



Regularization by Architecture: A Deep Prior Approach for Inverse Problems

Sören Dittmer¹ · Tobias Kluth¹ · Peter Maass¹ · Daniel Otero Bager¹

Received: 30 November 2018 / Accepted: 18 October 2019 / Published online: 30 October 2019
© The Author(s) 2019

Abstract

The present paper studies so-called deep image prior (DIP) techniques in the context of ill-posed inverse problems. DIP networks have been recently introduced for applications in image processing; also first experimental results for applying DIP to inverse problems have been reported. This paper aims at discussing different interpretations of DIP and to obtain analytic results for specific network designs and linear operators. The main contribution is to introduce the idea of viewing these approaches as the optimization of Tikhonov functionals rather than optimizing networks. Besides theoretical results, we present numerical verifications.

Keywords Inverse problems · Deep learning · Regularization by architecture · Deep inverse prior · Deep image prior

1 Introduction

Deep image priors (DIP) were recently introduced in deep learning for some tasks in image processing [19]. Usually, deep learning approaches to inverse problems proceed in two steps. In a first step (training), the parameters Θ of the deep neural network φ_{Θ} are optimized by minimizing a suitable loss function using large sets of training data. In a second step (application), new data are fed into the network for solving the desired task.

DIP approaches are radically different; they are based on unsupervised training using only a single data point y^{δ} . More

precisely, in the context of inverse problems, where we aim at solving ill-posed operator equations $Ax \sim y^{\delta}$, the task of DIP is to train a network $\varphi_{\Theta}(z)$ with parameters Θ by minimizing the simple loss function

$$\min_{\Theta} \|A\varphi_{\Theta}(z) - y^{\delta}\|^2. \quad (1.1)$$

The minimization is with respect to Θ , the random input z is kept fixed. After training, the solution to the inverse problem is approximated directly by $\hat{x} = \varphi_{\Theta}(z)$.

In image processing, common choices for A are the identity operator (denoising) or a projection operator to a subset of the image domain (inpainting). For these applications, it has been observed that minimizing the functional iteratively by gradient descent methods in combination with a suitable stopping criterion leads to amazing results [19].

Training with a single data point is the most striking property, which separates DIP from other neural network concepts. One might argue that the astonishing results [10,19,23,32] are only possible if the network architecture is fine-tuned to the specific task. This is true for obtaining optimal performance; nevertheless, the presented numerical results perform well even with somewhat generic network architectures such as autoencoders.

We are interested in analyzing DIP approaches for solving ill-posed inverse problems. As a side remark, we note that the applications (denoising, inpainting) mentioned above are modeled by either identity or projection operators, which are

Sören Dittmer, Tobias Kluth, Peter Maass and Daniel Otero Bager have contributed equally to this study.

The code of the analytic deep prior implementation is available at <https://github.com/oterobager/analytic-deep-prior>.

✉ Sören Dittmer
sdittmer@math.uni-bremen.de

Tobias Kluth
tkluth@math.uni-bremen.de

Peter Maass
pmaass@math.uni-bremen.de

Daniel Otero Bager
otero@math.uni-bremen.de

¹ Center for Industrial Mathematics (ZeTeM), University of Bremen, Bremen, Germany

not ill-posed in the functional analytical setting [13,21,28]. Typical examples of ill-posed inverse problems correspond to compact linear operators such as a large variety of tomographic measurement operators or parameter-to-state mappings for partial differential equations.

We aim at analyzing a specific network architecture φ_Θ and at interpreting the resulting DIP approach as a regularization technique in the functional analytical setting, and also at proving convergence properties for the minimizers of (1.1). In particular, we are interested in network architectures, which themselves can be interpreted as a minimization algorithm that solves a regularized inverse problem of the form

$$x(B) = \arg \min_x \frac{1}{2} \|Bx - y^\delta\|^2 + \alpha R(x), \tag{1.2}$$

where R is a given convex function and B a learned operator.

In general, deep learning approaches for inverse problems have their own characteristics, and naive applications of neural networks can fail for even the most simple inverse problems, as shown in [22]. However, there is a growing number of compelling numerical experiments using suitable network designs for some of the toughest inverse problems such as photo-acoustic tomography [17] or X-ray tomography with very few measurements [2,18]. Concerning networks based on deep prior approaches for inverse problems, first experimental investigations have been reported, as shown in [19,23,32]. Similar as for the above-mentioned tasks in image processing, DIPs for inverse problems rely on two ingredients:

1. A suitable network design, which leads to our phrase “regularization by architecture.”
2. Training algorithms for iteratively minimizing (1.1) with respect to Θ in combination with a suitable stopping criterion.

In this paper, we present different mathematical interpretations of DIP approaches, and we analyze two network designs in the context of inverse problems in more detail. It is organized as follows: In Sect. 2, we discuss some relations to existing results and make a short survey of the related literature. In Sect. 3, we then state different interpretations of DIP approaches and the network architectures that we use, as a basis for the subsequent analysis. We start with a first mathematical result for a trivial network design, which yields a connection to Landweber iterations. We then consider a fully connected feedforward network with L identical layers, which generates a proximal gradient descent for a modified Tikhonov functional. In Sect. 4, we use this last connection to define the notion of analytic deep prior networks, for which one can strictly analyze its regularization and convergence properties. The key to the theoretical findings is a change of

view, which allows for the interpretation of DIP approaches as optimizing families of Tikhonov functionals. Finally, we exemplify our theoretical findings with numerical examples for the standard linear integration operator.

