

Nr. 93
14. Apr. 2023

Preprint-Series: Department of Mathematics - Applied Mathematics

Uncertainty-Aware Null Space Networks for
Data-Consistent Image Reconstruction

C. Angermann, S. Göppel, M. Haltmeier

AppliedMathematics

Technikerstraße 13 - 6020 Innsbruck - Austria
Tel.: +43 512 507 53803 Fax: +43 512 507 53898
<https://applied-math.uibk.ac.at>

Uncertainty-Aware Null Space Networks for Data-Consistent Image Reconstruction

Christoph Angermann

VASCage – Centre on Clinical Stroke Research
6020 Innsbruck, Austria
E-mail: christoph.angermann@vascage.at

Simon Göppel

Department of Mathematics, University of Innsbruck
Technikerstrasse 13, 6020 Innsbruck, Austria
E-mail: simon.goepfel@uibk.ac.at

Markus Haltmeier

Department of Mathematics, University of Innsbruck
Technikerstrasse 13, 6020 Innsbruck, Austria
E-mail: markus.haltmeier@uibk.ac.at

April 14, 2023

Abstract

Reconstructing an image from noisy and incomplete measurements is a central task in several image processing applications. In recent years, state-of-the-art reconstruction methods have been developed based on recent advances in deep learning. Especially for highly underdetermined problems, maintaining data consistency is a key goal. This can be achieved either by iterative network architectures or by a subsequent projection of the network reconstruction. However, for such approaches to be used in safety-critical domains such as medical imaging, the network reconstruction should not only provide the user with a reconstructed image, but also with some level of confidence in the reconstruction. In order to meet these two key requirements, this paper combines deep null-space networks with uncertainty quantification. Evaluation of the proposed method includes image reconstruction from undersampled Radon measurements on a toy CT dataset and accelerated MRI

reconstruction on the fastMRI dataset. This work is the first approach to solving inverse problems that additionally models data-dependent uncertainty by estimating an input-dependent scale map, providing a robust assessment of reconstruction quality.

Keywords: Inverse problems, null space learning, neural network, uncertainty quantification, accelerated MRI, limited angle

1 Introduction

Inverse problems arise in a variety of sciences and play a central role in many engineering applications [17]. After discretization [30, 34], an inverse problem can be written in the form

$$y = \mathcal{A}x + \epsilon, \quad (1.1)$$

where $y, \epsilon \in \mathbb{R}^m$, $x \in \mathbb{R}^n$ and $\mathcal{A} \in \mathbb{R}^{m \times n}$. The goal is to recover the unknown signal x from the knowledge of the noisy data y and the measurement operator \mathcal{A} . In the deterministic case, the data perturbation ϵ is assumed to satisfy $\|\epsilon\|_2 \leq \delta$ for some noise level $\delta > 0$. Inverse problems are characterized by non-uniqueness and sensitivity with respect to data perturbations. In particular, for applications with incomplete data, the reduced number of measurements results in a high dimensional set of potential solutions. In order to obtain reliable reconstructions, it is therefore necessary to apply suitable regularization techniques to the problem at hand, which address non-uniqueness and instability [17, 46].

Prime examples of ill-posed inverse problems are found in image reconstruction, such as computed tomography (CT) and magnetic resonance imaging (MRI). Both of these techniques often encounter problems with limited data because measurements are only available for a limited subset. This is the case, for example, with limited-angle computed tomography, where angular projections are available only within a strict subset of the full angular range. These limitations may be due to the physical measurement setup, for example, and naturally occur in practical applications such as digital breast tomosynthesis, dental tomography or nondestructive testing. Due to the lack of available data, important features in the reconstructions can be obscured by artifacts caused by hard cut-offs in the measurement domain. While the characterization of these artifacts has been thoroughly investigated [11, 19, 20], the reliable correction of missing data is still a challenging task. Unlike computed tomography, MRI takes measurements without exposing the patient to ionizing radiation. MRI has a variety of applications for disease detection, general treatment, and prognosis. To save cost and time, it is common practice in MRI to reduce the acquisition time by taking an incomplete number of measurements. Due to the resulting undersampling, the patient image is no longer

uniquely determined by the available data. To overcome this limitation, methods of compressed sensing (CS) and ℓ^1 -regularization [10, 13, 15, 25, 41, 44] have been developed which allow accurate reconstruction from sparsely sampled data. As a result, CS and sparsity regularization became key tools in modern MRI imaging, and is still an active area of research.

In recent years, machine learning (ML) and neural networks (NN) have emerged as new paradigms for solving inverse problems [2, 4, 8, 33, 39, 43, 51], in which a trainable NN is adjusted to the available dataset. The fitting process, also referred to as network training, can be treated both as a supervised or unsupervised learning task. Supervised training attempts to infer a cross-domain correspondence between measurements and ground truth data, while unsupervised training attempts to estimate an overall image distribution from which reconstructions can be sampled [14, 24]. Note that any inverse problem can be considered a supervised ML task, as long as a sufficient amount of data is available and the forward operator is known explicitly. In many cases, data-driven solvers significantly reduce reconstruction time while improving accuracy. However, given the highly nonlinear structure of a NN, it remains unclear how accurately a network processes the provided data. This is of particular importance for medical applications, where reliable reconstruction is essential. In particular, for problems with limited data, it is desirable to explicitly enforce data consistency in order to obtain reliable predictions. The difficult-to-interpret nature of NNs combined with the lack of theoretical guarantees still limits the practical applicability of DL-based solvers in clinical trials [16, 22]. Null space networks that promote data-consistent reconstruction through architectural design have been introduced and analyzed in [48, 49] and shown to result in a convergent regularization technique.

