

Nr. 78
6. February 2021

Preprint-Series: Department of Mathematics - Applied Mathematics

Augmented NETT Regularization of Inverse Problems

D. Obmann, L. Nguyen, J. Schwab, M. Haltmeier



AppliedMathematics

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

Augmented NETT Regularization of Inverse Problems

Daniel Obmann¹, Linh Nguyen², Johannes Schwab¹, and Markus Haltmeier¹

¹Department of Mathematics, University of Innsbruck
Technikerstrasse 13, 6020 Innsbruck, Austria
{daniel.obmann,johannes.schwab,markus.haltmeier}@uibk.ac.at

²Department of Mathematics, University of Idaho
Moscow, ID 83844, lnguyen@uidaho.edu

February 06, 2021

Abstract

We propose aNETT (augmented NETwork Tikhonov) regularization as a novel data-driven reconstruction framework for solving inverse problems. An encoder-decoder type network defines a regularizer consisting of a penalty term that enforces regularity in the encoder domain, augmented by a penalty that penalizes the distance to the data manifold. We present a rigorous convergence analysis including stability estimates and convergence rates. For that purpose, we prove the coercivity of the regularizer used without requiring explicit coercivity assumptions for the networks involved. We propose a possible realization together with a network architecture and a modular training strategy. Applications to sparse-view and low-dose CT show that aNETT achieves results comparable to state-of-the-art deep-learning-based reconstruction methods. Unlike learned iterative methods, aNETT does not require repeated application of the forward and adjoint models, which enables the use of aNETT for inverse problems with numerically expensive forward models. Furthermore, we show that aNETT trained on coarsely sampled data can leverage an increased sampling rate without the need for retraining.

Keywords: inverse problems, regularization, stability guarantees, convergence rates, learned regularizer, computed tomography, neural networks

1 Introduction

Various applications in medical imaging, remote sensing and elsewhere require solving inverse problems of the form

$$y^\delta = \mathbf{K}x + \eta^\delta, \quad (1.1)$$

where $\mathbf{K}: \mathbb{X} \rightarrow \mathbb{Y}$ is an operator between Hilbert spaces modeling the forward problem, η^δ is the data perturbation, $y^\delta \in \mathbb{Y}$ is the noisy data and $x \in \mathbb{X}$ is the sought for signal. Inverse problems are well analyzed and several established approaches for its stable solution exist [13, 39]. Recently, neural networks and deep learning appeared as new paradigms for solving inverse problems [7, 17, 29, 31, 45]. Several approaches based on

deep learning have been developed, including post-processing networks [6, 21, 26, 30, 37], regularizing null-space networks [41, 42], plug-and-play priors [9, 38, 44], deep image priors [12, 43], variational networks [18, 23], network cascades [24, 40], learned iterative schemes [2, 4, 10, 20, 25, 47] and learned regularizers [5, 27, 28, 33].

Classical deep learning approaches may lack data consistency for unknowns very different from the training data. To address this issue, in [27] a deep learning approach named NETT (NETwork Tikhonov) regularization has been introduced which considers minimizers of the NETT functional

$$\mathcal{N}_{\alpha, y^\delta}(x) := \mathcal{D}(\mathbf{K}x, y^\delta) + \alpha \mathcal{Q}(\mathbf{E}(x)). \quad (1.2)$$

Here, \mathcal{D} is a similarity measure, $\mathbf{E}: \mathbb{X} \rightarrow \Xi$ is a trained neural network, $\mathcal{Q}: \Xi \rightarrow [0, \infty]$ a functional and $\alpha > 0$ a regularization parameter. In [27] it is shown that under suitable assumptions, NETT yields a convergent regularization method. This in particular includes provable stability guarantees and error estimates. Moreover, a training strategy has been proposed, where \mathbf{E} is trained such that $\mathcal{Q} \circ \mathbf{E}$ favors artifact-free reconstructions over reconstructions with artifacts.

1.1 The augmented NETT

One of the main assumptions for the analysis of [27] is the coercivity of the regularizer $\mathcal{Q} \circ \mathbf{E}$ which requires special care in network design and training. In order to overcome this limitation, we propose an augmented form of the regularizer for which we are able to rigorously prove coercivity. More precisely, for fixed $c > 0$, we consider minimizers x_α^δ of the augmented NETT functional

$$\mathcal{A}_{\alpha, y^\delta}(x) := \mathcal{D}(\mathbf{K}x, y^\delta) + \alpha \left(\mathcal{Q}(\mathbf{E}(x)) + \frac{c}{2} \|x - (\mathbf{D} \circ \mathbf{E})(x)\|_2^2 \right). \quad (1.3)$$

Here, $\mathcal{D}: \mathbb{Y} \times \mathbb{Y} \rightarrow [0, \infty]$ is a similarity measure and $\mathbf{D} \circ \mathbf{E}: \mathbb{X} \rightarrow \mathbb{X}$ is an encoder-decoder network trained such that for any signal x on a signal manifold we have $(\mathbf{D} \circ \mathbf{E})(x) \simeq x$ and that $\mathcal{Q}(\mathbf{E}(x))$ is small. We term this approach augmented NETT (aNETT) regularization. In this work we provide a mathematical convergence analysis for aNETT, present a novel modular training strategy and investigate its practical performance.

The term $\mathcal{Q}(\mathbf{E}(x))$ implements learned prior knowledge on the encoder coefficients, while smallness of $\|x - (\mathbf{D} \circ \mathbf{E})(x)\|_2^2$ forces x to be close to the signal manifold. The latter term also guarantees the coercivity of (1.3). In the original NETT version (1.2), coercivity of the regularizer requires coercivity conditions on the network involved. Indeed, in the numerical experiments, the authors of [27] observed a semi-convergence behaviour when minimizing (1.2), so early stopping of the iterative minimization scheme has been used as additional regularization. We attribute this semi-convergence behavior to the non-coercivity of the regularization term. In the present, paper we address this issue systematically by augmentation of the NETT functional which guarantees coercivity and allows a more stable minimization. Coercivity is also one main ingredient for the mathematical convergence analysis.

An interesting practical instance of aNETT takes \mathcal{Q} as a weighted ℓ^q -norm enforcing sparsity of the encoding coefficients [11, 16]. An important example for the similarity measure is given by the squared norm distance, which from a statistical viewpoint can

be motivated by a Gaussian white noise model. General similarity measures allow us to adapt to different noise models which can be more appropriate for certain problems.

1.2 Main contributions

The contributions of this paper are threefold. As described in more detail below, we introduce the aNETT framework, mathematically analyze its convergence, and propose a practical implementation that is applied to tomographic limited data problems.

- The first contribution is to introduce the structure of the aNETT regularizer $\mathcal{R}(x) = \mathcal{Q}(\mathbf{E}(x)) + (c/2) \|x - (\mathbf{D} \circ \mathbf{E})(x)\|_2^2$. The combination of the two terms in the regularizer, even in the case of a linear encoder seems to be new. The term $\mathcal{Q}(\mathbf{E}(x))$ enforces regularity of the analysis coefficients, which is an ingredient in most of existing variational regularization techniques. For example, this includes sparse regularization in frames or dictionaries, regularization with Sobolev norms or total variation regularization. On the other hand, the augmented term $\|x - (\mathbf{D} \circ \mathbf{E})(x)\|_2^2$ penalized distance to the signal manifold. It is the combination of these two terms that results in a stable reconstruction scheme without the need of strong assumptions on the involved networks.
- The second main contribution is the theoretical analysis of aNETT (1.3) in the context of regularization theory. We investigate the case where the image domain of the encoder is given by $\Xi = \ell^2(\Lambda)$ for some countable set Λ , and \mathcal{Q} is a coercive functional measuring the complexity of the encoder coefficients. The presented analysis is in the spirit of the analysis of NETT given in [27]. However, opposed to NETT, the required coercivity property is derived naturally for the class of considered regularizers. This supports the use of the regularizer \mathcal{R} also from a theoretical side. Moreover, the convergence rates results presented here uses assumptions significantly different from [27]. While we present our analysis for the transform domain $\Xi = \ell^2(\Lambda)$ we could replace the encoder space by a general Hilbert or Banach space.
- As a third main contribution we propose a modular strategy for training $\mathbf{D} \circ \mathbf{E}$ together with a possible network architecture. First, independent of the given inverse problem, we train a \mathcal{Q} -penalized autoencoder that learns representing signals from the training data with low complexity. In the second step, we train a task-specific network which can be adapted to specific inverse problem at hand. In our numerical experiments, we empirically found this modular training strategy to be superior to directly adapting the autoencoder to the inverse problem. For the \mathcal{Q} -penalized autoencoder, we train the modified version described in [35] of the tight frame U-Net of [19] in a way such that \mathcal{Q} poses additional constraints on the autoencoder during the training process.

1.3 Outline

In Section 2 we present the mathematical convergence analysis of aNETT. In particular, as an auxiliary result, we establish the coercivity of the regularization term. Moreover, we prove stability and derive convergence rates. Section 3 presents practical aspects for aNETT. We propose a possible architecture and training strategy for the networks,

and a possible ADMM based scheme to obtain minimizers of the aNETT functional. In Section 4, we present reconstruction results and compare aNETT with other deep learning based reconstruction methods. The paper concludes with a short summary and discussion. Parts of this paper were presented at the ISBI 2020 conference and the corresponding proceedings [34]. Opposed to the proceedings, this article treats a general similarity measure \mathcal{D} and considers a general complexity measure \mathcal{Q} . Further, all proofs and all numerical results presented in this paper are new.