2 Deep Prior and Related Research

We start with a description of general deep prior concepts. Afterward, we address similarities and differences to other approaches, such as LISTA [16], in more detail.

2.1 The Deep Prior Approach

Present results on deep prior networks utilize feedforward architectures. In general, a feedforward neural network is an algorithm that starts with input $x^0 = z$, computes iteratively

$$x^{k+1} = \phi \left(W_k x^k + b_k \right)$$

for $k = 0, \dots, L - 1$ and outputs

$$\varphi_\Theta(z) = x^L.$$

The parameters of this system are denoted by

$$\Theta = \{W_0, \dots, W_{L-1}, b_0, \dots, b_{L-1}\}$$

and ϕ denotes a nonlinear activation function.

In order to highlight one of the unique features of deep image priors, let us first refer to classical generative networks that require training on large data sets.

In this classical setting, we are given an operator $A : X \rightarrow Y$ between Hilbert spaces X, Y , as well as a set of training data (x_i, y_i^δ) , where y_i^δ is a noisy version of Ax_i satisfying $\|y_i^\delta - Ax_i\| \leq \delta$. Here, the usual deep learning approach is to use a network for direct inversion and the parameters Θ of the network are obtained by minimizing the loss function

$$\min_{\Theta} \sum_{i=1}^N \|\varphi_\Theta(y_i^\delta) - x_i\|^2. \tag{2.1}$$

After training, Θ is fixed and the network is used to approximate the solution of the inverse problem with new data y^δ by computing $x = \varphi_\Theta(y^\delta)$. For a recent survey on this approach and more general deep learning concepts for inverse problems see [5].

In general, this approach relies on the underlying assumption that complex distributions of suitable solutions x , e.g., the distribution of natural images, can be approximated by neural networks [6,8,33]. The parameters Θ are trained for the specific distribution of training data and are fixed after

training. One then expects that choosing a new data set as input, i.e., $z = y^\delta$ will generate a suitable solution to $Ax \sim y^\delta$ [7]. Hence, after training, the distribution of solutions is parametrized by the inputs z .

In contrast, DIP is an unsupervised approach using only a single data point for training. That means, for given data y^δ and fixed z , the parameters Θ of the network φ_Θ are obtained by minimizing the loss function (1.1). The solution to the inverse problem is then denoted by $\hat{x} = \varphi_\Theta(z)$. Hence, deep image priors keep z fixed and aim at parameterizing the solution with Θ . It has been observed in several works [10,19,23,32] that this approach indeed leads to remarkable results for problems such as inpainting or denoising.

To some extent, the success of deep image priors is rooted in the careful design of network architectures. For example, [19] uses a U-Net-like “hourglass” architecture with skip connections, and the amazing results show that such an architecture implicitly captures some statistics of natural images. However, in general, the DIP learning process may converge toward noisy images or undesirable reconstructions. The whole success relies on a combination of the architecture with a suitable optimization method and stopping criterion. Nevertheless, the authors claim the architecture has a positive impact on the exploration of the solution space during the iterative optimization of Θ . They show that the training process descends quickly to “natural-looking” images but requires much more steps to produce noisy images. This is also supported by the theoretical results of [29] and the observations of [35], which shows that deep networks can fit noise very well but need more training time to do so. Another paper that hints in this direction is [4], which analyzes whether neural networks could have a bias toward approximating low frequencies.

There are already quite a few works that deal with deep prior approaches. Following, we mention the most relevant ones to our work. The original deep image prior article [19] introduces the DIP concept and presents experimental evidence that today’s network architectures are in and of themselves conducive to image reconstruction. Another work [32] explores the applicability of DIP to problems in compressed sensing. Also, [23] discusses how to combine DIP with the regularization by denoising approach and [10] explores DIP in the context of stationary Gaussian processes. All of these introduce and discuss variants of DIP concepts; however, neither of them addresses the intrinsic regularizing properties of the network concerning ill-posed inverse problems.

2.2 Deep Prior and Unrolled Proximal Gradient Architectures

A major part of this paper is devoted to analyzing the DIP approach in combination with an unrolled proximal gradient

network φ_Θ . Hence, there is a natural connection to the well-established analysis of LISTA schemes. Before we sketch the state of research in this field, we highlight the two major differences (loss function, training data) to the present approach. LISTA is based on a supervised training using multiple data points (x_i, y_i^δ) , $i = 1, \dots, N$ where y_i^δ is a noisy representation of Ax_i . The loss function is (2.1). DIP, however, is based on unsupervised learning using the loss function (1.1) and a single data point y^δ . Hence, DIP with the unrolled proximal gradient network shares the architecture with LISTA, but its concept, as well as its analytic properties, is different. Nevertheless, the analysis we will present in Sect. 4 will exhibit structures similar to the ones appearing in the LISTA-related literature. Hence, we shortly review the major contributions in this field.

