



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Geometry of Deep Learning for Magnetic Resonance Fingerprinting

Citation for published version:

Golbabaee, M, Chen, D, Gómez, PA, Menzel, MI & Davies, ME 2019, Geometry of Deep Learning for Magnetic Resonance Fingerprinting. in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*., 18778383 , Institute of Electrical and Electronics Engineers, 44th IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2019, Brighton, United Kingdom, 12/05/19. <https://doi.org/10.1109/ICASSP.2019.8683549>

Digital Object Identifier (DOI):
[10.1109/ICASSP.2019.8683549](https://doi.org/10.1109/ICASSP.2019.8683549)

Link:
[Link to publication record in Edinburgh Research Explorer](#)

Document Version:
Peer reviewed version

Published In:
2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



GEOMETRY OF DEEP LEARNING FOR MAGNETIC RESONANCE FINGERPRINTING

Mohammad Golbabaee¹, Dongdong Chen², Pedro A. Gómez^{3,4}, Marion I. Menzel⁴ and Mike E. Davies².

¹Computer Science Department, University of Bath, ²School of Engineering, University of Edinburgh.

³Munich School of Bioengineering, Technische Universität München. ⁴GE Healthcare, Munich.

ABSTRACT

Current popular methods for Magnetic Resonance Fingerprint (MRF) recovery are bottlenecked by the heavy storage and computation requirements of a dictionary-matching (DM) step due to the growing size and complexity of the fingerprint dictionaries in multi-parametric quantitative MRI applications. In this paper we study a deep learning approach to address these shortcomings. Coupled with a dimensionality reduction first layer, the proposed MRF-Net is able to reconstruct quantitative maps by saving more than 60 times in memory and computations required for a DM baseline. Fine-grid manifold enumeration i.e. the MRF dictionary is only used for training the network and not during image reconstruction. We show that the MRF-Net provides a piece-wise affine approximation to the Bloch response manifold projection and that rather than memorizing the dictionary, the network efficiently clusters this manifold and learns a set of hierarchical *matched-filters* for affine regression of the NMR characteristics in each segment.

Index Terms— Magnetic resonance fingerprinting, inverse problem, deep learning, dictionary, manifold compressed sensing.

1. INTRODUCTION

Magnetic Resonance Fingerprinting (MRF) recently emerged to accelerate acquisition of the *quantitative* NMR characteristics such as the T1, T2 and T2* relaxation times, field inhomogeneity and perfusion [1, 2, 3, 4]. As opposed to mainstream qualitative assessments these *absolute* physical quantities can be used for tissue or pathology identification independent of the scanner or scanning sequences. Unlike conventional quantitative approaches MRF uses i) short and often complicated excitation pulses which encode many NMR parameters simultaneously, and ii) significantly undersampled k-space data. To overcome the lack of sufficient spatio-temporal information MRF incorporates a physical model based on exhaustively simulating a large dictionary of magnetic responses (fingerprints) for all combinations of the quantized NMR parameters. This dictionary is then used for matched-filtering in a model-based reconstruction scheme e.g. [5]. As occurs to any multi-parametric manifold enumeration, the main drawback of such approach is the size of this dictionary which grows exponentially in terms of the number of parameters and their quantization resolution; a serious (non-scalability) limitation of the current methods to be applicable in the emerging multi-parametric MRF applications. In conjunction with the widespread applications of machine learning methodologies, a number of recent empirical studies have proposed a dictionary-free deep learning (DL) approach to address this shortcoming [6, 7, 8, 9] the crux of which is to bypass the DM step by using compact deep

This work is funded by the EPSRC grant EP/M019802/1, ERC C-SENSE project (ERCADG-2015-694888), and the Scottish Research Partnership in engineering (SPRe) Award PECRE1718/18.