2 Mathematical analysis

In this section we prove the stability and convergence of aNETT as regularization method. Moreover, we derive convergence rates in the form of quantitative error estimates between exact solutions for noise-free data and aNETT regularized solutions for noisy data.

2.1 Assumptions and coercivity results

For our convergence analysis we make use of the following assumptions on the underlying spaces and operators involved.

Condition 2.1 (General assumptions for aNETT).

- (A1) \mathbb{X} and \mathbb{Y} are Hilbert spaces.
- (A2) $\Xi = \ell^2(\Lambda)$ for a countable Λ .
- (A3) $\mathbf{K}: \mathbb{X} \rightarrow \mathbb{Y}$ is weakly sequentially continuous.
- (A4) $\mathbf{E}: \mathbb{X} \rightarrow \Xi$ is weakly sequentially continuous.
- (A5) $\mathbf{D}: \Xi \rightarrow \mathbb{X}$ is weakly sequentially continuous.
- (A6) $\mathcal{Q}: \Xi \rightarrow [0, \infty]$ is coercive and weakly sequentially lower semi-continuous.

We set $\mathbf{N} := \mathbf{D} \circ \mathbf{E}$ and, for given $c > 0$, define

$$\mathcal{R}: \mathbb{X} \rightarrow [0, \infty]: x \mapsto \mathcal{Q}(\mathbf{E}(x)) + \frac{c}{2} \|x - (\mathbf{D} \circ \mathbf{E})(x)\|_2^2, \quad (2.1)$$

which we refer to as the aNETT (or augmented NETT) regularizer.

According to (A4)-(A6), the aNETT regularizer is weakly sequentially lower semi-continuous. As a main ingredient for our analysis we next prove its coercivity.

Theorem 2.2 (Coercivity of the aNETT regularizer \mathcal{R}). *If Condition 2.1 holds, then $\mathcal{R}: \mathbb{X} \rightarrow [0, \infty]$ is coercive.*

Proof. Let $(x_n)_{n \in \mathbb{N}}$ be some sequence in \mathbb{X} such that $(\mathcal{R}(x_n))_{n \in \mathbb{N}}$ is bounded. Then by definition of \mathcal{R} it follows that $(\mathcal{Q}(\mathbf{E}(x_n)))_{n \in \mathbb{N}}$ is bounded and by coercivity of \mathcal{Q} we have that $(\mathbf{E}(x_n))_{n \in \mathbb{N}}$ is also bounded. By assumption, \mathbf{D} is weakly sequentially continuous and thus $(\|\mathbf{N}(x_n)\|)_{n \in \mathbb{N}}$ must be bounded. Using that $c > 0$, we obtain the inequality $\|x_n\|^2 \leq 2\|x_n - \mathbf{N}(x_n)\|^2 + 2\|\mathbf{N}(x_n)\|^2 \leq (4/c)\mathcal{R}(x_n) + 2\|\mathbf{N}(x_n)\|^2$. This shows that $(x_n)_{n \in \mathbb{N}}$ is bounded and therefore that \mathcal{R} is coercive. \square

Example 2.3 (Sparse aNETT regularizer). To obtain a sparsity promoting regularizer we can choose $Q(\xi) = \|\xi\|_{1;w} := \sum_{\lambda \in \Lambda} w_\lambda |\xi_\lambda|^q$ where $q \in [1, 2]$ and $\inf_\lambda w_\lambda > 0$. Since $q \in [1, 2]$ we have $\|\cdot\|^2 \leq (\inf_\lambda w_\lambda)^{-1/q} \|\cdot\|_{1;w}$ and hence $\|\cdot\|_{1;w}$ is coercive. As a sum of weakly sequentially lower semi-continuous functionals it is also weakly sequentially lower semi-continuous [16]. Therefore Condition (A6) is satisfied for the weighted ℓ^q -norm. Together with Theorem 2.2, we conclude that the resulting weighted sparse aNETT regularizer $x \mapsto \|\mathbf{E}(x)\|_{1;w} + (c/2) \|x - (\mathbf{D} \circ \mathbf{E})(x)\|_2^2$ is a coercive and weakly sequentially lower semi-continuous functional.

For the further analysis we will make the following assumptions including the similarity measure $\mathcal{D}: \mathbb{Y} \times \mathbb{Y} \rightarrow [0, \infty]$.

Condition 2.4 (Similarity measure).

- (B1) $\forall y_0, y_1: \mathcal{D}(y_0, y_1) = 0 \Leftrightarrow y_0 = y_1$.
- (B2) \mathcal{D} is sequentially lower semi-continuous with respect to the weak topology in the first and the norm topology in the second argument.
- (B3) $\forall (y_n)_{n \in \mathbb{N}} \in \mathbb{Y}^{\mathbb{N}}: (\mathcal{D}(y, y_n) \rightarrow 0 \Rightarrow y_n \rightarrow y \text{ as } n \rightarrow \infty)$.
- (B4) $\mathcal{D}(y, y_n) \rightarrow 0 \text{ as } n \rightarrow \infty \Rightarrow (\forall z \in \mathbb{Y}: \mathcal{D}(z, y) < \infty \Rightarrow \mathcal{D}(z, y_n) \rightarrow \mathcal{D}(z, y))$.
- (B5) $\forall y \in \mathbb{Y} \forall \alpha > 0: \exists x \in \mathbb{X} \text{ with } \mathcal{D}(\mathbf{K}x, y) + \alpha \mathcal{R}(x) < \infty$.

While (B1)-(B4) restrict the choice of the similarity measure, (B5) is a technical assumption involving the forward operator, the regularizer and the similarity measure, that is required for the existence of minimizers. For a more detailed discussion of these assumptions we refer to [36].

Example 2.5 (Similarity measures using the norm). The classical example of a similarity measure satisfying (B1)-(B4) is given by $\mathcal{D}(y_0, y_1) = \|y_0 - y_1\|^p$ for some $p \geq 1$ and more generally by $\mathcal{D}(y_0, y_1) = \psi(\|y_0 - y_1\|)$, where $\psi: [0, \infty) \rightarrow [0, \infty)$ is a continuous and monotonically increasing function that satisfies $\forall t \geq 0: \psi(t) = 0 \Leftrightarrow t = 0$.

Taking into account Theorem 2.2, Conditions 2.1 and 2.4 imply that the aNETT functional $\mathcal{A}_{\alpha, y}$ defined by (1.3), (2.1) is proper, coercive and weakly sequentially lower semi-continuous. This in particular implies the existence of minimizers of $\mathcal{A}_{\alpha, y}$ for all data $y \in \mathbb{Y}$ and regularization parameters $\alpha > 0$.

2.2 Stability

Next we prove the stability of minimizing the aNETT functional $\mathcal{A}_{\alpha, y}$ regarding perturbations of the data y .

Theorem 2.6 (Stability). *Let Conditions 2.1 and 2.4 hold, $y \in \mathbb{Y}$ and $\alpha > 0$. Moreover, let $(y_n)_{n \in \mathbb{N}} \in \mathbb{Y}^{\mathbb{N}}$ be a sequence of perturbed data with $\mathcal{D}(y, y_n) \rightarrow 0$ and consider minimizers $x_n \in \arg \min \mathcal{A}_{\alpha, y_n}$. Then the sequence $(x_n)_{n \in \mathbb{N}} \in \mathbb{X}^{\mathbb{N}}$ has at least one weak accumulation point and weak accumulation points are minimizers of $\mathcal{A}_{\alpha, y}$. Moreover, for any weak accumulation point x^\dagger of $(x_n)_{n \in \mathbb{N}}$ and any subsequence $(x_{\tau(n)})_{n \in \mathbb{N}}$ with $x_{\tau(n)} \rightharpoonup x^\dagger$ we have $\mathcal{R}(x_{\tau(n)}) \rightarrow \mathcal{R}(x^\dagger)$.*

Proof. Let $x_* \in \mathbb{X}$ be such that $\mathcal{A}_{\alpha,y}(x_*) < \infty$. By definition of x_n we have $\mathcal{A}_{\alpha,y_n}(x_n) \leq \mathcal{A}_{\alpha,y_n}(x_*)$. Since by assumption $\mathcal{D}(y, y_n) \rightarrow 0$ and $\mathcal{D}(\mathbf{K}x_*, y) < \infty$, we have $\mathcal{D}(\mathbf{K}x_*, y_n) \rightarrow \mathcal{D}(\mathbf{K}x_*, y)$. This implies that $\mathcal{A}_{\alpha,y_n}(x_*)$ is bounded by some positive constant m_* for sufficiently large n . By definition of \mathcal{A}_{α,y_n} we have $\alpha\mathcal{R}(x_n) \leq \mathcal{A}_{\alpha,y_n}(x_n) \leq m_* + \alpha\mathcal{R}(x_*)$. Since \mathcal{R} is coercive it follows that $(x_n)_{n \in \mathbb{N}}$ is a bounded sequence and hence it has a weakly convergent subsequence.

