



Compressive MR Fingerprinting Reconstruction with Neural Proximal Gradient Iterations

Dongdong Chen¹, Mike E. Davies¹, and Mohammad Golbabaee²(✉)

¹ School of Engineering, University of Edinburgh, Edinburgh, UK
{d.chen,mike.davies}@ed.ac.uk

² Computer Science department, University of Bath, Bath, UK
mg2105@bath.ac.uk

Abstract. Consistency of the predictions with respect to the physical forward model is pivotal for reliably solving inverse problems. This consistency is mostly un-controlled in the current end-to-end deep learning methodologies proposed for the Magnetic Resonance Fingerprinting (MRF) problem. To address this, we propose PGD-NET, a learned proximal gradient descent framework that directly incorporates the forward acquisition and Bloch dynamic models within a recurrent learning mechanism. The PGD-NET adopts a compact neural proximal model for de-aliasing and quantitative inference, that can be flexibly trained on scarce MRF training datasets. Our numerical experiments show that the PGD-NET can achieve a superior quantitative inference accuracy, much smaller storage requirement, and a comparable runtime to the recent deep learning MRF baselines, while being much faster than the dictionary matching schemes. Code has been released at <https://github.com/edongdongchen/PGD-Net>.

Keywords: Magnetic resonance fingerprinting · Compressed Sensing · Deep learning · Learned proximal gradient descent

1 Introduction

Magnetic resonance fingerprinting (MRF) is an emerging technology that enables simultaneous quantification of multitudes of tissues' physical properties in short and clinically feasible scan times [20]. Iterative reconstruction methods based on Compressed Sensing (CS) have proven efficient to help MRF overcome the challenge of computing accurate quantitative images from the undersampled k-space measurements taken in aggressively short scan times [3, 10, 11, 30]. However, these methods require dictionary matching (DM) that is non-scalable and can create enormous storage and computational overhead. Further, such approaches often do not fully account for the joint spatiotemporal structures of the MRF data which can lead to poor reconstructions [14].

Deep learning methodologies have emerged to address DM’s computational bottleneck [9, 15, 24, 27], and in some cases to perform joint spatiotemporal MRF processing through using convolutional layers [4, 8, 12, 13, 17, 18, 26]. These models are trained in an end-to-end fashion *without* an explicit account for the known physical acquisition model (i.e. the forward operator) and a mechanism for explicitly enforcing measurement consistency according to this sampling model which can be crucial in the safety-first medical applications. Further, ignoring the structure of the forward model could lead to building unnecessary large inference models and possible overfitted predictions, especially for the extremely *scarce* labelled anatomical quantitative MRI datasets that are available for training.

Our Contributions: we propose PGD-NET a deep convolutional model that is able to learn and perform robust spatiotemporal MRF processing, and work with limited access to the ground-truth (i.e. labelled) quantitative maps. Inspired by iterative proximal gradient descent (PGD) methods for CS reconstruction [22], we adopt learnable, compact and shared convolutional layers within a data-driven proximal step, meanwhile explicitly incorporating the acquisition model as a non-trainable gradient step in all iterations. The proximal operator is an auto-encoder network whose decoder embeds the Bloch magnetic responses and its convolutional encoder embeds a de-aliasing projector to the tissue maps’ quantitative properties. Our work is inspired by recent general CS methodologies [1, 2, 7, 22, 25] that replace traditional hand-crafted image priors by deep data-driven models. To the best of our knowledge, this is the first work to adopt and investigate the feasibility of such an approach for solving the MRF inverse problem.

2 Methodology

MRF adopts a linear spatiotemporal compressive acquisition model:

$$\mathbf{y} = \overline{\mathbf{H}}(\overline{\mathbf{x}}) + \xi \quad (1)$$

where $\mathbf{y} \in \mathbb{C}^{m \times L}$ are the k-space measurements collected at L temporal frames and corrupted by some noise ξ , and $\overline{\mathbf{x}} \in \mathbb{C}^{n \times L}$ is the Time-Series of Magnetisation Images (TSMI) with n voxels across L timeframes. The forward operator $\overline{\mathbf{H}} : \mathbb{C}^{n \times L} \rightarrow \mathbb{C}^{m \times L}$ models Fourier transformations subsampled according to a set of temporally-varying k-space locations in each timeframe. Accelerated MRF acquisition implies working with heavily under-sampled data $m \ll n$, which makes $\overline{\mathbf{H}}$ becomes ill-posed for the inversion.