Due to the difficult-to-interpret structure, it is often argued that NNs suffer from the so-called black-box paradigm [12]. In general, no interpretable connection can be made between the learned network parameters and the actual inter-domain relationship. As a result, network behavior is almost unpredictable in the presence of out-of-distribution (OOD) data, which are not included in the training set [1, 35]. OOD data occurs, for example, when the scanner’s measurement quality suddenly degrades or unusual objects (such as tumors or metal parts in medical imaging) become visible in the patient. This is a significant limitation to the applicability of DL-based reconstruction methods in safety-critical applications. Dealing with data-dependent uncertainty and model uncertainty, also known as aleatoric (stochastic) and epistemic (systematic) uncertainty, respectively, has attracted much interest in computer vision over the last six years, along with effective tools for assessing the reliability of model predictions [1, 6, 21, 31, 35, 50]. In the context of inverse problems, uncertainty quantification has been addressed in a Bayesian framework by sampling the posterior distribution [9, 34], which captures either epistemic or aleatory uncertainty alone. A framework for deep learning based Bayesian inversion has been

established in [3]. Modeling uncertainty inherent in training data with respect to direct data consistent DL-based solvers requires further investigation.

In this work, we combine uncertainty estimation with data-consistent image reconstruction. For this purpose, we assume a parameterized Laplace distribution for the residuals between reconstructions and true images. Incorporating the uncertainty estimate into the reconstruction loss enables the simultaneous learning of a reconstruction and the corresponding uncertainty map. The uncertainty map can be used to infer the quality of a reconstructed image in the absence of ground truth. Data consistency is realized by null space networks. The approach is comprehensively tested using limited data problems in CT and MRI as practical use cases.

2 Methods

Consider the inverse problem (1.1) where $x \in \mathbb{R}^n$ stands for the vectorized image to be recovered and $\mathcal{A} \in \mathbb{R}^{m \times n}$ is the forward matrix. In this section, we introduce the scheme of null space networks and other networks architectures that will be used as benchmarks. We discuss what adjustments can be made in the loss function during network training to simultaneously derive an uncertainty map while reconstructing the signal.

2.1 Image reconstruction

A reconstruction method for (1.1) is a family $(\mathcal{M}_\alpha)_{\alpha \in \Lambda}$ of potentially nonlinear operators $\mathcal{M}_\alpha: \mathbb{R}^m \rightarrow \mathbb{R}^n: y \mapsto \mathcal{M}_\alpha(y)$ mapping noisy data $y \in \mathbb{R}^m$ to the reconstruction $\mathcal{M}_\alpha(y) \in \mathbb{R}^n$. The crucial property is the convergence of the method to a right inverse $\mathcal{M}_0: \text{ran}(\mathcal{A}) \subseteq \mathbb{R}^m \rightarrow \mathbb{R}^n$ of the forward matrix defined by the property $\mathcal{A} \circ \mathcal{M}_0 \circ \mathcal{A} = \mathcal{A}$; see [27, 42]. Convergence of the reconstruction method means that for all $x \in \mathbb{R}^n$ we have

$$\sup_y \|\mathcal{M}_{\hat{\alpha}(\delta, y)}(y) - \mathcal{M}_0(\mathcal{A}x)\| \rightarrow 0 \text{ as } \delta \rightarrow 0, \quad (2.1)$$

where the supremum is taken over all noisy data $y \in \mathbb{R}^m$ with $\|y - \mathcal{A}x\| \leq \delta$ and the parameter $\alpha = \hat{\alpha}(\delta, y)$ is chosen dependent on the noise level and the actual data. If $\Lambda = \mathbb{R}_{>0}$ then $\hat{\alpha}$ is called parameter choice and (2.1) means that $((\mathcal{M}_\alpha)_{\alpha > 0}, \hat{\alpha})$ is a regularization method for (1.1), see [17, 46].

Classical reconstruction methods converge to the Moore-Penrose inverse $\mathcal{M}_0 = \mathcal{A}^\dagger$ of \mathcal{A} . If \mathcal{A} has linearly independent columns, then $\mathcal{A}^\dagger := (\mathcal{A}^* \mathcal{A})^{-1} \mathcal{A}^*$ where \mathcal{A}^* is the conjugate transpose. Examples of such reconstruction methods include Tikhonov regularization, truncated SVD or Landweber regularization. Different examples are variational regularization, where the right inverse \mathcal{M}_0 is defined as a quasi-solution having minimal value

of a so-called regularization functional. More recent examples are neural network based methods as we discuss next.

2.2 Learned reconstruction

In learned reconstruction methods, $(\mathcal{M}_\theta)_{\theta \in \Theta}$ consist of a rich family of mappings that typically depends on a high dimensional parameter θ and includes neural networks. This allows to extract knowledge from available training data such as databases of medical images [23]. NNs provide state-of-the-art for a variety of image processing tasks such as segmentation [5, 40], image synthesis [6, 14] and reconstruction [26, 52]. The steady improvements in CNN architectures is evident in the increase in accuracy for medical imaging applications, with the U-net [45] architecture still considered a high-performance model due to its ability to learn reliably in a limited data environment.

In the context of image reconstruction, popular approaches use networks of the form [29, 33]

$$\mathcal{M}_{\alpha, \theta} := \mathcal{N}_\theta \circ \mathcal{B}_\alpha \quad (2.2)$$

where $(\mathcal{B}_\alpha)_{\alpha > 0}$ is a classical reconstruction method and $(\mathcal{U}_\theta)_{\theta \in \Theta}$ a neural network architecture. Here an initial reconstruction $\mathcal{B}_\alpha(y)$ is found based on data y and a network is trained to enhance this reconstruction according to the trained parameters. While \mathcal{B}_α serves as an approximate inversion of the forward operator, the learned part can be seen as a correction of the artifacts such aliasing artifacts for compressed sensing in MRI or streak artifacts in limited angle CT. Specifically, in our work we will use residual networks of the form

$$\mathcal{R}_\theta^{(1)} := \text{Id} + \mathcal{U}_{\theta_1}, \quad (2.3)$$

$$\mathcal{R}_\theta^{(2)} := (\text{Id} + \mathcal{U}_{\theta_2}) \circ (\text{Id} + \mathcal{U}_{\theta_1}), \quad (2.4)$$

defined by an arbitrary network architecture $(\mathcal{U}_\theta)_{\theta \in \Theta}$, as reference method. In this study, $(\mathcal{U}_\theta)_{\theta \in \Theta}$ is considered a basic U-net architecture unless stated otherwise.

