

# PINQI: An End-to-End Physics-Informed Approach to Learned Quantitative MRI Reconstruction

Felix F. Zimmermann , Christoph Kolbitsch , Patrick Schuenke , and Andreas Kofler 

**Abstract**—Quantitative Magnetic Resonance Imaging (qMRI) enables the reproducible measurement of biophysical parameters in tissue. The challenge lies in solving a nonlinear, ill-posed inverse problem to obtain the desired tissue parameter maps from acquired raw data. While various learned and non-learned approaches have been proposed, the existing learned methods fail to fully exploit the prior knowledge about the underlying MR physics, i.e. the signal model and the acquisition model. In this paper, we propose PINQI, a novel qMRI reconstruction method that integrates the knowledge about the signal, acquisition model, and learned regularization into a single end-to-end trainable neural network. Our approach is based on unrolled alternating optimization, utilizing differentiable optimization blocks to solve inner linear and non-linear optimization tasks, as well as convolutional layers for regularization of the intermediate qualitative images and parameter maps. This design enables PINQI to leverage the advantages of both the signal model and learned regularization. We evaluate the performance of our proposed network by comparing it with recently published approaches in the context of highly undersampled  $T_1$ -mapping, using both a simulated brain dataset, as well as real scanner data acquired from a physical phantom and in-vivo data from healthy volunteers. The results demonstrate the superiority of our proposed solution over existing methods and highlight the effectiveness of our method in real-world scenarios.

**Index Terms**—differentiable optimization, learned regularization, neural network,  $T_1$ -mapping, unrolled optimization, quantitative magnetic resonance imaging.

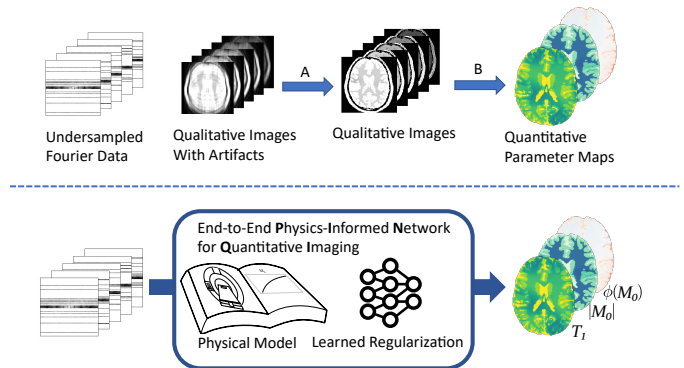
## I. INTRODUCTION

Magnetic Resonance Imaging (MRI) is a well-known method and an indispensable tool in clinical practice. However, the most commonly used protocols are qualitative, where the contrast in the images is determined by a mixture of different tissue parameters and acquisition properties. To overcome this issue, quantitative MRI (qMRI) techniques such as MR relaxometry have been proposed, which allow for the quantification of specific (bio)physical parameters of tissue such as spin-lattice ( $T_1$ ) and spin-spin relaxation times ( $T_2$ ). These quantitative measurements allow greater comparability between different devices at different sites and can be used to create more specific clinical guidelines. Typically, in qMRI series of qualitative images with different acquisition parameters

This paragraph of the first footnote will contain the date on which you submitted your paper for review

This work was supported in part by the Metrology for Artificial Intelligence for Medicine (M4AIM) project that is funded by the German Federal Ministry for Economic Affairs and Climate Action (BMWi) in the framework of the QI-Digital initiative. This work was funded in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), Grant 372486779.

F. F. Zimmermann (e-mail: felix.zimmermann@ptb.de), C. Kolbitsch, P. Schuenke, and A. Kofler are with Physikalisch-Technische Bundesanstalt (PTB), Braunschweig and Berlin, Germany



**Figure 1:** The problem to be solved in quantitative MRI is obtaining maps of physical parameters from undersampled measurements in Fourier space. Most previous methods consider the two steps of A) reconstructing artifact-free qualitative images and B) obtaining the parameter maps as two disjoint and independent steps (see text for details). We propose a novel end-to-end method making use of prior knowledge about the physics of data acquisition and learned regularization.

are recorded and a non-linear inverse problem is solved to obtain the tissue parameters. The clinical application of MRI and especially qMRI is challenging due to the relatively long measurement times. Hence, the acquired data is typically undersampled in Fourier space ( $k$ -space) to reduce the scan time at the cost of making the problem more ill-posed. This leads to artifacts that need to be accounted for by adopting appropriate regularization methods. Both for the reconstruction of purely qualitative images, as well as quantitative parameter maps, different approaches utilizing parallel data recording with multiple receiver coils [1], compressed sensing, and varying regularization schemes [2]–[5], model-based reconstruction [6]–[8], and combinations thereof have been proposed. More recently, neural networks were introduced as learned image reconstructions and as data-driven regularization methods [9]–[13], or for mapping reconstructed images to quantitative maps [14]–[16].

The existing data-driven approaches to solve the qMRI inverse problem, i.e. to obtain the parameter maps from the acquired  $k$ -space data, can be broadly categorized as follows. The first type of approach splits the problem into two disjoint tasks: A) image reconstruction and B) parameter regression (see Figure 1). The second type jointly solves both tasks to obtain the parameter maps directly from the  $k$ -space data. For the first task of qualitative image reconstruction, data-driven methods have made great progress: State-of-the-art results can be achieved by incorporating deep learning into

an unrolled model-based reconstruction by iteratively applying a neural network as a form of learned regularization as well as enforcing consistency with the recorded k-space data and the linear signal model [9], [10], [17]–[19]. The second task, parameter regression, is then either carried out by classical regression using non-linear solvers [20]–[23], dictionary matching [24], or by learned methods. Within the latter, a further distinction can be made between pixelwise mappings, such as MyoMapNet for  $T_1$ -mapping [15], [16] and methods using the information of multiple pixels via convolutional neural networks (CNNs), thereby implicitly learning a suitable spatial regularization [12], [14], [25]. MANTIS [14], for example, uses a UNet [26] to predict the parameter maps (originally  $T_2$  maps) from qualitative magnitude images with severe undersampling artifacts. DeepT1 [12], proposed for myocardial  $T_1$ -mapping, consists of two separate data-driven parts. First, an iterative reconstruction using recurrent CNNs [18] and data-consistency layers for image reconstruction is trained to reconstruct artifact-free qualitative images. Second, a UNet is trained to predict the quantitative maps from the magnitude of the reconstructed images. A similar approach is used in the recently proposed CoRRRECT [13] for motion-corrected  $R_2^*$ -mapping. It also uses an unrolled, CNN-regularized image reconstruction, followed by a parameter mapping UNet, with the major difference that both parts are trained simultaneously end-to-end. However, most learning-based methods entirely ignore prior knowledge about the physics of the signal and acquisition model relating the parameters with the quantitative images or use it only during training as part of a loss function [14], [25], [27]. Few learned methods are trained in an end-to-end manner, i.e. from k-space data to parameter maps, and fully incorporate the physical model at inference time. PGD-Net [28] unrolls a proximal gradient descent scheme with an approximated signal function, which is implemented as a pre-trained neural network and used as a differentiable proxy of the true MR fingerprinting signal function. DOPAMINE [11] unrolls a first-order gradient descent scheme that includes an implicitly learned regularizer where small residual CNNs operating on the different parameter maps are used to learn the gradient of a regularizer. Neither employs image-space regularization nor makes use of the particular form of the forward model and both only use shallow networks.

Combining the knowledge about the physics of the acquisition model by data-consistency layers with deep learning regularization has greatly improved qualitative image reconstruction [10], [17], [29]. Thus, we want to investigate a novel approach to incorporate the full knowledge of the physics of the signal model into a learned qMRI reconstruction, while employing learned regularization in image- and parameter-space.

#### A. Our Contributions

Our main contributions to the field of deep-learning-based quantitative imaging consist of

- Introduction of non-linear optimization as a differentiable layer into qMRI reconstruction.
- A novel and general end-to-end **Physics-Informed Network for Quantitative Imaging**, *PINQI*, based on

unrolled half quadratic splitting and differentiable optimization.

- Validation of the proposed approach on the task of  $T_1$ -mapping. We show the transferability of the network trained solely on synthetic data to in-vivo measurements.

We also demonstrate the usage of implicit differentiation for efficient evaluation of gradients with respect to all inputs of the commonly data-consistency layers of linear inverse problems as well as the non-linear optimization layer.

## II. METHODOLOGY

First, we introduce the notation and formulate the inverse problem to be solved in quantitative imaging. Next, we give a short introduction to differentiable optimization. We present implicit differentiation as a technique to obtain the Jacobian of solutions of optimization problems with respect to the parameters of the problems. Finally, we introduce our proposed PINQI solution to the quantitative imaging problem using differentiable optimization.