Similarities are most visible when considering algorithms and convergence analysis for sparse coding applications [20,25,30,31,34]. The field of sparse coding makes heavy use of proximal splitting algorithms and, since the advent of LISTA, of trained architectures inspired by truncated versions of these algorithms. In the broadest sense, all of these methods are expressions of “Learning to learn by gradient descent” [3]. Once more, we would like to emphasize that these results utilize multiple data points while DIP does not require any training data but only one measurement. Another key difference is that we approach the topic from an ill-posed inverse problem perspective, which (a) grounds our approach in the functional analytic realm and (b) considers ill-posed (not only ill-conditioned) problems in the Nashed sense, i.e., allows the treatment of unstable inverses [13]. These two points fundamentally differentiate the present approach from traditional compressed sensing considerations which usually deal with (a) finite dimensional formulations and (b) forward operators given by well-conditioned, carefully hand-crafted settings or dictionaries, which are optimized using large sets of training data [30].

Coming back to LISTA for sparse coding applications, there are many excellent papers [15,24,25] which are devoted to a strict mathematical analysis of different aspects of LISTA-like approaches. In [25], the authors show under which conditions sparse coding can benefit from LISTA-like trained structures and ask how good trained sparsity estimators can be given a computational budget. The article [15] deals with a similar trade-off proposing the quite exciting, “inexact proximal gradient descent.” The paper [9] proposes, based on theoretically founded considerations, a sibling architecture to LISTA. Moreover, [27] argues that deep learning architectures, in general, can be interpreted as multistage proximal splitting algorithms.

Finally, we want to point at publications, which address deep learning with only a few data points for training, see, e.g., [14] and the references therein. However, they do not address the architectures relevant for our publication, and




Fig. 1 A simple network with scalar input, a single layer and no activation function. For any arbitrary input z one obtains $\varphi_\Theta(z) = \Theta$ (Color figure online)

they do not refer to the specific complications of inverse problems.

3 Deep Prior Architectures and Interpretations

In this section, we discuss different perspectives on deep prior networks, which open the path to provable mathematical results. The first two subsections are devoted to special network architectures, and the last two subsections deal with more general points of view.

3.1 A Trivial Architecture

We aim at solving ill-posed inverse problems. For a given operator A , the general task in inverse problems is to recover an approximation for x^\dagger from measured noisy data

$$y^\delta = Ax^\dagger + \tau,$$

where τ , with $\|\tau\| \leq \delta$, describes the noise in the measurement.

The deep image prior approach to inverse problems asks to train a network $\varphi_\Theta(z)$ with parameters Θ and fixed input z by minimizing $\|A\varphi_\Theta(z) - y^\delta\|^2$ with an optimization method such as gradient descent with early stopping. After training, a final run of the network computes $\hat{x} = \varphi_\Theta(z)$ as an approximation to x^\dagger .

We consider a trivial single-layer network without activation function, as shown in Fig. 1. This network simply outputs Θ , i.e., $\varphi_\Theta(z) = \Theta$. In this case, the network parameter Θ is a vector, which is chosen to have the same dimension as x . That means, that training the network by gradient descent of $\|A\varphi_\Theta(z) - y^\delta\|^2 = \|A\Theta - y^\delta\|^2$ with respect to Θ is equivalent to the classical Landweber iteration, which is a gradient descent method for $\|Ax - y^\delta\|^2$ with respect to x .

Landweber iterations are slowly converging. However, in combination with a suitable stopping rule, they are optimal regularization schemes for diminishing noise level $\delta \rightarrow 0$, [13,21,28]. Despite the apparent trivialization of the neural network approach, this shows that there is potential in training such networks with a single data point for solving ill-posed inverse problems.




Fig. 2 Unrolled proximal gradient network with $L = 2$ (Color figure online)

3.2 Unrolled Proximal Gradient Architecture

In this section, we aim at rephrasing DIP, i.e., the minimization of (1.1) with respect to Θ , as an approach for learning optimized Tikhonov functionals for inverse problems. This change of view, i.e., regarding deep inverse priors as optimization of functionals rather than networks, opens the way for analytic investigations in Sect. 4.

We use the particular architecture, which was introduced in [16], i.e., a fully connected feedforward network with L layers of identical size,

$$\varphi_\Theta(z) = x^L, \tag{3.1}$$

where

$$x^{k+1} = \phi(Wx^k + b) \tag{3.2}$$

The affine linear map $\Theta = (W, b)$ is the same for all layers. The matrix W is restricted to obey $I - W = \lambda B^* B$ (I denotes the identity operator) for some B and the bias is determined via $b = \lambda B^* y^\delta$, as shown in Fig. 2. If the activation function of the network is chosen as the proximal mapping of a regularizing functional $\lambda \alpha R$, then $\varphi_\Theta(z)$ is identical to the L th iterate of a proximal gradient descent method for minimizing

$$J_B(x) = \frac{1}{2} \|Bx - y^\delta\|^2 + \alpha R(x), \tag{3.3}$$

see [12] or “Appendix 1”.