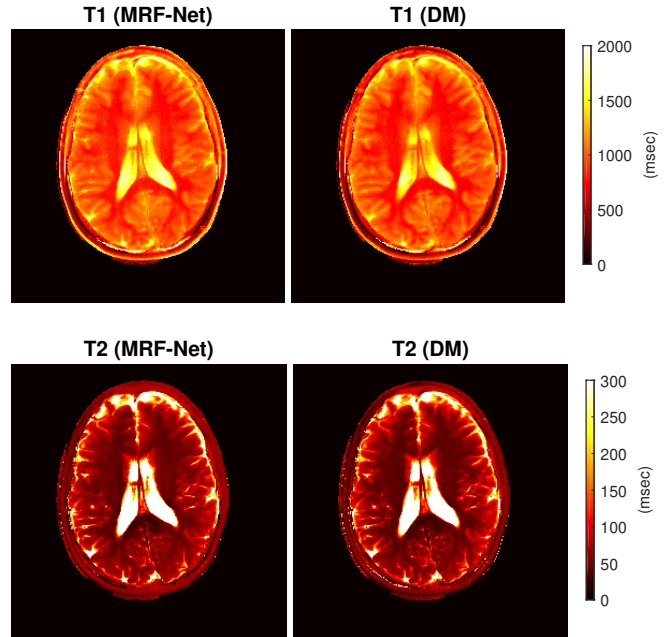


Fig. 1: Reconstructed T1 and T2 maps using the proposed dictionary-less MRF-Net and dictionary matching (DM) baseline.

neural networks. However, reasons that DL works so well for this problem are poorly understood.

This paper aims at uncovering the underlying mechanisms by which DL achieves such progress for the MRF framework from a geometrical point of view. We show that the MRF-Net provides a piece-wise affine approximation to the Bloch response manifold projection and that rather than memorizing the dictionary, the network efficiently clusters this manifold layer-by-layer and implicitly learns a set of hierarchical *matched-filters* for affine parameter regression in each segment. Besides, we propose a competitive architecture to the current DL baselines. The proposed MRF-Net features a (unsupervisedly learned) dimensionality reduction first layer which promotes a low-rank subspace prior during image reconstruction, and results in less units and training resources as required for the uncompressed DL approaches earlier proposed for this problem. Our *in-vivo* experiment for estimating two NMR quantities (i.e. a small-size MRF problem) indicates that the MRF-Net is capable of saving more than 60x in storage and model-fitting computations as required for a dimension-reduced DM baseline (Figure 1).

2. PROBLEM STATEMENT

MRF acquisitions follow a linear spatio-temporal model:

$$Y \approx \mathcal{A}(X), \quad (1)$$

where $Y \in \mathbb{C}^{m \times L}$ denotes noisy k-space measurements collected at $t = 1, \dots, L$ temporal frames after each excitation. The MRF image sequence is a complex-valued matrix X of spatio-temporal resolution $n \times L$ i.e. n spatial voxels and L temporal frames. The forward operator $\mathcal{A} := P_{\Omega} F S(\cdot)$ models multi-coil sensitivity maps S and a sub-sampled Fourier operator $P_{\Omega} F$ which represents the k-space acquisition with respect to a set of *temporally-varying* locations $\Omega = \bigcup_{t=1}^L \Omega_t$ where $\text{Card}(\Omega_t) = m \ll n$.

The main source of quantitative measurements are the per-voxel magnetization response of proton dipoles obtained from dynamic rotations of the external magnetic field i.e. a sequence of Flip Angles (FA) $\{\alpha_t\}_{t=1}^L$ applying at certain repetition times $\{TR_t\}$. *Tissues with different NMR characteristics respond distinctively to these excitations.* The MRF framework relies on this principle to regularize the under-determined problem (1) by a temporal model and enable parameter estimation. Magnetization trajectories (responses) —denoted by $\mathcal{B}(\Theta; TR, \alpha) \in \mathbb{C}^L$ —are distinct solutions of the *Bloch differential equations* for a given set of intrinsic NMR parameters $\Theta \in \mathbb{R}^P$ and excitation sequence $\{\alpha_t, TR_t\}$ [10]. Current MRF approaches discretize through a dense sampling the parameter space $[\Theta] = [T1] \times [T2] \times \dots$ and simulate a large dictionary of normalized fingerprints $D = \{D_j\}_{j=1}^d$ where,

$$D_j := \mathcal{B}([\Theta_j]; TR, \alpha) \quad \forall j = 1, \dots, d, \quad (2)$$

for all d combinations of the quantized parameters. Under the *voxel purity* assumption each spatial voxel of the MRF image corresponds to a unique NMR parameter and would approximately match to a temporal trajectory in the fingerprint dictionary: $X_v \in D \quad \forall v = 1, \dots, n$, where X_v denotes the normalized v -th row of X i.e. a multi-dimensional spatial voxel.