#### A. Problem Statement

By  $\mathbf{p}(\mathbf{r}) = [p_1(\mathbf{r}), \dots, p_{N_P}(\mathbf{r})]^\top$  we denote for each location  $\mathbf{r}$  the  $N_P$  relevant physical parameters to be determined, such as relaxation times. For notational brevity, we will write  $\mathbf{p} \in \mathbb{R}^{N_P \cdot N}$  for the vector representation of the parameters at  $N = N_x \cdot N_y$  discrete 2D positions. In qMRI, one considers the forward model given by

$$\mathbf{k} = (\mathbf{A} \circ q)(\mathbf{p}) + \mathbf{e}, \quad (1)$$

where  $\mathbf{k} \in \mathbb{C}^{N_r \cdot N_c \cdot N_k}$  is the vector of recorded k-space data,  $\mathbf{A}$  is a linear acquisition operator,  $q$  a (non-linear) signal model, and  $\mathbf{e}$  random noise. As the acquisition is performed  $N_r$  times with different sampling parameters and changes in the MR sequence influencing the signal model, the forward model is given by  $\mathbf{A} \circ q : \mathbb{R}^{N_P \cdot N} \rightarrow \mathbb{C}^{N_r \cdot N_c \cdot N_k}$  with  $\mathbf{p} \mapsto [(\mathbf{A}_1 \circ q_1)(\mathbf{p})^\top, \dots, (\mathbf{A}_{N_r} \circ q_{N_r})(\mathbf{p})^\top]^\top$ . Here, the encoding operators at each acquisition  $\mathbf{A}_i : \mathbb{C}^N \rightarrow \mathbb{C}^{N_c \cdot N_k}$  typically are undersampled Fourier operators. These can be written as  $\mathbf{A}_i = \mathbf{S}_i \mathbf{F} \mathbf{C}$  with  $\mathbf{F}$  a being two-dimensional Fourier transform and  $\mathbf{S}_i$  undersampling operators, masking out all but  $N_k$  discrete points in k-space. The undersampling masks are typically varied between different acquisitions. For  $N_c$  receiver coils acquiring data in parallel, the coil sensitivity operator  $\mathbf{C}$  expands each image to  $N_c$  different views acquired by the receiver coils by pixelwise multiplying with their respective spatially varying complex-valued sensitivity maps  $\mathbf{c}_c \in \mathbb{C}^N$ , often normalized such that for each pixel  $\sum_i^N |c_i|^2 = 1$ .

A qMRI reconstruction aims to obtain the tissue parameters  $\mathbf{p}$  from the acquired data  $\mathbf{k}$ . For uncorrelated Gaussian noise  $\mathbf{e}$ , the maximum likelihood estimate is

$$\mathbf{p}^* = \arg \min_{\mathbf{p}} \|\mathbf{A}q(\mathbf{p}) - \mathbf{k}\|^2. \quad (2)$$

As the inverse problem of obtaining  $\mathbf{p}^*$  is typically ill-posed, instead the regularized problem

$$\min_{\mathbf{p}} \|\mathbf{A}q(\mathbf{p}) - \mathbf{k}\|^2 + \mathcal{R}(\mathbf{p}) \quad (3)$$

with a regularizer  $\mathcal{R}$  has to be considered.

## B. Differentiable Optimization

In many computer vision tasks, data-driven approaches with differentiable layers solving an inner optimization problem within a larger neural network have been used successfully [30]–[32]. Yet so far, in medical imaging, optimization of an inner problem is mainly only used as data-consistency layers in unrolled reconstruction networks [10] for linear problems. In general, to incorporate optimization problems as a layer into a larger network, which can then be trained end-to-end with gradient descent algorithms, the gradients of the solution of the inner problem with respect to the inputs must be calculated. For data-consistency layers in linear problems, this can, for example, be achieved by automatic differentiation through the steps performed by the inner optimizer if each operation is differentiable [9]. The downsides are 1) all intermediate results have to be kept available, resulting in a linear dependence of the memory required during training on the number required iterations, and 2) algorithms for solving non-linear problems typically contain non-differentiable operations [23]. Alternatively, for linear problems, the gradient of the output with respect to the previous solution estimate can be efficiently calculated with matrix calculus [10], [33]. A more general approach to obtain gradients through an optimization layer uses a technique known as *implicit differentiation* [34]. We first revise the concept before applying it in Section II-D.

Let  $\mathcal{F}_\alpha : \mathbb{R}^n \rightarrow \mathbb{R}$  be the twice continuously differentiable objective of the inner optimization problem solved by the layer,  $\alpha \in \mathbb{R}^p$  the parameters to backpropagate the gradient for and  $\mathbf{f} : \mathbb{R}^p \times \mathbb{R}^n \rightarrow \mathbb{R}^n$  with  $\mathbf{f}(\alpha, \mathbf{x}) = \nabla_{\mathbf{x}} \mathcal{F}_\alpha(\mathbf{x})$  be the gradient of the objective function. For some  $\mathbf{x}^* \in \mathbb{R}^n$  to be a minimizer of  $\mathcal{F}$ , the condition

$$\mathbf{0} = \mathbf{f}(\alpha, \mathbf{x}^*(\alpha)) \quad (4)$$

has to be fulfilled. The well-known implicit function theorem (IFT) in a suitable notation [35], [36] states:

**Theorem 1 (Implicit Function Theorem):** Let  $\mathbf{f} : \mathbb{R}^p \times \mathbb{R}^n \rightarrow \mathbb{R}^n$  be a continuous differentiable function,  $\alpha_0 \in \mathbb{R}^p$  and  $\mathbf{x}_0 \in \mathbb{R}^n$  such that  $\mathbf{f}(\alpha_0, \mathbf{x}_0) = \mathbf{0}$  with non singular Jacobian  $\frac{\partial \mathbf{f}}{\partial \mathbf{x}}(\alpha_0, \mathbf{x}_0) \in \mathbb{R}^{n \times n}$ . Then, there exist an open set  $S \subset \mathbb{R}^n$  with  $\alpha_0 \in S$  and a unique continuously differentiable function  $\mathbf{x}^* : S \rightarrow \mathbb{R}^n$  such that  $\mathbf{x}^*(\alpha_0) = \mathbf{x}_0$  and  $\mathbf{f}(\alpha, \mathbf{x}^*(\alpha)) = \mathbf{0}$  for all  $\alpha \in S$ .

Hence, Equation (4) implicitly defines the minimizer  $\mathbf{x}^*$  as a function of  $\alpha$ . Differentiating both sides of Equation (4) yields

$$\mathbf{0} = \frac{\partial \mathbf{f}(\alpha, \mathbf{x}^*(\alpha))}{\partial \alpha} = \frac{\partial \mathbf{f}}{\partial \mathbf{x}^*} \frac{\partial \mathbf{x}^*(\alpha)}{\partial \alpha} + \frac{\partial \mathbf{f}}{\partial \alpha}, \quad (5)$$

where  $\mathbf{x}^*$  without explicit dependency on  $\alpha$  it is treated as a fixed value. By rearranging we obtain an expression for the Jacobian of the solution mapping, i.e

$$\frac{\partial \mathbf{x}^*(\alpha)}{\partial \alpha} = - \left( \frac{\partial \mathbf{f}}{\partial \mathbf{x}^*} \right)^{-1} \frac{\partial \mathbf{f}}{\partial \alpha}. \quad (6)$$

Therefore, by the chain rule, given the gradient of an outer loss  $\frac{\partial \mathcal{L}}{\partial \mathbf{x}}$  as a row vector at  $\mathbf{x} = \mathbf{x}^*$ , the row vector  $\frac{\partial \mathcal{L}}{\partial \alpha}$  at

$\alpha = \alpha_0$  can be written as

$$\frac{\partial \mathcal{L}}{\partial \alpha}(\alpha_0) = \frac{\partial \mathcal{L}}{\partial \mathbf{x}}(\mathbf{x}^*) \frac{\partial \mathbf{x}^*(\alpha)}{\partial \alpha}(\alpha_0, \mathbf{x}^*) \quad (7)$$

$$= - \frac{\partial \mathcal{L}}{\partial \mathbf{x}}(\mathbf{x}^*) \left( \frac{\partial^2 \mathcal{F}}{\partial \mathbf{x} \partial \mathbf{x}^T}(\alpha_0, \mathbf{x}^*) \right)^{-1} \frac{\partial^2 \mathcal{F}}{\partial \alpha \partial \mathbf{x}}(\alpha_0, \mathbf{x}^*). \quad (8)$$

Finally, dropping the explicit points of evaluation in the notation, we obtain the column vector

$$\left( \frac{\partial \mathcal{L}}{\partial \alpha} \right)^T = - \frac{\partial^2 \mathcal{F}}{\partial \alpha \partial \mathbf{x}} \left( \frac{\partial^2 \mathcal{F}}{\partial \mathbf{x} \partial \mathbf{x}^T} \right)^{-1} \left( \frac{\partial \mathcal{L}}{\partial \mathbf{x}} \right)^T, \quad (9)$$

where we have used the symmetry of the Hessian  $\frac{\partial^2 \mathcal{F}}{\partial \mathbf{x} \partial \mathbf{x}^T}$ .

This gives the general blueprint to implement the back-propagation step for optimization layers in deep-learning frameworks. Either the analytical Hessian or the vector-Hessian-product functionality of the framework can be used in conjunction with an iterative solver for linear problems to approximately obtain  $(\frac{\partial^2 \mathcal{F}}{\partial \mathbf{x} \partial \mathbf{x}^T})^{-1} (\frac{\partial \mathcal{L}}{\partial \mathbf{x}})^T$  in Equation (9). For an approximate solution with an error with norm  $\epsilon$ , it has been shown that the error of the estimated gradient is  $\mathcal{O}(\epsilon)$  [37]. The derivative with respect to the parameters of the gradient of the inner optimization loss used in Equation (9),  $\frac{\partial^2 \mathcal{F}}{\partial \alpha \partial \mathbf{x}}$ , can also either be calculated analytically or by autograd functionality.

