the convex function $f = \frac{1}{2} \|X - X_g\|_2^2$ has a Lipschitz continuous gradient with constant L [13]. We refer the readers to original GFBS paper [13] for details concerning the proof of the convergence of the algorithm.

4 Experiments

Experimental Setup: We evaluate our method on two different types of perfusion MRI datasets. We use three DSC-MRI brain perfusion sequences ($128 \times 128 \times 60$) and one in-vivo breath-hold cardiac perfusion sequence¹ ($192 \times 192 \times 40$) from [4] with normalized intensities. All the perfusion datasets used in the experiments are acquired with full-sampling and the fully-sampled sequences are artificially corrupted by multiplying its corresponding k-space representation with a binary undersampling mask and subsequently adding Gaussian white noise. To simulate undersampling, we retrospectively apply a time-varying variable density Cartesian mask in our experiments (see Fig. 1). The sampling ratio is set to 1/4 for brain sequences and 1/6 for cardiac sequence. We compare our method with three state-of-the-art reconstruction methods: the dynamic total variation (DTV) [4], (k,t)-space via low-rank plus sparse prior (kt-RPCA) [14], and fast total variation and nuclear norm regularization (FTVNRR) [17]. To ensure fair comparison, similar to the experiments presented in [3], we empirically fine-tune the optimal regularization parameters for all methods individually for each dataset and depending on noise level. For our proposed method, we specifically set $\lambda_2 = 0.25$ for all noise levels and set $\lambda_1 = 0.025$ for relatively high level noise and $\lambda_1 = 0.001$ for low noise levels. We test the following noise levels and report the results: $\sigma = \{10^{-1}, 5 \times 10^{-2}, 10^{-2}, 5 \times 10^{-3}, 10^{-3}\}$. For the proposed method, we set $N_w = 7 \times 7 \times 7$, $N_p = 5 \times 5 \times 5$, and $w_1 = w_2 = 0.5$ for all sequences. We consider using small cubic neighborhoods for N_w and N_p since large neighborhoods drastically increase the computation time. To reduce the computational burden, we also employ an optimized blockwise version of the non-local means filter that was proposed by Coupé et al. [7] for 3D medical data. We adopt the Peak Signal-to-Noise Ratio (PSNR) as the metric for quantitative evaluation. Instead of directly calculating PSNR on a whole image or 3D sequence, we employ a region-based analysis by calculating the PSNR on randomly selected 100 image blocks (50×50) in 2D frames. This allows us to test for statistical differences using paired t-test when comparing different methods.

Results: Fig. 1 and 2 demonstrate a single reconstructed frame of the first and third brain perfusion dataset, respectively, and the estimation of perfusion time profiles averaged over voxels inside a small region of interest. The results in Fig. 1 reveal that kt-RPCA and FTVNRR show quite similar performances, and the DTV yields both better reconstruction and estimation of perfusion signal compared to these two methods. Compared with all three methods, our proposed method can achieve the best reconstruction and very accurate estimation of perfusion time profiles even when the k-space measurements are contaminated with a relatively high level noise ($\sigma = 5 \times 10^{-2}$). The reconstruction results of our method are also statistically significant (p -value $< 10^{-5}$) when compared with all other methods. Moreover, both kt-RPCA and FTVNRR result in over spatial smoothing (see close-up views in Fig. 1) and along time as well, which can be clearly seen from smoothening of the perfusion peaks in the third-fourth column