Bloch Response Model: Per-voxel TSMI temporal signal evolution is related to the quantitative NMR parameters/properties such as $\{T1_v, T2_v\}$ relaxation times, through the solutions of the *Bloch differential equations* $\overline{\mathbf{x}}_v \approx \rho_v \overline{\mathbf{B}}(T1_v, T2_v)$, scaled by the ρ_v proton density (PD) in each voxel v [19, 20].

The Subspace Dimension-Reducing Model: In many MRF applications (including ours) a low $s \ll L$ dimensional subspace $V \in \mathbb{C}^{L \times s}$ embeds the Bloch solutions $\mathcal{B}(\cdot) \approx VV^H \mathcal{B}(\cdot)$. This subspace can be computed through

PCA decomposition of the MRF dictionary [23], and enables re-writing (1) in a compact form that is beneficial to the storage, runtime and accuracy of the reconstruction [3, 15]:

$$\mathbf{y} = \mathbf{H}(\mathbf{x}) + \xi \quad \text{where, for each voxel } \mathbf{x}_v \approx \rho_v \mathcal{B}(T1_v, T2_v), \quad (2)$$

and $\mathbf{x} \in \mathbb{C}^{n \times s}$ is the dimension-reduced TSMI, $\mathbf{H} := \overline{\mathbf{H}} \circ V$ and $\mathcal{B} = V^H \overline{\mathcal{B}}$ denotes the subspace-compressed Bloch solutions (for more details see [14]).

Tissue Quantification: Given the compressed measurements \mathbf{y} , the goal of MRF is to solve the inverse problem (2) and to compute the underlying multi-parametric maps $\mathbf{m} = \{T1, T2, \rho\}$ (and \mathbf{x} as a bi-product). Such problems are typically casted as an optimisation problem of the form:

$$\arg \min_{\mathbf{x}, \mathbf{m}} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \phi(\mathbf{x}, \mathbf{m}), \quad (3)$$

and solved iteratively by the proximal gradient descent (PGD):

$$\text{PGD} : \begin{cases} \mathbf{g}^{(t+1)} = \mathbf{x}^{(t)} + \alpha^{(t)} \mathbf{H}^H (\mathbf{y} - \mathbf{H}\mathbf{x}^{(t)}) \rightarrow \text{Gradient with step size } \alpha \\ \{\mathbf{x}^{(t+1)}, \mathbf{m}^{(t+1)}\} = \text{Prox}_\phi(\mathbf{g}^{(t+1)}) \rightarrow \text{Proximal update} \end{cases} \quad (4)$$

where the gradient updates encourage k-space fidelity (the first term of (3)), and the *proximal operator* $\text{Prox}_\phi(\cdot)$ enforces image structure priors through a *regularisation* term $\phi(\cdot)$ that makes the inverse problem well-posed. The Bloch dynamics in (2) place an important temporal constraint (prior) for per-voxel trajectories of \mathbf{x} . Projecting onto this model (i.e. a temporal Prox model) has been suggested via iterative dictionary search schemes [3, 10]. This approach boost MRF reconstruction accuracy compared to the non-iterative DM [20], however, DM is non-scalable and can create enormous storage and computational overhead. Further, such approach processes data independently per voxel and neglects important spatial domain regularities in the TSIMIs and quantitative maps.

3 PGD-Net for MRF Quantification

We propose to learn a data-driven proximal operator within the PGD mechanism for solving the MRF problem. Implemented by compact networks with convolutional layers, the *neural Prox* improves the storage overhead and the sluggish runtime of the DM-based PGD by orders of magnitudes. Further, trained on quantitative MR images, the neural Prox network learns to simultaneously enforce spatial- and temporal-domain data structures within PGD iterations.