## C. Proposed Reconstruction Network PINQI

For solving the qMRI inverse problem, we propose the following, physics-informed end-to-end approach *PINQI*: First, we choose the regularization in Equation (3) such that the objective is finding

$$\mathbf{p}^* = \arg \min_{\mathbf{p}} \|\mathbf{A}q(\mathbf{p}) - \mathbf{k}\|_2^2 + \lambda_{\tilde{\mathbf{p}}} \|\mathbf{p} - \mathbf{p}_{\text{reg}}\|_2^2, \quad (10)$$

with  $\mathbf{p}_{\text{reg}}$  a regularizing prediction for the parameters. We introduce an auxiliary variable  $\mathbf{y} := q(\mathbf{p})$  and include equality by a quadratic penalty constraint as well as an additional regularizing prediction for the images  $\mathbf{y}_{\text{reg}}$ ,

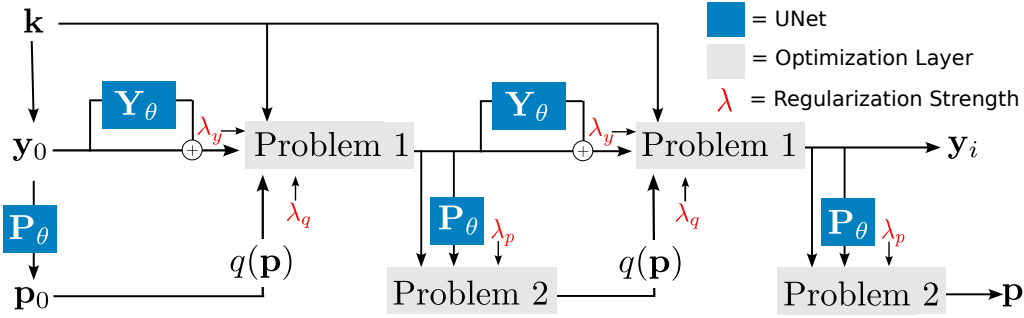
$$\min_{\mathbf{y}, \mathbf{p}} \|\mathbf{A}\mathbf{y} - \mathbf{k}\|_2^2 + \lambda_{\tilde{\mathbf{p}}} \|\mathbf{p} - \mathbf{p}_{\text{reg}}\|_2^2 + \lambda_q \|\mathbf{q}(\mathbf{p}) - \mathbf{y}\|_2^2 + \lambda_y \|\mathbf{y} - \mathbf{y}_{\text{reg}}\|_2^2, \quad (11)$$

where all  $\lambda$  are positive regularization strengths. Next, we split Problem (11) into two subproblems which we solve in an alternating fashion [38] within our unrolled reconstruction network.

1) *Linear Subproblem: Optimization in y*: The subproblem of obtaining artifact-free qualitative images

$$\mathbf{y}^* = \arg \min_{\mathbf{y}} \|\mathbf{A}\mathbf{y} - \mathbf{k}\|_2^2 + \lambda_y \|\mathbf{y} - \mathbf{y}_{\text{reg}}\|_2^2 + \lambda_q \|\mathbf{q}(\mathbf{p}) - \mathbf{y}\|_2^2 \quad (12)$$

is similar to the problem commonly solved in MR image reconstruction networks [10] by the data-consistency blocks, but extended by the model-based [6]–[8] last term which penalizes a discrepancy between reconstructed qualitative images and predictions based on the estimate for the quantitative parameters. As we consider multi-coil imaging, the minimizer of Equation (12) does not have a closed-form analytical solution. Instead, we solve a problem of the form  $\mathbf{H}\mathbf{x} = \mathbf{b}$  with



**Figure 2:** Schematic of the unrolled physics-informed network to solve Equation (10) by quadratic splitting as used by our proposed PINQI. We alternate between solving two subproblems, Problem 1 is a linear data-consistency problem and solved by a differentiable conjugate-gradient block. Subproblem 2 is solved by a differentiable non-linear optimization block.  $\mathbf{Y}_\theta$  denotes a residual UNet operating on qualitative images,  $\mathbf{P}_\theta$  the parameter prediction UNet. The predictions of these subnetworks are used as regularizers with (learnable) strength  $\lambda_y$  and  $\lambda_p$ , respectively. Consistency between both subproblems is relaxed to a quadratic penalty weighted by  $\lambda_q$ . For more details regarding the formulations of the two subproblems, see the main text.

$\mathbf{H} = \mathbf{A}^H \mathbf{A} + (\lambda_q + \lambda_y) \mathbf{I}_N$  and  $\mathbf{b} = \mathbf{A}^H \mathbf{k} + \lambda_q q(\mathbf{p}) + \lambda_y \mathbf{y}_{\text{reg}}$  approximately, for example with the conjugate gradient (CG) algorithm [23].

2) *Non-Linear Subproblem: Optimization in  $\mathbf{p}$ :* By introducing  $\lambda_p = \lambda_{\bar{p}}/\lambda_q$ , the second subproblem can be written as finding

$$\mathbf{p}^* = \arg \min_{\mathbf{p}} \|q(\mathbf{p}) - \mathbf{y}\|_2^2 + \lambda_p \|\mathbf{p} - \mathbf{p}_{\text{reg}}\|_2^2. \quad (13)$$

Due to the non-linear signal function  $q$ , this subproblem is non-linear. Hence, we introduce the non-linear optimization layer depicted in Figure 3. This layer uses L-BFGS [22], [23] to approximately solve the problem in the forward pass. Depending on the concrete signal model, within this layer, different solvers and preconditioning techniques might be chosen [39] instead.


Finally, we construct our proposed PINQI as shown in Figure 2 by alternating between both subproblems for a fixed number of iterations. In each iteration  $i = 1 \dots T$ , we obtain predictions for the qualitative, artifact-free images  $\mathbf{y}^i$  and for the quantitative parameters  $\mathbf{p}^i$ . For regularization, we use predictions obtained by trained subnetworks with parameters  $\theta \in \mathbb{R}^n$  as  $\mathbf{y}_{\text{reg}}^i = \mathbf{y}_{\theta}^i = \mathbf{Y}_{\theta}(\mathbf{y}^{i-1}, i)$  and  $\mathbf{p}_{\text{reg}}^i = \mathbf{p}_{\theta}^i = \mathbf{P}_{\theta}(\mathbf{y}^i, i)$ , respectively. All learnable parameters of the UNet subnetworks are shared between iterations.

#### D. Gradients of the Subproblem Solutions

We propose to train PINQI end-to-end, i.e. we construct an objective function  $\mathcal{L}$  which depends on the final predicted real-valued quantitative parameter maps. Also, we aim to use gradient descent-based algorithms to optimize the learnable network weights  $\theta$  and use backpropagation to find the direction of the steepest descent. Thus, we need to be able to calculate the gradients of the solutions found by the solvers of both subproblems with respect to all variables depending on  $\theta$ . We achieve this by implementing the solvers as differentiable optimization layers as presented in Section II-B.

First, we specialize the general concept to a regularized non-linear regression problem, i.e. an inner problem of obtaining  $\mathbf{p}^* = \arg \min_{\mathbf{p}} \mathcal{F}(\mathbf{p})$  with an objective

$$\mathcal{F}(\mathbf{p}) = \|q(\mathbf{p}) - \mathbf{y}\|_2^2 + \lambda \|\mathbf{p} - \mathbf{p}_{\text{reg}}\|_2^2, \quad (14)$$



**Figure 3:** Proposed non-linear optimization layer, finding  $\mathbf{p}^* = \arg \min_{\mathbf{p}} \mathcal{F}(\mathbf{p})$  with an off-the-shelf solver while allowing backpropagation of the gradients (red) to the regularization parameters,  $\lambda$  and  $\mathbf{p}_{\text{reg}}$ , and data  $\mathbf{y}$ .

as used in subproblem 2. Suppose the gradient of the inner objective at  $\mathbf{x}^*$  is continuously differentiable and its Jacobian is invertible. Then, the required gradient of  $\mathcal{L}$  with respect to the trainable weights can simply be found by applying Equation (9). The resulting equations for the propagation of gradients to  $\lambda$ ,  $\mathbf{p}_{\text{reg}}$  and  $\mathbf{y}$  are shown in Figure 3, which summarizes our proposed non-linear regression layer for this subproblem. Within this layer, we use CG to approximate  $\mathbf{g}$ , which denotes the product of the inverse Hessian of  $\mathcal{F}$  with the gradient of  $\mathcal{L}$  with respect to the output of the solver. We employ automatic differentiation to compute the Hessian-vector product applied inside the CG algorithm, as well as to compute the derivative (with respect to  $\theta$ ) of the gradient (with respect to  $\mathbf{p}$ ) of the inner loss. In Section III-D we investigate the influence of including a single instance of our differentiable non-linear optimization layer at the end of an otherwise unchanged qualitative image reconstruction network.

Similarly, backpropagation through the solver of the linear image reconstruction subproblem 1 can also be efficiently calculated, as demonstrated in [10] for gradients with respect to  $\mathbf{y}_{\text{reg}}$ . Extending upon these results, we use implicit differentiation to derive the gradients with respect to all inputs: For an inner optimization objective of the form

$$\mathcal{F}(\mathbf{y}) = \|\mathbf{A}\mathbf{y} - \mathbf{k}\|^2 + \sum_i \lambda_i \|\mathbf{y} - \mathbf{y}_i\|^2, \quad (15)$$

and a solution  $\mathbf{y}^* \in \mathbb{C}^N$ , that fulfills the necessary condition

$\nabla F(\mathbf{y}^*) = 0$ , application of Equation (9) and defining

$$\mathbf{g} := (\mathbf{A}^H \mathbf{A} + \sum_i \lambda_i \mathbf{I}_N)^{-1} \left( \frac{\partial \mathcal{L}(\mathbf{y}^*)}{\partial \mathbf{y}^{*H}} \right)^\top \quad (16)$$

gives the gradients of the outer loss (in Wirtinger calculus [40]):

$$\frac{\partial \mathcal{L}}{\partial \mathbf{y}_i^H} = \lambda_i \mathbf{g}, \quad \frac{\partial \mathcal{L}}{\partial \lambda_i} = \text{Re}\{(\mathbf{y}_i - \mathbf{y})^H \mathbf{g}\}, \quad \frac{\partial \mathcal{L}}{\partial \mathbf{k}^H} = \mathbf{A} \mathbf{g}. \quad (17)$$

For multi-coil MRI, the linear operator defining the data recorded in a single acquisition by the  $j$ -th coil is of the form  $\mathbf{A}_j = \mathbf{SFC}_j$ , and,

$$\frac{\partial \mathcal{L}}{\partial \mathbf{c}_j^H} = (\mathbf{F}^H (\mathbf{A}_j \mathbf{y}^* - \mathbf{k}_j)) \odot \bar{\mathbf{g}} + (\mathbf{F}^H \mathbf{A}_j \mathbf{g}) \odot \bar{\mathbf{y}}^* \quad (18)$$

can be used to backpropagate the gradient of the outer loss to the  $j$ -th estimated unnormalized sensitivity map  $\mathbf{c}_j$ . Here  $\odot$  denotes the Hadamard product,  $\bar{\mathbf{g}}$  the complex conjugate of  $\mathbf{g}$ , and  $\mathbf{k}_j$  the data recorded by the  $j$ -th coil.