Let $(x_{\tau(n)})_{n \in \mathbb{N}}$ be a weakly convergent subsequence of $(x_n)_{n \in \mathbb{N}}$ and denote its limit by x^\dagger . By the lower semi-continuity we get $\mathcal{D}(\mathbf{K}x^\dagger, y) \leq \liminf_{n \rightarrow \infty} \mathcal{D}(\mathbf{K}x_{\tau(n)}, y_{\tau(n)})$ and $\mathcal{R}(x^\dagger) \leq \liminf_{n \rightarrow \infty} \mathcal{R}(x_{\tau(n)})$. Thus for all $x \in \mathbb{X}$ with $\mathcal{A}_{\alpha,y}(x) < \infty$ we have

$$\begin{aligned} \mathcal{A}_{\alpha,y}(x^\dagger) &\leq \liminf_{n \rightarrow \infty} \mathcal{D}(\mathbf{K}x_{\tau(n)}, y_{\tau(n)}) + \alpha \liminf_{n \rightarrow \infty} \mathcal{R}(x_{\tau(n)}) \\ &\leq \limsup_{n \rightarrow \infty} \mathcal{D}(\mathbf{K}x_{\tau(n)}, y_{\tau(n)}) + \alpha \mathcal{R}(x_{\tau(n)}) \leq \limsup_{n \rightarrow \infty} \mathcal{A}_{\alpha,y_{\tau(n)}}(x) = \mathcal{A}_{\alpha,y}(x). \end{aligned}$$

This shows that $x^\dagger \in \arg \min \mathcal{A}_{\alpha,y}$ and, by considering $x = x^\dagger$ in the above displayed equation, that $\mathcal{A}_{\alpha,y_{\tau(n)}}(x_{\tau(n)}) \rightarrow \mathcal{A}_{\alpha,y}(x^\dagger)$. Moreover, we have

$$\begin{aligned} \limsup_{n \rightarrow \infty} \alpha\mathcal{R}(x_{\tau(n)}) &\leq \limsup_{n \rightarrow \infty} \mathcal{A}_{\alpha,y_{\tau(n)}}(x_{\tau(n)}) - \liminf_{n \rightarrow \infty} \mathcal{D}(\mathbf{K}x_{\tau(n)}, y_{\tau(n)}) \\ &\leq \mathcal{A}_{\alpha,y}(x^\dagger) - \mathcal{D}(\mathbf{K}x^\dagger, y) = \alpha\mathcal{R}(x^\dagger). \end{aligned}$$

This shows $\mathcal{R}(x_{\tau(n)}) \rightarrow \mathcal{R}(x^\dagger)$ as $n \rightarrow \infty$ and concludes the proof. \square

In the following we say that the similarity measure \mathcal{D} satisfies the quasi triangle-inequality if there is some $q \geq 1$ such that

$$\forall y_0, y_1, y_2 \in \mathbb{Y}: \quad \mathcal{D}(y_0, y_1) \leq q \cdot (\mathcal{D}(y_0, y_2) + \mathcal{D}(y_2, y_1)). \quad (2.2)$$

While this property is essential for deriving convergence rate results, we will show below that it is not enough to guarantee stability of minimizing the augmented NETT functional in the sense of Theorem 2.6. Note that [27] assumes the quasi triangle-inequality (2.2) instead of Condition (B4). The following remarks shows that (2.2) is not sufficient for the stability result of Theorem 2.6 to hold and therefore Condition (B4) has to be added to the list of assumptions in [27] required for the stability.

Example 2.7 (Instability in the absence of Condition (B4)). Consider the similarity measure $\mathcal{D}: \mathbb{Y} \times \mathbb{Y} \rightarrow [0, \infty]$ defined by

$$\mathcal{D}(y_0, y_1) := H(y_1) \|y_0 - y_1\|^2, \quad (2.3)$$

where $H: \mathbb{Y} \rightarrow [0, 1]$ is defined by $H(y_1) = 1$ if $\|y_1\| \leq 1$ and $H(y_1) = 2$ otherwise. Moreover, choose $\mathbb{X} = \mathbb{Y}$, let $\mathbf{K} = \text{Id}$ be the identity operator and suppose the regularizer takes the form $\mathcal{R} = \|\cdot\|^2$.

- The similarity measure defined in (2.3) satisfies (B1)-(B3): Convergence with respect to \mathcal{D} is equivalent to convergence in norm which implies that (B3) is satisfied. Moreover, we have $\forall y_0, y_1: \mathcal{D}(y_0, y_1) = 0 \Leftrightarrow y_0 = y_1$, which is (B1). Consider sequences $z_n \rightarrow z$ and $y_n \rightharpoonup y$. The sequential lower semi-continuity stated in (B2) can be derived by separately looking at the cases $\|z\| \leq 1$ and $\|z\| > 1$. In the first case, by the continuity of the norm we have $\forall n \in \mathbb{N}: H(z) \leq H(z_n)$. In the second case, we have $\|y_n\| > 1$ for n sufficiently large. In both cases, the lower semi-continuity property follows from the weak lower semi-continuity property of the norm.

- The similarity measure defined by (2.3) does not satisfy (B4): To see this, we define $y_n := (1 + \delta_n)y$ where $\|y\| = 1$ and $\delta_n > 0$ is taken as a non-increasing sequence converging to zero. We have $y_n \rightarrow y$ and hence also $\mathcal{D}(y, y_n) \rightarrow 0$ as $n \rightarrow \infty$. For any $z \in \mathbb{Y}$ we have $\mathcal{D}(z, y) < \infty$ and $\mathcal{D}(z, y_n) = 2\|z - y_n\|^2 \rightarrow 2\|z - y\|^2 = 2\mathcal{D}(z, y)$ as $n \rightarrow \infty$. In particular, $\mathcal{D}(z, y_n)$ does not converge to $\mathcal{D}(z, y)$ if $z \neq y$ and therefore (B5) does not hold. In summary, all requirements for Theorem 2.6 are satisfied, except of the continuity assumption (B4).

- We have $\|y_0 - y_1\|^2 \leq \mathcal{D}(y_0, y_1) \leq 2\|y_0 - y_1\|^2$ which implies that the similarity measure satisfies the quasi triangle-inequality (2.2). However as shown next, this is not sufficient for stable reconstruction in the sense of Theorem 2.6. To that end, let $y_n := (1 + \delta_n)y$ with $\|y\| = 1$ and $\delta_n \downarrow 0$ be as above and let $\alpha > 0$. In particular, $\mathcal{D}(x, y_n) = 2\|x - y_n\|^2$ and $\mathcal{D}(x, y) = \|x - y\|^2$. Therefore the minimizer of $\arg \min \mathcal{D}(x, y_n) + \alpha\|x\|^2$ with perturbed data y_n is given by $x_n = y_n/(1 + \alpha/2)$ and the minimizer of $\arg \min \mathcal{D}(x, y_n) + \alpha\|x\|^2$ for data y is given by $x^\dagger = y/(1 + \alpha)$. We see that $x_n \rightarrow y/(1 + \alpha/2)$, which is clearly different from x^\dagger . In particular, minimizing $\mathcal{D}(\cdot, y) + \alpha\|\cdot\|^2$ does not stably depend on data y . Theorem 2.6 states that stability holds if (B4) is satisfied.

While the above example may seem somehow constructed, it shows that one has to be careful when choosing the similarity measure in order to obtain a stable reconstruction scheme.

2.3 Convergence

In this subsection we prove that minimizers of the aNETT functional for noisy data converge to \mathcal{R} -minimizing solution of the equation $\mathbf{K}x = y$ as the noise level goes to zero and the regularization parameter is chosen properly. Here and below, we use the following notation.

Definition 2.8 (\mathcal{R} -minimizing solutions). For $y \in \mathbb{Y}$, we call an element $x^\dagger \in \text{dom}(\mathcal{R})$ an \mathcal{R} -minimizing solution of the equation $\mathbf{K}x = y$ if

$$x^\dagger \in \arg \min \{ \mathcal{R}(x) : x \in \text{dom}(\mathcal{R}) \wedge \mathbf{K}x = y \}.$$

An \mathcal{R} -minimizing solution always exists provided that data satisfies $y \in \mathbf{K}(\text{dom}(\mathcal{R}))$, which means that the equation $\mathbf{K}x = y$ has at least one solution with finite value of \mathcal{R} . To see this, consider a sequence of solutions $(x_n)_{n \in \mathbb{N}}$ with $\mathcal{R}(x_n) \rightarrow \inf \{ \mathcal{R}(x) : \mathbf{K}x = y \wedge \mathcal{R}(x) < \infty \}$. Since \mathcal{R} is coercive there exists a weakly convergent subsequence $(x_{\tau(n)})_{n \in \mathbb{N}}$ with weak limit x^\dagger . Using the weak sequential lower semi-continuity of \mathcal{R} one concludes that x^\dagger is an \mathcal{R} -minimizing solution.

We first show weak convergence.

Theorem 2.9 (Weak convergence of aNETT). *Suppose Conditions 2.1 and 2.4 are satisfied. Let $y \in \mathbf{K}(\text{dom}(\mathcal{R}))$, $(\delta_n)_{n \in \mathbb{N}} \in (0, \infty)^{\mathbb{N}}$ with $\delta_n \rightarrow 0$ and let $(y_n)_{n \in \mathbb{N}} \in \mathbb{Y}^{\mathbb{N}}$ satisfy $\mathcal{D}(y, y_n) \leq \delta_n$. Choose $\alpha_n > 0$ such that $\lim_{n \rightarrow \infty} \delta_n/\alpha_n = \lim_{n \rightarrow \infty} \alpha_n = 0$ and let $x_n \in \arg \min \mathcal{A}_{\alpha_n, y_n}$. Then the following hold:*

- (a) $(x_n)_{n \in \mathbb{N}}$ has at least one weakly convergent subsequence.
- (b) All accumulation points of $(x_n)_{n \in \mathbb{N}}$ are \mathcal{R} -minimizing solutions of $\mathbf{K}x = y$.