2.3 Data consistency

Given an initial reconstruction $\mathcal{B}_\alpha(y)$ and any network \mathcal{U}_θ , the two step method (2.2) does not promote data-consistent solutions. This means that even if $\mathcal{A}(\mathcal{B}_\alpha(y))$ is close to the data y this is not necessarily the case for $\mathcal{A}(\mathcal{U}_\theta \circ \mathcal{B}_\alpha)(y)$. To address this issue, one approach is to consider variational or iterative networks [2, 22, 28, 37, 47, 53]. However, models with an iterative architecture increase training time and controlling the data discrepancy $\|\mathcal{A}x - y\|^2$ is still challenging.

In order to derive a data-consistent solution from any reconstruction $x_{\mathcal{M}} := \mathcal{M}(y)$ we therefore consider the orthogonal projection

$$\mathcal{P}_{\mathbb{L}}(x_{\mathcal{M}}) = \arg \min_x \{\|x^{\mathcal{M}} - x\| \mid \mathcal{A}x = y\} \quad (2.5)$$

onto the solution space $\mathbb{L}(\mathcal{A}, y) := \{x \mid \mathcal{A}x = y\}$. An iterative way to calculate the orthogonal projection is given by the Landweber iteration [18, 38]

$$\forall j \in \mathbb{N}: \quad x_{j+1} = x_j - \lambda_j \mathcal{A}^*(\mathcal{A}x_j - y), \quad (2.6)$$

with initial value $x_0 = x_{\mathcal{M}}$, where $\lambda_j > 0$ is the stepsize for the j -th iteration. Note that the Landweber iterations (2.6) will be applied to the output of an already trained network. This allows for a large number of iterations, which ensures a good approximation of the projection onto \mathbb{L} and does not prolong the training process. In the ill-posed case, regularization can be integrated into (2.6) by early stopping, by replacing the forward model \mathcal{A} in with an operator having a stable pseudoinverse, or by replacing the data y with $(\mathcal{A} \circ \mathcal{B}_{\alpha})(y)$ in a two-step method $\mathcal{M}_{\alpha, \theta} = \mathcal{N}_{\theta} \circ \mathcal{B}_{\alpha}$.

2.4 Null space networks

In order to obtain solutions that are guaranteed to be data consistent, the concept of null space networks has been introduced in [48]. Null space networks basically denote residual networks that only modify components in the kernel of the forward problem. More formally, let $(\mathcal{U}_{\theta})_{\theta \in \Theta}$ be any network architecture and \mathcal{P}_0 denote the orthogonal projection onto the null space of \mathcal{A} . We then call

$$\Psi_{\theta}^{(1)} := (\text{Id} + \mathcal{P}_0 \circ \mathcal{U}_{\theta}), \quad (2.7)$$

$$\Psi_{\theta}^{(2)} := (\text{Id} + \mathcal{P}_0 \circ \mathcal{U}_{\theta_2}) \circ (\text{Id} + \mathcal{P}_0 \circ \mathcal{U}_{\theta_1}), \quad (2.8)$$

null space network associated to operator \mathcal{A} with architecture $(\mathcal{U}_{\theta})_{\theta \in \Theta}$ and cascade length one and two, respectively. An illustration of a null space network (with cascade length one) is shown in Figure 1.

We use null space networks in the context of the two-step reconstruction (2.2). Note that any null space network is a particular form of a residual network, where the residual correction only operates in the null space of the forward operator \mathcal{A} and by that ensuring data consistency $\mathcal{A} \circ \Psi_{\theta}^{(2)} = \mathcal{A} \circ \Psi_{\theta}^{(1)} = \mathcal{A}$. One readily verifies that $(\text{Id} + \mathcal{P}_0 \circ \mathcal{U}_{\theta})(\mathcal{A}^{\ddagger}y) = \mathcal{P}_{\mathbb{L}}(\text{Id} + \mathcal{U}_{\theta})(\mathcal{A}^{\ddagger}y)$. Therefore two-step reconstructions using either a null space network $\text{Id} + \mathcal{P}_0 \circ \mathcal{U}_{\theta}$ or a projected residual network $\mathcal{P}_{\mathbb{L}} \circ (\text{Id} + \mathcal{U}_{\theta})$ for given parameters coincide [7]. However, training a null-space network is clearly different from training the residual network followed by a projection onto $\mathcal{P}_{\mathbb{L}}$.




Figure 1: **Null space network.** Illustration of a null space network $\text{Id} + \mathcal{P}_0 \mathcal{U}_\theta$. The last layer projects the output of the residual part \mathcal{U}_θ on the null space of forward operator \mathcal{A} .

2.5 Network training

Let $(\mathcal{M}_\theta)_{\theta \in \Theta}$ be a reconstruction method whose parameter vector θ is selected based on training data $(\mathcal{A}x_i, x_i)_{i=1}^N$. Note that the subscript α for the classical reconstruction $(\mathcal{B}_\alpha)_{\alpha > 0}$ in (2.2) is omitted here for simplicity. The common training procedure minimizes the empirical risk

$$\mathcal{L}(\theta) := \frac{1}{N} \sum_{i=1}^N \|\mathcal{M}_\theta(\mathcal{A}x_i) - x_i\|_1 \quad (2.9)$$

According to [6, 50], the underlying assumption can be interpreted that every component of the residual image $(\epsilon_{\theta,p})_{p=1}^n := \mathcal{M}_\theta(\mathcal{A}x) - x$ follows a Laplace distribution $\epsilon_{\theta,p} \sim \text{Laplace}(0, \sigma)$ with density $\exp(-|\epsilon_{\theta,p}|/\sigma)/(2\sigma)$. Maximum likelihood optimization yields

$$\begin{aligned} \hat{\theta} &= \arg \max_{\theta} \frac{1}{N} \prod_{i,p} \frac{1}{2\sigma} \exp(-|\epsilon_{\theta,p}|/\sigma) \\ &= \arg \min_{\theta} \frac{1}{N} \sum_{i,p} (|\mathcal{M}_\theta(\mathcal{A}x_i) - x_i|_p)/\sigma + \frac{1}{N} \log(2\sigma) \\ &= \arg \min_{\theta} \frac{1}{N} \sum_i \|\mathcal{M}_\theta(\mathcal{A}x_i) - x_i\|_1/\sigma + \frac{1}{N} \log(2\sigma). \end{aligned}$$

Hence, for pixel independent scale parameter σ , maximum likelihood optimization minimization recovers minimization of the empirical risk (2.9).