### III. APPLICATION TO $T_1$ -MAPPING

The proposed PINQI can be used for many quantitative MR imaging techniques by adapting  $q$  to the respective signal model and  $\mathbf{A}$  to the acquisition operator.

For our experiments, we chose a  $T_1$ -mapping of the brain using a saturation recovery sequence as a typical, yet conceptually simple, qMRI problem. Here, the signal model  $q_i$  at the  $i$ -th acquisition is given for each pixel by

$$q_i(\text{Re } M_0, \text{Im } M_0, T_1) = M_0(1 - \exp(-\tau_i/T_1)) \quad (19)$$

with  $\text{Re } M_0$  and  $\text{Im } M_0$  denoting the real and imaginary part of the complex initial magnetization  $M_0$ ,  $T_1$  the longitudinal relaxation time, and  $\tau_i$  the  $i$ -th saturation recovery time, i.e. the delay between the magnetization preparation pulse and the data acquisition. As encoding operator  $\mathbf{A}$ , we chose a Fourier transform with Cartesian sampling, 4-8-fold undersampling, and 8 receiver coils.

All code relating to a PyTorch implementation of PINQI, the optimization layer, our implementations of the reference methods, the synthetic data generation as well as the MR sequences will be made available after peer-review at [github.com/fzimmermann89/PINQI](https://github.com/fzimmermann89/PINQI).

#### A. Training Dataset and Data Acquisition

We utilized synthetic data for training and validation and conducted testing on both synthetic data and real scanner data.

1) *Synthetic Data*: The synthetic data was generated from the BrainWeb dataset [41], which consists of three-dimensional segmentation masks of 20 healthy human brains, which we split into 16/1/3 for training/validation/testing. During the training process, we randomly assigned anatomically plausible values [21] for  $M_0$  and  $T_1$  to each of the 11 tissue classes in every sample on-the-fly. These values were combined with axial slices of the masks, resampled to  $192 \times 192$  pixels, into initial synthetic parameter maps. To increase spatial variability and prevent overfitting to piecewise constant images, we augmented the  $T_1$ -maps by multiplying them with low-variance 2D polynomials. Similarly, we augmented the

$M_0$ -maps by a random spatially slowly varying complex phase and bias field. As additional augmentations during training, we performed flips and rotations  $< 10^\circ$ . The resulting parameter maps were considered ground truth labels. In addition, we generated sample-specific masks indicating the presence of brain tissue within the maps. During training, these masks were used to weigh down the losses outside the brain region. During the testing phase, all quantitative measures used to report the performance of the methods were restricted to the relevant brain tissue.

The saturation recovery times were set as either 0.5 s, 1 s, 1.5 s, 2 s, and 8 s (synthetic comparison, phantom, and in-vivo experiments), or as 0.5 s, 0.7 s, 0.9 s, 1.1 s, 1.3 s, 1.6 s, 2 s, and 8 s (ablation study). We used variable density 1D-undersampling with 8 randomly generated [42] bird-cage-like receiver coils with random orientation. For each coil, we used a Gaussian amplitude profile with randomly varying half-width in  $[0.2, 0.5]$  times the field-of-view and a random phase modulated by a slowly varying random 2D polynomial. The undersampling patterns were randomly chosen for each recovery time and for each sample. Finally, for each sample, we added complex Gaussian noise with randomly chosen standard variation  $\sigma \in (0.001, 0.04)$  to simulate noisy measurements.

Our synthetic validation and testing datasets consisted of a fixed set of generated labels and simulated noisy undersampled k-space measurements. Here we used the ground truth sensitivity maps, thus assume a fully known forward model.

2) *Data Acquisition*: For further validation of our proposed method, we acquired two sets of scanner data on a 3T MRI system (Siemens Verio): 1) data from a physical phantom and 2) data from the brains of healthy volunteers. The examination was approved by the local ethics committee and is in accordance with the relevant guidelines and regulations. Written informed consent was received from all subjects prior to the examination. The phantom consisted of 9 tubes filled with liquids prepared to have  $T_1$  times in the range of approx. 250 ms-1750 ms. We used adiabatic saturation pulses, 5 saturation recovery times, spoiled GRE readouts with inside-out order and  $1 \text{ mm} \times 1 \text{ mm} \times 8 \text{ mm}$  resolution,  $192 \times 192$  matrix size, 6 deg flip angle, a 32-channel head coil. For each saturation time, the 1D Cartesian undersampling was chosen randomly, analogously to our training. As a reference, we performed fully sampled saturation recovery measurements with 18 delay times and obtained the  $T_1$  values by pixel-wise regression on the reconstructed images. Coil sensitivity maps were either estimated from the autocalibration region, i.e. the fully sampled 12/10/8 central lines, of the  $4 \times 6 \times 8 \times$  undersampled image with the longest  $\tau$ , or, for the reference, from the fully-sampled image [43]–[45]. All MR sequences were implemented using the vendor-agnostic Pulseseq framework [46].

#### B. PINQI: Implementation Details

We set the number of alternations between the subproblems in our implementation of PINQI as  $T = 5$  (empirically chosen). In the non-linear subproblem, we used the L-BFGS algorithm [22], [23] with a trigonometric transformation of the variables [39] to enforce the bounds  $\text{Re}(M_0) \in (-2, 2)$ ,

$\text{Im}(M_0) \in (-2, 2)$ , and  $R_1 := 1/T_1 \in (-1s^{-1}, 20s^{-1})$ . These bounds are much wider than any plausible predictions and only served to increase stability at the beginning of training. We initialized the solver in the first iteration with  $\mathbf{p}_{\text{reg}}$ , in subsequent iterations with the result obtained in the previous iteration. The linear problems in Equation (12) and Equation (9) were approximately solved by CG [23] with the norm of the residual  $< 10^{-6}$  as stopping criterion. Both regularizing networks,  $\mathbf{P}_\theta$  and  $\mathbf{Y}_\theta$ , are UNets [26] with residual blocks, each consisting of two SiLU [47] activated convolutions (window size 3) and group normalizations (group size 16) [48]. In  $\mathbf{Y}_\theta$ , the two convolutions of each block handle all temporal points as batched samples, and a third convolution operates along the temporal direction, handling all spatial points as batched samples [49]. In both UNets, we condition on the iteration of the unrolled reconstruction by performing learned projections to scale and shift values for each feature map [50], [51] after the first convolution in each encoder and decoder block. The downscaling in the UNet is done with stride 2, kernel 2 max-pooling operations in the spatial dimensions. Each upscaling is performed with 2x bilinear interpolation in the spatial dimensions followed by a 3x3 convolution. The number of output features at the different layers in  $\mathbf{Y}_\theta$  is empirically chosen as (16, 32, 48, 64), and as (32, 64, 96, 128) in  $\mathbf{P}_\theta$ . This results in 597,323 and 2,235,308 learnable weights, respectively. We found initializing the network such that each block mainly operates pixel-wise [52] improved training stability.

All regularization strengths  $\lambda_{p,i}$ ,  $\lambda_{y,i}$ , and  $\lambda_{q,i}$  are iteration-dependent learnable parameters. We enforced the positivity of these parameters by setting  $\lambda_i = \log \frac{1}{5} (1 + \exp\{5\tilde{\lambda}_i\})$  and optimizing for each  $\tilde{\lambda}_i$ . We empirically chose an initialization corresponding to  $\lambda_{p,i} = 3$ ,  $\lambda_{y,i} = 0.1$ , and  $\lambda_{q,i} = 0.1 + 0.05i$ . The training was performed by minimizing the MSE between the predicted quantitative parameters,  $R_1 = 1/T_1$ ,  $\text{Re}(M_0)$ ,  $\text{Im}(M_0)$ , and the corresponding target parameters. Deep supervision [53], [54] was applied by also incorporating the MSE loss of the predicted parameters during all previous iterations, strongly weighted down by a factor of 0.05. We pretrained  $\mathbf{P}_\theta$  for one epoch with random linear interpolation between the zero-filled reconstructions of the noisy k-space data and the ground truth (noise- and artifact-free) qualitative images as input with MSE loss on the estimated parameters. The training was performed for 80 epochs. Both training phases were performed using the Adam optimizer [55] with a maximum learning rate 0.004 for the UNet parameters and 0.001 for all  $\tilde{\lambda}$ 's, a linear warmup and cosine learning rate schedule [56], weight decay [57] 0.01, and a batch size of 16.