- (c) For every convergent subsequence $(x_{\tau(n)})_{n \in \mathbb{N}}$ it holds $\mathcal{R}(x_{\tau(n)}) \rightarrow \mathcal{R}(x^\dagger)$.
(d) If the \mathcal{R} -minimizing solution x^\dagger is unique then $x_n \rightharpoonup x^\dagger$.

Proof. (a): Because $y \in \mathbf{K}(\text{dom}(\mathcal{R}))$, there exists an \mathcal{R} -minimizing solution of the equation $\mathbf{K}x = y$ which we denote by x^\dagger . Because $x_n \in \arg \min \mathcal{A}_{\alpha_n, y_n}$ we have

$$\begin{aligned} \alpha_n \mathcal{R}(x_n) &\leq \mathcal{D}(\mathbf{K}x_n, y_n) + \alpha_n \mathcal{R}(x_n) \leq \mathcal{D}(\mathbf{K}x^\dagger, y_n) + \alpha_n \mathcal{R}(x^\dagger) \\ &= \mathcal{D}(y, y_n) + \alpha_n \mathcal{R}(x^\dagger) \leq \delta_n + \alpha_n \mathcal{R}(x^\dagger). \end{aligned} \quad (2.4)$$

Because $\alpha_n, \delta_n \rightarrow 0$ this shows that $(\mathcal{R}(x_n))_{n \in \mathbb{N}}$ is bounded. Due to the coercivity of the aNETT regularizer (see Theorem 2.2), this implies that $(x_n)_{n \in \mathbb{N}}$ has a weakly convergent subsequence.

(b), (c): Let $(x_{\tau(n)})_{n \in \mathbb{N}}$ be a weakly convergent subsequence of $(x_n)_{n \in \mathbb{N}}$ with limit x_* . From the weak lower semi-continuity we get $\mathcal{D}(\mathbf{K}x_*, y) \leq \liminf_{n \rightarrow \infty} \mathcal{D}(\mathbf{K}x_{\tau(n)}, y_{\tau(n)}) + \alpha_{\tau(n)} \mathcal{R}(x_{\tau(n)}) \leq \liminf_{n \rightarrow \infty} \mathcal{D}(\mathbf{K}x^\dagger, y_{\tau(n)}) + \alpha_{\tau(n)} \mathcal{R}(x^\dagger) = 0$, which shows that x_* is a solution of $\mathbf{K}x = y$. Moreover,

$$\mathcal{R}(x^\dagger) \leq \mathcal{R}(x_*) \leq \liminf_{n \rightarrow \infty} \mathcal{R}(x_{\tau(n)}) \leq \liminf_{n \rightarrow \infty} \frac{\delta_{\tau(n)}}{\alpha_{\tau(n)}} + \mathcal{R}(x^\dagger) = \mathcal{R}(x^\dagger)$$

where for the second last inequality we used (2.4) and for the last equality we used that $\delta_n/\alpha_n \rightarrow 0$. Therefore, x_* is an \mathcal{R} -minimizing solution of the equation $\mathbf{K}x = y$. In a similar manner we derive $\mathcal{R}(x_*) \leq \liminf_{n \rightarrow \infty} \mathcal{R}(x_{\tau(n)}) \leq \limsup_{n \rightarrow \infty} \delta_{\tau(n)}/\alpha_{\tau(n)} + \mathcal{R}(x^\dagger) = \mathcal{R}(x^\dagger)$ which shows $\mathcal{R}(x_{\tau(n)}) \rightarrow \mathcal{R}(x_*)$.

(d): If $\mathbf{K}x = y$ has a unique \mathcal{R} -minimizing solution x^\dagger , then every subsequence of $(x_n)_{n \in \mathbb{N}}$ has itself a subsequence weakly converging to x^\dagger , which implies that $(x_n)_{n \in \mathbb{N}}$ weakly converges to the \mathcal{R} -minimizing solution. \square

Next we derive strong convergence of the regularized solutions. To this end we recall the absolute Bregman distance, the modulus of total nonlinearity and the total nonlinearity, defined in [27].

Definiton 2.10 (Absolute Bregman distance). Let $\mathcal{F}: \mathbb{X} \rightarrow [0, \infty]$ be Gâteaux differentiable at $x_* \in \mathbb{X}$. The absolute Bregman distance $\Delta_{\mathcal{F}}(\cdot, x_*): \mathbb{X} \rightarrow [0, \infty]$ at x_* with respect to \mathcal{F} is defined by

$$\forall x \in \mathbb{X}: \quad \Delta_{\mathcal{F}}(x, x_*) := |\mathcal{F}(x) - \mathcal{F}(x_*) - \mathcal{F}'(x)(x - x_*)|.$$

Here and below $\mathcal{F}'(x_*)$ denotes the Gâteaux derivative of \mathcal{F} at x_* .

Definiton 2.11 (Modulus of total nonlinearity and total nonlinearity). Let $\mathcal{F}: \mathbb{X} \rightarrow [0, \infty]$ be Gâteaux differentiable at $x_* \in \mathbb{X}$. We define the modulus of total nonlinearity of \mathcal{F} at x_* as $\nu_{\mathcal{F}}(x_*, \cdot): [0, \infty) \rightarrow [0, \infty): t \mapsto \inf\{\Delta_{\mathcal{F}}(x, x_*): \|x - x_*\| = t\}$. We call \mathcal{F} totally nonlinear at x_* if $\nu_{\mathcal{F}}(x_*, t) > 0$ for all $t \in (0, \infty)$.

Using these definitions we get the following convergence result in the norm topology.

Theorem 2.12 (Strong convergence of aNETT). *Let Conditions 2.1 and 2.4 hold, $y \in \mathbf{K}(\text{dom}(\mathcal{R}))$ and let \mathcal{R} be totally nonlinear at all \mathcal{R} -minimizing solutions of $\mathbf{K}x = y$. Let $(y_n)_{n \in \mathbb{N}}, (x_n)_{n \in \mathbb{N}}, (\alpha_n)_{n \in \mathbb{N}}$ be as in Theorem 2.9. Then there is a subsequence $(x_{\tau(n)})_{n \in \mathbb{N}}$ which converges in norm to an \mathcal{R} -minimizing solution x^\dagger of $\mathbf{K}x = y$. If the \mathcal{R} -minimizing solution is unique, then $x_n \rightarrow x^\dagger$ as $n \rightarrow \infty$.*

Proof. In [27, Proposition 2.9] it is shown that the total nonlinearity of \mathcal{R} implies that for every bounded sequence $(z_n)_{n \in \mathbb{N}}$ with $\Delta_{\mathcal{R}}(z_n, z) \rightarrow 0$ it holds that $z_n \rightarrow z$. Theorem 2.9 gives us a weakly converging subsequence $(x_{\tau(n)})_{n \in \mathbb{N}}$ of $(x_n)_{n \in \mathbb{N}}$ with weak limit x^\dagger and $\mathcal{R}(x_{\tau(n)}) \rightarrow \mathcal{R}(x^\dagger)$. By the definition of the absolute Bregman distance it follows that $\Delta_{\mathcal{R}}(x_{\tau(n)}, x^\dagger) \rightarrow 0$ and hence, together with [27, Proposition 2.9], that $x_{\tau(n)} \rightarrow x^\dagger$. If the \mathcal{R} -minimizing solution of $\mathbf{K}x = y$ is unique, then every subsequence has a subsequence converging to x^\dagger and hence the claim follows. \square

2.4 Convergence rates

We will now prove convergence rates by deriving quantitative estimates for the absolute Bregman distance between \mathcal{R} -minimizing solutions for exact data and regularized solutions for noisy data. The convergence rates will be derived under the additional assumption that \mathcal{D} satisfies the quasi triangle-inequality (2.2).

Proposition 2.13 (Convergence rates for aNETT). *Let the assumptions of Theorem 2.12 be satisfied and suppose that \mathcal{D} satisfies the quasi triangle-inequality (2.2) for some $q \geq 1$. Let $x^\dagger \in \mathbb{X}$ be an \mathcal{R} -minimizing solution of $\mathbf{K}x = y$ such that \mathcal{R} is Gâteaux differentiable at x^\dagger and assume there exist $\epsilon, c > 0$ with*

$$\forall x \in \mathbb{X}: \quad \|x - x^\dagger\| \leq \epsilon \Rightarrow \Delta_{\mathcal{R}}(x, x^\dagger) \leq \mathcal{R}(x) - \mathcal{R}(x^\dagger) + c\sqrt{\mathcal{D}(\mathbf{K}x, \mathbf{K}x^\dagger)}. \quad (2.5)$$

For any $\delta > 0$, let $y^\delta \in \mathbb{Y}$ be noisy data satisfying $\mathcal{D}(y^\delta, \mathbf{K}x^\dagger)^{1/2} \leq \delta$, $\mathcal{D}(\mathbf{K}x^\dagger, y^\delta)^{1/2} \leq \delta$, and write $x_\alpha^\delta \in \arg \min \mathcal{A}_{\alpha, y^\delta}$. Then the following hold:

- (a) For sufficiently small α , it holds $\Delta_{\mathcal{R}}(x_\alpha^\delta, x^\dagger) \leq \delta^2/\alpha - c\delta\sqrt{q} + c^2\alpha q/4$.
- (b) If $\alpha \asymp \delta$, then $\Delta_{\mathcal{R}}(x_\alpha^\delta, x^\dagger) = \mathcal{O}(\delta)$ as $\delta \rightarrow 0$.