Remark 3.1 Restricting activation functions to be proximal mappings is not as severe as it might look at first glance. For example, ReLU is the proximal mapping for the indicator function of positive real numbers, and soft shrinkage is the proximal mapping for the modulus function.

This allows the interpretation that every weight update, i.e., every gradient step for minimizing (1.1) with respect to Θ or B , changes the functional J_B . Hence, DIP can be regarded as optimizing a functional, which in-turn is minimized by the network. This view is the starting point for investigating convergence properties in Sect. 4.

3.3 Two Perspectives Based on Regression

The following subsections address more general concepts, which open the way to further analytic investigations, which, however, are not considered further in this paper. The reader interested in the regularization properties for DIP approaches for inverse problems only may jump directly to Sect. 4.

In this subsection, we present two different perspectives on solving inverse problems with the DIP via the minimization of a functional as discussed in the subsection above. The first perspective is based on a reinterpretation of the minimization of the functional (1.1) in the finite, real setting, i.e., $A \in \mathbb{R}^{m \times n}$. This setting allows us to write

$$\min_{\Theta} \|A\varphi_{\Theta}(z) - y^{\delta}\|^2 = \min_{x \in \mathcal{R}(\varphi_{\Theta}(z))} \|Ax - y^{\delta}\|^2 \tag{3.4}$$

$$= \min_{x \in \mathcal{R}(\varphi_{\Theta}(z))} \sum_{i=1}^m (x^* a_i - y_i^{\delta})^2, \tag{3.5}$$

where $\mathcal{R}(\varphi_{\Theta}(z))$ denotes the range of the network with regard to Θ for a fixed z and a_i the rows of the matrix A as well as y_i^{δ} the entries of the vector y^{δ} . This setting allows for the interpretation that we are solving a linear regression, parameterized by x , which is constrained by a deep learning hypothesis space and given by data pairs of the form (a_i, y_i^{δ}) .

The second perspective is based on the rewriting of the optimization problem via the method of Lagrange multipliers. We start by considering the constrained optimization problem

$$\min_{x \in X, \Theta} \|Ax - y^{\delta}\|^2 \text{ s.t. } \|x - \varphi_{\Theta}(z)\|^2 = 0. \tag{3.6}$$

If we now assume that φ has continuous first partial derivatives with regard to Θ , the Lagrange functional

$$\mathcal{L}(\Theta, x, \lambda) = \|Ax - y^{\delta}\|^2 + \lambda \|x - \varphi_{\Theta}(z)\|^2, \tag{3.7}$$

with the correct Lagrange multiplier $\lambda = \lambda_0$, has a stationary point at each minimum of the original constraint optimization problem. This gives us a direct connection to unconstrained variational approaches like Tikhonov functionals.

3.4 The Bayesian Point of View

The Bayesian approach to inverse problems focuses on computing MAP (maximum a posteriori probability) estimators, i.e., one aims for

$$\hat{x} = \arg \max_{x \in X} p(x|y^{\delta}), \tag{3.8}$$

where $p : X \times Y \rightarrow \mathbb{R}_+ \cup \{0\}$ is a conditional PDF. From standard Bayesian theory, we obtain

$$\hat{x} = \arg \min_{x \in X} \{-\log[p(y^{\delta}|x)] - \log[p(x)]\}. \tag{3.9}$$

The setting for inverse problems, i.e., $Ax + \tau = y^{\delta}$ with $\tau \sim \text{Normal}(0, \sigma^2 \mathbb{1}_Y)$, yields $(\lambda = 2\sigma^2)$

$$\hat{x} =: \arg \min_{x \in X} \|Ax - y^{\delta}\|^2 - \lambda \log[p(x)].$$

We now decompose x into $x_{\perp} := P_{\mathcal{N}(A)^{\perp}}(x)$, and $x_{\mathcal{N}} := P_{\mathcal{N}(A)}(x)$, where $\mathcal{N}(A)$ denotes the nullspace of A and where $P_{\mathcal{N}(A)}(x)$, resp. $P_{\mathcal{N}(A)^{\perp}}(x)$, denotes the orthogonal projection onto $\mathcal{N}(A)$, resp. $\mathcal{N}(A)^{\perp}$. Setting $\hat{x} = (x_{\mathcal{N}}, x_{\perp})$ yields

$$\begin{aligned} \hat{x} &= \arg \min_{x \in X} \|Ax_{\perp} - y^{\delta}\|^2 - \lambda \log p(x_{\mathcal{N}}, x_{\perp}) \\ &= \arg \min_{x \in X} \|Ax_{\perp} - y^{\delta}\|^2 - \lambda \log p(x_{\perp}) - \lambda \log p(x_{\mathcal{N}}|x_{\perp}). \\ &= \arg \min_{x \in X} \underbrace{\|Ax_{\perp} - y^{\delta}\|^2 - \lambda \log p(x_{\perp})}_{(I)} - \lambda \log p(x_{\mathcal{N}}|x_{\perp}). \end{aligned} \tag{II}$$

The data y^{δ} only contain information about x_{\perp} , which in classical regularization is exploited by restricting any reconstruction to $\mathcal{N}(A)^{\perp}$.