### C. Methods of Comparison

Besides our proposed PINQI approach, we provide results of our implementations of four recently published learned reconstruction methods trained and tested on the same dataset. As the general superiority of the selected learned comparison methods to non-learned methods has been established by the respective authors [10]–[14], we abstain from comparing against non-learned baselines.

a) *DeepT1* [12]: Our reimplement of DeepT1 tries to stay as close as possible to the provided description in the original publication while adapting to our synthetic data and the saturation recovery signal model. The task of magnitude-only image reconstruction by an unrolled reconstruction uses a recurrent CNN (297,538 learnable weights). We replaced the original data-consistency layers, which were only suitable for a single coil setting, with CG-based ones. The parameter estimation from the reconstructed magnitude images was done with a UNet (20,547,906 learnable weights).

b) *CoRECT* [13]: We implemented the network of CoRECT as described in the original publication, using a CNN (187,010 learnable parameters) in the unrolled image reconstruction and a UNet (46,741,698 parameters) for the parameter estimation. We adapted the method to our setting. Besides a supervised reconstruction loss and a self-supervised parameter estimation loss, we also included a supervised MSE loss of the quantitative parameters. This addition was motivated by the availability of ground truth labels in our setting and improved the results on our validation set, ensuring a fair comparison. For combining the loss terms, we empirically found weighting factors of 1, 0.1, and 1, respectively.

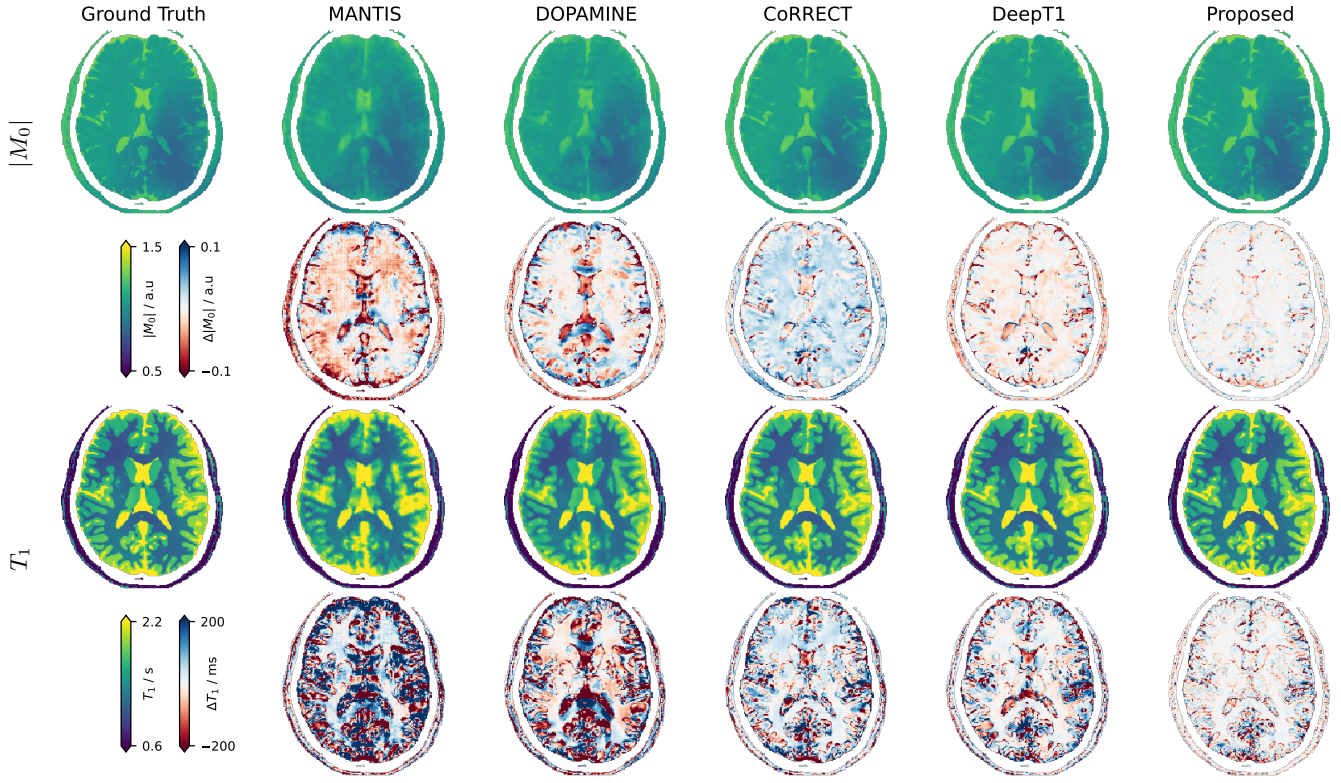
c) *MANTIS* [14]: We adopted the MANTIS method for  $T_1$ -mapping and training on our simulated dataset. As MANTIS takes zero-filled *magnitude* images as input, it cannot predict the phase of  $M_0$ . The particular UNet architecture proposed for MANTIS uses 29,248,258 learnable parameters. We used the proposed combination of the MSE of the quantitative parameters and in k-space as loss and tried different weighting of the latter. Although we noticed only a minimal influence on the performance on the validation set, we report results for a network trained with the optimal weighting found in our experiments of 0.01.

d) *DOPAMINE* [11]: We modified the DOPAMINE method for our signal model and supervised training. Our implementation performs 10 iterations of unrolled gradient descent. Each overall stepsize and each weighting of the neural network predicted step were scalar trainable parameters with softplus enforced positivity. The two regularizing CNNs and the network for starting point prediction combined had 317,894 trainable parameters in total. The training was performed on MSE of  $\mathbf{p}$ . As the training was highly unstable, we slightly modified DOPAMINE by soft-clipping each update in the unrolled gradient descent via a scaled tanh. A clipping threshold higher than the parameters' dynamic range over the training data was sufficient to stabilize training. We emphasize that during inference on the validation set, each update step proposed by the network was much smaller than allowed by the clipping, thus our modification should not negatively influence test performance.

### D. Ablation Study

Ablations of different parts of our proposed unrolled network are used to demonstrate the incremental benefits of the major components of our method. For each ablation, the respective learning rate of the parameters of the UNet was tuned based on the validation set MSE. The study was performed for eight acquisition times and 8-fold undersampling.





**Figure 4:** Exemplary results of the different methods for one simulated measurement from the test dataset at 8-fold acceleration. More details regarding the different methods are provided in the text. For each method, the magnitude of  $M_0$  and deviation from the ground truth are shown in the top two rows. The calculated  $T_1$  and deviation from ground truth are shown in the bottom two rows.

*a) No Signal Function:* To highlight the benefit of incorporating the knowledge about the signal function into the network through the proposed unrolling, we instead consider, similar to DeepT1 and CoRECT, solving the two problems of image reconstruction and parameter estimation only once. We use an unrolled model-based image reconstruction with a UNet-predicted regularization (5 iterations of network application and data-consistency) [9], [10]. The quantitative parameters are predicted by a separate UNet. Both UNets have the same architecture as the corresponding ones in PINQI. We performed a pretraining of the reconstruction network for 5 epochs (optimizing on MSE of the complex-valued qualitative images) followed by end-to-end optimization for 60 epochs. This baseline approach does not use the non-linear signal model  $q$  at inference time but benefits from architectural improvements to the CNNs over the comparison methods.

*b) No Image-Space Regularization:* We set  $\lambda_y = 0$  and remove  $\mathbf{Y}_\theta$  from PINQI to investigate the importance of the learned regularization in image-space while retaining the regularization in the non-linear subproblem. The training setup and all other parameters remain the same.

*c) No Parameter-Space Regularization:* We set  $\lambda_p = 0$  and remove  $\mathbf{P}_\theta$  from PINQI, effectively removing the learned regularization in parameter-space. Instead, we only perform a non-regularized regression on the signal model in the non-linear subproblem 2. All other parameters remain unchanged.

*d) No Non-Linear Solver:* We remove the proposed differentiable non-linear optimization layer for subproblem 2

from our network and instead directly consider the network prediction of the parameters as the solution of Equation (13). Note that we still iterate between the subproblems and train end-to-end. This ablation results in a learned reconstruction scheme similar to, for example, PGD-Net [28].

*e) Gradient Descent:* Instead of the proposed L-BFGS-based optimization with implicit gradient calculation, we perform 5 steps of gradient descent on Equation (13) in each alternation. We use standard backward-mode autograd to obtain the gradients of the steps and make the stepsizes for each alternation trainable parameters. The memory consumption during training is approximately 10% higher than in PINQI.

*f) Fixed  $\mathbf{p}_{\text{reg}}$  and  $\mathbf{y}_{\text{reg}}$  across all iterations:* We only apply the image-space and parameter-space UNets once, i.e.  $\mathbf{p}_{\text{reg}} = \mathbf{P}_\theta(\mathbf{y}_0)$  and  $\mathbf{y}_{\text{reg}} = \mathbf{Y}_\theta(\mathbf{y}_0)$ , instead of updating the UNet predictions in each iteration of the unrolled reconstruction. While the number of training parameters stays constant, this significantly reduces the computation necessary in a single forward pass. We keep the number of iterations  $T = 5$ , iterating between both subproblems to solve Equation (11).

*g) Single Iteration / Two Iterations:* We set either  $T = 1$  or  $T = 2$ , performing either a single iteration or two iterations of PINQI, respectively. Note, that  $T = 1$  differs from the experiment in *f*), where we still alternate between solving two subproblems.

