On the Spatial and Temporal Influence for the Reconstruction of Magnetic Resonance Fingerprinting

Fabian Balsiger[1,2,3] FABIAN.BALSIGER@ARTORG.UNIBE.CH
[1] *Insel Data Science Center, Inselspital, Bern University Hospital, University of Bern, Switzerland*
[2] *NMR Laboratory, Institute of Myology, Neuromuscular Investigation Center, France*
[3] *NMR Laboratory, CEA, DRF, IBFJ, MIRCen, France*

Olivier Scheidegger[2,4,5] OLIVIER.SCHEIDEGGER@INSEL.CH
[4] *Department of Neurology, Inselspital, Bern University Hospital, University of Bern, Switzerland*
[5] *Support Center for Advanced Neuroimaging (SCAN), Institute for Diagnostic and Interventional Neuroradiology, Inselspital, Bern University Hospital, University of Bern, Switzerland*

Pierre G. Carlier[2,3] P.CARLIER@INSTITUT-MYOLOGIE.ORG
Benjamin Marty[2,3] B.MARTY@INSTITUT-MYOLOGIE.ORG
Mauricio Reyes[1] MAURICIO.REYES@MED.UNIBE.CH

Abstract

Magnetic resonance fingerprinting (MRF) is a promising tool for fast and multiparametric quantitative MR imaging. A drawback of MRF, however, is that the reconstruction of the MR maps is computationally demanding and lacks scalability. Several works have been proposed to improve the reconstruction of MRF by deep learning methods. Unfortunately, such methods have never been evaluated on an extensive clinical data set, and there exists no consensus on whether a fingerprint-wise or spatiotemporal reconstruction is favorable. Therefore, we propose a convolutional neural network (CNN) that reconstructs MR maps from MRF-WF, a MRF sequence for neuromuscular diseases. We evaluated the CNN's performance on a large and highly heterogeneous data set consisting of 95 patients with various neuromuscular diseases. We empirically show the benefit of using the information of neighboring fingerprints and visualize, via occlusion experiments, the importance of temporal frames for the reconstruction.

Keywords: Magnetic Resonance Fingerprinting, Image Reconstruction, Convolutional Neural Network, Quantitative Magnetic Resonance Imaging, Neuromuscular Diseases.

1. Introduction

Neuromuscular diseases impose a high burden to patients and society regarding severity and economic costs (Larkindale et al., 2014). For such diseases, noninvasive outcome measures to monitor disease progression and treatment effects are desperately needed. A promising tool for obtaining noninvasive outcome measures is magnetic resonance imaging (MRI) (Carlier et al., 2016). However, MRI is rarely performed for outcome measures due to the long acquisition times associated with quantitative MRI, which renders it clinically not applicable. Magnetic resonance fingerprinting (MRF) (Ma et al., 2013) is a novel and promising approach for quantitative MRI, which allows acquiring multiple quantitative MR parameters in one fast scan. The principle of MRF is to pseudo-randomly vary the MR sequence parameters to obtain a unique signal evolution, or fingerprint, per voxel and tissue type. Each fingerprint is then compared to a dictionary of pre-computed fingerprints

to estimate the quantitative MR parameters at interest. For instance, we proposed MRF-WF (Marty and Carlier, 2018), a MRF sequence for water T1 relaxation time ($T1_{H2O}$) and fat fraction (FF) quantification in skeletal muscles with an acquisition time of only 50 seconds. But while the acquisition time is short, the reconstruction of the $T1_{H2O}$ and FF maps requires approximately 20 hours per scan using a standard dictionary matching algorithm.

Several methods attempting to accelerate the reconstruction of MRF have been proposed lately. Fast group matching (Cauley et al., 2015) and search window reduction (Gómez et al., 2016) aim both at reducing the number of fingerprint comparisons during the dictionary matching. The dictionary matching has also been investigated within the frame of tree data structures and iterative reconstruction (Cline et al., 2017; Golbabaee et al., 2018b). Further, the compression of the fingerprints has also been investigated (McGivney et al., 2014; Assländer et al., 2018). However, these approaches suffer mainly by a over-discretized reconstruction, an approximation by compression, and most importantly a poor scalability with more MR parameters. To cope with the mentioned problems, several works have focused on replacing the dictionary matching by learning the mapping between fingerprints and MR parameters by deep learning. The deep learning-based MRF reconstruction can mainly be grouped into fingerprint-wise and spatial-temporal reconstruction. The fingerprint-wise reconstruction takes a single fingerprint as input and estimates the MR parameter(s) at the fingerprint's voxel. The proposed architectures vary from fully-connected networks (Cohen et al., 2018; Golbabaee et al., 2018a; Barbieri et al., 2018), to fully-convolutional with 1-D convolutions (Hoppe et al., 2017), the combination of both (Virtue et al., 2017), and recurrent neural networks (Oksuz et al., 2019). A limitation of the fingerprint-wise reconstruction is that they do not take advantage of the information between neighboring fingerprints. Recently, spatiotemporal reconstruction by inputting a neighborhood of fingerprints to a neural network has been shown to be beneficial (Balsiger et al., 2018; Fang et al., 2018). The proposed approaches reconstruct MR parameters from 5×5 patches via a convolutional neural network (CNN) (Balsiger et al., 2018) and MR maps slice-wise with an encoder-decoder CNN architecture (Fang et al., 2018).