3. PARAMETER ESTIMATION

A popular approach for parameter estimation is to perform back-projection (adjoint operator) on the k-space data $\hat{X} := \mathcal{A}^H(Y) \in \mathbb{C}^{n \times L}$ followed by dictionary matching to identify the highest correlated atom and its corresponding NMR parameters for each (normalized) voxel of the highly aliased back-projected image \hat{X} :

$$[\Theta_v] = \text{NNS}_D(\hat{X}_v), \quad \forall v = 1, \dots, n. \quad (3)$$

Here $\text{NNS}_D(x) := \text{argmin}_j \|x - D_j\|_2$ denotes the nearest neighbour search which serves as a Euclidean projection onto the discrete set of fingerprints i.e. the manifold of Bloch Eq solutions. A temporal (subspace) compression can be used to shrink the search dimension i.e. $\tilde{X}_v := V_s^H \hat{X}_v, \tilde{D}_j := V_s^H D_j$ across the $s \leq L$ dominant principal components of $DD^H \approx V_s \Lambda V_s^H$ [11], and to promote a low-rank subspace prior during image reconstruction [12, 13]. However, enumerating the multi-parametric MRF manifold in order for (3) to be an accurate projection introduces an exponentially growing complexity (in terms of P) to the storage and computations needed for conducting NNS. A recent line of research [14, 15] shows that certain tree search strategies can benefit from the *low intrinsic dimensionality* of the MRF manifold and significantly accelerate the matching step. However storage of the dictionary or the corresponding tree still remains a big challenge for fine-grid enumerations.

3.1. MRF-Net

In this study we propose training a 4-layer fully connected feed-forward network dubbed as the MRF-Net for approximating the




Fig. 2: Illustration of the MRF-Net: pre-trained linear dimension reduction used for the first layer. Three last layers use ReLU nonlinearities (orange) and are trained by the standard back-propagation. MRF manifold projection by a continuous mapping $\mathcal{F} : \mathbb{C}^L \rightarrow \mathbb{R}^P$:

$$\Theta_v = \mathcal{F}(\hat{X}_v), \quad (4)$$

where $\mathcal{F}(x) \approx \text{argmin}_{\Theta} \|x - \mathcal{B}(\Theta)\|_2$. The first layer of MRF-Net unsupervisedly learns the best linear projection onto the subspace of clean fingerprints through principal component analysis, and it is kept fixed during training other layers. Three other layers include nonlinear ReLU activations in order to approximate the dimension-reduced projection function $\text{argmin}_{\Theta} \|V_s^H x - V_s^H \mathcal{B}(\Theta)\|_2$. The size of MRF-Net is $1000 - 10 - 200 - 30 - 2$ including an input layer $L = 1000$ fed with voxel sequences from the back-projected images, and 4 hidden layers as shown in Figure 2. Dimensions of the input/output and hidden units are customized here for the Steady State Precession (FISP) sequence [2] which encodes $P = 2$ NMR characteristics i.e. $\Theta = \{T1, T2\}$ relaxation times. The MRF dictionary corresponding to the FISP sequence is shown to be well represented by very few principal components [11] e.g. $s = 10$ here, which determines first layer’s dimension accordingly. Thanks to this dimensionality-reduction, MRF-Net requires far less units and training resources compared to the uncompressed DL approaches proposed earlier in [6, 7].

3.1.1. Training MRF-Net

Fine-grid manifold enumeration i.e. the MRF dictionary is only used for training and not during image reconstruction. To avoid losing discrimination between fingerprints —by the magnitude-only data treatment proposed in [6]—we adopt a phase-alignment heuristic used in practice [16, 17] to align dictionary atoms (for training) and back-projected images (the input). With this treatment we can assume that without loss of generality the MRF-Net consists of real-valued parameters and approximates a real-valued mapping.