For further investigation of the advantage of incorporating differentiable non-linear regression into the reconstruction, we used the same UNet architecture as used in PINQI for  $\mathbf{Y}_\theta$  also

in an unrolled reconstruction of the qualitative images  $\mathbf{y}$  and compare the following two cases:

*h) Neural Network Reconstruction and Separate Regression:* We optimized the network by minimizing the MSE of the qualitative images. At inference time, we used the resulting reconstructed images to perform a non-linear least-squares regression using BFGS to obtain  $\mathbf{p}$ .

*i) Neural Network with Non-linear Optimization Layer for End-to-End Regression:* We include the non-linear regression by means of the proposed optimization layer (Figure 3 in the network. As a result, we were able to train the network in an end-to-end fashion, minimizing the MSE of the quantitative parameters.

#### IV. RESULTS

##### A. Comparison with Reference Methods

The results of training PINQI and the comparison methods on the synthetic dataset and application on the test set are shown in Figure 5. Our proposed PINQI method achieves at the highest acceleration factor we investigated, 8-fold, a normalized root mean squared error (nRMSE) of the  $T_1$  maps of less than 0.10 whereas the lowest nRMSE achieved by one of the comparison methods is 0.13 (DeepT1 and the similar performing CoRECT) and both MANTIS and DOPAMINE have nRMSE exceeding 0.2. The mean absolute error (MAE) of  $T_1$  is again lowest for PINQI (0.05) compared to 0.07, 0.15, and 0.13 for DeepT1, MANTIS, and DOPAMINE, respectively. In terms of the Structural Similarity Index (SSIM) [58] (calculated with  $7 \times 7$  pixel uniform windows, only considering windows fully inside the brain), PINQI achieves the best, i.e. highest, result of 0.939 compared to the other methods. Examples of parameter maps for  $T_1$  as well as the magnitude of  $M_0$  obtained for a random slice of the test dataset are shown in Figure 4. Here, only PINQI was able to resolve the fine details. At lower undersampling factors, PINQI also achieves superior results compared to all comparison methods. For example, at 4-fold undersampling it yields 0.073/0.042/0.961 in terms of nRMSE/MAE/SSIM compared to the second best results of 0.091/0.054/0.938.

##### B. Ablation Study

The performance results obtained from the ablations are summarized in terms of nRMSE and MAE of  $T_1$  in Figure 6. The greatest increase in both MAE and nRMSE compared to PINQI as proposed is observed if, instead of the proposed differentiable optimization layer, gradient descent steps with a learned step size is used within the network (Figure 6, *e*). This showcases the importance of the proposed optimization layer for PINQI. The next biggest increases are if either the UNets are only applied once to obtain  $\mathbf{p}_{\text{reg}}$  and  $\mathbf{y}_{\text{reg}}$  independent of the iteration of the unrolled scheme (Figure 6, *f*), the iteration number is drastically decreased to  $T = 1$  (Figure 6, *g*), or the parameter regularization network is completely removed (Figure 6, *b*). These observations highlight the importance of learned regularization. The removal of the L-BFGS solver in the non-linear subproblem 13, corresponding to setting  $\lambda_p^i = \infty$ , resulting in  $\mathbf{p}^i = \mathbf{p}_0^i$ , also severely degrades the

**Table I:**  $T_1$ -values (in s) and standard deviations over ROIs obtained for the physical phantom by the fully-sampled reference method with 18 saturations delays compared to PINQI with 8-fold undersampling and 5 delays.

Reference	0.28	0.38	0.39	0.45	0.59	0.60	1.15	1.42	1.76
Std. Dev.	0.01	0.01	0.03	0.01	0.02	0.02	0.04	0.06	0.13
PINQI	0.29	0.37	0.42	0.46	0.60	0.59	1.10	1.47	1.76
Std. Dev.	0.03	0.02	0.06	0.05	0.02	0.02	0.05	0.16	0.14
Difference	4 %	-1 %	7 %	3 %	2 %	-2 %	-4 %	4 %	1 %

performance. In this ablation, the composition of parameter estimation UNet and  $q$  can be interpreted as a learned proximal mapping, as used in other unrolled reconstruction methods [28]. We observed further degradation by full removal of the explicit knowledge about the signal function  $q$ , i.e. only solving the linear image reconstruction problem with Fourier-space data-consistency and using a UNet for parameter prediction, similar as in other recent methods [12], [13].

The comparison between a learned reconstruction of qualitative images with the *separate* parameter regression as a post-processing step (Figure 6, *h*), and the inclusion of the regression in the network with *end-to-end* training utilizing the proposed differentiable optimization layer (Figure 6, *i*), shows a reduction of the mean nRMSE of the  $T_1$ -maps of  $>20\%$ .

##### C. Phantom and In-Vivo Application


We applied the final trained network to the data acquired from the physical phantom and used the reference measurement to obtain the mean  $T_1$  for each region-of-interest (ROI) for comparison. The RMS deviation between our proposed method and these reference values was 35 ms, the full results for all nine ROIs are shown in Table I.

To qualitatively demonstrate the applicability of the proposed method to real, unseen in-vivo measurements, we present the results of the volunteer study in Figure 7. Even though the acquisition was severely undersampled, the network predicts  $T_1$ -maps in agreement with the fully-sampled references and only minor artifacts remain.


#### V. DISCUSSION

We have demonstrated the superiority of PINQI for saturation recovery  $T_1$ -mapping compared to four current state-of-the-art qMRI reconstruction techniques, as evidenced by the lower  $T_1$ -errors on the synthetic test dataset (Figure 5). Given that only  $T_1$  holds clinical significance among the parameters obtainable in a saturation recovery sequence, and considering that most of the compared methods do not provide access to the phase of  $M_0$ , our quantitative comparison here was primarily focused on  $T_1$ . The methods used for comparison, MANTIS, DeepT1, CoRECT, and DOPAMINE all have their own characteristics and limitations. MANTIS utilizes a fully learned mapping from magnitude qualitative images with artifacts to quantitative parameter maps, but only incorporates the physical model during training. DeepT1 and CoRECT enforce data-consistency between qualitative images and recorded k-space data, but do not enforce consistency between the final predicted parameter maps and recorded

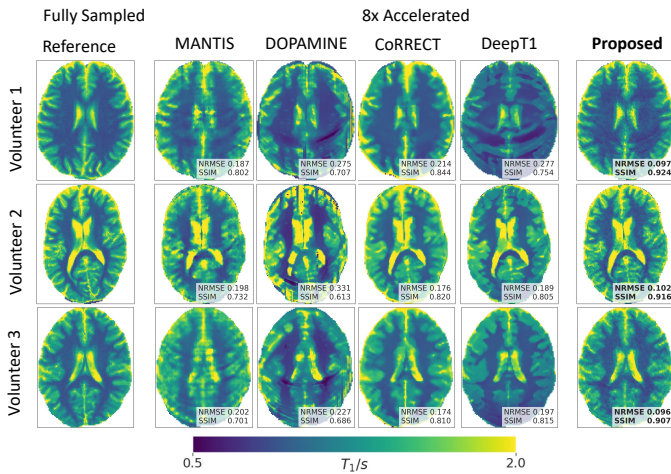




**Figure 5:** Comparison of our proposed PINQI with four different state-of-the-art learned qMRI reconstruction methods [11]–[14] in terms of nRMSE, MAE, and SSIM of  $T_1$  for each sample of the test set at 4-fold (darkest, bottom), 6-fold and 8-fold (brightest, top) undersampling. The mean values over all samples at 4-fold/6-fold/8-fold undersampling are provided as labels. PINQI improves upon all of the comparison methods in all three metrics.



**Figure 6:** Results of the ablation study for  $T_1$ , highlighting the importance of iteratively applying both regularizing UNets as well as the inclusion of the signal function via a non-linear regression layer. Note, here  $N_r = 8$  and the saturation times  $\tau$  differ from those used in Figure 5. For further details, see the full descriptions of the experiments in Section III-D with corresponding captions.



**Figure 7:** Examples of application of our proposed PINQI and comparison methods to 8-fold accelerated saturation recovery scans of the brains of three volunteers. All networks were solely trained on synthetic data. Only PINQI can successfully remove most of the artifacts caused by the severe undersampling and results in  $T_1$  closely matching the fully sampled saturation recovery measurements used as reference.

data at inference time. DOPAMINE explicitly uses knowledge about the non-linear signal model during inference, but does not employ image-space regularization and can use only shallow CNNs for parameter-space regularization due to the high number of iterations necessary. In contrast, PINQI incorporates data-consistency for the non-linear signal model through non-linear optimization layers, allowing for the inclusion of prior knowledge about the underlying physics at each iteration of the unrolled reconstruction. Additionally, PINQI makes use of both image- and parameter-space regularization and explicitly formulates the influence of all UNets as a learned regularizer. In total, these distinctions position PINQI as a unique method in qMRI reconstruction. Note that for a fair comparison, we reimplemented and adjusted the comparison methods to our dataset and signal model, as well as optimized hyperparameters rather than relying on the choices of the respective works. We were unable to use DOPAMINE as published, as training was highly unstable. We speculate, that the combination of an iteration-dependent but data-independent stepsize factor and a first-order gradient step without any thresholds was causing high susceptibility to even minor outliers. We had to modify the training to overcome this issue. During training with a batch size of 8, our PyTorch implementation of PINQI utilizes approximately 36 GB of memory in our specific configuration. This is comparable to our implementations of CoRECT (34 GB), DeepT1 (23 GB), and DOPAMINE (17 GB), but significantly higher than MANTIS (2 GB). It is crucial to note that these figures are highly dependent on the implementation details and the nature of the problem. Therefore, they should be interpreted as rough guidance for estimating memory consumption.