Despite the increasing research being conducted on deep learning-based reconstruction of MRF, there exist several limitations. In general, all studies conducted experiments either on phantom or volunteer data. Furthermore, the data sets were limited in size with the maximum being six scans (Balsiger et al., 2018). Further, while the studies on spatiotemporal reconstruction have shown a potential advantage over fingerprint-wise reconstruction, it remains unclear to what extent the spatial dimension should be considered in the reconstruction. Last, none of the studies investigated the influence of the temporal frames on the reconstruction. Therefore, we propose a CNN that reconstructs quantitative MR maps from MRF. First, we show its performance on a large and highly heterogeneous data set with 95 scans of patients suffering from various neuromuscular diseases who were imaged at two anatomical regions. Second, we report on three different strategies to incorporate spatial information: i) fingerprint-wise reconstruction (Cohen et al., 2018), ii) slice-wise reconstruction with spatial pooling operations (Fang et al., 2018), and iii) our proposed reconstruction with varying sizes of receptive fields. Last, to gain insights into the MR parameter reconstruction, we visualize the importance of the temporal frames via occlusion experiments applied on the temporal domain.

2. Materials and Methods

2.1. MRF Acquisition and Reconstruction

We acquired the MRF-WF sequence at the legs and thighs levels in patients with various neuromuscular diseases, resulting in a highly heterogeneous data set with 95 scans (43 female, 52 male; age = 53.7 $\pm$ 19.5 years). The MRF-WF acquisition consisted of a non-selective inversion followed by a 1400 radial spokes FLASH echo train (golden angle scheme) with varying echo time (TE), repetition time (TR), and nominal flip angle (FA). We used a field of view of 350 $\times$ 350 mm^2 with a voxel size of 1.0 $\times$ 1.0 $\times$ 8.0 mm^3 and five slices per scan, which resulted in a total acquisition time of 50 s. All experiments were performed on a 3 Tesla Siemens MAGNETOM Prisma^{fit} scanner (Siemens Healthineers, Erlangen, Germany).

We reconstructed a MRF image space series consisting of 175 temporal frames from the raw k-space data after the acquisition. The reconstruction involved view sharing with a k-space weighted image contrast filter with 55 spokes, and compressed sensing with total variation (Marty et al., 2018a,b). This reconstruction resulted in a complex-valued MRF image space series of size 350 $\times$ 350 $\times$ 175 voxels per slice. For each scan, $T1_{H2O}$, FF, and transmit field efficacy (B1) reference MR maps were reconstructed from the MRF image space series using dictionary matching (Marty and Carlier, 2018). The total reconstruction time was approximately 20 hours per scan using a standard desktop computer (2.6 GHz Intel Xenon E5-2630, 48 GB memory).

2.2. CNN Map Reconstruction

We propose a CNN architecture that reconstructs MR maps patch-wise from MRF image space series. Let us consider a 2-D+time MRF image space series $I \in \mathbb{C}^{X \times Y \times T}$ and its corresponding reference MR maps $Q \in \mathbb{R}^{X \times Y \times M}$ with $X \times Y = 350 \times 350$ being the spatial dimension, $T = 175$ being the number of temporal frames, and $M = 3$ being the number of MR maps ($T1_{H2O}$, FF, B1). Our CNN learns the mapping $\mathscr{M} : I_P \rightarrow Q_P$, i.e. to reconstruct a patch $Q_P \in \mathbb{R}^{QP_X \times QP_Y \times M} \subset Q$ of the MR maps from a patch $I_P \in \mathbb{C}^{IP_X \times IP_Y \times T} \subset I$ of the 2-D+time MRF image space series. In our experiments, we reconstructed non-overlapping output patches of a size $QP_X \times QP_Y = 32 \times 32$ with the input patch size being set to $IP_X \times IP_Y = 42 \times 42$ according to the spatial receptive field of the CNN (11×11).[1]

2.2.1. PRE-PROCESSING

We normalized the real and imaginary parts of each MRF image space series I to zero mean and unit variance. Each MR map has been normalized to the range $[0, 1]$ using the minimum and maximum values of the entire data set.