2.6 Uncertainty quantification

Obviously, the assumption of a fixed scale for all residuals $\epsilon_{\theta,p} := (\mathcal{M}_\theta(\mathcal{A}x) - x)_p$ is quite strong. In the context of uncertainty quantification it can be relaxed considering input-dependent scale parameter. To do so, the scale is now taken as a whole image $\sigma = (\sigma_p)_{p=1}^n$ modeled as a function of data $y \in \mathbb{R}^m$. More precisely, the network now simultaneously predicts a reconstruction $x_\theta = \mathcal{M}_\theta^x(y)$ and its corresponding scale map $\sigma_\theta = \mathcal{M}_\theta^\sigma(y)$. In our implementation, the reconstruction network \mathcal{M}_θ^x will be taken as a two-step architecture (2.2) combined with the null space network (2.7) or (2.8). The uncertainty estimating network $\mathcal{M}_\theta^\sigma$ has the similar architecture and nearly shares all parameters. More precisely, we split the network architecture \mathcal{U}_θ in the last null space block to obtain two output-branches, where one gives the data-consistent reconstruction and one the uncertainty map. For further implementation details we refer to the available github repository <https://github.com/anger-man/cascaded-null-space-learning>. The networks \mathcal{M}_θ^x and $\mathcal{M}_\theta^\sigma$ are trained by simultaneously minimizing the uncertainty-aware loss

$$\mathcal{L}^{\text{unc}}(\theta) := \frac{1}{N} \sum_{i,p} \frac{|(\mathcal{M}_\theta^x(\mathcal{A}x_i) - x_i)_p|}{(\mathcal{M}_\theta^\sigma(\mathcal{A}x_i))_p} + \frac{1}{N} \sum_{i,p} \log(2 \cdot (\mathcal{M}_\theta^\sigma(\mathcal{A}x_i))_p). \quad (2.10)$$

By taking $(\mathcal{M}_\theta^\sigma(\mathcal{A}x_i))_p = \sigma$ as a constant, (2.10) would reduce to (2.9). However simultaneous minimization allows both reconstructing the image and the associated uncertainty. For regions with large absolute residuals we obtain large values in the scale map corresponding to high uncertainty. At the same time, the logarithmic term penalizes the model to avoid predicting high uncertainty for all pixel regions.

The benefits of uncertainty-aware models are versatile. The training process is more robust against OOD data in the training set, insufficient measurement quality is established during prediction, and unwanted objects or artifacts that may effect the reconstruction are detected and localized in the absence of ground truth data.

2.7 Implemented methods

In our simulations we compare a total of eight different learned reconstruction methods against the ground truth and the Moore-Penrose inverse:

- x_0 : Ground truth used for error evaluation.
- $x^\dagger = \mathcal{A}^\dagger(y)$: Pseudoinverse reconstruction.
- $x_{\mathcal{R},1} = \mathcal{R}_\theta^{(1)}(x^\dagger)$: Two-step reconstruction (2.2) + (2.3).
- $x_{\mathcal{R},2} = \mathcal{R}_\theta^{(2)}(x^\dagger)$: Two-step reconstruction (2.2) + (2.4).

- $\mathcal{P}_{\mathbb{L}}(x_{\mathcal{R},1})$: Projection of $x_{\mathcal{R},1}$ onto $\mathbb{L}(\mathcal{A}, y)$ (2.5).
- $\mathcal{P}_{\mathbb{L}}(x_{\mathcal{R},2})$: Projection of $x_{\mathcal{R},2}$ onto $\mathbb{L}(\mathcal{A}, y)$ (2.5).
- $x_{\Psi,1} = \Psi_{\theta}^{(1)}(x^{\dagger})$: Null space reconstruction (2.2) + (2.7).
- $x_{\Psi,2} = \Psi_{\theta}^{(2)}(x^{\dagger})$: Null space reconstruction (2.2) + (2.8).
- $x_{\Psi,1}^{\text{unc}}$: Uncertainty-aware reconstruction using (2.7).
- $x_{\Psi,2}^{\text{unc}}$: Uncertainty-aware reconstruction using (2.8).

The underlying networks $\mathcal{R}_{\theta}^{(1)}$, $\mathcal{R}_{\theta}^{(2)}$, $\Psi_{\theta}^{(1)}$, $\Psi_{\theta}^{(2)}$ are trained using the MAE loss function (2.9), except for the uncertainty-aware networks $x_{\Psi,1}^{\text{unc}}$, $x_{\Psi,2}^{\text{unc}}$, which are trained by minimizing the uncertainty aware loss (2.10). For the reconstructions $\mathcal{P}_{\mathbb{L}}(x_{\mathcal{R},1})$ and $\mathcal{P}_{\mathbb{L}}(x_{\mathcal{R},2})$ the projection onto the space of all solutions of $\mathcal{A}x = y$ is calculated using 15 iterations of the Landweber iteration. Reconstruction quality is assessed by peak-signal-to-noise ratio (PSNR) and the structural-similarity index (SSIM).

3 Results

In this section we present the numerical results for the discussed reconstruction methods. To this end, we will consider two examples of inverse problem of the form (1.1): (A) Reconstruction from subsampled Fourier measurements on fastMRI dataset [36, 54]; (B) reconstruction of limited angle Radon measurements on a self-acquired toy dataset. Note that it is not the aim of this study to propose a new state-of-the-art model to solve specific inverse problems. With these experiments, we show that simple adjustments to the network architecture and slight modifications to the loss function result in a significant increase in reconstruction quality and the benefit of pixel-based confidence prediction for each reconstruction. Note also that these modifications are implemented and evaluated considering the basic U-network architecture widely used in image processing tasks. However, our rationale can be applied to any network architecture.