Prox Auto-encoder: We implement $\text{Prox} : \mathbf{g} \rightarrow \{\mathbf{x}, \mathbf{m}\}$ through a deep convolutional auto-encoder network:

$$\text{Prox} := \text{BLOCH} \circ \mathcal{G}, \quad (5)$$

consisting of an encoder $\mathcal{G} : \mathbf{g} \rightarrow \mathbf{m}$ and a decoder $\text{BLOCH} : \mathbf{m} \rightarrow \mathbf{x}$ subnetworks. The *information bottleneck* in the (neural) Prox auto-encoder corresponds to

projecting multichannel TSMIs to the low-dimensional manifold of the tissues' intrinsic (quantitative) property maps [14].

Decoder Network: Creates a differentiable model for *generating* the Bloch magnetic responses. This network uses 1×1 filters to process image time-series in a voxel-wise manner. Given quantitative properties $\mathbf{m}_v = \{T1_v, T2_v, \gamma_v\}$, the decoder *approximates* (dimension-reduced) Bloch responses in voxel v i.e. $\text{BLOCH}(\mathbf{m}_v) \approx \rho_v \mathcal{B}(T1_v, T2_v)$. This network is trained separately from the encoder. Training uses physical (Bloch) simulations for many combinations of the T1, T2 and PD values which can flexibly produce a rich training dataset [14].

Encoder Network: Projects \mathbf{g} the gradient-updated TSMIs in each iteration (i.e. the first line of (4)) to the quantitative property maps \mathbf{m} . Thus, \mathcal{G} must simultaneously (i) learn to incorporate spatial-domain regularities to de-alias TSMIs from the undersampling artefacts, and (ii) resolve the temporal-domain *inverse* mapping from the (noisy) TSMIs to the quantitative property maps. For this, and unlike BLOCH which applies pixel-wise temporal-only processing, \mathcal{G} uses multichannel convolution filters with wider receptive fields to learn/enable spatiotemporal processing of the TSMIs.




Fig. 1. Overview of the proposed proximal gradient descent network (PGD-Net) for tissue quantification in the compressive MR fingerprinting.

PGD-Net: Fig. 1a shows the *recurrent* architecture of the proposed *learned* PGD algorithm, coined as the PGD-NET. The trainable parameters within the PGD-NET are those of the encoder network \mathcal{G} (Fig. 1b) and the step sizes α_t . Other operators such as \mathbf{H} , \mathbf{H}^H and BLOCH (pre-trained separately, Fig. 1c) are kept frozen during training. Further, \mathcal{G} 's parameters are *shared* through all iterations. In practice, a truncated $T \geq 1$ recurrent iterations is used for training. Supervised training requires the MRF measurements, TSMIs, and the ground truth property maps to form the training input \mathbf{y} and target \mathbf{x}, \mathbf{m} samples.

Note there are many arts of engineering to determine the optimal network architecture, including different ways to encode temporal [17] or spatial-temporal information [5], these aspects are somewhat orthogonal to the model consistency question. Indeed, such mechanisms could also be incorporated in PGD-NET.

Training Loss: Given a training set $\{\mathbf{x}_i, \mathbf{m}_i, \mathbf{y}_i\}_{i=1}^N$, and $T \geq 1$ recurrent iterations of the PGD-NET (i.e. iterations used in PGD), the loss is defined as

$$\mathcal{L} = \gamma \sum_{i=1}^N \ell(\mathbf{x}_i, \mathbf{x}_i^{(T)}) + \sum_{j \in \{T_1, T_2, \rho\}} \beta_j \sum_{i=1}^N \ell(\mathbf{m}_{ij}, \mathbf{m}_{ij}^{(T)}) + \lambda \sum_{t=1}^T \sum_{i=1}^N \ell(\mathbf{y}_i, \mathbf{H}(\mathbf{x}_i^{(t)})), \quad (6)$$

where ℓ is the MSE loss defined with appropriate weights γ, β_j, λ on the reconstructed TSMIs \mathbf{x} (which measures the Bloch dynamic consistency) and tissue property maps \mathbf{m} , as well as on \mathbf{y} to maximise k-space data consistency with respect to the (physical) forward acquisition model. In this paper, the scaling between parameters γ, β_j and λ were initialized based on the physics (see 4.3).