2.2.2. CNN ARCHITECTURE

Our CNN architecture consists of two types of blocks: a temporal block that leverages the temporal information of the fingerprints, and a spatial block that leverages the spatial correlation between neighboring fingerprints. After each temporal block follows a spatial block, with this sequence being repeated five times. We treat the temporal dimension T as feature channels, and therefore,

1. The size of Q_P is mainly restricted by GPU memory. We could not determine a statistically significant difference to other patch sizes.

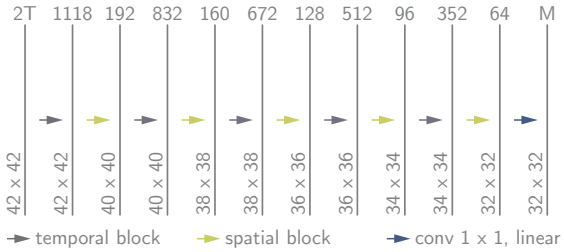

Figure 1: The proposed CNN architecture. The numbers at the top of the bars denote the number of feature channels ($T = 175$ and $M = 3$ for MRF-WF), and numbers at the lower left of the bars denote the spatial patch size.

the concatenated real and imaginary parts of the complex-valued I_P serve as input, i.e., we used a real-valued input $I_{Preal} \in \mathbb{R}^{IP_X \times IP_Y \times 2T}$. The output Q_P is reconstructed after a final 2-D convolution with kernel size of 1×1, M feature channels, and a linear activation (Figure 1).

Our temporal blocks learn features along the fingerprints by applying 1×1 convolutions and are inspired by dense blocks (Huang et al., 2017). The temporal blocks contain four layers, and each layer is connected to the output feature channels of every preceding layer. The number of feature channels, or growth rate, n is equal for the four layers in a temporal block, and decreases within the CNN from 192, 160, 128, 92, to 64. In order to reuse features and to further facilitate the gradient flow, we additionally concatenate the input of a temporal block to its output arriving at $4 \cdot n + n_{in}$ features channels after each temporal block, with n_{in} being the number of input feature channels to the temporal block. Each layer in our temporal block consists of a 2-D convolution with kernel size of 1×1, rectified linear unit (ReLu) activation function (Glorot et al., 2011), and batch normalization (Ioffe and Szegedy, 2015). After each temporal block, our spatial blocks extract spatial features by performing a 2-D convolution with a kernel size of 3×3 (valid padding) followed by a ReLu activation function, and batch normalization. The number of feature channels of the spatial blocks decreases equally to the growth rate of the temporal blocks within the CNN. Code is available online[2].

2.2.3. TRAINING

We trained our CNN for 100 epochs with a batch size of 50 randomly selected patches. We used an Adam optimizer (Kingma and Ba, 2015) to minimize a mean squared error loss with a learning rate of 0.001, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. The CNN was implemented with TensorFlow 1.10.0 (Google, Mountain View, CA, U.S.) and Python 3.6.5 (Python Software Foundation, Wilmington, DE, U.S.).

2.3. Comparison and Evaluation

We compared our CNN to three other deep learning-based MRF reconstruction methods. The first method is a neural network with two hidden fully-connected layers (Cohen et al., 2018). The method

2. https://github.com/fabianbalsiger/mrf-reconstruction-midl2019

requires single fingerprint as input and is, therefore, an extreme variant where no neighboring fingerprint information is considered. The second method is an encoder-decoder CNN with two spatial pooling operations (Fang et al., 2018). The method reconstructs entire MR map slices and hence represents the opposite extreme variant due to the spatial pooling operations, which increase the receptive field drastically compared to our method. We implemented both methods as originally described with slight modifications for our MRF sequence and we used the same pre-processing and training scheme as for our method. Last, we also show the reconstructions of our previous architecture with a receptive field of 5×5 (Balsiger et al., 2018).