Proof. By definition of x_α^δ we have $\mathcal{D}(\mathbf{K}x_\alpha^\delta, y^\delta) + \alpha\mathcal{R}(x_\alpha^\delta) - \alpha\mathcal{R}(x^\dagger) \leq \mathcal{D}(\mathbf{K}x^\dagger, y^\delta) \leq \delta^2$. By Theorem 2.12 for sufficiently small α we can assume that $\|x_\alpha^\delta - x^\dagger\| \leq \epsilon$ and hence

$$\begin{aligned} \alpha\Delta_{\mathcal{R}}(x_\alpha^\delta, x^\dagger) &\leq \alpha\mathcal{R}(x_\alpha^\delta) - \alpha\mathcal{R}(x^\dagger) + c\alpha\sqrt{\mathcal{D}(\mathbf{K}x_\alpha^\delta, \mathbf{K}x^\dagger)} \\ &= \mathcal{R}(x_\alpha^\delta) - \mathcal{D}(\mathbf{K}x_\alpha^\delta, y^\delta) - \left(\mathcal{R}(x_\alpha^\delta) - \mathcal{D}(\mathbf{K}x_\alpha^\delta, y^\delta)\right) + c\alpha\sqrt{\mathcal{D}(\mathbf{K}x_\alpha^\delta, \mathbf{K}x^\dagger)} \\ &\leq \delta^2 - \mathcal{D}(\mathbf{K}x_\alpha^\delta, y^\delta) + c\alpha\delta\sqrt{q} + c\alpha\sqrt{q\mathcal{D}(\mathbf{K}x_\alpha^\delta, y^\delta)}. \end{aligned}$$

Together with the inequality of arithmetic and geometric means $(a + b)/2 \geq \sqrt{ab}$ for $a = \mathcal{D}(\mathbf{K}x_\alpha^\delta, y^\delta)$ and $b = c^2\alpha^2q/4$ this implies $\alpha\Delta_{\mathcal{R}}(x_\alpha^\delta, x^\dagger) \leq \delta^2 - c\alpha\delta\sqrt{q} + c^2\alpha^2q/4$ which shows (a). Item (b) is an immediate consequence of (a). \square

The following results is our main convergence rates result. It is similar to Proposition [27, Theorem 3.1], but uses different assumptions.

Theorem 2.14 (Convergence rates for finite rank operators). *Let the assumptions of Theorem 2.12 be satisfied, take $\mathcal{D}(y_1, y_2) = \|y_1 - y_2\|^2$, assume \mathbf{K} has finite dimensional range and that \mathcal{R} is Lipschitz continuous and Gâteaux differentiable. For any $\delta > 0$, let $y^\delta \in \mathbb{Y}$ be noisy data satisfying $\|y^\delta - \mathbf{K}x^\dagger\| \leq \delta$ and write $x_\alpha^\delta \in \arg \min \mathcal{A}_{\alpha, y^\delta}$. Then for the parameter choice $\alpha \asymp \delta$ we have the convergence rates result $\Delta_{\mathcal{R}}(x_\alpha^\delta, x^\dagger) = \mathcal{O}(\delta)$ as $\delta \rightarrow 0$.*

Proof. According to Proposition 2.13, it is sufficient to show that (2.5) holds with $\|\mathbf{K}x - \mathbf{K}x^\dagger\|$ in place of $\mathcal{D}(\mathbf{K}x, \mathbf{K}x^\dagger)^{1/2}$. For that purpose, let \mathbf{P} denote the orthogonal projection onto the null-space $\ker(\mathbf{K})$ and let L be a Lipschitz constant of \mathcal{R} . Since \mathbf{K} restricted to $\ker(\mathbf{K})^\perp$ is injective with finite dimensional range, we can choose a constant $a > 0$ such that $\forall z \in \ker(\mathbf{K})^\perp: \|\mathbf{K}z\| \geq a\|z\|$.

We first show the estimates

$$\forall x \in \mathbb{X}: \quad \mathcal{R}(x^\dagger) - \mathcal{R}(x) \leq (L/a) \|\mathbf{K}x - \mathbf{K}x^\dagger\| \quad (2.6)$$

$$\forall x \in \mathbb{X}: \quad |\langle \mathcal{R}'(x^\dagger), x - x^\dagger \rangle| \leq (\|\mathcal{R}'(x^\dagger)\|/a) \|\mathbf{K}x - \mathbf{K}x^\dagger\|. \quad (2.7)$$

To that end, let $x \in \mathbb{X}$ and write $x_0 := (x^\dagger - \mathbf{P}x^\dagger) + \mathbf{P}x$. Then $\mathbf{K}x_0 = \mathbf{K}x^\dagger$. Since x^\dagger is an \mathcal{R} -minimizing solution, we have $\mathcal{R}(x^\dagger) - \mathcal{R}(x) \leq \mathcal{R}(x_0) - \mathcal{R}(x) \leq L\|x_0 - x\|$. Since $x_0 - x \in \ker(\mathbf{K})^\perp$, we have $\|\mathbf{K}x^\dagger - \mathbf{K}x\| = \|\mathbf{K}(x_0 - x)\| \geq a\|x_0 - x\|$. The last two estimates prove (2.6). Because x^\dagger is an \mathcal{R} -minimizing solution, we have $\langle \mathcal{R}'(x^\dagger), x^\dagger - x \rangle = 0$ whenever $x^\dagger - x \in \ker(\mathbf{K})$. On the other hand, using that \mathcal{R} is Gâteaux differentiable and that \mathbf{K} has finite rank, shows $|\langle \mathcal{R}'(x^\dagger), x^\dagger - x \rangle| \leq \|\mathcal{R}'(x^\dagger)\| a^{-1} \|\mathbf{K}(x^\dagger - x)\|$ for $x^\dagger - x \in \ker(\mathbf{K})^\perp$. This proves (2.7).

Inequality (2.6) implies $|\mathcal{R}(x) - \mathcal{R}(x^\dagger)| \leq \mathcal{R}(x) - \mathcal{R}(x^\dagger) + 2(L/a) \|\mathbf{K}x - \mathbf{K}x^\dagger\|$. Together with (2.7) this yields

$$\begin{aligned} \Delta_{\mathcal{R}}(x, x^\dagger) &= |\mathcal{R}(x) - \mathcal{R}(x^\dagger) - \langle \mathcal{R}'(x^\dagger), x - x^\dagger \rangle| \\ &\leq |\mathcal{R}(x) - \mathcal{R}(x^\dagger)| + |\langle \mathcal{R}'(x^\dagger), x^\dagger - x \rangle| \\ &\leq \mathcal{R}(x) - \mathcal{R}(x^\dagger) + (2L + \|\mathcal{R}'(x^\dagger)\|) a^{-1} \|\mathbf{K}x^\dagger - \mathbf{K}x\|, \end{aligned}$$

which proves (2.5) with $c = (2L + \|\mathcal{R}'(x^\dagger)\|)/a$. \square

Note that the theoretical results stated remain valid, if we replace \mathcal{R} by a general coercive and weakly lower semi-continuous regularizer $\mathcal{R}: \mathbb{X} \rightarrow [0, \infty]$.

3 Practical realization

In this section we investigate practical aspects of aNETT. We present a possible network architecture together with a possible training strategy in the discrete setting. Further we discuss minimization of aNETT using the ADMM algorithm.

For the sake of clarity we restrict our discussion to the finite dimensional case where $\mathbb{X} = \mathbb{R}^{N \times N}$ and $\ell^2(\Lambda) = \mathbb{R}^\Lambda$ for a finite index set Λ .

3.1 Proposed modular aNETT training

To find a suitable network $\mathbf{D} \circ \mathbf{E}$ defining the aNETT regularizer $\mathcal{R}(x) = \mathcal{Q}(\mathbf{E}(x)) + (c/2) \|x - \mathbf{D} \circ \mathbf{E}(x)\|^2$, we propose a modular data driven approach that comes in two separate steps. In a first step, we train a \mathcal{Q} -regularized denoising autoencoder $\mathbf{D}^\mathcal{Q} \circ \mathbf{E}$ independent of the forward problem \mathbf{K} , whose purpose is to well represent elements of a training data set by low complexity encoder coefficients. In a second step, we train a task-specific network that increases the ability of the aNETT regularizer to distinguish between clean images and images containing problem specific artifacts.

Let $x_1, \dots, x_m \in \mathcal{M}$ denote the given set of artifact-free training phantoms.

- **Q-REGULARIZED AUTOENCODER:**

First, an autoencoder $\mathbf{D}^{\mathcal{Q}} \circ \mathbf{E}$ is trained such that x_i is close to $\mathbf{D}^{\mathcal{Q}} \circ \mathbf{E}(x_i)$ and that $\mathcal{Q}(\mathbf{E}(x_i))$ is small for the given training signals. For that purpose, let $(\mathbf{D}_{\theta} \circ \mathbf{E}_{\theta})_{\theta \in \Theta}$ be a family of autoencoder networks, where $\mathbf{E}_{\theta}: \mathbb{R}^{N \times N} \rightarrow \mathbb{R}^{\Lambda}$ are encoder and $\mathbf{D}_{\theta}: \mathbb{R}^{\Lambda} \rightarrow \mathbb{R}^{N \times N}$ decoder networks, respectively.