4 Numerical Experiments

4.1 Anatomical Dataset

We construct a dataset of brain scans acquired using the 1.5T GE HDxT scanner with 8-channel receive-only head RF coil. For setting ground-truth (GT) values for the T1, T2 and PD parameters, gold standard anatomical maps were acquired using MAGIC quantification protocol [21]. Ground-truth quantitative maps were acquired from 8 healthy volunteers (16 axial brain slices each, at the spatial resolution of 128×128 pixels). From these parametric maps, we then construct the TSMIs and MRF measurements using the MRF acquisition protocol mentioned below to form the training/testing tuples $(\mathbf{m}_i, \mathbf{x}_i, \mathbf{y}_i)$. Data from 7 subjects were used for training our models, and one subject was kept for performance testing. We augmented training data into total 224 samples using random rotations (uniform angles in $[-8^\circ, 8^\circ]$), and left-right flipping of the GT maps. Training batches at each learning epoch were corrupted by i.i.d Gaussian noises of 30 dB SNR added to \mathbf{y} (we similarly add noise to the k-space test data).

4.2 MRF Acquisition

Our experiments use an excitation sequence of $L = 200$ repetitions which jointly encodes T1 and T2 values using an inversion pulse followed by a flip angle schedule that linearly ramps up from 1° to 40° , i.e. $\times 4$ truncated sequence than [14, 16]. Following [16], we set acquisition parameters $T_{\text{inv}} = 18$ ms (inversion time), fixed $TR = 10$ ms (repetition time), and $TE = 0.46$ ms (echo time). Spiral readouts subsample the k-space frequencies (the 128×128 Cartesian FFT grid) across 200 repetition times. We sample spatial frequencies $k(t) \propto (1.05)^{\frac{16\pi\tau}{1000}} e^{j\frac{16\pi\tau}{1000}}$ for $\tau = 1, 2, \dots, 1000$, which after quantisation to the nearest FFT grid, results in $m = 654$ samples per timeframe. In every repetition, similar to [20], this spiral pattern rotates by 7.5° in order to sub-sample new k-space frequencies. Given the anatomical T1, T2 and PD maps, we simulate magnetic responses using the Extended Phase Graph (EPG) formalism [29] and construct TSMIs and k-space measurements datasets, and use them for training and retrospective validations.


4.3 Reconstruction Algorithms

Two DM baselines namely, the non-iterative Fast Group Matching (FGM) [6] and the model-based iterative algorithm BLIP empowered by the FGM’s fast searches, were used for comparisons. For this, a MRF dictionary of 113’781 fingerprints was simulated over a dense grid of $(T1, T2) = [100:10:4000] \times [20:2:600]$ ms values. We implemented FGM searches on GPU using 100 groups for clustering this dictionary. The BLIP algorithm uses backtracking step size search and runs for maximum 20 iterations if is not convergent earlier. Further, we compared against related deep learning MRF baselines MRFCNN [8] and SCQ [12]. In particular, MRFCNN is a fully convolutional network and SCQ mainly uses 3 U-nets to separately infer T1, T2 and PD maps. The input to these networks is the dimension-reduced back-projected TSMIs $\mathbf{H}^H(\mathbf{y})$, and their training losses only consider quantitative maps consistency i.e. the second term in (6).

We trained PGD-NET with recurrent iterations $T = 2$ and 5 to learn appropriate proximal encoder \mathcal{G} and the step sizes $\alpha^{(t)}$. The architectures of \mathcal{G} and BLOCH networks are illustrated in Fig. 1. Similar to [14], the MRF dictionary was used for pre-training the BLOCH decoder that embeds a differentiable model for generating Bloch magnetic responses. A compact shallow network with one hidden layer and 1×1 filters (for pixel-wise processing) implements our BLOCH model [14]. On the other hand, our encoder \mathcal{G} has two residual blocks with 3×3 filters (for de-aliasing) followed by three convolutional layers with 1×1 filters for quantitative inference. The final hyper-parameters were $\beta = [1, 20, 2.5]$, $\gamma = 10^{-3}$ and $\lambda = 10^{-2}$ selected via a multiscale grid search to minimize error w.r.t. the ground truth. The inputs were normalized such that PD ranged in $[0, 1]$; smaller weights were used for \mathbf{x} and \mathbf{y} since they have higher energy than PD; we set $\lambda > \gamma$ since \mathbf{x} ’s norm is larger than \mathbf{y} ; $T1/T2$ values typically exhibit different ranges with $T1 \gg T2$, justifying their relative weightings in β to balance these terms. We used ADAM optimiser with 2000 epochs, mini-batch size 4 and learning rate 10^{-4} . We pre-trained our encoder \mathcal{G} using back-projected TSMIs to initialise the recurrent training, and also to compare the *encoder alone* predictions to the PGD-NET. All algorithms use a $s = 10$ dimensional MRF subspace representation for temporal-domain dimensionality reduction. The input and output channels are respectively 10 and 3 for MRFCNN, SCQ and \mathcal{G} . All networks were implemented in PyTorch, and trained and tested on NVIDIA 2080Ti GPUs.