However, if available, $p(x_{\mathcal{N}}|x_{\perp})$ is a measure on how to extend x_{\perp} with an $x_{\perp} \in \mathcal{N}(A)^{\perp}$ to a suitable $x = (x_{\mathcal{N}}, x_{\perp})$. The classical regularization of inverse problems uses the trivial extension by zero, i.e., $x = (0, x_{\perp})$, which is not necessarily optimal. If we accept the interpretation that a network can be a meaningful parametrization of the set of suitable solutions x , then $p(x) \equiv 0$ for all x not in the range of the network and optimizing the network will indeed yield a non-trivial completion $x = (x_{\mathcal{N}}, x_{\perp})$. More precisely (I) can be interpreted to be a deep prior on the measurement and (II) to be a deep prior on the nullspace part of the problem.

4 Deep Priors and Tikhonov Functionals

In this section, we consider the particular network architecture given by unrolled proximal gradient schemes, as shown in Sect. 3.2. We aim at embedding this approach into the classical regularization theory for inverse problems. For a strict mathematical analysis, we will introduce the notion of an analytic deep prior network, which then allows interpreting the training of the deep prior network as an optimization of a Tikhonov functional. The main result of this section is Theorem 4.2, which states that analytic deep priors in combination with a suitable stopping rule are indeed order optimal regularization schemes. Numerical experiments in Sect. 4.2 demonstrate that such deep prior approaches lead to

smaller reconstruction errors when compared with standard Tikhonov reconstructions. The superiority of this approach can be proved, however, only for the rather unrealistic case, that the solution coincides with a singular function of A .

4.1 Unrolled Proximal Gradient Networks as Deep Priors for Inverse Problems

In this section, we consider linear operators A and aim at rephrasing DIP, i.e., the minimization of (1.1) with respect to Θ , as a constrained optimization problem. This change of view, i.e., regarding deep inverse priors as an optimization of a simple but constrained functional, rather than networks, opens the way for analytic investigations. We will use an unrolled proximal gradient architecture for the network $\varphi_\Theta(z)$ in (1.1). The starting point for our investigation is the common observation, as shown in [11,16] or “Appendix 1”, that an unrolled proximal gradient scheme as defined in Sect. 3.2 approximates a minimizer $x(B)$ of (3.3). Assuming that a unique minimizer $x(B)$ exists as well as neglecting the difference between $x(B)$ and the approximation $\varphi_\Theta(z)$ achieved by the unrolled proximal gradient motivates the following definition of analytic deep priors.

Definition 4.1 Let us assume that measured data $y^\delta \in Y$, a fixed $\alpha > 0$, a convex penalty functional $R : X \rightarrow \mathbb{R}$ and a measurement operator $A \in \mathcal{L}(X, Y)$ are given. We consider the minimization problem

$$\min_B F(B) = \min_B \frac{1}{2} \|Ax(B) - y^\delta\|^2, \tag{4.1}$$

subject to the constraint

$$\begin{aligned} x(B) &= \arg \min_x J_B(x) \\ &= \arg \min_x \frac{1}{2} \|Bx - y^\delta\|^2 + \alpha R(x). \end{aligned} \tag{4.2}$$

We assume that for every $B \in \mathcal{L}(X, Y)$, there is a unique minimizer $x(B)$. We call this constrained minimization problem an analytic deep prior and denote by $x(B)$ the resulting solution to the inverse problems posed by A and y^δ .

We can also use this technical definition as the starting point of our consideration and retrieve the neural network architecture by considering the following approach for solving the minimization problem stated in the above definition. Assuming that R has a proximal operator, we can compute $x(B)$, given B , via proximal gradient method. That is, via the (for a suitable choice of $\lambda > 0$ and an arbitrary $x^0 = z \in X$) converging iteration

$$x^{k+1} = \text{Prox}_{\lambda\alpha R} \left(x^k - \lambda B^*(Bx^k - y^\delta) \right). \tag{4.3}$$

Following this iteration for L steps can be seen as the forward pass of a particular architecture of a fully connected feedforward network with L layers of identical size as described in (3.1) and (3.2). The affine linear map given by $\Theta = (W, b)$ is the same for all layers. Moreover, the activation function of the network is given by the proximal mapping of $\lambda\alpha R$, the matrix W is given via $I - W = \lambda B^* B$ (I denotes the identity operator), and the bias is determined by $b = \lambda B^* y^\delta$.

From now on we will assume that the difference between x^L and $x(B)$ is negligible, i.e.,

$$x^L = x(B). \tag{4.4}$$

Remark 4.1 The task in the DIP approach is to find Θ (network parameters). Analogously, in the analytic deep prior, we try to find the operator B .

We now examine the analytic deep image prior utilizing the proximal gradient descent approach to compute $x(B)$. Therefore, we will focus on the minimization of (4.1) with respect to B for given data y^δ by means of gradient descent.