To achieve that unperturbed images are sparsely represented by \mathbf{E} , whereas disrupted images are not, we apply the following training strategy. We randomly generate images $x_i + a_i \epsilon_i$ where ϵ_i is additive Gaussian white noise with a standard deviation proportional to the mean value of x_i , and $a_i \in \{0, 1\}$ is a binary random variable that takes each value with probability 0.5. For the numerical results below we use a standard deviation of 0.05 times the mean value of x_i . To select the particular autoencoder based on the training data, we consider the following training strategy

$$\theta^* \in \arg \min_{\theta} \frac{1}{m} \sum_{i=1}^m \|(\mathbf{D}_{\theta} \circ \mathbf{E}_{\theta})(x_i + a_i \epsilon_i) - x_i\|_2^2 + \nu(1 - a_i) \mathcal{Q}(\mathbf{E}_{\theta}(x_i)) + \beta \|\theta\|_2^2, \quad (3.1)$$

and set $[\mathbf{D}^{\mathcal{Q}}, \mathbf{E}] := [\mathbf{D}_{\theta^*}, \mathbf{E}_{\theta^*}]$. Here $\nu, \beta > 0$ are regularization parameters.

Including perturbed signals $x_i + \epsilon_i$ in (3.1) increases robustness of the \mathcal{Q} -regularized autoencoder. To enforce regularity for the encoder coefficients only on the noise-free images, the penalty $\mathcal{Q}(\mathbf{E}(\cdot))$ is only used for the noise-free inputs, reflected by the pre-factor $1 - a_i$. Using auto-encoders, regularity for a signal class could also be achieved by means of dimensionality reduction techniques, where \mathbb{R}^{Λ} is used as a bottleneck in the network architecture. However, in order to get a regularizer that is able to distinguish between perturbed and undisturbed perturbed signals we use \mathbb{R}^{Λ} to be of sufficiently high dimensionality.

- **TASK-SPECIFIC NETWORK:**

Numerical simulations showed that the \mathcal{Q} -regularized autoencoder alone was not able to sufficiently well distinguish between artifact-free training phantoms and images containing problem specific artifacts. In order to address this issue, we compose the operator independent network with another network \mathbf{U} , that is trained to distinguish between image with and without problem specific artifacts.

For that purpose, we consider randomly generated images z_1, \dots, z_m where either $z_i = (\mathbf{D}^{\mathcal{Q}} \circ \mathbf{E})(x_i)$ or $z_i = (\mathbf{D}^{\mathcal{Q}} \circ \mathbf{E})(\mathbf{K}^{\dagger}(\mathbf{K}x_i + \eta_i))$ with equal probability. Here \mathbf{K}^{\dagger} is an approximate right inverse and η_i are error terms. We choose a network architecture $(\mathbf{U}_{\theta})_{\theta \in \Theta}$ and select $\mathbf{U} = \mathbf{U}_{\theta^*}$, where

$$\theta^* \in \arg \min_{\theta} \frac{1}{m} \sum_{i=1}^m \|\mathbf{U}_{\theta}(z_i) - x_i\|_2^2 + \gamma \|\theta\|_2^2, \quad (3.2)$$

for some regularization parameter $\gamma > 0$. In particular, the image residuals $(\mathbf{D}^{\mathcal{Q}} \circ \mathbf{E})(\mathbf{K}^{\dagger}(\mathbf{K}x_i + \eta_i)) - (\mathbf{D}^{\mathcal{Q}} \circ \mathbf{E})(x_i)$ now depend on the specific inverse problem and we can consider them to consist of operator and training signal specific artifacts.

The above training procedure ensures that the network \mathbf{U} adapts to the inverse problem at hand as well as to the \mathcal{Q} -regularized autoencoder. Training the network \mathbf{U} independently of $\mathbf{D}^{\mathcal{Q}} \circ \mathbf{E}$, or directly training the auto-encoder to distinguish between

images with and without problem specific artifacts, we empirically found to perform considerably worse.

The final autoencoder is then given as $\mathbf{N} = \mathbf{D} \circ \mathbf{E}$ with modular decoder $\mathbf{D} := \mathbf{U} \circ \mathbf{D}^{\mathcal{Q}}$. For the numerical results we take $(\mathbf{U}_{\theta})_{\theta \in \Theta}$ as the tight frame U-Net of [19]. Moreover, we choose $(\mathbf{D}_{\theta} \circ \mathbf{E}_{\theta})_{\theta \in \Theta}$ as modified tight frame U-Net proposed in [35] for deep synthesis regularization. In particular, opposed to the original tight frame U-net, the modified tight frame U-Net does not involve skip connections.

3.2 Possible aNETT minimization

For minimizing the aNETT functional (1.3) we use the alternating direction method of multipliers (ADMM) with scaled dual variable [8, 14, 15]. For that purpose, the aNETT minimization problem is rewritten as the following constraint minimization problem

$$\begin{cases} \arg \min_{x, \xi} & \mathcal{D}(\mathbf{K}x, y^{\delta}) + \alpha \mathcal{Q}(\xi) + \alpha c/2 \|x - \mathbf{N}(x)\|^2 \\ \text{s. t.} & \mathbf{E}(x) = \xi. \end{cases}$$

The resulting ADMM update scheme with scaling parameter $\rho > 0$ initialized by $\xi_0 = \mathbf{E}(\mathbf{K}^{\dagger}y^{\delta})$ and $\eta_0 = 0$ then reads as follows:

$$\begin{aligned} \text{(S1)} \quad x_{k+1} &= \arg \min_x \mathcal{D}(\mathbf{K}x, y^{\delta}) + (\alpha c/2) \|x - \mathbf{N}(x)\|_2^2 + (\rho/2) \|\mathbf{E}(x) - \xi_k + \eta_k\|_2^2. \\ \text{(S2)} \quad \xi_{k+1} &= \arg \min_{\xi} \alpha \mathcal{Q}(\xi) + (\rho/2) \|\mathbf{E}(x_{k+1}) - \xi + \eta_k\|_2^2. \\ \text{(S3)} \quad \eta_{k+1} &= \eta_k + (\mathbf{E}(x_{k+1}) - \xi_{k+1}). \end{aligned}$$

One interesting feature of the above approach is that the signal update (S1) is independent of the possibly non-smooth penalty \mathcal{Q} . Moreover, the encoder update (S2) uses the proximal mapping of \mathcal{Q} which in important special cases can be evaluated explicitly and therefore fast and exact. Moreover, it guarantees regular encoder coefficients during each iteration. For example, if we choose the penalty as the ℓ^1 -norm, then (S2) is a soft-thresholding step which results in sparse encoder coefficients. Step (S1) in typical cases has to be computed iteratively via an inner iteration. To find an approximate solution for (S1) for the results presented below we use gradient descent with at most 10 iterations. We stop the gradient descent updates early if the difference of the functional evaluated at two consecutive iterations is below our predefined tolerance of 10^{-5} .

The concrete implementation of the aNETT minimization requires specification of the similarity measure, the total number of outer iterations N_{iter} , the step-size γ for the iteration in (S1) and the parameters defining the aNETT functional. These specifications are selected dependent of the inverse problem at hand. Table 3.1 lists the particular choices for the reconstruction scenarios considered in the following section.

The ADMM scheme for aNETT minimization shares similarities with existing iterative neural network based reconstruction methods. In particular, ADMM inspired plug-and-play priors [9, 38, 44] may be most closely related. However, opposed the plug and play

	α	c	γ	N_{iter}	N_{φ}	noise model	\mathcal{D}
Sparse view	10^{-4}	10^2	$5 \cdot 10^{-1}$	50	40	Gaussian, $\sigma = 0.02$	$\mathcal{D}_{\ \cdot\ ^2}$
Low dose	$5 \cdot 10^{-3}$	10^2	10^{-3}	20	1138	Poisson, $p = 10^4$	\mathcal{D}_{KL}
Universality	10^{-4}	10^2	$5 \cdot 10^{-1}$	50	160	Gaussian, $\sigma = 0.02$	$\mathcal{D}_{\ \cdot\ ^2}$

Table 3.1: Parameter specifications for proposed aNETT functional and its numerical minimization.

approach we can deduce convergence from existing results for ADMM for non-convex problems [46]. While convergence of (S1)-(S3) and relations with plug and play priors are interesting and relevant, they are beyond the scope of this work. This also applies to the comparison with other iterative minimization schemes for minimizing aNETT.




Figure 3.1: Left: Example image randomly drawn from the dataset. Middle: Corresponding sparse view sinogram (40 directions). Right: low dose sinogram.

4 Application to sparse view and low dose CT

In this section we apply aNETT regularization to sparse view and low-dose computed tomography (CT). For the experiments we always choose \mathcal{Q} to be the ℓ^1 -norm. The parameter specifications for the proposed aNETT functional and its numerical minimization are given in Table 3.1. For quantitative evaluation, we use the peak-signal-to-noise-ratio (PSNR) defined by

$$\text{PSNR}(x, x_{\text{rec}}) := 20 \log_{10} \left(\frac{\max x}{\|x - x_{\text{rec}}\|_2} \right).$$

Here $x \in \mathbb{R}^{N \times N}$ is the ground truth image and $x_{\text{rec}} \in \mathbb{R}^{N \times N}$ its numerical reconstruction. Higher value of PSNR indicates better reconstruction.

4.1 Discretization and dataset

For sparse view CT as well as for low dose CT we work with a discretization of the Radon transform $(\mathbf{R}f)(\varphi, s) := \int_{\mathbb{R}} f(s \cos(\varphi) - t \sin(\varphi), s \sin(\varphi) + t \cos(\varphi)) dt$. The values $(\mathbf{R}f)(\varphi, s)$ are integrals of the function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ over lines orthogonal to

PSNR	FBP	LPD	Post	aNETT
Sparse view	23.8 ± 1.3	37.9 ± 1.2	37.1 ± 0.9	37.1 ± 1.0
Low dose	36.9 ± 1.6	43.6 ± 1.3	44.1 ± 1.5	43.9 ± 1.3
Universality	32.4 ± 1.6	na	37.7 ± 0.8	38.3 ± 1.0

Table 4.1: Overview of metric results evaluated on the test-set. The values shown are the average of the PSNR \pm the standard deviation calculated over the test dataset. The values in bold show the best results. The na entry means that LPD cannot be applied after changing the angular sampling pattern.