4.4 Results and Discussions

Table 1 and Fig. 2 compare the performances of the MRF baselines against our proposed PGD-NET using $T = 2$ and 5 recurrent iterations. We also include inference results using the proposed *encoder alone* \mathcal{G} , without proximal iterations. Reconstruction performances were measured by the Normalised RMSE = $\frac{\|T1 - T1^{GT}\|}{\|T1^{GT}\|}$, MAE = $|T1 - T1^{GT}|$, Structural Similarity Index Metric (SSIM) [28], the required storage for the MRF dictionary (in DM methods) or the networks, and the algorithm runtimes averaged over the test image slices.



Ground truth (T1, T2, PD) anatomical maps acquired by the MAGIC gold standard [21]




Fig. 2. A slice of the true T1, T2 and PD maps acquired by the gold standard MAGIC (top), and the corresponding MRF reconstructions using (from top to bottom) FGM, BLIP+FGM, MRFCNN, SCQ, encoder alone \mathcal{G} , and the PROXNET with $T = 2$ and $T = 5$ algorithms.

Table 1. Average errors (NRMSE, SSIM, MAE), memory (for storing a dictionary or a network) and runtimes (per image slice) required for computing T1, T2 and PD maps using the MRF baselines and our PGD-NET algorithm.

| | NRMSE | | | SSIM | | | MAE (msec) | | Time (sec) | Memory (MB) |
|-------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|-------------|------------|--------------|-------------|
| | T1 | T2 | PD | T1 | T2 | PD | T1 | T2 | | |
| FGM | 0.475 | 0.354 | 1.12 | 0.614 | 0.652 | 0.687 | 350.0 | 14.6 | 1.29 | 8.81 |
| BLIP+FGM | 0.230 | 0.545 | 0.073 | 0.886 | 0.880 | 0.984 | 91.7 | 8.0 | 79.28 | 8.81 |
| MRFCNN | 0.155 | 0.158 | 0.063 | 0.943 | 0.972 | 0.987 | 80.3 | 5.4 | 0.083 | 4.72 |
| SCQ | 0.172 | 0.177 | 0.064 | 0.929 | 0.967 | 0.984 | 91.7 | 6.1 | 0.132 | 464.51 |
| \mathcal{G} (encoder alone) | 0.142 | 0.155 | 0.065 | 0.948 | 0.973 | 0.987 | 77.1 | 5.6 | 0.067 | 0.55 |
| PGD-NET ($T = 2$) | 0.104 | 0.138 | 0.050 | 0.973 | 0.979 | 0.991 | 59.9 | 5.0 | 0.078 | 0.57 |
| PGD-NET ($T = 5$) | 0.100 | 0.132 | 0.045 | 0.975 | 0.981 | 0.992 | 50.8 | 4.6 | 0.103 | 0.57 |