¹ Available at: <http://web.engr.illinois.edu/~cchen156/SSMRI.html>

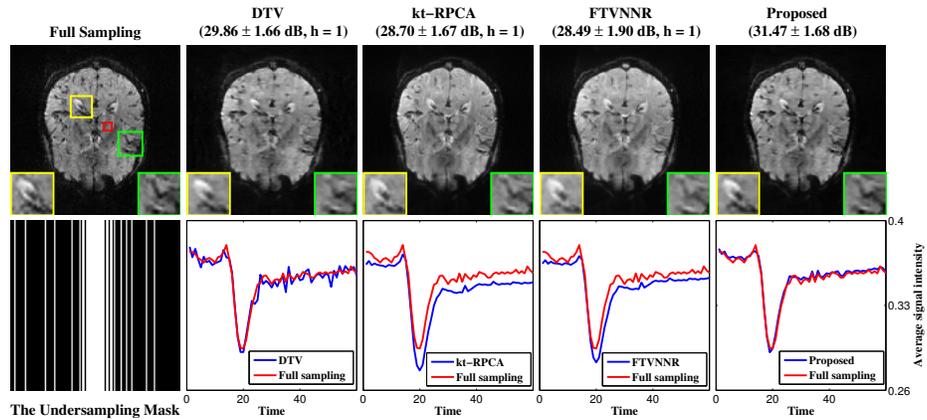


Fig. 1. (Top) Results (mean \pm std, h-value) of the 22nd frame of the first brain dataset and close-up views of two regions of interest (yellow and green square). h=1 specifies the statistical significance between the results of proposed and compared method, (Bottom) An exemplary undersampling mask and for each method, estimation of perfusion time profiles averaged over the voxels inside the red square shown in top-left figure. The standard deviation of added Gaussian noise is $\sigma = 5 \times 10^{-2}$. Our method achieves both the best frame-based reconstruction and the most accurate estimation of peaks and temporal pattern of perfusion signal.

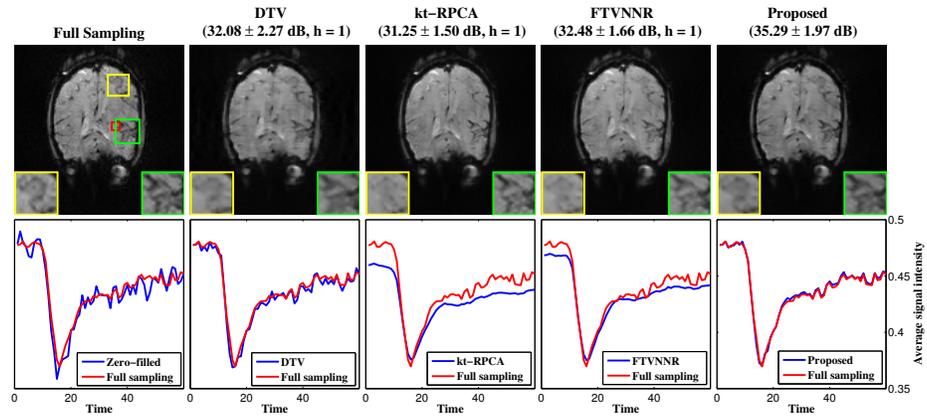


Fig. 2. (Top) Results (mean \pm std, h-value) of the 15th frame of the third brain dataset and close-up views of two regions of interest (yellow and green square), (Bottom) For each method, estimation of perfusion time profiles averaged over the voxels inside the red square shown in top-left figure. The standard deviation of added Gaussian noise is $\sigma = 10^{-3}$. Our method again achieves both the best frame-based reconstruction and the most accurate estimation of peaks and temporal pattern of perfusion signal.

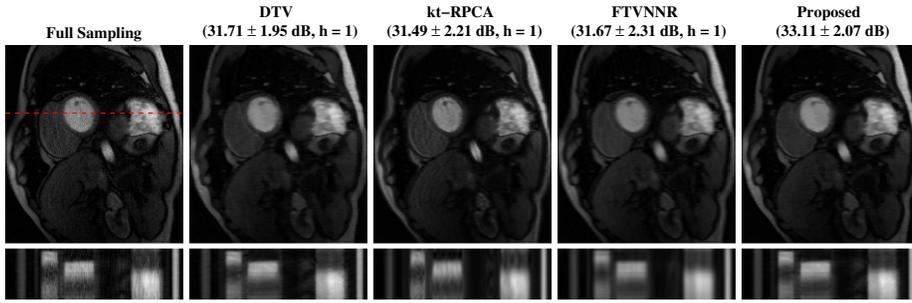


Fig. 3. (Top) Results of the 18th frame of the cardiac dataset with added noise $\sigma = 10^{-2}$, (Bottom) Temporal cross sections by the red dashed line. All methods can produce high quality spatial frames, however, our method yields less noisy and blurry temporal profiles, and the aliasing artifacts are also mostly removed.