We evaluated the performance of our CNN and the two compared methods by randomly splitting our data set into training/validation/testing sets (n=55/20/20). Note that we purposely did not apply any stratification regarding neuromuscular disease or anatomical region to evaluate the robustness to the heterogeneous data set. The best performing model on the validation set, in terms of the normalized root mean squared error (NRMSE), was selected and applied to the testing set. We calculated the NRMSE, peak signal-to-noise ratio (PSNR), and structural similarity index measure (SSIM) (Wang et al., 2004) between the predicted and the reference maps on a scan level. Additionally, we calculated the coefficient of determination (R2) between the predicted and reference mean values of manually segmented region of interests within the major muscles across all scans. For the calculations of the NRMSE and PSNR, background voxels were excluded using a mask segmented by thresholding an anatomical image (out-of-phase). The mask for the $T1_{H2O}$ map was further processed to exclude voxels with a FF higher than 0.65 because the confidence of $T1_{H2O}$ values is highly decreased at high FFs. For the SSIM, we used a window size of 7×7, $K_1 = 0.01$, $K_2 = 0.03$, and L was set to the maximum value of the reference map.

3. Experiments and Results

We first present the comparison between our method and (Cohen et al., 2018; Fang et al., 2018; Balsiger et al., 2018) for the deep learning-based reconstruction of MRF-WF. Second, we investigate the influence of the CNN's spatial receptive field on the reconstruction. Third, we visualize the importance of the temporal frames on the reconstruction.

3.1. Deep Learning-based Reconstruction of MRF-WF

Figure 2 shows the reconstruction of $T1_{H2O}$, FF, and B1 maps of a representative case. Compared to our method, the fingerprint-wise method of (Cohen et al., 2018) produced noisier and the slice-wise with spatial pooling operations method of (Fang et al., 2018) smoothed map reconstructions. Further, we are able to improve over our previous architecture (Balsiger et al., 2018). Quantitatively, our method yielded the best reconstructed MR maps for all evaluation metrics (Table 1).

3.2. Influence of the Receptive Field

To this date, it is unclear how the information of neighboring fingerprints affects the deep learning-based MR map reconstruction. We think that an optimal solution lies between the extreme variants of fingerprint-wise reconstruction and reconstruction including spatial pooling operations with large receptive fields. Therefore, we relied on a patch-wise MR map reconstruction and determined the effect of the receptive field by experimentation. We modified the spatial blocks in our CNN by performing convolutions with a kernel size of 1×1 instead of 3×3 to obtain CNNs with receptive

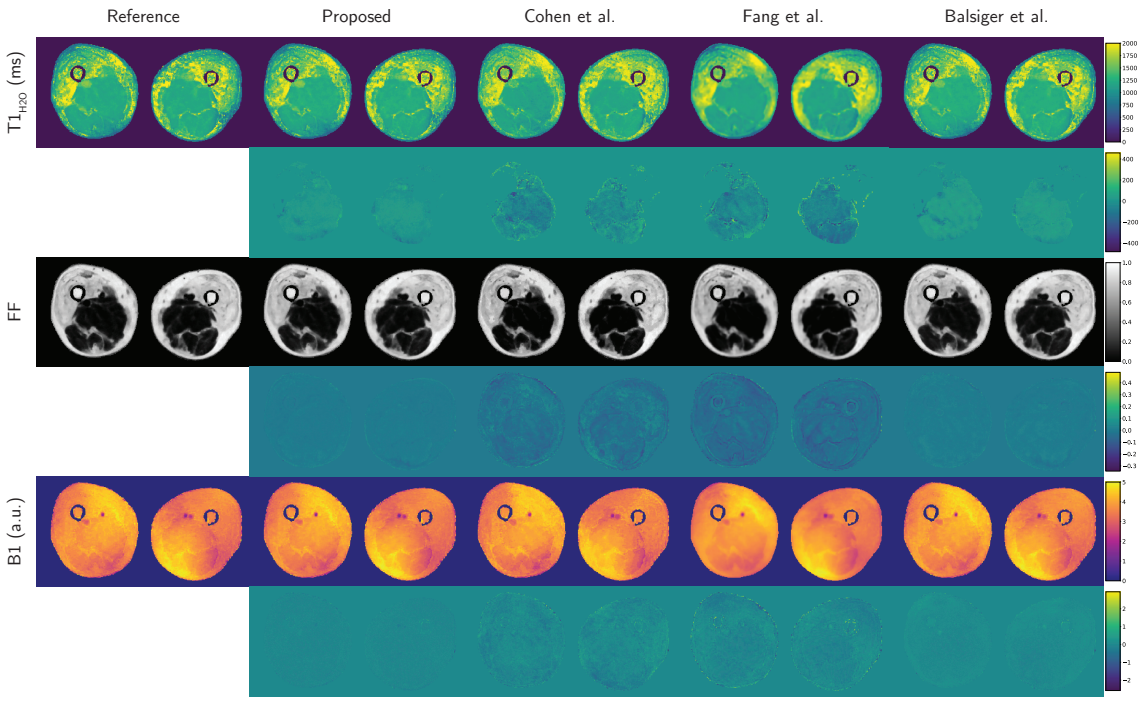

Figure 2: $T1_{H2O}$, FF, and B1 maps of a 73 years old male patient with inclusion body myositis. The reference maps, the reconstructions, and the reconstruction errors (reference minus reconstruction) of our and the compared methods are shown. a.u.: arbitrary unit.