A fully connected feed-forward network is composed of N (here $N = 4$) layers each applying an affine transform followed by nonlinear activation functions. The network can be modelled as $\mathcal{F} \equiv h^{(N)}(x) : \mathbb{R}^L \rightarrow \mathbb{R}^P$ through a hierarchy of hidden mappings $h^{(i)}(x) : \mathbb{R}^L \rightarrow \mathbb{R}^{\dim(i)}$ initialized by inputs $h^{(0)}(x) = x$:

$$h^{(i)}(x) = f_i \left(W^{(i)} h^{(i-1)}(x) + \beta^{(i)} \right) \quad \text{for } i = 1, 2, \dots, N. \quad (5)$$

Here $W^{(i)} \in \mathbb{R}^{\dim(i) \times \dim(i-1)}$ and $\beta^{(i)} \in \mathbb{R}^{\dim(i)}$ are the weights and biases at the i -th layer and $f_i(\cdot)$ is an element-wise nonlin-

ear activation function. The first Layer of MRF-Net consists of identity activation $f_1(u) = u$, and unsupervised trained parameters $W^{(1)} = V_s^H, \beta^{(1)} = \mathbf{0}$. The remaining layers use ReLU activations $f_i(u) = \max(u, 0)$ for $i = 2, 3, 4$. To learn these layers, MRF-Net is supervised trained to minimize the mean-squared regression loss between $h_k^{(4)}, \Theta_k$.¹ Training inputs are dimension-reduced (phase-aligned) atoms of the fine-grid MRF dictionary \tilde{D}_k corrupted by zero-mean independent Gaussian noises ξ_k with SNRs randomly selected between 40-60 dB. We use the Extended Phase Graph framework [19] to simulate Bloch Eq responses to the FISP sequence for all combinations of $T1=[100:10:4000]$ (msec) and $T2=[20:2:600]$ (msec), and build a dictionary with $d = 113781$ atoms for training. After noise corruption (i.e. data augmentation by factor 100) we perform NNS searches to find correct training labels Θ_k (and not those originally generated the fingerprints):

$$\Theta_k := \text{NNS}_{\tilde{D}}(\tilde{D}_k + \xi_k), \quad (6)$$

which enables learning a projection mapping rather than a denoiser.

4. GEOMETRY OF DEEP LEARNING FOR BLOCH RESPONSE MANIFOLD PROJECTION

In this part we show that the MRF-Net provides a piece-wise affine approximation to the Bloch response manifold projection. Our analysis is inspired by the recent work [20] and is in relation to the authors' previous works [14, 16] on adopting cover tree data structures to cluster dictionary atoms in hierarchical segments and using fast approximate NNS search for Dictionary Matching (DM). We show that MRF-Net also clusters the input space, however as opposed to the cover trees, the network does not memorize the MRF dictionary but rather efficiently encodes a set of deep (hierarchical) *matched-filters* for affine regression of the NMR quantities in each segment. Further, segments here have piece-wise affine boundaries.

4.1. Affine spline function approximation

For a network composed of piecewise linear activation functions such as ReLU and the (linear) identity activation, we have that:

Remark 1. Each layer's output $h^{(i)}$ is a piece-wise affine transformation of its direct input $h^{(i-1)}$. Composition of such layers gives mappings $h^{(i)}(x) : \mathbb{R}^L \rightarrow \mathbb{R}^{\dim(i)}$ which $\forall i$ are piece-wise affine transformations of the input $h^{(0)} = x$ (see e.g. [21]). Further, using continuous activation functions (as above) and for bounded $\{W^{(i)}, \beta^{(i)}, i\}$, we have that $h^{(i)}$ is *Lipschitz continuous*.