In our ablation study, we highlighted the importance of both image- and parameter-space regularization networks, as well as the utilization of the non-linear optimization layer and the signal model. Removal of any of these major components resulted in severe degradation, as summarized in Figure 6. By employing an unrolling approach and alternating between the two subproblems, we were able to update the predictions used as learned regularizations in each iteration, which proved to be highly beneficial compared to a single application of

the UNets. This suggests, that the UNets were able to make use of the gradually restored information within the unrolled algorithm. Moreover, we have demonstrated that the addition of a single non-linear optimization, as a final layer, to an otherwise unchanged learned reconstruction, can significantly reduce the error in the resulting parameter maps compared to performing a separate regression as a post-processing step outside the network (see Figure 6, *End-to-End Regression (i)* vs. *Separate Regression (h)*). This approach can be seamlessly integrated into existing reconstruction methods, enabling training with an objective that aligns more closely with the true objective in quantitative MRI. Essentially, it can be viewed as a task-specific loss function, which effectively penalizes reconstruction errors of the intermediate qualitative images that contribute to poor final regression results. Unlike a simple  $L_p$  loss on the qualitative intermediate images, this approach allows for more efficient learning by focusing on important features [59], [60]. We emphasize the potential of adding a single differentiable regression layer to both existing and future quantitative imaging methods as a straightforward means of incorporating knowledge about the signal model. As the implicit differentiation-based numerical gradient calculation only provides an approximation in a neighborhood of the possibly local, minimum obtained by the inner optimization, the usage may necessitate additional regularization, as employed in PINQI. Additionally, gradient noise might be an issue for certain problems.

Finally, we demonstrated that PINQI, even when trained solely on synthetic data, can be successfully applied to real data. This was confirmed both quantitatively, with low deviations from the reference measurements of the physical phantom in Table I, and qualitatively, through the reconstruction of nearly artifact-free maps from undersampled data in Figure 7. In contrast, the comparison methods cannot faithfully reconstruct the parameter maps in this setting with significantly higher undersampling than in the original publications and without specifically acquired in-vivo training data. In our synthetic evaluation, the data-consistent reconstruction methods were able to utilize the ground truth forward model, specifically  $q$  and the true sensitivity maps used the acquisition operator  $\mathbf{A}$ , whereas in the tests on real scanner data, these are not available. Here, we also acknowledge the possibility of self-supervised fine-tuning of our method for adaptation from the assumed forward model during training to the partially unknown and potentially different true forward model during testing, or a shifted distribution of the parameter maps [61]–[64]. In particular, the sensitivity maps estimated from undersampled data could be enhanced by an additional subnetwork [9]. However, for the sake of simplicity and comparability, we refrained from implementing such modules in our network or in any of the comparison methods. Nevertheless, Equation (18), obtained through implicit differentiation, would facilitate an efficient extension.

Although the formulation of an optimization objective with learned regularization provides some interpretability compared to a direct learned mapping of the quantitative parameters, as in, for example, CoRECT, MANTIS, or DeepT1, the inner workings of PINQI remain largely a black-box. The unrolling,

the repeated application of the subnetworks, and the inner optimizers increase the memory requirements during training as well as the computational requirements at inference time of PINQI compared to the comparison methods, potentially limiting its application in settings without suitable accelerator hardware or in 3D acquisitions. Notably, in our synthetic setting, we demonstrated that the results of the proposed PINQI at an acceleration factor of 8 are superior to those obtained by all other comparison methods at 4-fold acceleration. Despite the increased reconstruction time (2.3 s for 8 slices, using an Nvidia A6000) compared to, for example, DeepT1 (0.6 s) and CoRECT (0.3 s), opting for faster acquisition at lower errors can be an acceptable trade-off. Finally, PINQI requires a known, twice-differentiable closed-form signal model, which may limit its direct application to techniques such as magnetic resonance fingerprinting.

While we have only demonstrated the application of PINQI to Cartesian undersampled  $T_1$ -mapping, the splitting of PINQI can be trivially adapted to many inverse problems of the form described by Equation (3) by adjusting the non-linear signal model and the linear acquisition operator. Thus, while its effectiveness for other signal models remains to be further explored, PINQI represents a general physics-informed network for many quantitative imaging inverse problems.

## VI. CONCLUSION

Our proposed method, PINQI, presents a novel approach to quantitative image reconstruction by combining unrolled optimization, differentiable optimization layers, and learned regularization to fully utilize prior knowledge regarding the underlying physics. On a representative qMRI task,  $T_1$  mapping of the brain, we demonstrated superiority compared to three established data-driven reconstruction methods on a synthetic test dataset. Finally, we showed that our model, while trained on synthetic data alone, is transferable to in-vivo data.

## REFERENCES

- [1] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: Sensitivity encoding for fast MRI," *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 1999.
- [2] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed Sensing MRI," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 72–82, 2008.
- [3] M. Doneva, P. Börner, H. Eggers, C. Stehning, J. S  n  gas, and A. Mertins, "Compressed sensing reconstruction for magnetic resonance parameter mapping," *Magnetic Resonance in Medicine*, vol. 64, no. 4, pp. 1114–1120, 2010.
- [4] X. Wang, V. Roeloffs, J. Klosowski, Z. Tan, D. Voit, M. Uecker, and J. Frahm, "Model-based  $T_1$  mapping with sparsity constraints using single-shot inversion-recovery radial FLASH," *Magnetic Resonance in Medicine*, vol. 79, no. 2, pp. 730–740, 2018.
- [5] B. Zhao, F. Lam, and Z. P. Liang, "Model-based MR parameter mapping with sparsity constraints: Parameter estimation and performance bounds," *IEEE Transactions on Medical Imaging*, vol. 33, no. 9, pp. 1832–1844, 2014.
- [6] K. T. Block, M. Uecker, and J. Frahm, "Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint," *Magnetic Resonance in Medicine*, vol. 57, no. 6, pp. 1086–1098, 2007.
- [7] J. Tran-Gia, D. St  b, T. Wech, D. Hahn, and H. K  stler, "Model-based Acceleration of Parameter mapping (MAP) for saturation prepared radially acquired data," *Magnetic Resonance in Medicine*, vol. 70, no. 6, pp. 1524–1534, 2013.