of Fig. 1. In contrast, the proposed method reconstructs a perfusion pattern that is in good agreement with the pattern of the fully sampled data (see Fig. 1 bottom fifth column), and produces less blurry image regions and sharper edges. The perfusion time profiles obtained from the third dataset (see Fig. 2 bottom plots) also demonstrate the success of our proposed method. Considering the spatial outputs, when looking at details in close-up views of Fig. 2, the reconstructions obtained by kt-RPCA and FTVNNR are more blurry and thus lacking some finer details in yellow region, whereas the reconstruction obtained by proposed method involves more finer information in yellow region and provides sharper edges in green region.

We also validate our method on a cardiac perfusion data from [4] and the results are presented in Fig. 3. All methods here can produce high quality images, however, when looking temporal cross sections at bottom, it can be observed that our method gives less noisy and with lower aliasing artifacts reconstruction on

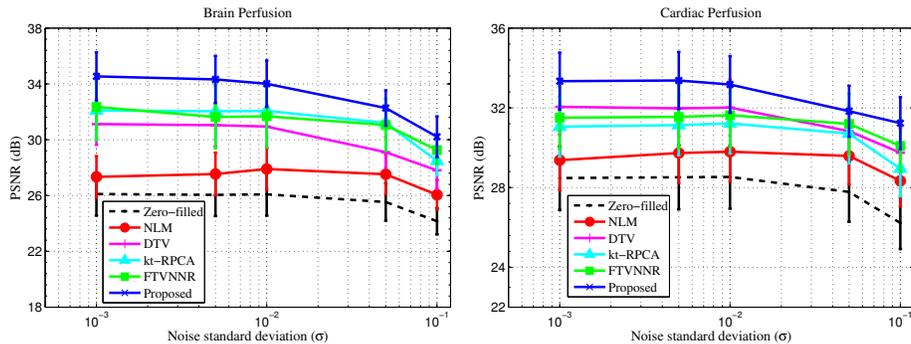


Fig. 4. PSNR results versus noise std (σ) for (left) Brain, (right) Cardiac datasets. Our joint local and nonlocal regularization based method performs the best.

myocardium surface while FTVNMR provides more blurry result. The reason is that our method can utilize both local consistency in temporal differences and nonlocal similarities between spatio-temporal neighborhoods of MR sequence while the FTVNMR does not explicitly exploit sparsity in temporal domain.

Quantitative results of different reconstruction methods on both brain and cardiac perfusion datasets are shown in Fig. 4. Note that the NLM only solves the \mathcal{G}_2 -subproblem of Sect. 3. From the figure, one can clearly see that our proposed method achieves the highest mean PSNR for all noise levels applied. The running time of all methods on the brain and cardiac datasets is provided in Table 1. Compared with the other three methods, our method needs the highest amount of processing time. However, due to its faster convergence, our method can achieve the best reconstruction accuracy within the first 3-4 iterations on average, which approximately takes 4.5 minutes for cardiac dataset on a desktop with Intel Xeon CPU E3-1226 v3 Processor.

Table 1. The time cost of different reconstruction methods.

Time (Seconds)	DTV	kt-RPCA	FTVNMR	Proposed
Brain ($128 \times 128 \times 60$)	54.5	194.5	74.3	304.6
Cardiac ($192 \times 192 \times 40$)	46.2	263.9	81.7	278.1

5 Conclusion

We have presented a robust reconstruction model for perfusion MRI, which is based on a joint regularization of pixel-wise and patch-wise spatio-temporal constraints. Numerical experiments validate the efficiency of our method over the state-of-the-art methods in terms of reconstruction accuracy and estimation of perfusion time profiles in varying noise conditions. We also introduce an iterative algorithm that efficiently solves convex optimization problems with mixtures of regularizers. Our algorithm provides fast convergence and can be easily extended to other medical image applications, in particular denoising and super-resolution. The proposed method can be also extended to parallel MR imaging [8] and be applied to multi-coil data. Future work will aim at expanding our work with the fitting of pharmacokinetic models and quantitative analysis of perfusion parameters on 3D+t brain perfusion data using partial k-space measurements.