Table 1: Results for the MR map reconstruction.

Metric	**MR map**	**Proposed**	**Cohen et al.**	**Fang et al.**	**Balsiger et al.**
NRMSE	$T1_{H2O}$	**0.024 ± 0.016**	0.065 ± 0.026	0.064 ± 0.028	0.032 ± 0.017
	FF	**0.017 ± 0.008**	0.042 ± 0.005	0.056 ± 0.005	0.021 ± 0.007
	B1	**0.024 ± 0.007**	0.051 ± 0.006	0.062 ± 0.006	0.035 ± 0.006
PSNR	$T1_{H2O}$	**36.16 ± 4.211**	26.53 ± 1.869	26.85 ± 2.557	33.27 ± 3.485
	FF	**35.84 ± 3.038**	27.50 ± 0.969	25.00 ± 0.799	33.90 ± 2.564
	B1	**34.27 ± 1.856**	27.59 ± 0.864	25.88 ± 0.842	31.08 ± 1.135
SSIM	$T1_{H2O}$	**0.978 ± 0.011**	0.931 ± 0.029	0.915 ± 0.040	0.972 ± 0.013
	FF	**0.990 ± 0.007**	0.955 ± 0.024	0.953 ± 0.033	0.985 ± 0.012
	B1	**0.977 ± 0.009**	0.934 ± 0.025	0.926 ± 0.030	0.974 ± 0.008
R2	$T1_{H2O}$	**0.964**	0.888	0.909	0.949
	FF	**0.999**	0.998	0.996	**0.999**
	B1	**0.998**	0.967	0.950	0.993

fields lower than 11×11, and added additional 3×3 convolutions with 64 channels before the last 1×1 convolutional layer of our CNN to obtain CNNs with receptive fields up to 21×21.[3] Figure 3 shows the NRMSE for the $T1_{H2O}$ map reconstruction, which is the most difficult, with varying receptive fields of our CNN. The optimal receptive field for MRF-WF is 11×11, with a statistically significant better reconstruction compared to receptive fields of size 5×5 and lower. Figures for the FF and B1 map reconstruction can be found in the Appendix A.

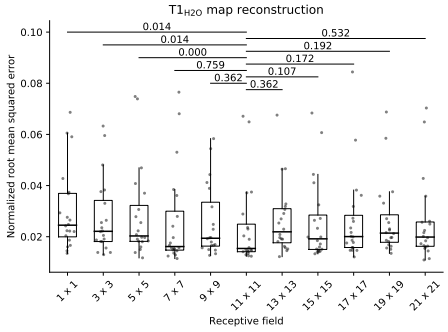

Figure 3: Effect of the receptive field on the $T1_{H2O}$ map reconstruction. Paired Wilcoxon signed rank test with significance level of 0.05 and Bonferroni correction for multiple comparisons.

3.3. Importance of the Temporal Frames

We adopted the occlusion experiments proposed by (Zeiler and Fergus, 2014) for MRF to examine the importance of the temporal frames of the fingerprint for the MR map reconstruction. We applied the occlusion in the temporal domain, which allows for simple yet effective visualization of the importance of the temporal frames. Therefore, we blacked out each t-th temporal frame and reconstructed the MR maps with this occluded MRF image space series. We then considered the absolute difference of the NRMSE to the non-occluded MRF image space series as the importance of the t-th temporal image for the reconstruction. The importance of each temporal frame t is shown in Figure 4(a). We observe that the first temporal frames, after the non-selective inversion pulse, have a high importance for the $T1_{H2O}$ and FF map reconstruction. Additionally, the importance correlates with changes of the MRF-WF sequence parameters over the course of the acquisition (cf. Figure 4(b)). Note that we could also observe the same pattern for the importance when running the occlusion experiment on modified versions of our network.

4. Discussion

We evaluated a CNN consisting of temporal and spatial blocks for the MR map reconstruction from MRF-WF, and compared it to the current state-of-the-art deep learning-based MRF reconstruction

3. We remark that this results in a slightly different number of learnable parameters, which could be compensated. However, we decided not to do so as this would change the CNN architecture.