$(\cos(\varphi), \sin(\varphi))^T$ for angle $\varphi \in [0, \pi)$ and signed distance $s \in \mathbb{R}$. We discretize the Radon transform using the ODL library [1] where we assume that the function has compact support in $[-1, 1]^2$ and sampled on an equidistant grid. We use N_φ equidistant samples of $\varphi \in [0, \pi)$ and N_s equidistant samples of $s \in [-1.5, 1.5]$. In both cases, we end up with an inverse problem of the form (1.1), where $\mathbf{K}: \mathbb{R}^{N \times N} \rightarrow \mathbb{R}^{N_\varphi \times N_s}$ is the discretized linear forward operator. Elements $x \in \mathbb{R}^{N \times N}$ will be referred to as CT images and the elements $y \in \mathbb{R}^{N_\varphi \times N_s}$ as sinograms.

or all results presented below we work with image size 512×512 and use $N_s = 768$. The number of angular samples N_φ is taken 40 for low the dose CT and $N_\varphi = 1138$ for the low dose example. In both cases we use the CT images from the Low Dose CT Grand Challenge dataset [32] provided by the Mayo Clinic. The dataset consists of 512×512 grayscale images of 10 different patients, where for each patient there are multiple CT scanning series available. We use the split 7/2/1 for training, validation and testing which corresponds to 4267/1143/526 CT images in the respective sets. We use the validation set to select networks which achieve the minimal loss on the validation set. The test set is used to evaluate the final performance. Note that by splitting the dataset according to patient we avoid validation and testing on images patients that have already be seen during training time. An example image and the corresponding simulated sparse view and low-dose data are shown in Figure 3.1.

4.2 Numerical results

We compare results of aNETT to the learned primal-dual algorithm (LPD) [3], the tight frame U-Net [19] applied as post-processing network (CNN), and the filtered back-projection (FBP) as a base-line. Minimization of the loss-function for all methods was done using Adam [22] for 100 epochs, cosine decay learning rate $\eta_t = (\eta_0/2) \cdot (1 + \cos(\pi t/100))$ with $\eta_0 = 10^{-3}$ in the t -th epoch, and a batch-size of 4. For the LPD we take hyper-parameters $N_{\text{primal}} = N_{\text{dual}} = 5$ and $N = 7$ network iterations and train according to [3]. For training of the tight frame U-Net we do not follow the patch approach of [19] but instead use full images obtained with FBP as CNN inputs. Training of all the networks was done on a GTX 1080 Ti with an Intel Xeon Bronze 3104 CPU.

- **SPARSE VIEW CT:** To simulate sparse view data we evaluate the Radon transform for $N_\varphi = 40$ directions. We generate noisy data $y^\delta = \mathbf{K}x + \eta^\delta$ by adding Gaussian white noise with standard deviation taken as 0.02 times the mean value of $\mathbf{K}x$. We use the ℓ^2 -norm distance as the similarity measure. Quantitative results evaluated on the test




Figure 4.1: Reconstructions for sparse view CT data from $N_\varphi = 40$ angular directions. The intensity range of all images is $[-500, 500]$ HU.

set are shown in Table 4.1. All learning-based methods yield comparable performance in terms of PSNR and clearly outperform FBP. The reconstructions shown in Figure 4.1 indicate that aNETT reconstructions are less smooth than CNN reconstructions and less blocky than LPD reconstructions.

- **Low Dose CT:** For the low dose problem, we use a fully sampled sinogram with $N_\varphi = 1138$ and add Poisson noise corresponding to 10^4 incident photons per pixel bin. The Kullback-Leibler divergence \mathcal{D}_{KL} is a more appropriate than the squared ℓ^2 -norm distance in case of Poisson noise and the reported values and reconstructions use the Kullback-Leibler divergence as the similarity measure. Quantitative results are shown in Table 4.1. Again, all learning-based methods give similar results and significantly outperform FBP. Visual comparison of the reconstructions in Figure 4.2 shows that CNN yields cartoon like images and the LPD reconstruction again looks blocky. The aNETT reconstruction shows more texture than the CNN reconstruction and at the same time is less blocky than the LPD reconstruction.

- **UNIVERSALITY:** In practical applications, we may not have a fixed sampling pattern. If we have many different sampling patterns, then training network for each sampling pattern is infeasible and hence reconstruction methods should be applicable to different sampling scenarios. Additionally, it is desirable that an increased number of samples indeed increases performance. In order to test this issue, we consider the sparse view CT problem but with an increased number of angular samples without retraining the networks. LPD is not applicable in this case, as the changing the forward operator changes the network architecture. Quantitative evaluation for this scenario is given in Table 4.1. We see that aNETT are better than CNN in terms of PSNR. The advantage




Figure 4.2: Reconstructions results from low dose CT data. The intensity range of all images is $[-200, 200]$ HU.

of aNETT over CNN, however, is best observed in Figure 4.3. One observes that CNN yields a similar reconstructions for both angular sampling patterns. On the other hand, aNETT is able to synergistically combine the increased sampling rate of the sinogram with the network trained on coarsely sampled data. Despite using the network, aNETT with 160 angular samples reconstructs small details which are not present in the reconstruction from 40 angular samples.

4.3 Discussion

The results show that the proposed aNETT regularization is competitive with prominent deep-learning methods such as LPD and post-processing CNNs. We found that the aNETT does not suffer as much from over-smoothing which is often observed in other deep-learning reconstruction methods. This can for example be seen in Figure 4.2 where the CNN yields an over-smoothed reconstruction and the aNETT reconstruction shows more texture. Besides this, aNETT reconstructions are less blocky than LPD reconstructions. Moreover, aNETT is able to leverage higher sampling rates to reconstruct small details while other deep-learning methods fail to do so. We conjecture that this is due the synergistic interplay of the aNETT regularizer with the data-consistency term in (1.3). In some scenarios, it may not be possible to retrain networks. Especially for learned iterative schemes network training is a time-consuming task. Training aNETT on the other hand is straightforward and, as demonstrated, yields a method which is robust to changes of the forward problem during testing time.

Finally, we note that aNETT relies on minimizing (1.3) iteratively. With the use of

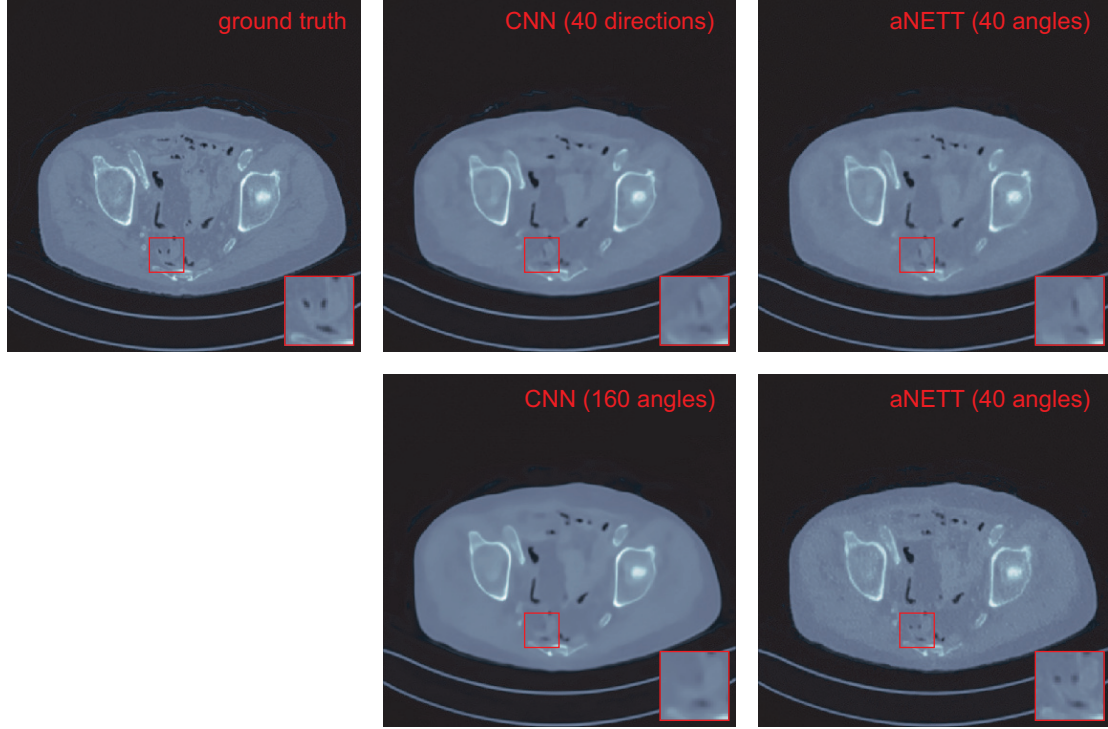


Figure 4.3: Universality of aNETT due to change of angular sampling pattern. Top row: Ground truth and reconstructions from 40 angular directions. Bottom: Reconstructions from 160 angular directions. All reconstruction use the networks trained with 40 angular directions. While aNETT shows increased resolution for ncreased angular sampling, CNN does not.

the ADMM minimization scheme presented in this article, aNETT is slower than the methods used for comparison in this article. Designing faster optimization schemes for (1.3) is beyond the scope of this work, but is an important and interesting aspect.