In MRF-NET the last non-linearity is mainly used to impose non-negativity of the estimated parameters, and therefore most of the prediction task is done by the preceding layers. We denote by

$$z^{(N)}(x) := W^{(N)}h^{(N-1)} + \beta^{(N)} \quad (7)$$

as the *weighted outputs* before the last non-linearity. We have the following *affine spline* representation for the weighted outputs [20]:

$$z^{(N)}(x) = \mathbf{A}[x]x + [x] := \sum_r (A_r x + b_r) \mathbf{1}_{\Omega_r}(x), \quad (8)$$

where $\mathbf{1}_{\Omega_r}(x)$ is the indicator function with respect to a segment $\Omega_r \in \mathbb{R}^L$, returning x if it belongs to the segment and 0 otherwise

¹We use Adam optimizer [18] where the gradient updates are computed by the standard back-propagation on a moderate CPU desktop. Optimization parameters are as follows: batch size 50, 30 epochs and the step-size 10^{-2} decaying at the rate of 0.8 after each epoch.

—segments form a disjoint partitioning of the input space. Matrices $A_r \in \mathbb{R}^{P \times L}$ and vectors $b_r \in \mathbb{R}^P$ define the corresponding input-output affine mapping for each segment. We use the short-hands $\mathbf{A}[x], [x]$ to represent the input-dependent (piece-wise affine) mapping of $z^{(N)}(x)$: P input-dependent offsets represented by $[x]$ and similarly, an input-dependent $P \times L$ matrix $\mathbf{A}[x]$ where each row represents a *matched-filter* (acting on x and measures the mutual correlation) corresponding to a certain output coordinate $p = 1, \dots, P$.

Match-filters and offsets here are used for *regressing* the outputs —not to be confused with the MRF dictionary matching (DM). In other word, during training the network $h^{(N)}(x)$ learns $\{W^{(i)}, \beta^{(i)}\}$ or equivalently $\{\mathbf{A}[x], [x]\}$ to provide a continuous and piece-wise affine *approximation* between input and output e.g. the $\text{NNS}_D(\cdot)$ function (Bloch manifold projection) for the MRF reconstruction problem. The universal approximation theorem [22] states that a 2-layer shallow network with large enough units can provide an arbitrarily close approximation to any Borel-measurable function. Deeper networks however are often more favourable in practice to efficiently reduce the number of hidden units [23]. For certain manifold embedding tasks this has been proven e.g. in [24, 25, 26, 27].

4.2. Visualizing MRF-Net's segments on Bloch manifold

Finite-sized networks with bounded weights and biases introduce piecewise affine boundaries for their corresponding segments Ω_r [21]. It is easy to verify that each layer i in (5) with aforementioned non-linearities introduces segments with piecewise affine boundaries in its direct input space $h^{(i-1)} \in \mathbb{R}^{\dim(i-1)}$. Further each segment corresponds to an affine transformation of $h^{(i-1)}$. Therefore, composition of such layers results in piecewise affine segments Ω_r in the input space for the function $z^{(N)}(x)$.

Remark 2. Continuity of the mapping $z^{(N)}(x)$ implies that adjacent segments $\Omega_r, \Omega_{r'}$ correspond to distinct $A_r, A_{r'}$. Indeed, if $A_r = A_{r'}$ and the only difference is in the offsets $b_r \neq b_{r'}$, then $\Omega_r, \Omega_{r'}$ won't intersect on boundaries. Therefore they are not adjacent segments unless contradicting the continuity assumption.

This remark gives an idea for visualizing the segments as follows: for densely sampled input signals x , we compute gradients of the weighted output with respect to the input. The gradients determine input-dependant slopes in affine spline formulation (8) i.e. rows of $\mathbf{A}[x]$ at a point x are populated as follows $\forall p = 1, 2, \dots, P$:²

$$\begin{aligned} \mathbf{A}[x]_{(p,\cdot)} &= \left(\nabla_x z_p^{(N)}(x) \right)^T \\ &:= \left[\frac{\partial z_p^{(N)}(x)}{\partial x_1}, \frac{\partial z_p^{(N)}(x)}{\partial x_2}, \dots, \frac{\partial z_p^{(N)}(x)}{\partial x_L} \right]. \end{aligned} \quad (9)$$

For a given input x the gradients in (9) can be efficiently calculated using *back-propagation*. We feed forward x to identify all activations $z^{(i)}, h^{(i)}$ and then follow the recursion starting at the vector $y = \mathbf{e}_p$:

$$\begin{aligned} y &\leftarrow f'_i(z^{(i-1)}) \odot \left(\left(W^{(i)} \right)^T y \right) \quad i = N, N-1, \dots, 2. \\ \nabla_x z_p^{(N)}(x) &= \left(W^{(1)} \right)^T y, \end{aligned} \quad (10)$$

²For a matrix X we denote by $X_{(p,\cdot)}$ as its p -th row. For a vector x , x_p denotes its p -th element. We also later denote by \mathbf{e}_p the coordinate vector whose p -th element is one and zero elsewhere.