The non-iterative FGM results in incorrect maps due to the severe under-sampling artefacts. The model-based BLIP iterations improve this, however, due to lacking spatial regularisation, BLIP has limited accuracy and cannot fully remove aliasing artefacts (e.g. see T2 maps in Fig. 2) despite 20 iterations and very long runtime. In contrast, all deep learning methods outperform BLIP not only in accuracy but also in having 2 to 3 orders of magnitude faster reconstruction times—an important advantage of the learning-based methods. The proposed PGD-NET consistently outperforms all baselines, including DM and learning-based methods, over all defined accuracy metrics. This is achieved due to learning an effective spatiotemporal model (only) for the proximal operator i.e. the \mathcal{G} and BLOCH networks, directly incorporating the physical acquisition model \mathbf{H} into the recurrent iterations to avoid over-parameterisation of the overall inference model, as well as enforcing reconstructions to be consistent with the Bloch dynamics and the k-space data through the multi-term training loss (6). The MRFCNN and SCQ over-parametrise the inference by 1 and 3 orders of magnitude larger model sizes (the SCQ requires larger memory than DM) and are unable to achieve PGD-NET’s accuracy e.g. see the corresponding over-smoothed T2 maps in Fig. 2. Finally, we observe that despite having roughly the same model size (storage), the *encoder alone* \mathcal{G} predictions are not as accurate as the results of the PGD-NET’s recurrent iterations. By increasing the number of iterations T we observe that the PGD-NET’s accuracy consistently improves despite having an acceptable longer inference time.

5 Conclusions

In this work we showed that the consistency of the computed quantitative maps with respect to the physical forward acquisition model and the Bloch dynamics is important for reliably solving the MRF inverse problem using compact deep neural networks. For this, we proposed PGD-NET, a learned model-based iterative reconstruction framework that directly incorporates the forward acquisition and Bloch dynamic models within a recurrent learning mechanism with

a multi-term training loss. The PGD-NET adopts a data-driven neural proximal model for spatiotemporal processing of the MRF data, TSMI de-aliasing and quantitative inference. A chief advantage of this model is its compactness (a small number of weights/biases to tune), which might makes it particularly suitable for supervised training using scarce quantitative MRI datasets. Through our numerical validations we showed that the proposed PGD-NET achieves a superior quantitative inference accuracy, much smaller storage requirement, and a comparable runtime to the recent deep learning MRF baselines, while being much faster than the MRF fast dictionary matching schemes. In future work, we plan to evaluate the non-simulated scanner datasets with higher diversities and possible pathologies to further validate the method’s potential for clinical usage.

Acknowledgements. We thank Pedro Gómez, Carolin Prikl and Marion Menzel from GE Healthcare for the quantitative anatomical maps dataset. DC and MD are supported by the ERC C-SENSE project (ERCADG-2015-694888).

References

1. Adler, J., Öktem, O.: Solving ill-posed inverse problems using iterative deep neural networks. *Inverse Prob.* **33**(12), 124007 (2017)
2. Aggarwal, H.K., Mani, M.P., Jacob, M.: MoDL: model-based deep learning architecture for inverse problems. *IEEE Trans. Med. Imaging* **38**(2), 394–405 (2018)
3. Assländer, J., Cloos, M.A., Knoll, F., Sodickson, D.K., Hennig, J., Lattanzi, R.: Low rank alternating direction method of multipliers reconstruction for MR fingerprinting. *Magn. Reson. Med.* **79**(1), 83–96 (2018)
4. Balsiger, F., et al.: Magnetic resonance fingerprinting reconstruction via spatiotemporal convolutional neural networks. In: Knoll, F., Maier, A., Rueckert, D. (eds.) *MLMIR 2018*. LNCS, vol. 11074, pp. 39–46. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00129-2_5
5. Balsiger, F., Scheidegger, O., Carlier, P.G., Marty, B., Reyes, M.: On the spatial and temporal influence for the reconstruction of magnetic resonance fingerprinting. In: *International Conference on Medical Imaging with Deep Learning*, pp. 27–38 (2019)
6. Cauley, S.F., et al.: Fast group matching for MR fingerprinting reconstruction. *Magn. Reson. Med.* **74**(2), 523–528 (2015)
7. Chen, D., Davies, M.E.: Deep decomposition learning for inverse imaging problems. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2020)
8. Chen, D., Golbabae, M., Gómez, P.A., Menzel, M.I., Davies, M.E.: Deep fully convolutional network for MR fingerprinting. In: *International Conference on Medical Imaging with Deep Learning (MIDL)*, London, United Kingdom, 08–10 July 2019
9. Cohen, O., Zhu, B., Rosen, M.S.: MR fingerprinting deep reconstruction network (drone). *Magn. Reson. Med.* **80**(3), 885–894 (2018)
10. Davies, M., Puy, G., Vandergheynst, P., Wiaux, Y.: A compressed sensing framework for magnetic resonance fingerprinting. *SIAM J. Imaging Sci.* **7**(4), 2623–2656 (2014)
11. Doneva, M., Amthor, T., Koken, P., Sommer, K., Börnert, P.: Matrix completion-based reconstruction for undersampled magnetic resonance fingerprinting data. *Magn. Reson. Imaging* **41**, 41–52 (2017)