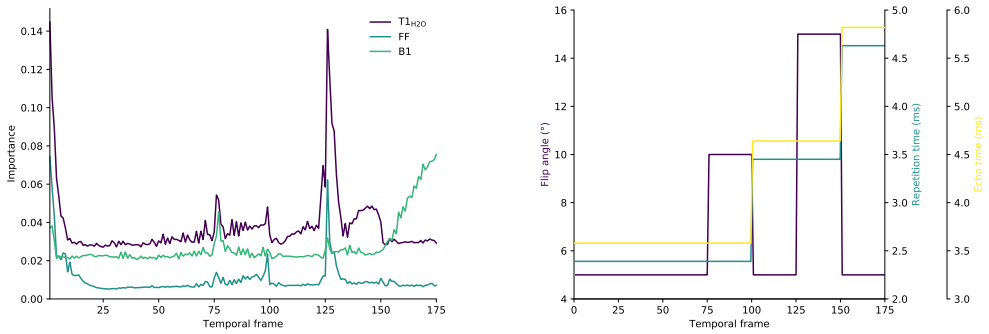

(*a*) The importance for the $T1_{H2O}$, FF, and B1 map reconstruction.

(*b*) The MRF-WF sequence parameters.

Figure 4: The importance of the temporal frames for the MR map reconstruction.

on a large and highly heterogeneous data set consisting only of patient scans. Our proposed CNN outperformed the compared methods quantitatively and qualitatively. We attribute the good performance of our CNN to the trade-off between temporal and spatial feature learning. An indication for the superiority of such a spatiotemporal trade-off is also the noisier or over-smoothed reconstructions of the compared methods.

We conducted an experiment to show the influence of the spatial receptive field on the reconstruction. It is evident that the reconstruction benefits from the information of neighboring fingerprints, however, only to some extent. For MRF-WF, a receptive field larger than 11×11 voxels is not beneficial, which we also attribute to the highly diverse appearance of neuromuscular pathologies. Especially incorporating spatial features from multiple scales as with the method of (Fang et al., 2018) yielded worse reconstructions. For future developments, it is therefore essential to consider the application of the MRF sequence. MRF scans with small or highly diverse pathologies might benefit from the information of neighboring fingerprints but only to a certain extent. We also note that the optimal receptive field is certainly depended on the in-plane voxel size and the *k*-space sampling.

Visualizing the importance of the temporal frames might be a useful tool for further developments of MRF reconstruction, and the MRF sequence development itself. From a MR physical point of view, the first temporal frames should be sensitive to $T1_{H2O}$ due to the applied inversion pulse at the beginning of the MRF-WF acquisition. We could observe a high importance of the first few temporal frames for the $T1_{H2O}$ but also for the FF reconstruction by our occlusion experiment. Further, the high correlation of the importance to changes in the MRF sequence parameters might suggest using a smoother transition of parameter values. However, further work needs to be done to better understand the reconstruction. For instance, it remains unclear from a MR physical point of view, why the last 25 temporal frames are of high importance for the B1 map reconstruction.

5. Conclusion

We proposed a spatiotemporal CNN to reconstruct MR maps from MRF, and showed empirically that a deep learning-based reconstruction benefits from the information between neighboring fingerprints and learns features that can make sense from a MR physical point of view.

Acknowledgments

This research was partially supported by the Swiss National Science Foundation (SNSF) under grant number 184273. The authors thank the NVIDIA corporation for the donation of a GPU and appreciate the valuable discussions with Pierre-Yves Baudin and Alain Jungo.

References

Jakob Assländer, Martijn A. Cloos, Florian Knoll, Daniel K. Sodickson, Jürgen Hennig, and Riccardo Lattanzi. Low Rank Alternating Direction Method of Multipliers Reconstruction for MR Fingerprinting. *Magnetic Resonance in Medicine*, 79(1):83–96, 2018. ISSN 07403194. doi: 10.1002/mrm.26639.

Fabian Balsiger, Amaresha Shridhar Konar, Shivaprasad Chikop, Vimal Chandran, Olivier Scheidegger, Sairam Geethanath, and Mauricio Reyes. Magnetic Resonance Fingerprinting Reconstruction via Spatiotemporal Convolutional Neural Networks. In Florian Knoll, Andreas Maier, and Daniel Rueckert, editors, *Machine Learning for Medical Image Reconstruction*, volume 11074 of *Lecture Notes in Computer Science*, pages 39–46. Springer, Cham, 2018. doi: 10.1007/978-3-030-00129-2_5.

Marco Barbieri, Leonardo Brizi, Enrico Giampieri, Francesco Solera, Gastone Castellani, Claudia Testa, and Daniel Remondini. Circumventing the Curse of Dimensionality in Magnetic Resonance Fingerprinting through a Deep Learning Approach. *arXiv preprint arXiv:1811.11477*, 2018.