Fig. 3: MRF-Net’s segments on the manifold of Bloch responses to the FISP sequence depicted across the three dominant principal components (left), and the corresponding segments on the T_1, T_2 grid used for generating this dictionary (right).

where \odot denotes the element-wise vector product, and $f'_i(\cdot)$ denotes the (element-wise) derivative of the activation function. For an identity activation $f'_i(\cdot)$ is an all-one vector and for the ReLU activation it returns one for the positive (direct) input coordinates and zero elsewhere. By vector quantization (e.g. k-means clustering) we cluster regions of x which output distinct slopes A_r and identify Ω_r . We note that for a classification problem [20] used a similar idea to separately identify each hidden layer’s segments and then intersect them in a hierarchical (layer-by-layer) fashion to get Ω_r .

Fine-sampling could be used to visualize intersection of the MRF-Net’s segments with the Bloch response manifold. For this purpose we compute (9) for inputs corresponding to a dense sample of T_1, T_2 grid i.e. the MRF dictionary. Figure 3 visualizes dominant MRF-Net’s segments on the manifold of Bloch responses to the FISP sequence used in our experimental validations.

4.3. Deep matched-filtering

The identity (8) interprets how data is treated by our network. Ignoring the offsets, the (two) rows of $\mathbf{A}[x]$ correlate with inputs belonging to a segment and predict T_1 and T_2 . Each segment Ω_r of the input space has a distinct set of (two) matched-filters i.e. rows of A_r , whose correlations with input data (belonging to that segment) will linearly regress the outputs. The end-to-end matched-filtering parameters (i.e. $\{A_r, \Omega_r, b_r\}$) are implicitly learned during training the network and learning the corresponding layer-by-layer affine transformations. In Figure 4 we choose two input segments that include standard T_1, T_2 values measured for the White (WM) and Gray (GM) Matters in healthy volunteers’ brains [28]. We show the clean magnetic responses (i.e. fingerprints) associated with that region together with the matched-filters used for predicting T_1 and T_2 quantities. As can be seen, *matched-filters peak at discriminant parts of the fingerprints* that is where the sequence encodes a significant amount of information (i.e. sensitivity) about the underlying NMR characteristics in that segment. These peaks are visible in the beginning of the FISP sequence due to using an Inversion Recovery (i.e. rotating the external magnetic field by 180°) and they repeat because of the periodic pattern of flip angles used for FISP excitations (see [2, Figure 1.b]).

5. IN-VIVO EXPERIMENT

An *in-vivo* MRF dataset was acquired using the Steady State Precession (FISP) sequence in [2] and spiral readouts which sample $m = 732$ k-space locations in each of the $L = 1000$ time-frames in order to reconstruct $n = 256 \times 256$ resolution parametric T_1 and T_2 maps. Other scanning parameters are $TE/T_{inv}=2/18$




Fig. 4: FISP dictionary fingerprints (left column) for two segments of the MRF-Net including the standard characteristics measured for Gray (top row) and White (second row) Matters in healthy volunteers’ brains. Right column shows the implicit matched-filters that MRF-Net builds for regressing T_1, T_2 quantities in each segment.