12. Fang, Z., et al.: Deep learning for fast and spatially-constrained tissue quantification from highly-accelerated data in magnetic resonance fingerprinting. *IEEE Trans. Med. Imaging* **38**(10), 2364–2374 (2019)
13. Fang, Z., Chen, Y., Nie, D., Lin, W., Shen, D.: RCA-U-Net: residual channel attention U-Net for fast tissue quantification in magnetic resonance fingerprinting. In: Shen, D., et al. (eds.) *MICCAI 2019. LNCS*, vol. 11766, pp. 101–109. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32248-9_12
14. Golbabaee, M., et al.: Compressive MRI quantification using convex spatiotemporal priors and deep auto-encoders. arXiv preprint [arXiv:2001.08746](https://arxiv.org/abs/2001.08746) (2020)
15. Golbabaee, M., Chen, D., Gómez, P.A., Menzel, M.I., Davies, M.E.: Geometry of deep learning for magnetic resonance fingerprinting. In: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7825–7829 (2019)
16. Gómez, P.A., et al.: Rapid three-dimensional multiparametric MRI with quantitative transient-state imaging. arXiv preprint [arXiv:2001.07173](https://arxiv.org/abs/2001.07173) (2020)
17. Hoppe, E., et al.: Deep learning for magnetic resonance fingerprinting: a new approach for predicting quantitative parameter values from time series. *Stud. Health Technol. Inf.* **243**, 202 (2017)
18. Hoppe, E., et al.: RinQ fingerprinting: recurrence-informed quantile networks for magnetic resonance fingerprinting. In: Shen, D., et al. (eds.) *MICCAI 2019. LNCS*, vol. 11766, pp. 92–100. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32248-9_11
19. Jiang, Y., Ma, D., Seiberlich, N., Gulani, V., Griswold, M.A.: MR fingerprinting using fast imaging with steady state precession (FISP) with spiral readout. *Magn. Reson. Med.* **74**(6), 1621–1631 (2015)
20. Ma, D., et al.: Magnetic resonance fingerprinting. *Nature* **495**(7440), 187 (2013)
21. Marcel, W., AB, S.: New technology allows multiple image contrasts in a single scan. *SPRING*, 6–10 (2015). [GESIGNAPULSE.COM/MR](https://www.gesignapulse.com/mr)
22. Mardani, M., et al.: Neural proximal gradient descent for compressive imaging. In: *Advances in Neural Information Processing Systems*, pp. 9573–9583 (2018)
23. McGivney, D.F., et al.: SVD compression for magnetic resonance fingerprinting in the time domain. *IEEE Trans. Med. Imaging* **33**(12), 2311–2322 (2014)
24. Oksuz, I., et al.: Magnetic resonance fingerprinting using recurrent neural networks. In: *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1537–1540 (2019)
25. Rick Chang, J., Li, C.L., Poczos, B., Vijaya Kumar, B., Sankaranarayanan, A.C.: One network to solve them all-solving linear inverse problems using deep projection models. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5888–5897 (2017)
26. Song, P., Eldar, Y.C., Mazor, G., Rodrigues, M.R.: HYDRA: hybrid deep magnetic resonance fingerprinting. *Med. Phys.* **46**(11), 4951–4969 (2019)
27. Virtue, P., Yu, S.X., Lustig, M.: Better than real: complex-valued neural nets for MRI fingerprinting. arXiv preprint [arXiv:1707.00070](https://arxiv.org/abs/1707.00070) (2017)
28. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**(4), 600–612 (2004)
29. Weigel, M.: Extended phase graphs: dephasing, RF pulses, and echoes-pure and simple. *J. Magn. Reson. Imaging* **41**(2), 266–295 (2015)
30. Zhao, B., et al.: Improved magnetic resonance fingerprinting reconstruction with low-rank and subspace modeling. *Magn. Reson. Med.* **79**(2), 933–942 (2018)