Pierre G Carlier, Benjamin Marty, Olivier Scheidegger, Paulo Loureiro de Sousa, Pierre-Yves Baudin, Eduard Snezhko, and Dmitry Vlodavets. Skeletal Muscle Quantitative Nuclear Magnetic Resonance Imaging and Spectroscopy as an Outcome Measure for Clinical Trials. *Journal of Neuromuscular Diseases*, 3(1):1–28, 2016. ISSN 2214-3599. doi: 10.3233/JND-160145.

Stephen F. Cauley, Kawin Setsompop, Dan Ma, Yun Jiang, Huihui Ye, Elfar Adalsteinsson, Mark A. Griswold, and Lawrence L. Wald. Fast Group Matching for MR Fingerprinting Reconstruction. *Magnetic Resonance in Medicine*, 74(2):523–528, 2015. ISSN 07403194. doi: 10.1002/mrm.25439.

Christopher C. Cline, Xiao Chen, Boris Mailhe, Qiu Wang, Josef Pfeuffer, Mathias Nittka, Mark A. Griswold, Peter Speier, and Mariappan S. Nadar. AIR-MRF: Accelerated iterative reconstruction for magnetic resonance fingerprinting. *Magnetic Resonance Imaging*, 41:29–40, 2017. ISSN 0730-725X. doi: 10.1016/J.MRI.2017.07.007.

Ouri Cohen, Bo Zhu, and Matthew S. Rosen. MR fingerprinting Deep RecOnstruction NEtwork (DRONE). *Magnetic Resonance in Medicine*, 80(3):885–894, 2018. ISSN 07403194. doi: 10.1002/mrm.27198.

Zhenghan Fang, Yong Chen, Mingxia Liu, Yiqiang Zhan, Weili Lin, and Dinggang Shen. Deep Learning for Fast and Spatially-Constrained Tissue Quantification from Highly-Undersampled Data in Magnetic Resonance Fingerprinting (MRF). In Yinghuan Shi, Heung-Il Suk, and Mingxia Liu, editors, *Machine Learning in Medical Imaging*, volume 11046 of *Lecture Notes in Computer Science*, pages 398–405. Springer, Cham, 2018. doi: 10.1007/978-3-030-00919-9_46.

Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep Sparse Rectifier Neural Networks. In Geoffrey Gordon, David Dunson, and Miroslav Dudík, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 315–323, Fort Lauderdale, 2011. PMLR.

Mohammad Golbabaee, Dongdong Chen, Pedro A. Gómez, Marion I. Menzel, and Mike E. Davies. Geometry of Deep Learning for Magnetic Resonance Fingerprinting. *arXiv preprint arXiv:1809.01749*, 2018a. URL http://arxiv.org/abs/1809.01749.

Mohammad Golbabaee, Zhouye Chen, Yves Wiaux, and Mike Davies. CoverBLIP: accelerated and scalable iterative matched-filtering for Magnetic Resonance Fingerprint reconstruction. *arXiv preprint arXiv:1810.01967*, 2018b.

Pedro A. Gómez, Miguel Molina-Romero, Cagdas Ulas, Guido Bounincontri, Jonathan I. Sperl, Derek K. Jones, Marion I. Menzel, and Bjoern H. Menze. Simultaneous Parameter Mapping, Modality Synthesis, and Anatomical Labeling of the Brain with MR Fingerprinting. In Sebastien Ourselin, Leo Joskowicz, Mert R. Sabuncu, Gozde Unal, and William Wells, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, pages 579–586, Cham, 2016. Springer International Publishing. ISBN 978-3-319-46726-9. doi: 10.1007/978-3-319-46726-9_67.

Elisabeth Hoppe, Gregor Körzdörfer, Tobias Würfl, Jens Wetzl, Felix Lugauer, Josef Pfeuffer, and Andreas Maier. Deep Learning for Magnetic Resonance Fingerprinting: A New Approach for Predicting Quantitative Parameter Values from Time Series. In R. Röhrig, A. Timmer, H. Binder, and U. Sax, editors, *German Medical Data Sciences: Visions and Bridges*, volume 243, pages 202–206, Oldenburg, 2017. Oldenburg. doi: 10.3233/978-1-61499-808-2-202.

Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely Connected Convolutional Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269. IEEE, 2017. ISBN 978-1-5386-0457-1. doi: 10.1109/CVPR.2017.243.

Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on International Conference on Machine Learning*, Proceedings of Machine Learning Research, pages 448–456, Lille, 2015. PMLR.