3.1 Study A: 4-fold accelerated MRI

The first case study is devoted to a real-world scenario of compressed sensing in MRI. To this end, we use the public fastMRI dataset, which consists of 1594 multi-coil knee MRI scans (<https://fastmri.med.nyu.edu/>). The experiments are based on the subset of 571 randomly selected diagnostic cases without fat-suppression. For a practice-oriented evaluation on unseen individuals, 5% of the scans have been randomly selected to represent the test cases. This yields nearly 20000 samples for training and 1000 samples for

testing. For the training set $(x_i)_{i=1}^N$, we draw magnitude images $x_i \in \mathbb{R}^{320 \times 320}$ obtained by fully sampled multi-coil data. Similar to [22], our model corresponds to the simpler modality of single-coil MRI. The fastMRI challenge also provides emulated single-coil data, that is drawn in a retrospective way from the multi-coil measurements. However, we decided to sample from the multi-coil reconstructions in the favor of higher image quality and noise-reduced measurements.




Figure 2: **Accelerated MRI reconstruction results.** Ground truth image x_0 (central slice of a randomly selected volume from the test set), pseudoinverse reconstruction x^\ddagger , projected two-step reconstructions $\mathcal{P}_L(x_{\mathcal{R},1})$, $\mathcal{P}_L(x_{\mathcal{R},2})$ and null-space reconstructions $x_{\Psi,1}$, $x_{\Psi,2}$.

Implementation details

Here, the forward operator takes the form $\mathcal{A} = \mathcal{S} \circ \mathcal{F}$, where $\mathcal{F}: \mathbb{C}^{320 \times 320} \rightarrow \mathbb{C}^{320 \times 320}$ is the two-dimensional discrete Fourier transform and $\mathcal{S}: \mathbb{C}^{320 \times 320} \rightarrow \mathbb{C}^{320 \times 320}$ is a subsampling operator, implemented via a binary mask omitting Fourier measurement lines in the phase-encoding direction. The amount of dropped k-space lines depends on the acceleration factor and in our case is set to 75 %, following the subsampling scheme in [32]. The Fourier transform is a unitary operator. Therefore, the Moore-Penrose inverse of \mathcal{A} is given by the conjugate transpose of the forward operator, $\mathcal{A}^\ddagger = \mathcal{A}^* = \mathcal{F}^* \circ \mathcal{S}$. Note that

we actually have complex-valued data. Thus before feeding the instances to a neural network function, we concatenate real and imaginary part such that they are treated as separate channels, $x^\ddagger \in \mathbb{R}^{2(320 \times 320)}$. The data processing pipeline and the models are implemented in Python using PyTorch library (<https://pytorch.org/>) for GPU-accelerated computation. For any details on training and hyperparameter selection, we refer to our github repository <https://github.com/anger-man/cascaded-null-space-learning>.

Table 1: **Accelerated MRI**. Quantitative evaluation on nearly 1000 test slices. The reported metrics are PSNR and $100 \times \text{SSIM}$ (higher is better).

benchmark				
	$x_{\mathcal{R},1}$	$\mathcal{P}_{\mathbb{L}}(x_{\mathcal{R},1})$	$x_{\mathcal{R},2}$	$\mathcal{P}_{\mathbb{L}}(x_{\mathcal{R},2})$
PSNR	31.29	32.12	31.5	32.35
$100 \times \text{SSIM}$	85.13	86.17	85.44	86.5
null space networks				
	$x_{\Psi,1}$	$x_{\Psi,1}^{\text{unc}}$	$x_{\Psi,2}$	$x_{\Psi,2}^{\text{unc}}$
PSNR	32.25	32.28	33.39	33.44
$100 \times \text{SSIM}$	86.31	86.4	88.2	88.29

Reconstruction results

The quantitative results shown in Table 1 demonstrate the superior accuracy of DL based solvers. The benchmark $x_{\mathcal{R},1}$ already achieves a PSNR of 31.29 and a SSIM of 0.851 for the set of nearly 1k unseen test slices. The result is quite impressive when we consider an acceleration rate of 4 in the compressed sensing scenario, i.e., only 25% of the lines in phase encoding directions have been kept. The projection of the previously trained basic method onto the solution space \mathbb{L} leads to approximately data-consistent solutions. Its impact is confirmed by a slight increase for both test metrics. The extension of $x_{\Psi,1}$ to a cascaded scheme $x_{\Psi,2}$ has a positive impact on the entire optimization procedure as indicated by a significant increase to 33.39 and 0.882 for PSNR and SSIM metric, respectively.

In Figure 2 we observe the prominent aliasing artifacts in the pseudoinverse reconstruction $x^\ddagger = \mathcal{A}^\ddagger(y)$. These artifacts are obviously removed in all learned reconstructions. The examination of the magnified areas shows that the contrast in the high-frequency details is higher for the null space networks than for the residual networks. Using the null space approach for joint reconstruction and uncertainty prediction does not significantly change the reconstruction quality. Therefore, it is possible to incorporate uncertainty awareness into the cascade of null space blocks without affecting the re-

covery performance at all. Benefits of an uncertainty estimate are investigated in the following.




Figure 3: **Position-related uncertainty in accelerated MRI.** The presence of an atypical irregularity in the knee causes artifacts on the entire horizontal trace of the pseudo-inverse solution x^\ddagger (first column) and leads to an unpredictable behavior of the reconstruction $x_{\Psi,2}^{\text{unc}}$ at the corresponding locations (second column). In the simultaneously predicted uncertainty maps (third column), these irregular objects are clearly indicated by higher uncertainty values.

Uncertainty quantification

The uncertainty reconstruction provides a pixel-wise confidence value additionally to the reconstruction, which can be used to detect suspicious objects that are not covered by the training distribution. Furthermore, uncertainty evaluation based on these region-based confidence maps can serve as a quality assessment of the overall prediction. In the following, both advantageous properties are investigated numerically.