The stationary points are characterized by $\partial F(B) = 0$, and gradient descent iterations with stepsize η are given by

$$B^{\ell+1} = B^\ell - \eta \partial F(B^\ell). \tag{4.5}$$

Hence, we need to compute the derivative of F with respect to B .

Lemma 4.1 Consider an analytic deep prior with the proximal gradient descent approach as described above. We define

$$\psi(x, B) = \text{Prox}_{\lambda\alpha R} (x - \lambda B^*(Bx - y^\delta)) - x. \tag{4.6}$$

Then,

$$\partial F(B) = \partial x(B)^* A^* (Ax(B) - y^\delta) \tag{4.7}$$

with

$$\partial x(B) = -\psi_x(x(B), B)^{-1} \psi_B(x(B), B), \tag{4.8}$$

which leads to the gradient descent

$$B^{\ell+1} = B^\ell - \eta \partial F(B^\ell). \tag{4.9}$$

This lemma allows to obtain an explicit description of the gradient descent for B , which in turn leads to an iteration of functionals J_B and minimizers $x(B)$. We will now exemplify this derivation for a rather academic example, which however highlights in particular the differences between a classical Tikhonov minimizer, i.e.,

$$x(A) = \arg \min_x \frac{1}{2} \|Ax - y^\delta\|^2 + \frac{\alpha}{2} \|x\|^2,$$

and the solution of the DIP approach.

4.1.1 Example

In this example, we examine analytic deep priors for linear inverse problems $A : X \rightarrow Y$, i.e., $A, B \in \mathcal{L}(X, Y)$, and

$$R(x) = \frac{1}{2} \|x\|^2. \tag{4.10}$$

The rather abstract characterization of the previous section can be made explicit for this setting. Since $J_B(x)$ is the classical Tikhonov regularization, which can be solved by

$$x(B) = (B^*B + \alpha I)^{-1} B^* y^\delta, \tag{4.11}$$

we can rewrite the analytic deep prior reconstruction as $x(B)$, where B is minimizing

$$F(B) = \frac{1}{2} \|A(B^*B + \alpha I)^{-1} B^* y^\delta - y^\delta\|^2. \tag{4.12}$$

Lemma 4.2 *Following Lemma 4.1, assuming $B^0 = A$ and computing one step of gradient descent to minimize the functional with respect to B , yields*

$$B_1 = A - \eta \partial F(A) \tag{4.13}$$

with

$$\begin{aligned} \partial F(A) &= \partial x(A)^* A^* (Ax(A) - y^\delta) \\ &= \alpha AA^* y^\delta (y^\delta)^* A (A^* A + \alpha I)^{-3} \end{aligned} \tag{4.14}$$

$$\begin{aligned} &+ \alpha A (A^* A + \alpha I)^{-3} A^* y^\delta (y^\delta)^* A \\ &- \alpha y^\delta (y^\delta)^* A (A^* A + \alpha I)^{-2}. \end{aligned} \tag{4.15}$$

This expression nicely collapses if $y^\delta (y^\delta)^*$ commutes with AA^* . For illustration, we assume the rather unrealistic case that $x^+ = u$, where u is a singular function for A with singular value σ . The dual singular function is denoted by v , i.e., $Au = \sigma v$ and $A^*v = \sigma u$ and we further assume that the measurement noise in y^δ is in the direction of this singular function, i.e., $y^\delta = (\sigma + \delta)v$, as shown in Fig. 3. In this case, the problem is indeed one-dimensional and we obtain an iteration restricted to the span of u , resp. the span of v .

Lemma 4.3 *The setting described above yields the following gradient step for the functional in (4.12):*

$$B^{\ell+1} = B^\ell - c_\ell v u^* \tag{4.16}$$

with

$$c_\ell = c(\alpha, \delta, \sigma, \eta) = \eta \sigma (\sigma + \delta)^2 (\alpha + \beta_\ell^2 - \sigma \beta_\ell) \frac{\beta_\ell^2 - \alpha}{(\beta_\ell^2 + \alpha)^3},$$




Fig. 3 Example of $y^\delta = (\sigma + \delta)v$ where v is a singular function of A (integral operator) (Color figure online)




Fig. 4 Comparison of the Tikhonov reconstruction (orange broken line), the result obtained in (4.17) (blue continuous line) and the direct inverse. In this example, we considered $\alpha = 10^{-3}$ (Color figure online)

and the iteration (4.16) in-turn results in the sequence $x(B^\ell)$ with the unique attractive stationary point

$$x = \begin{cases} \frac{1}{2\sqrt{\alpha}} (\sigma + \delta)u, & \sigma < 2\sqrt{\alpha} \\ \frac{1}{\sigma} (\sigma + \delta)u, & \text{otherwise.} \end{cases} \tag{4.17}$$

For comparison, the classical Tikhonov regularization would yield $\frac{\sigma}{\sigma^2 + \alpha} (\sigma + \delta)u$. This is depicted in Fig. 4.