Diederik P. Kingma and Jimmy Lei Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations 2015*, pages 1–15, 2015. ISBN 9781450300728. doi: 10.1145/1830483.1830503.

Jane Larkindale, Wenya Yang, Paul F. Hogan, Carol J. Simon, Yiduo Zhang, Anjali Jain, Elizabeth M. Habeeb-Louks, Annie Kennedy, and Valerie A. Cwik. Cost of illness for neuromuscular

diseases in the United States. *Muscle & Nerve*, 49(3):431–438, 2014. ISSN 0148639X. doi: 10.1002/mus.23942.

Dan Ma, Vikas Gulani, Nicole Seiberlich, Kecheng Liu, Jeffrey L Sunshine, Jeffrey L Duerk, and Mark A Griswold. Magnetic resonance fingerprinting. *Nature*, 495(7440):187–192, 2013. ISSN 0028-0836. doi: 10.1038/nature11971.

Benjamin Marty and Pierre G Carlier. Quantification of water T1 and fat fraction in skeletal muscle tissue using an optimal MR fingerprinting radial sequence (MRF-WF). In *International Society for Magnetic Resonance in Medicine*, 2018.

Benjamin Marty, B. Coppa, and Pierre G. Carlier. Fast, Precise, and Accurate Myocardial T1 Mapping Using a Radial MOLLI Sequence With FLASH Readout. *Magnetic Resonance in Medicine*, 79(3):1387–1398, 2018a. ISSN 07403194. doi: 10.1002/mrm.26795.

Benjamin Marty, Bertrand Coppa, and Pierre G. Carlier. Monitoring skeletal muscle chronic fatty degenerations with fast T1-mapping. *European Radiology*, 28(11):4662–4668, 2018b. ISSN 0938-7994. doi: 10.1007/s00330-018-5433-z.

Debra F. McGivney, Eric Pierre, Dan Ma, Yun Jiang, Haris Saybasili, Vikas Gulani, and Mark A. Griswold. SVD Compression for Magnetic Resonance Fingerprinting in the Time Domain. *IEEE Transactions on Medical Imaging*, 33(12):2311–2322, 2014. ISSN 0278-0062. doi: 10.1109/TMI.2014.2337321.

Ilkay Oksuz, Gastao Cruz, James Clough, Aurelien Bustin, Nicolo Fuin, Rene M. Botnar, Claudia Prieto, Andrew P. King, and Julia A. Schnabel. Magnetic Resonance Fingerprinting using Recurrent Neural Networks. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*. IEEE, 2019. URL http://arxiv.org/abs/1812.08155.

Patrick Virtue, Stella X. Yu, and Michael Lustig. Better than real: Complex-valued neural nets for MRI fingerprinting. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 3953–3957. IEEE, 2017. ISBN 978-1-5090-2175-8. doi: 10.1109/ICIP.2017.8297024.

Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. ISSN 1057-7149. doi: 10.1109/TIP.2003.819861.

Matthew D. Zeiler and Rob Fergus. Visualizing and Understanding Convolutional Networks. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, volume 8689 of *Lecture Notes in Computer Science*, pages 818–833. Springer International Publishing, Cham, 2014. doi: 10.1007/978-3-319-10590-1_53.

Appendix A. Influence of the Receptive Field on the FF and B1 Map Reconstruction

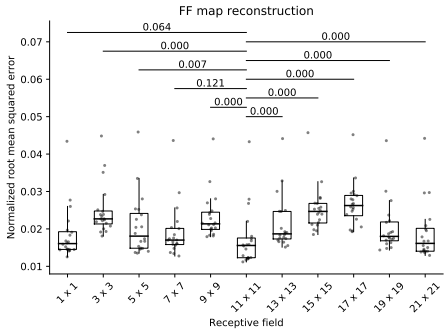

Figure 5: Effect of the receptive field on the FF map reconstruction. Paired Wilcoxon signed rank test with significance level of 0.05 and Bonferroni correction for multiple comparisons.

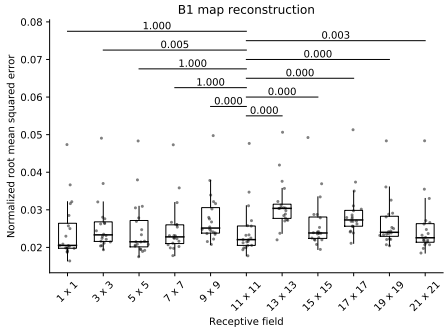

Figure 6: Effect of the receptive field on the B1 map reconstruction. Paired Wilcoxon signed rank test with significance level of 0.05 and Bonferroni correction for multiple comparisons.