Unknown local properties of a patient’s knee under investigation can lead to unpredictable behavior of the reconstruction method in that region. Unknown properties here refer to the occurrence of non-ubiquitous objects such as cancer or bone fractures. We assume here that the presence of these objects is true, i.e., the irregularity is not caused by the measurement process itself. For the simulations, we select two magnitude images from the test cases and add salt-and-pepper noise to small subregions. We then

apply the forward operator to obtain measurement data y_{cor} describing the actual measured data if such an irregularity would be present in the examined knee. The results of applying the Moore-Penrose inverse \mathcal{A}^\ddagger and the uncertainty aware null space approach to y_{cor} are shown in Figure 3. As can be seen, the predicted uncertainty maps show large values at the corresponding location, visualized by a brighter color intensity. We claim that this characteristic of the uncertainty-aware null space networks clearly identifies unreliable regions in the reconstruction caused by the presence of abnormal features in an examined knee.




Figure 4: **Prediction error versus mean uncertainty for accelerated MRI.** A clear correlation is found between the prediction error and the average value of the modeled uncertainty maps. Perturbations of the measurement data simultaneously increases prediction error and mean uncertainty score.

The experiments in this study are based on the root-sum-of-squares reconstruction of multi-coil measurements. As a consequence, the slices in the training set $(x_i)_{i=1}^N$ denote noise-reduced samples. We select 100 test slices and perturb the measured data by additive noise, i.e. we generate measurements $y_i = \mathcal{A}x_i + \epsilon$ for $i = 1, \dots, N$, where $\|\epsilon\|_2 \leq \delta$ for $\delta \in \{0, 5 \times 10^3, 9 \times 10^3\}$. For each test slice, we define the average pixel value of the uncertainty map as the image-based uncertainty score and plot it against the corresponding mean absolute residual between reconstruction and noise-free ground truth in Figure 4. In the noise-reduced case $\epsilon = 0$, we infer a fairly clear positive correlation between the mean uncertainty and the mean absolute residuals. For $\|\epsilon\|_2 \leq 5 \times 10^3$ the prediction error as well as the uncertainty values increase. The effect is amplified by increasing the scale of the additive Gaussian noise to $\|\epsilon\|_2 \leq 9 \times 10^3$. The perturbation of the measured data during testing leads to OOD data, which is clearly indicated by an increase in uncertainty scores. We conclude that the proposed method

can be used to assess reconstruction quality (correlation with MAE) and to detect out-of-distribution data (larger scale for noise-perturbed data).




Figure 5: **Limited view CT reconstruction results.** Ground truth image x_0 (randomly selected), pseudoinverse reconstruction x^\ddagger , projected two-step reconstructions $\mathcal{P}_L(x_{\mathcal{R},1})$, $\mathcal{P}_L(x_{\mathcal{R},2})$ and null-space reconstructions $x_{\Psi,1}$, $x_{\Psi,2}$.

3.2 Study B: limited angle CT

In our second study, we address limited angle CT, where the angular range of the available data is limited to a strict subset. The experiments are based on artificial phantoms. The structure of the phantoms $(x_i)_{i=1}^N$ of size 192×192 pixels is based on the Shepp-Logan phantom. Each sample consists of a randomly rotated phantom disc that contains smaller ellipses and rectangles drawn in a random fashion. Furthermore, some instances also contain high-frequency details (cf. Figure 5). All in all we generated 2000 toy samples for training and another 200 slices for testing.

Implementation details

The forward operator takes the form $\mathcal{A} = \mathcal{S} \circ \mathcal{R}$, where \mathcal{R} is the discretized Radon transform with full angular range $[0^\circ, 180^\circ)$. Here, $\mathcal{S} \in \{0, 1\}^{m \times m}$ is a binary mask for

removing angular projections such that we obtain a total of 60 angular projections with angular directions in $[0^\circ, 120^\circ]$. We again make use of the Moore-Penrose inverse as initial reconstruction method. In this case, the inverse is approximated by the singular value decomposition (SVD) [18]. Python code for the discretization of \mathcal{A} and a stable inexact pseudoinverse using truncated SVD is provided in our github repository <https://github.com/anger-man/cascaded-null-space-learning>.

Table 2: **Limited view CT**. Quantitative evaluation on 200 test phantoms. The reported metrics are PSNR and SSIM $\times 100$ (higher is better).

benchmark				
	$x_{\mathcal{R},1}$	$\mathcal{P}_{\mathbb{L}}(x_{\mathcal{R},1})$	$x_{\mathcal{R},2}$	$\mathcal{P}_{\mathbb{L}}(x_{\mathcal{R},2})$
PSNR	34.88	34.9	35.97	35.98
$100 \times \text{SSIM}$	98.38	98.36	98.71	98.7
null space networks				
	$x_{\Psi,1}$	$x_{\Psi,1}^{\text{unc}}$	$x_{\Psi,2}$	$x_{\Psi,2}^{\text{unc}}$
PSNR	35.87	35.8	37.87	37.14
$100 \times \text{SSIM}$	97.47	97.59	98.78	98.68

Reconstruction results

Quantitative results are summarized in Table 2. The baseline method $x_{\mathcal{R},1}$ yields a PSNR of 34.88 and SSIM of 0.984. Both metrics increase to 35.97 and 0.987 if using the cascaded U-net architecture $x_{\mathcal{R},2}$. Interestingly, subsequent projection onto \mathbb{L} for $x_{\mathcal{R},1}$ and $x_{\mathcal{R},2}$ does not enhance results at all. This is mainly due to the small stepsize of $\lambda = 0.003$ for the Landweber iterations. However, increasing the stepsize leads to instabilities and thus insufficient solutions. Again, leveraging the methods $x_{\mathcal{R},1}$ and $x_{\mathcal{R},2}$ to the null space schemes $x_{\Psi,1}$ and $x_{\Psi,2}$ significantly increases the PSNR. Best results have been achieved by the cascaded null space network $x_{\Psi,2}$, which yields a PSNR and SSIM of 37.87 and 0.988, respectively. Better reconstruction performance is only observed in the PSNR while the SSIM does not seem to improve. However, for the toy dataset, the SSIM can be observed to already be rather small. Analogously to study A, incorporating the uncertainty-aware risk functional in (2.10) does not show any impact on the recovery performance of the null space networks.