msec, 8 head-coils, 3 Tesla GE HDx MRI system (GE Medical Systems, Milwaukee, WI), variable density spiral sampling with 89 interleaves, $22.5 \times 22.5 \text{cm}^2$ FOV, 256×256 voxel spatial resolution with 5mm slice thickness. As discussed in Section 3.1.1, we simulate a dictionary of $d = 113781$ atoms which finely samples the grid $T_1 \in [100, 4000]$ (msec), $T_2 \in [20, 600]$ (msec). The baseline DM scheme requires direct access to this dictionary for NNS searches whereas the proposed DL approach only uses it for data augmentation and training the MRF-Net. Figure 1 compares the reconstructed parametric maps using DM with brute-force searches and the proposed MRF-Net. Note that the computation-memory complexity of a dimension-reduced dictionary matching —without a fast tree search—is $\mathcal{O}(snL + snd)$ which in this example is more than 60 times higher than the requirements of the (dimension-reduced) MRF-Net. This comparison is on a moderate-size MRF dictionary encoding only two parameters and we expect that for the emerging applications and dictionaries encoding a large number of intrinsic NMR characteristics e.g. T_2^* , field inhomogeneity, perfusion, diffusion, etc, this gap substantially grows. We leave this direction for further future investigations.

6. CONCLUSION

In this paper we study a dictionary-less deep learning approach for the MRF reconstruction problem. Featuring a subspace compression in its first layer, the proposed MRF-Net is compact, easily trained and is capable of achieving comparable estimation accuracy to a dimension-reduced DM baseline, however, with 60 times less storage and computations. The MRF dictionary is only used for training and not during image reconstruction. We show that the MRF-Net provides a piece-wise affine approximation to the Bloch response manifold projection through which, the network efficiently clusters the input space and learns hierarchical *matched-filters* for affine regression of the quantitative parameters in each segment. Future directions could extend this work to applications with a large number of intrinsic NMR characteristics as well as incorporating

spatial regularities by e.g. using convolutional networks [9].