- [8] K. M. Becker, J. Schulz-Menger, T. Schaeffter, and C. Kolbitsch, "Simultaneous high-resolution cardiac T1 mapping and cine imaging using model-based iterative image reconstruction," *Magnetic Resonance in Medicine*, vol. 81, no. 2, pp. 1080–1091, 2019.
- [9] A. Sriram, J. Zbontar, T. Murrell, A. Defazio, C. L. Zitnick, N. Yakubova, F. Knoll, and P. Johnson, "End-to-End Variational Networks for Accelerated MRI Reconstruction," *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020. Lecture Notes in Computer Science*, vol. 12262 LNCS, pp. 64–73, 2020.
- [10] H. K. Aggarwal, M. P. Mani, and M. Jacob, "ModL: Model-Based Deep Learning Architecture for Inverse Problems," *IEEE Transactions on Medical Imaging*, vol. 38, no. 2, pp. 394–405, 2019.
- [11] Y. Jun, H. Shin, T. Eo, T. Kim, and D. Hwang, "Deep model-based magnetic resonance parameter mapping network (DOPAMINE) for fast T1 mapping using variable flip angle method," *Medical Image Analysis*, vol. 70, p. 102017, 2021.
- [12] H. Jeelani, Y. Yang, R. Zhou, C. M. Kramer, M. Salerno, and D. S. Weller, "A Myocardial T1-Mapping Framework with Recurrent and U-Net Convolutional Neural Networks," *Proceedings - International Symposium on Biomedical Imaging*, vol. 2020-April, pp. 1941–1944, 2020.
- [13] X. Xu, W. Gan, S. V. Kothapalli, D. A. Yablonskiy, and U. S. Kamilov, "CoRECT: A Deep Unfolding Framework for Motion-Corrected Quantitative R2\* Mapping," 2022.
- [14] F. Liu, L. Feng, and R. Kijowski, "MANTIS: Model-Augmented Neural Network with Incoherent k-space Sampling for efficient MR parameter mapping," *Magnetic Resonance in Medicine*, vol. 82, no. 1, pp. 174–188, 2019.
- [15] R. Guo, H. El-Rewaady, S. Assana, X. Cai, A. Amyar, K. Chow, X. Bi, T. Yankama, J. Cirillo, P. Pierce, B. Goddu, L. Ngo, and R. Nezafat, "Accelerated cardiac T1 mapping in four heartbeats with inline MyoMapNet: a deep learning-based T1 estimation approach," *Journal of Cardiovascular Magnetic Resonance*, vol. 24, no. 1, pp. 1–15, 2022.
- [16] 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.
- [17] J. Schlemper, J. Caballero, J. V. Hajnal, A. N. Price, and D. Rueckert, "A Deep Cascade of Convolutional Neural Networks for Dynamic MR Image Reconstruction," *IEEE Transactions on Medical Imaging*, vol. 37, no. 2, pp. 491–503, 2018.
- [18] C. Qin, J. Schlemper, J. Caballero, A. N. Price, J. V. Hajnal, and D. Rueckert, "Convolutional recurrent neural networks for dynamic MR image reconstruction," *IEEE Transactions on Medical Imaging*, vol. 38, no. 1, pp. 280–290, 2019.
- [19] Y. Yang, J. Sun, H. Li, and Z. Xu, "Deep ADMM-Net for compressive sensing MRI," *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [20] K. Chow, J. A. Flewitt, J. D. Green, J. J. Pagano, M. G. Friedrich, and R. B. Thompson, "Saturation recovery single-shot acquisition (SASHA) for myocardial T1 mapping," *Magnetic Resonance in Medicine*, vol. 71, no. 6, pp. 2082–2095, 2014.
- [21] J. Z. Bojorquez, S. Bricq, C. Acquitier, F. Brunotte, P. M. Walker, and A. Lalande, "What are normal relaxation times of tissues at 3 T?" *Magnetic Resonance Imaging*, vol. 35, pp. 69–80, 2017.
- [22] D. C. Liu and J. Nocedal, "On the limited memory BFGS method for large scale optimization," *Mathematical Programming*, vol. 45, no. 1-3, pp. 503–528, 8 1989.
- [23] J. Nocedal and S. Wright, *Numerical Optimization*, ser. Springer Series in Operations Research and Financial Engineering. Springer New York, 2006.
- [24] D. Ma, V. Gulani, N. Seiberlich, K. Liu, J. L. Sunshine, J. L. Duerk, and M. A. Griswold, "Magnetic resonance fingerprinting," *Nature*, vol. 495, no. 7440, pp. 187–192, 3 2013.
- [25] H. Li, M. Yang, J. H. Kim, C. Zhang, R. Liu, P. Huang, D. Liang, X. Zhang, X. Li, and L. Ying, "SuperMAP: Deep ultrafast MR relaxometry with joint spatiotemporal undersampling," *Magnetic Resonance in Medicine*, vol. 89, no. 1, pp. 64–76, 2023.
- [26] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 9351, pp. 234–241, 2015.
- [27] K. Hammernik, T. Kustner, B. Yaman, Z. Huang, D. Rueckert, F. Knoll, and M. Akcakaya, "Physics-Driven Deep Learning for Computational Magnetic Resonance Imaging: Combining physics and machine learning for improved medical imaging," *IEEE Signal Processing Magazine*, vol. 40, no. 1, pp. 98–114, 2023.
- [28] D. Chen, M. E. Davies, and M. Golbabaee, "Deep Unrolling for Magnetic Resonance Fingerprinting," *Proceedings - International Symposium on Biomedical Imaging*, no. 2, 2022.
- [29] A. Kofler, M. Haltmeier, T. Schaeffter, and C. Kolbitsch, "An end-to-end-trainable iterative network architecture for accelerated radial multi-coil 2d cine mr image reconstruction," *Medical Physics*, vol. 48, no. 5, pp. 2412–2425, 2021.
- [30] B. Amos and J. Z. Kolter, "OptNet: Differentiable optimization as a layer in neural networks," *34th International Conference on Machine Learning, ICML 2017*, vol. 1, pp. 179–191, 2017.
- [31] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Zico Kolter, "Differentiable convex optimization layers," *Advances in Neural Information Processing Systems*, vol. 32, no. NeurIPS, 2019.
- [32] Z. Lv, F. Dellaert, J. M. Rehg, and A. Geiger, "Taking a deeper look at the inverse compositional algorithm," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2019-June, pp. 4576–4585, 2019.
- [33] K. Hammernik, T. Kustner, and D. Rueckert, "Machine Learning for MRI Reconstruction," in *Magnetic Resonance Image Reconstruction*, C. Prieto, M. I. Doneva, and M. Akcakaya, Eds. Elsevier, 2022, ch. 11, pp. 281–317.
- [34] S. Gould, B. Fernando, A. Cherian, P. Anderson, R. S. Cruz, and E. Guo, "On differentiating parameterized argmin and argmax problems with application to bi-level optimization," 2016.
- [35] M. Crowder, J. R. Magnus, and H. Neudecker, *Matrix Differential Calculus with Applications in Statistics and Econometrics*. John Wiley & Sons, 1989.
- [36] O. de Oliveira, "The implicit and the inverse function theorems: Easy proofs," *Real Analysis Exchange*, vol. 39, no. 1, pp. 207–218, 2013.
- [37] F. Pedregosa, "Hyperparameter optimization with approximate gradient," *33rd International Conference on Machine Learning, ICML 2016*, vol. 2, pp. 1150–1159, 2016.
- [38] M. V. Afonso, J. M. Bioucas-Dias, and M. A. Figueiredo, "Fast image recovery using variable splitting and constrained optimization," *IEEE Transactions on Image Processing*, vol. 19, no. 9, pp. 2345–2356, 2010.
- [39] P. E. Gill, W. Murray, and M. H. Wright, *Practical Optimization*. Academic Press, 1981.
- [40] K. Kreutz-Delgado, "The complex gradient operator and the CR-calculus," *arXiv preprint*, 2009.
- [41] B. Aubert-Broche, M. Griffin, B. B. Pike, A. C. Evans, and D. L. Collins, "Twenty new digital brain phantoms for creation of validation image data bases," *IEEE Transactions on Medical Imaging*, vol. 25, no. 11, pp. 1410–1416, 2006.
- [42] M. J. Muckley, R. Stern, T. Murrell, and F. Knoll, "mrisensesim.py," TorchKbNufft: A High-Level, Hardware-Agnostic Non-Uniform Fast Fourier Transform, ISMRM Workshop on Data Sampling & Image Reconstruction. Online github.com/mmuckley/torchkbnufft/tree/v0.3.4/torchkbnufft/mri.
- [43] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin, "Adaptive reconstruction of phased array MR imagery," *Magnetic Resonance in Medicine*, vol. 43, no. 5, pp. 682–690, 2000.
- [44] S. J. Inati, M. S. Hansen, and P. Kellman, "A Fast Optimal Method for Coil Sensitivity Estimation and Adaptive Coil Combination for Complex Images," *Proceedings of the 22nd Annual Meeting of ISMRM*, 2014.
- [45] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, and M. Lustig, "ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA," *Magnetic Resonance in Medicine*, vol. 71, no. 3, pp. 990–1001, 2014.
- [46] K. J. Layton, S. Kroboth, F. Jia, S. Littin, H. Yu, J. Leupold, J.-F. Nielsen, T. Stöcker, and M. Zaitsev, "Pulseq: A rapid and hardware-independent pulse sequence prototyping framework," *Magnetic Resonance in Medicine*, vol. 77, no. 4, pp. 1544–1552, 2017.
- [47] B. Zoph and Q. V. Le, "Searching for activation functions," *6th International Conference on Learning Representations, ICLR 2018 - Workshop Track Proceedings*, no. 1, pp. 1–12, 2018.
- [48] Y. Wu and K. He, "Group Normalization," *International Journal of Computer Vision*, vol. 128, no. 3, pp. 742–755, 2020.
- [49] Z. Qiu, T. Yao, and T. Mei, "Learning Spatio-Temporal Representation with Pseudo-3D Residual Networks," *Proceedings of the IEEE International Conference on Computer Vision*, vol. 2017-October, pp. 5534–5542, 10 2017.
- [50] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, "FiLM: Visual reasoning with a general conditioning layer," *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pp. 3942–3951, 2018.
- [51] F. F. Zimmermann and A. Kofler, "NoSENSE: Learned unrolled cardiac MRI reconstruction without explicit sensitivity maps," *International*

*Workshop on Statistical Atlases and Computational Models of the Heart (STACOM)*, 2023.

- [52] T. Bachlechner, B. P. Majumder, H. Mao, G. Cottrell, and J. McAuley, "ReZero is All You Need: Fast Convergence at Large Depth," *37th Conference on Uncertainty in Artificial Intelligence, UAI 2021*, no. 1, pp. 1352–1361, 2021.
- [53] B. Zhou and S. Kevin Zhou, "Dudornet: Learning a dual-domain recurrent network for fast MRI reconstruction with deep T1 prior," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 4272–4281, 2020.
- [54] L. Wang, C.-Y. Lee, Z. Tu, and S. Lazebnik, "Training deeper convolutional networks with deep supervision," 2015.
- [55] D. P. Kingma and J. L. Ba, "Adam: A method for stochastic optimization," *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, pp. 1–15, 2015.
- [56] T. He, Z. Zhang, H. Zhang, Z. Zhang, J. Xie, and M. Li, "Bag of tricks for image classification with convolutional neural networks," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2019-June, pp. 558–567, 2019.
- [57] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- [58] 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*, vol. 13, no. 4, pp. 600–612, 2004.
- [59] X. Zhang, Q. Duchemin, K. Liu\*, C. Gultekin, S. Flassbeck, C. Fernandez-Granda, and J. Assländer, "Cramér–Rao bound-informed training of neural networks for quantitative MRI," *Magnetic Resonance in Medicine*, vol. 88, no. 1, pp. 436–448, 2022.
- [60] A. Kendall, Y. Gal, and R. Cipolla, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7482–7491.
- [61] D. Gilton, G. Ongie, and R. Willett, "Model Adaptation for Inverse Problems in Imaging," *IEEE Transactions on Computational Imaging*, vol. 7, no. 2, pp. 661–674, 2021.
- [62] K. Hammernik, J. Schlemper, C. Qin, J. Duan, R. M. Summers, and D. Rueckert, "Systematic evaluation of iterative deep neural networks for fast parallel mri reconstruction with sensitivity-weighted coil combination," *Magnetic Resonance in Medicine*, vol. 86, no. 4, pp. 1859–1872, 2021.
- [63] F. F. Zimmermann, A. Kofler, C. Kolbitsch, and P. Schuenke, "Semi-Supervised Learning for Spatially Regularized Quantitative MRI Reconstruction - Application to Simultaneous T1, B0, B1 Mapping," 2023, 1166, ISMRM Annual Meeting.
- [64] Y. Jun, J. Cho, X. Wang, M. Gee, P. E. Grant, B. Bilgic, and B. Gagoski, "SSL-QALAS: Self-Supervised Learning for rapid multiparameter estimation in quantitative MRI using 3D-QALAS," *Magnetic Resonance in Medicine*, vol. 90, no. 5, pp. 2019–2032, 2023.