In Figure 5(b) we see the characteristic limited angle streak artifacts by in the pseudoinverse reconstruction. Again, all learned reconstruction strategies are able to remove these artifacts completely. The single U-net architectures shown in Figure 5 (c) and (e) have problems when reconstructing high-frequency details. This can be clearly observed

when we look at the enlarged patches in the red boxes. Qualitative reconstruction of these patches is more satisfying for the cascaded architectures in Figure 5 (d) and (f), while the latter yields the best visual reconstruction compared to the ground truth.

Uncertainty quantification




Figure 6: **Position-related uncertainty in limited view CT.** The presence of an unknown material in the synthetic phantom causes artifacts classical reconstruction x^\ddagger (first column). While the reconstruction method $\Psi_\theta^{(2)}(x^\ddagger)$ reliably reconstructs the phantom (second column) the predicted uncertainty maps (third column) clearly identifies this irregular object and unreliable regions in the prediction by a higher uncertainty values.

Similar to Study A, the benefits of simultaneously modeling the uncertainty maps for the toy phantom dataset are qualitatively explored. In this case, we simulate OOD data by adding a square object to the phantom. This may be interpreted as a very simple simulation of for example metal being present inside the human body, which can be the case due to dentures or artificial joints. Metal artifacts in CT can be particularly problematic in the areas surrounding the metal objects, where the image quality can be severely degraded. Again, behavior of the Moore-Penrose inverse and the cascaded null space network is investigated on two test phantoms shown in Figure 6. The inserted square perturbation yields strong streak artifacts for the model-based reconstruction x^\ddagger (first row) and is visible in the DL-based reconstruction of $\Psi_\theta^{(2)}(x^\ddagger)$ (second row). The position of the square is clearly marked with high intensity in the corresponding uncer-

tainty maps (third row). Also in the setting of limited angle CT the uncertainty-aware null space network clearly identifies unreliable regions during reconstruction caused by presence of unknown metal artifacts.

4 Conclusion

In this paper we presented a learning-based method for simultaneous image reconstruction and uncertainty estimation. The proposed uncertainty estimation procedure introduces a second output branch into the reconstruction network, which can be seen as scale map for the Laplace distributed residuals between reconstructed image and ground truth. This leads to small uncertainty values for regions that are easy to learn based on the training data. On the other hand, rapid changes in pixel intensities and out of distribution data create new regions of high frequency detail, a property for which the second output branch has already been developed to predict high uncertainty. Our experiments demonstrate that simple modifications to the network architecture and slight modifications to the risk functional yield uncertainty information and increased reconstruction quality. The approach is implemented and evaluated based on the standard U-net architecture commonly used in image processing tasks. However, our framework can also be combined with more complex architectures or integrated into state-of-the-art image reconstruction methods.

Acknowledgment

The contribution of C. A. is supported by VASCage – Centre on Clinical Stroke Research. VASCage is a COMET Centre within the Competence Centers for Excellent Technologies (COMET) programme and funded by the Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology, the Federal Ministry of Labour and Economy, and the federal states of Tyrol, Salzburg and Vienna. COMET is managed by the Austrian Research Promotion Agency (Österreichische Forschungsförderungsgesellschaft). The contribution of S. G. is part of a project that has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 847476. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

References

- [1] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, et al. A review of uncertainty

- quantification in deep learning: Techniques, applications and challenges. *Information fusion*, 76:243–297, 2021.
- [2] J. Adler and O. Öktem. Solving ill-posed inverse problems using iterative deep neural networks. *Inverse Problems*, 33(12):124007, 2017.
- [3] J. Adler and O. Öktem. Deep bayesian inversion. *arXiv:1811.05910*, 2018.
- [4] F. Altekrüger, A. Denker, P. Hagemann, J. Hertrich, P. Maass, and G. Steidl. Patchnr: Learning from small data by patch normalizing flow regularization. *CoRR*, abs/2205.12021, 2022.
- [5] C. Angermann and M. Haltmeier. Deep structure learning using feature extraction in trained projection space. *Computers & electrical engineering*, 92:107097, 2021.
- [6] C. Angermann, M. Haltmeier, and A. R. Siyal. Unsupervised joint image transfer and uncertainty quantification using patch invariant networks. In *Computer vision—ECCV 2022 workshops: Tel Aviv, Israel, October 23–27, 2022, proceedings, part VIII*, pages 61–77. Springer, 2023.
- [7] S. Antholzer, J. Schwab, and M. Haltmeier. Deep learning versus ℓ^1 -minimization for compressed sensing photoacoustic tomography. In *2018 IEEE International Ultrasonics Symposium (IUS)*, pages 206–212. IEEE, 2018.
- [8] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb. Solving inverse problems using data-driven models. *Acta Numerica*, 28:1–174, 2019.
- [9] J. M. Bardsley. *Computational Uncertainty Quantification for Inverse Problems: An Introduction to Singular Integrals*. SIAM, 2018.
- [10] 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: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 57(6):1086–1098, 2007.
- [11] L. Borg, J. Friel, J. S. Jørgensen, and E. T. Quinto. Analyzing reconstruction artifacts from arbitrary incomplete x-ray ct data. *SIAM journal on imaging sciences*, 11(4):2786–2814, 2018.
- [12] V. Buhrmester, D. Münch, and M. Arens. Analysis of explainers of black box deep neural networks for computer vision: A survey. *Machine learning and knowledge extraction*, 3(4):966–989, 2021.
- [13] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE transactions on information theory*, 52(2):489–509, 2006.