5 Conclusion

We have proposed the aNETT (augmented NETwork Tikhonov) for which derived coercivity of the regularizer under quite mild assumptions on the networks involved. Using this coercivity we presented a convergence analysis of aNETT with a general similarity measure \mathcal{D} . We proposed a modular training strategy in which we first train an \mathcal{Q} -regularized autoencoder independent of the problem at hand and then a network which is adapted to the problem and first autoencoder. Experimentally we found this training strategy to be superior to directly training the autoencoder on the full task. Lastly, we conducted numerical simulations demonstrating the feasibility of aNETT.

The experiments show that aNETT is able to keep up with the classical post-processing CNNs and the learned primal-dual approach for to sparse view and low dose CT. Typical deep learning methods work well for a fixed sampling pattern on which they have been trained on. However, reconstruction methods are expected to perform better if we use an increased sampling rate. We have experimentally shown that aNETT is able to leverage

higher sampling rates to reconstruct small details in the images which are not visible in the other reconstructions. This universality can be advantageous in applications where one is not fixed to one sampling pattern or is not able to train a network for every sampling pattern.

Acknowledgments

D.O. and M.H. acknowledge support of the Austrian Science Fund (FWF), project P 30747-N32. The research of L.N. has been supported by the National Science Foundation (NSF) Grants DMS 1212125 and DMS 1616904.

References

- [1] J. Adler, H. Kohr, and O. Öktem. ODL-a python framework for rapid prototyping in inverse problems. *Royal Institute of Technology*, 2017.
- [2] J. Adler and O. Öktem. Solving ill-posed inverse problems using iterative deep neural networks. *Inverse Probl.*, 33:124007, 2017.
- [3] J. Adler and O. Öktem. Learned primal-dual reconstruction. *IEEE Trans. Med. Imag.*, 37(6):1322–1332, 2018.
- [4] H. K. Aggarwal, M. P. Mani, and M. Jacob. MoDL: model-based deep learning architecture for inverse problems. *IEEE Trans. Med. Imag.*, 38(2):394–405, 2018.
- [5] S. Antholzer and M. Haltmeier. Discretization of learned NETT regularization for solving inverse problems. *arXiv:2011.03627*, 2020.
- [6] S. Antholzer, M. Haltmeier, and J. Schwab. Deep learning for photoacoustic tomography from sparse data. *Inverse Probl. Sci. and Eng.*, 27(7):987–1005, 2018.
- [7] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb. Solving inverse problems using data-driven models. *Acta Numer.*, 28:1–174, 2019.
- [8] S. Boyd, N. Parikh, and E. Chu. *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc, 2011.
- [9] S. H. Chan, X. Wang, and O. A. Elgendy. Plug-and-play admm for image restoration: Fixed-point convergence and applications. *IEEE Trans. Comput. Imag.*, 3(1):84–98, 2016.
- [10] J. R. Chang, C.-L. Li, B. Póczos, and B. V. Kumar. One network to solve them all—solving linear inverse problems using deep projection models. In *IEEE International Conference on Computer Vision (ICCV)*, pages 5889–5898, 2017.
- [11] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Comm. Pure Appl. Math.*, 57(11):1413–1457, 2004.

- [12] S. Dittmer, T. Kluth, P. Maass, and D. O. Baguer. Regularization by architecture: A deep prior approach for inverse problems. *J. Math. Imaging Vis.*, 62(3):456–470, 2020.
- [13] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of inverse problems*, volume 375 of *Mathematics and its Applications*. Kluwer Academic Publishers Group, Dordrecht, 1996.
- [14] D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Comput. Math. Appl.*, 2(1):17–40, 1976.
- [15] R. Glowinski and A. Marroco. Sur l’approximation, par elements finis d’ordre un, et la resolution, par penalisation-dualite, d’une classe de problemes de dirichlet non lineares. *RAIRO Anal. Numer.*, 9:41–76, 1975.
- [16] M. Grasmair, M. Haltmeier, and O. Scherzer. Sparse regularization with l^q penalty term. *Inverse Probl.*, 24(5):055020, 2008.
- [17] M. Haltmeier and L. V. Nguyen. Regularization of inverse problems by neural networks. *arXiv:2006.03972*, 2020.
- [18] 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.
- [19] Y. Han and J. C. Ye. Framing U-Net via deep convolutional framelets: Application to sparse-view CT. *IEEE Trans. Med. Imag.*, 37:1418–1429, 2018.
- [20] A. Hauptmann, J. Adler, S. R. Arridge, and O. Oktem. Multi-scale learned iterative reconstruction. *IEEE Trans. Comput. Imag.*, page to appear, 2021.
- [21] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE Trans. Image Process.*, 26(9):4509–4522, 2017.
- [22] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2014. *arXiv:1412.6980*.
- [23] E. Kobler, T. Klatzer, K. Hammernik, and T. Pock. Variational networks: connecting variational methods and deep learning. In *German Conference on Pattern Recognition*, pages 281–293. Springer, 2017.
- [24] A. Kofler, M. Haltmeier, C. Kolbitsch, M. Kachelrieß, and M. Dewey. A U-Nets cascade for sparse view computed tomography. In *International Workshop on Machine Learning for Medical Image Reconstruction (MLIMR)*, pages 91–99. Springer, 2018.
- [25] 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. *arXiv:2102.00783*, 2021.

- [26] D. Lee, J. Yoo, and J. C. Ye. Deep residual learning for compressed sensing MRI. In *IEEE 14th International Symposium on Biomedical Imaging*, pages 15–18, 2017.
- [27] H. Li, J. Schwab, S. Antholzer, and M. Haltmeier. NETT: Solving inverse problems with deep neural networks. *Inverse Probl.*, 36(6):065005, 2020.
- [28] S. Lunz, O. Öktem, and C.-B. Schönlieb. Adversarial regularizers in inverse problems. In *Advances in Neural Information Processing Systems*, pages 8507–8516, 2018.
- [29] A. Maier, C. Syben, T. Lasser, and C. Riess. A gentle introduction to deep learning in medical image processing. *Zeitschrift für Medizinische Physik*, 29(2):86–101, 2019.
- [30] A. K. Maier, C. Syben, B. Stimpel, T. Würfl, M. Hoffmann, F. Schebesch, W. Fu, L. Mill, L. Kling, and S. Christiansen. Learning with known operators reduces maximum error bounds. *Nat. Mach. Intell.*, 1(8):373–380, 2019.
- [31] M. T. McCann, K. H. Jin, and M. Unser. Convolutional neural networks for inverse problems in imaging: A review. *IEEE Signal Process. Mag.*, 34(6):85–95, 2017.
- [32] C. McCollough. TU-FG-207A-04: overview of the low dose CT grand challenge. *Med. Phys.*, 43(6):3759–3760, 2016.
- [33] S. Mukherjee, S. Dittmer, Z. Shumaylov, S. Lunz, O. Öktem, and C.-B. Schönlieb. Learned convex regularizers for inverse problems. *arXiv:2008.02839*, 2020.
- [34] D. Obmann, L. Nguyen, J. Schwab, and M. Haltmeier. Sparse aNETT for solving inverse problems with deep learning. In *2020 IEEE 17th International Symposium on Biomedical Imaging Workshops (ISBI Workshops)*, pages 1–4, 2020.
- [35] D. Obmann, J. Schwab, and M. Haltmeier. Deep synthesis regularization of inverse problems. *Inverse Probl.*, 37(1):015005, 2021.
- [36] C. Pöschl. *Tikhonov regularization with general residual term*. PhD thesis, University of Innsbruck, 2008.
- [37] Y. Rivenson, Y. Zhang, H. Günaydin, D. Teng, and A. Ozcan. Phase recovery and holographic image reconstruction using deep learning in neural networks. *Light Sci. Appl.*, 7(2):17141–17141, 2018.
- [38] Y. Romano, M. Elad, and P. Milanfar. The little engine that could: Regularization by denoising (red). *SIAM J. Imaging Sci.*, 10(4):1804–1844, 2017.
- [39] O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen. *Variational methods in imaging*, volume 167 of *Applied Mathematical Sciences*. Springer, New York, 2009.
- [40] J. Schlemper, J. Caballero, J. V. Hajnal, A. Price, and D. Rueckert. A deep cascade of convolutional neural networks for mr image reconstruction. In *Proc. Inf. Process. Med. Imaging*, pages 647–658. Springer, 2017.

- [41] J. Schwab, S. Antholzer, and M. Haltmeier. Deep null space learning for inverse problems: convergence analysis and rates. *Inverse Probl.*, 35(2):025008, 2019.
- [42] J. Schwab, S. Antholzer, and M. Haltmeier. Big in Japan: Regularizing networks for solving inverse problems. *J. Math. Imaging Vis.*, 62:445–455, 2020.
- [43] D. Ulyanov, A. Vedaldi, and V. Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018.
- [44] S. V. Venkatakrisnan, C. A. Bouman, and B. Wohlberg. Plug-and-play priors for model based reconstruction. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 945–948. IEEE, 2013.
- [45] G. Wang. A perspective on deep imaging. *IEEE Access*, 4:8914–8924, 2016.
- [46] Y. Wang, W. Yin, and J. Zeng. Global convergence of ADMM in nonconvex nonsmooth optimization. *Journal of Scientific Computing*, 78(1):29–63, 2019.
- [47] Y. Yang, J. Sun, H. Li, and Z. Xu. Deep ADMM-net for compressive sensing MRI. In *Proc. 30th International Conference on Neural Information Processing Systems*, pages 10–18, 2016.