4.1.2 Constrained System of Singular Functions

In the previous example, we showed that if we do gradient descent starting from $B_0 = A$ and assume the rather simple case $y^\delta = (\sigma + \delta)v$, we obtain the iteration $B^{\ell+1} = B^\ell - c_\ell v u^*$, i.e., $B^{\ell+1}$ has the same singular functions as A and only one of the singular values is different.

We now analyze the optimization from a different perspective. Namely, we focus on finding directly a minimizer of (4.1) for a general $y^\delta \in Y$; however, we restrict B to be an operator such that B^*B commutes with A^*A , i.e., A and B share a common system of singular functions. Hence, B has the following representation.

$$B = \sum_i \beta_i v_i u_i^*, \quad \beta_i \in \mathbb{R}_+ \cup \{0\}, \tag{4.18}$$

where $\{u_i, \sigma_i, v_i\}$ is the singular value decomposition of A . That means, we restrict the problem to finding optimal singular values β_i for B . In this case, we show that a global minimizer exists and that it has interesting properties.

Theorem 4.1 For any $y^\delta \in Y$, there exist a global minimizer (in the constrained singular functions setting) of (4.1) given by $B_\alpha = \sum \beta_i^\alpha v_i u_i^*$ with

$$\beta_i^\alpha(\sigma) = \begin{cases} \frac{\sigma_i}{2} + \sqrt{\frac{\sigma_i^2}{4} - \alpha} & \sigma \geq 2\sqrt{\alpha} \\ \sqrt{\alpha} & \sigma < 2\sqrt{\alpha} \end{cases} \quad (4.19)$$

Remark 4.2 The singular values obtained in Theorem 4.1 match the ones obtained in the previous section for general B but simple $y^\delta = (\sigma + \delta)v$.

Remark 4.3 The minimizer from Theorem 4.1 does not depend on y^δ , i.e., $\forall : y^\delta \in Y$ it holds that B_α is a minimizer of (4.1). The solution to the inverse problem does still depend on y^δ since

$$x(B_\alpha) = \arg \min_x \frac{1}{2} \|B_\alpha x - y^\delta\|^2 + \alpha R(x). \quad (4.20)$$

Remark 4.4 In the original DIP approach, some of the parameters of the network may be similar for different y^δ , for example, the parameters of the first layers of the encoder part of the UNet. Other parameters may strongly depend on y^δ . In this particular case of the analytic deep prior (constrained system of singular functions), we have a explicit separation of which parameters ($b = \lambda B^* y^\delta$) depend on y^δ and which do not ($W = I - \lambda B^* B$).

From now on we consider the notation $x(B, y^\delta)$ to incorporate the dependency of $x(B)$ on y^δ . Following the classical filter theory for order optimal regularization schemes, [13,21,28], we obtain the following theorem.

Theorem 4.2 The pseudoinverse $K_\alpha : Y \rightarrow X$ defined as

$$K_\alpha(y^\delta) := x(B_\alpha, y^\delta) \quad (4.21)$$

is an order optimal regularization method given by the filter functions

$$F_\alpha(\sigma) = \begin{cases} 1 & \sigma \geq 2\sqrt{\alpha} \\ \frac{\sigma}{2\sqrt{\alpha}} & \sigma < 2\sqrt{\alpha} \end{cases} \quad (4.22)$$

The regularized pseudoinverse K_α is quite similar to the truncated singular value decomposition (TSVD) but is a softer version because it does not have a jump (see Fig. 5). We call this method *Soft TSVD*.

The disadvantage of Tikhonov, in this case, is that it damps all singular values, and the disadvantage of TSVD is that it throws away all the information related to small singular values. On the other hand, the Soft TSVD does not damp the higher singular values (similar to TSVD) and does not throw away the information related to smaller singular values but does damp it (similar to Tikhonov). For a comparison of the filter functions, see Table 1. Moreover, what is interesting




Fig. 5 Filter response of TSVD, Tikhonov and the Soft TSVD (Color figure online)

Table 1 Values of ν for which TSVD, Tikhonov and the Soft TSVD are order optimal

| Method | $F_\alpha(\sigma)$ | γ | ν |
|-----------|---|----------|---------------|
| Tikhonov | $\frac{\sigma^2}{\sigma^2 + \alpha}$ | 1/2 | $2 > \nu > 0$ |
| TSVD | $\begin{cases} 1 & \sigma \geq \alpha \\ 0 & \sigma < \alpha \end{cases}$ | 1 | $\nu > 0$ |
| Soft TSVD | $\begin{cases} 1 & \sigma \geq 2\sqrt{\alpha} \\ \frac{\sigma}{2\sqrt{\alpha}} & \sigma < 2\sqrt{\alpha} \end{cases}$ | 1/2 | $\nu > 0$ |

For more details see the Proof of Theorem 4.2 in “Appendix 2”

is how this method comes out from Definition 4.1, which is stated in terms of the Tikhonov pseudoinverse, and that the optimal singular values do not depend on y^δ .