- [14] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1):53–65, 2018.
- [15] D. Donoho. Compressed sensing. *IEEE transactions on information theory*, 52(4):1289–1306, 2006.
- [16] M. Elad. Deep, deep trouble: deep learning’s impact on image processing, mathematics, and humanity. *SIAM news*, 50(4), 2017.
- [17] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of inverse problems*, volume 375. Springer Science & Business Media, 1996.
- [18] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of inverse problems*, volume 375. Springer Science & Business Media, 1996.
- [19] J. Friel and E. T. Quinto. Characterization and reduction of artifacts in limited angle tomography. *Inverse problems*, 29(12):125007, nov 2013.
- [20] J. Friel and E. T. Quinto. Limited data problems for the generalized radon transform in \mathbb{R}^n . *SIAM journal on mathematical analysis*, 48(4):2301–2318, 2016.
- [21] Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International conference on machine learning*, pages 1050–1059. PMLR, 2016.
- [22] M. Genzel, J. Macdonald, and M. Marz. Solving inverse problems with deep neural networks-robustness included. *IEEE transactions on pattern analysis and machine intelligence*, 2022.
- [23] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.
- [24] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, editors, *Advances in neural information processing systems*, volume 27. Curran Associates, Inc., 2014.
- [25] M. Grasmair, M. Haltmeier, and O. Scherzer. Sparse regularization with lq penalty term. *Inverse Problems*, 24(5):055020, 2008.
- [26] H. Gupta, K. H. Jin, H. Q. Nguyen, M. T. McCann, and M. Unser. Cnn-based projected gradient descent for consistent ct image reconstruction. *IEEE transactions on medical imaging*, 37(6):1440–1453, 2018.
- [27] M. Haltmeier and L. Nguyen. *Regularization of Inverse Problems by Neural Networks*, pages 1065–1093. Springer International Publishing, Cham, 2023.

- [28] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll. Learning a variational network for reconstruction of accelerated mri data. *Magnetic resonance in medicine*, 79(6):3055–3071, 2018.
- [29] Y. S. Han, J. Yoo, and J. C. Ye. Deep residual learning for compressed sensing ct reconstruction via persistent homology analysis. *arXiv:1611.06391*, 2016.
- [30] P. C. Hansen. *Discrete inverse problems: insight and algorithms*. SIAM, 2010.
- [31] R. Hu, Q. Huang, S. Chang, H. Wang, and J. He. The MBPEP: a deep ensemble pruning algorithm providing high quality uncertainty prediction. *Applied intelligence*, 49:2942–2955, 2019.
- [32] C. M. Hyun, H. P. Kim, S. M. Lee, S. Lee, and J. K. Seo. Deep learning for undersampled MRI reconstruction. *Physics in medicine & biology*, 63(13):135007, 2018.
- [33] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE transactions on image processing*, 26(9):4509–4522, 2017.
- [34] J. Kaipio and E. Somersalo. *Statistical and computational inverse problems*, volume 160. Springer Science & Business Media, 2006.
- [35] A. Kendall and Y. Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30, 2017.
- [36] F. Knoll, J. Zbontar, A. Sriram, M. J. Muckley, M. Bruno, A. Defazio, M. Parente, K. J. Geras, J. Katsnelson, H. Chandarana, et al. fastMRI: A publicly available raw k-space and dicom dataset of knee images for accelerated mr image reconstruction using machine learning. *Radiology. artificial intelligence*, 2(1), 2020.
- [37] A. Kofler, M. Haltmeier, C. Kolbitsch, M. Kachelrieß, and M. Dewey. A u-nets cascade for sparse view computed tomography. In *Machine Learning for Medical Image Reconstruction: MLMIR 2018*, pages 91–99. Springer, 2018.
- [38] L. Landweber. An iteration formula for fredholm integral equations of the first kind. *American journal of mathematics*, 73(3):615–624, 1951.
- [39] H. Li, J. Schwab, S. Antholzer, and M. Haltmeier. Nett: Solving inverse problems with deep neural networks. *Inverse Problems*, 36(6):065005, 2020.
- [40] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.

- [41] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. Compressed sensing mri. *IEEE signal processing magazine*, 25(2):72–82, 2008.
- [42] M. Z. Nashed. Inner, outer, and generalized inverses in banach and hilbert spaces. *Numer. func. anal. opt.*, 9(3-4):261–325, 1987.
- [43] G. Ongie, A. Jalal, C. A. Metzler, R. G. Baraniuk, A. G. Dimakis, and R. Willett. Deep learning techniques for inverse problems in imaging. *IEEE journal on selected areas in information theory*, 1(1):39–56, 2020.
- [44] B. Popilka, S. Setzer, and G. Steidl. Signal recovery from incomplete measurements in the presence of outliers. *Inverse Problems and Imaging*, 1(4):661, 2007.
- [45] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [46] O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen. *Variational methods in imaging*. Springer, 2009.
- [47] J. Schlemper, J. Caballero, J. V. Hajnal, A. Price, and D. Rueckert. A deep cascade of convolutional neural networks for mr image reconstruction. In *Information Processing in Medical Imaging: 25th International Conference, IPMI 2017, Boone, NC, USA, June 25-30, 2017, Proceedings 25*, pages 647–658. Springer, 2017.
- [48] J. Schwab, S. Antholzer, and M. Haltmeier. Deep null space learning for inverse problems: convergence analysis and rates. *Inverse problems*, 35(2):025008, 2019.
- [49] J. Schwab, S. Antholzer, and M. Haltmeier. Big in japan: Regularizing networks for solving inverse problems. *Journal of mathematical imaging and vision*, 62(3):445–455, 2020.
- [50] U. Upadhyay, Y. Chen, and Z. Akata. Robustness via uncertainty-aware cycle consistency. *Advances in neural information processing systems*, 34:28261–28273, 2021.
- [51] G. Wang, J. C. Ye, and B. De Man. Deep learning for tomographic image reconstruction. *Nature Machine Intelligence*, 2(12):737–748, 2020.
- [52] H. Yao, F. Dai, S. Zhang, Y. Zhang, Q. Tian, and C. Xu. Dr2-net: Deep residual reconstruction network for image compressive sensing. *Neurocomputing*, 359:483–493, 2019.
- [53] G. Yiasemis, J.-J. Sonke, C. Sánchez, and J. Teuwen. Recurrent variational network: A deep learning inverse problem solver applied to the task of accelerated mri

reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 732–741, 2022.

- [54] J. Zbontar, F. Knoll, A. Sriram, T. Murrell, Z. Huang, M. J. Muckley, A. Defazio, R. Stern, P. Johnson, M. Bruno, et al. fastMRI: An open dataset and benchmarks for accelerated mri. *arXiv preprint arXiv:1811.08839*, 2018.