References

- [1] D. Ma, V. Gulani, N. Seiberlich, K. Liu, J. Sunshine, J. Durek, and M. Griswold, "Magnetic resonance fingerprinting," *Nature*, vol. 495, no. 7440, pp. 187–192, 2013.
- [2] N. Jiang Y, D. Ma, N. Seiberlich, V. Gulani, and M. Griswold, "MR fingerprinting using fast imaging with steady state precession (fisp) with spiral readout," *Magnetic resonance in medicine*, vol. 74, no. 6, pp. 1621–1631, 2015.
- [3] K. L. Wright, Y. Jiang, D. Ma, D. C. Noll, M. A. Griswold, V. Gulani, and L. Hernandez-Garcia, "Estimation of perfusion properties with mr fingerprinting arterial spin labeling," *Magnetic resonance imaging*, vol. 50, pp. 68–77, 2018.
- [4] B. Rieger, F. Zimmer, J. Zapp, S. Weingärtner, and L. R. Schad, "Magnetic resonance fingerprinting using echo-planar imaging: Joint quantification of T1 and T2* relaxation times," *Magnetic resonance in medicine*, vol. 78, pp. 1724–1733, 2017.
- [5] M. Davies, G. Puy, P. Vandergheynst, and Y. Wiaux, "A compressed sensing framework for magnetic resonance fingerprinting," *SIAM Journal on Imaging Sciences*, vol. 7, no. 4, pp. 2623–2656, 2014.
- [6] O. Cohen, B. Zhu, and M. S. Rosen, "MR fingerprinting deep reconstruction network (DRONE)," *Magnetic resonance in medicine*, vol. 80, no. 3, pp. 885–894, 2018.
- [7] P. Virtue, X. Y. Stella, and M. Lustig, "Better than real: Complex-valued neural nets for mri fingerprinting," in *Image Processing (ICIP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 3953–3957.
- [8] E. Hoppe, G. Kördörfer, T. Würfl, J. Wetzl, F. Lugauer, J. Pfeuffer, and A. Maier, "Deep learning for magnetic resonance fingerprinting: A new approach for predicting quantitative parameter values from time series." *Studies in health technology and informatics*, vol. 243, p. 202, 2017.
- [9] F. Balsiger, A. S. Konar, S. Chikop, V. Chandran, O. Scheidegger, S. Geethanath, and M. Reyes, "Magnetic resonance fingerprinting reconstruction via spatiotemporal convolutional neural networks," in *International Workshop on Machine Learning for Medical Image Reconstruction*. Springer, 2018, pp. 39–46.
- [10] E. Jaynes, "Matrix treatment of nuclear induction," *Physical Review*, vol. 98, no. 4, p. 1099, 1955.
- [11] D. F. McGivney, E. Pierre, D. Ma, Y. Jiang, H. Saybasili, V. Gulani, and M. A. Griswold, "SVD compression for magnetic resonance fingerprinting in the time domain," *IEEE transactions on medical imaging*, vol. 33, no. 12, pp. 2311–2322, 2014.
- [12] J. Assländer, M. A. Cloos, F. Knoll, D. K. Sodickson, J. Hennig, and R. Lattanzi, "Low rank alternating direction method of multipliers reconstruction for mr fingerprinting," *Magnetic resonance in medicine*, vol. 79, no. 1, pp. 83–96, 2018.
- [13] B. Zhao, K. Setsompop, E. Adalsteinsson, B. Gagoski, H. Ye, D. Ma, Y. Jiang, P. Ellen Grant, M. A. Griswold, and L. L. Wald, "Improved magnetic resonance fingerprinting reconstruction with low-rank and subspace modeling," *Magnetic resonance in medicine*, vol. 79, no. 2, pp. 933–942, 2018.
- [14] M. Golbabaee, Z. Chen, Y. Wiaux, and M. E. Davies, "Cover tree compressed sensing for fast MR fingerprint recovery," in *2017 IEEE 27th International Workshop on Machine Learning for Signal Processing (MLSP)*, Sept 2017, pp. 1–6.
- [15] M. Golbabaee and M. E. Davies, "Inexact gradient projection and fast data driven compressed sensing," *IEEE Transactions on Information Theory*, vol. 64, no. 10, pp. 6707–6721, 2018.
- [16] M. Golbabaee, Z. Chen, Y. Wiaux, and M. Davies, "Coverblip: accelerated and scalable iterative matched-filtering for magnetic resonance fingerprint reconstruction," *arXiv preprint arXiv:1810.01967*, 2018.
- [17] C. C. Cline, X. Chen, B. Mailhe, Q. Wang, J. Pfeuffer, M. Nitka, M. A. Griswold, P. Speier, and M. S. Nadar, "Air-mrf: Accelerated iterative reconstruction for magnetic resonance fingerprinting," *Magnetic resonance imaging*, vol. 41, pp. 29–40, 2017.
- [18] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [19] M. Weigel, "Extended phase graphs: Dephasing, RF pulses, and echoes-pure and simple," *Journal of Magnetic Resonance Imaging*, vol. 41, no. 2, pp. 266–295, 2015.
- [20] R. Balestriero and richard baraniuk, "A spline theory of deep learning," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 10–15 Jul 2018, pp. 374–383.
- [21] G. F. Montufar, R. Pascanu, K. Cho, and Y. Bengio, "On the number of linear regions of deep neural networks," in *Advances in neural information processing systems*, 2014, pp. 2924–2932.
- [22] G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Mathematics of control, signals and systems*, vol. 2, no. 4, pp. 303–314, 1989.
- [23] O. Delalleau and Y. Bengio, "Shallow vs. deep sum-product networks," in *Advances in Neural Information Processing Systems*, 2011, pp. 666–674.
- [24] R. Basri and D. Jacobs, "Efficient representation of low-dimensional manifolds using deep networks," *arXiv preprint arXiv:1602.04723*, 2016.
- [25] U. Shaham, A. Cloninger, and R. R. Coifman, "Provable approximation properties for deep neural networks," *Applied and Computational Harmonic Analysis*, 2016.
- [26] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: going beyond euclidean data," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 18–42, 2017.
- [27] D. Chen, J. Lv, and Y. Zhang, "Unsupervised multi-manifold clustering by learning deep representation," in *AAAI Workshops*, 2017.

- [28] J. P. Wansapura, S. K. Holland, R. S. Dunn, and W. S. Ball Jr, "Nmr relaxation times in the human brain at 3.0 tesla," *Journal of Magnetic Resonance Imaging*, vol. 9, no. 4, pp. 531–538, 1999.