Acknowledgements. The research leading to these results has received funding from the European Union’s H2020 Framework Programme (H2020-MSCA-ITN-2014) under grant agreement no 642685 MacSeNet. We also thank Dr. Christine Preibisch (Neuroradiology, Klinikum rechts der Isar der TU München) for providing brain perfusion datasets.

References

1. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. on Imag. Sci.* 2(1), 183–202 (2009)
2. Buades, A., Coll, B., Morel, J.M.: A non-local algorithm for image denoising. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. vol. 2, pp. 60–65 (2005)
3. Caballero, J., Price, A.N., Rueckert, D., Hajnal, J.: Dictionary learning and time sparsity for dynamic MR data reconstruction. *IEEE Trans. Med. Imag.* 33(4), 979–994 (2014)
4. Chen, C., Li, Y., Axel, L., Huang, J.: Real time dynamic MRI with dynamic total variation. In: *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 138–145. Springer (2014)
5. Chen, C., et al.: Preconditioning for accelerated iteratively reweighted least squares in structured sparsity reconstruction. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2713–2720 (2014)
6. Conturo, T.E., et al.: Arterial input functions for dynamic susceptibility contrast MRI: Requirements and signal options. *J. of Mag. Reson. Imag.* 22, 697–703 (2005)
7. Coupé, P., Yger, P., Prima, S., Hellier, P., Kervrann, C., Barillot, C.: An optimized blockwise nonlocal means denoising filter for 3-D magnetic resonance images. *IEEE Trans. Med. Imag.* 27(4), 425–441 (2008)
8. Deshmane, A., Gulani, V., Griswold, M.A., Seiberlich, N.: Parallel MR imaging. *J. Mag. Reson. Imag.* 36(1), 55–72 (2012)
9. Huang, J., Zhang, S., Metaxas, D.N.: Efficient MR image reconstruction for compressed MR imaging. *Med. Imag. Anal.* 15(5), 670–679 (2011)
10. Lingala, S.G., Hu, Y., DiBella, E., Jacob, M.: Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR. *IEEE Trans. Med. Imag.* 30(5), 1042–1054 (2011)
11. Marks, R.J.: Alternating projections onto convex sets. In: Jansson PA, editor. *Deconvolution of images and spectra*, 2nd ed. Orlando, FL: Academic Press (1996)
12. Protter, M., Elad, M., Takeda, H., Milanfar, P.: Generalizing the nonlocal-means to super-resolution reconstruction. *IEEE Trans. on Imag. Proc.* 18(1), 36–51 (2009)
13. Raguét, H., Fadili, J., Peyré, G.: A generalized forward-backward splitting. *SIAM J. on Imag. Sci.* 6(3), 1199–1226 (2013)
14. Trémouhéac, B., Dikaios, N., Atkinson, D., Arridge, S.R.: Dynamic MR image reconstruction-separation from undersampled (k-t)-space via low-rank plus sparse prior. *IEEE Trans. Med. Imag.* 33(8), 1689–1701 (2014)
15. Ulas, C., Gómez, P., Sperl, J.I., Preibisch, C., Menze, B.H.: Spatio-temporal MRI reconstruction by enforcing local and global regularity via dynamic total variation and nuclear norm minimization. In: *IEEE 13th International Symposium on Biomedical Imaging (ISBI)*. pp. 306–309 (2016)
16. Yang, Z., Jacob, M.: Nonlocal regularization of inverse problems: A unified variational framework. *IEEE Trans. on Imag. Proc.* 22(8), 3192–3203 (2013)
17. Yao, J., et al.: Accelerated dynamic MRI reconstruction with total variation and nuclear norm regularization. In: *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 635–642. Springer (2015)