At this point, the relation to the original DIP approach becomes more abstract. We considered a simplified network architecture where all layers share the same weights that come from an iterative algorithm for solving inverse problems. That means, we let the solution to the original inverse problem be the solution of another problem with different operator B . The DIP approach in this case is transformed to finding an optimal B and allows us to do the analysis in the functional analysis setting. What we learn from the previous results is that we can establish interesting connections between the DIP approach and the classical inverse problems theory. This is important because it shows that deep inverse priors can be used to solve really ill-posed inverse problems.

Remark 4.5 In the original DIP, the input z to the network is chosen arbitrarily and is of minor importance. However, once the weights have been trained for a given y^δ , z cannot be changed because it would affect the output of the network, i.e., it would change the obtained reconstruction. In the analytic deep prior, the input to the unrolled proximal gradient method is completely irrelevant (assuming an infinite number of layers). After finding the “weights” B , a different input will still produce the same solution $\hat{x} = x(B) = \varphi_\Theta(z)$.

Remark 4.5 tells us that there is still a gap between the original DIP and the analytic one. This was expected because of the obvious trivialization of the network architecture but serves as motivation for further research.




Fig. 6 Example of y^δ for $x^\dagger = u$ (singular function) with a SNR of 17.06 db (Color figure online)




Fig. 7 Example of a more general y^δ with a SNR of 18.97 db (Color figure online)

4.2 Numerical Experiments

We now use the analytic deep inverse prior approach for solving an inverse problem with the following integration operator $A : L^2([0, 1]) \rightarrow L^2([0, 1])$

$$(Ax)(t) = \int_0^t x(s) ds. \tag{4.23}$$

A is a linear and compact operator, hence the inverse problem is ill-posed. Let $A_n \in \mathbb{R}^{n \times n}$ be a discretization of A and $x^\dagger \in \mathbb{R}^n$ to be one of its discretized singular vectors u . We set the noisy data $y^\delta = A_n x^\dagger + \delta \tau$ with $\tau \sim \text{Normal}(0, \mathbb{1}_n)$, as shown in Fig. 6. A more general example, i.e., where x^\dagger is not restricted to be a singular function, is also included (Fig. 7).

We aim at recovering x^\dagger from y^δ considering the setting established in Definition 4.1 for $R(\cdot) = \frac{1}{2} \|\cdot\|^2$. That means that the solution x is parametrized by the operator B . Solving the inverse problem is now equivalent to finding optimal B that minimizes the loss function (1.1) for the single data point (z, y^δ) .

To find such a B , we go back to the DIP and the neural network approach. We write $x(B)$ as the output of the network φ_Θ defined in (3.1) with some randomly initialized input z . We optimize with respect to B , which is a matrix in the discretized setting, and obtain a minimizer B_{opt} of (1.1). For more details, please refer to ‘‘Appendix 3.’’

In Fig. 8, we show some reconstruction results. The first plot of each row contains the true solution x^\dagger , the standard

Tikhonov solution $x(A)$ and the reconstruction obtained with the analytic deep inverse approach $x(B_{\text{opt}})$ after B converged. For each case, we provide additional plots depicting:

- The true error of the network’s output $x(B)$ after each update of B in a logarithmic scale.
- The squared Frobenius norm of $B_k - B_{k+1}$ after each update of B .
- The matrix B_{opt} .

For all choices of α , the training of B converges to a matrix B_{opt} , such that $x(B_{\text{opt}})$ has a smaller true error than $x(A)$. In the third plot of each row, one can check that B indeed converges to some matrix B_{opt} , which is shown in the last plot. The networks were trained using gradient descent with 0.05 as learning rate.

The theoretical findings of the previous subsections allow us to compute, either the exact update (4.16) for B in the rather unrealistic case that $y^\delta = (\sigma + \delta)v$, or the exact solution $x(B_\alpha, y^\delta)$ if we restrict B to have the same system of singular functions as A (Theorem 4.1). In the numerical experiments, we do not consider any of these restrictions, and therefore, we cannot directly apply our theoretical results. Instead, we implement the network approach (see ‘‘Appendix 3’’) to be able to find B_{opt} in a more general scenario. Nevertheless, as it can be observed in the last plot of each row in Fig. 8, B_{opt} contains some patterns that reflect, to some extent, that B keeps the same singular system but with different singular values. Namely, B is updated in a similar way as in (4.16). With the current implementation, we could also use more complex regularization functionals R , in order to reduce the gap between our analytic approach and the original DIP. This is also a motivation for further research.

5 Summary and Conclusion

In this paper, we investigated the concept of deep inverse priors/regularization by architecture. This approach neither requires massive amounts of ground truth/surrogate data, nor pretrained models/transfer learning. The method is based on a single measurement. We started by giving different qualitative interpretations of what regularization is and specifically how regularization by architecture fits into this context.

We followed up with the introduction of the analytic deep prior by explicitly showing how unrolled proximal gradient architectures, allow for a somewhat transparent regularization by architecture. Specifically, we showed that their results can be interpreted as solutions of optimized Tikhonov functionals and proved precise equivalences to regularization techniques. We further investigated this point of view with an academic example, where we implemented the analytic deep inverse prior and tested its numerical applicability.

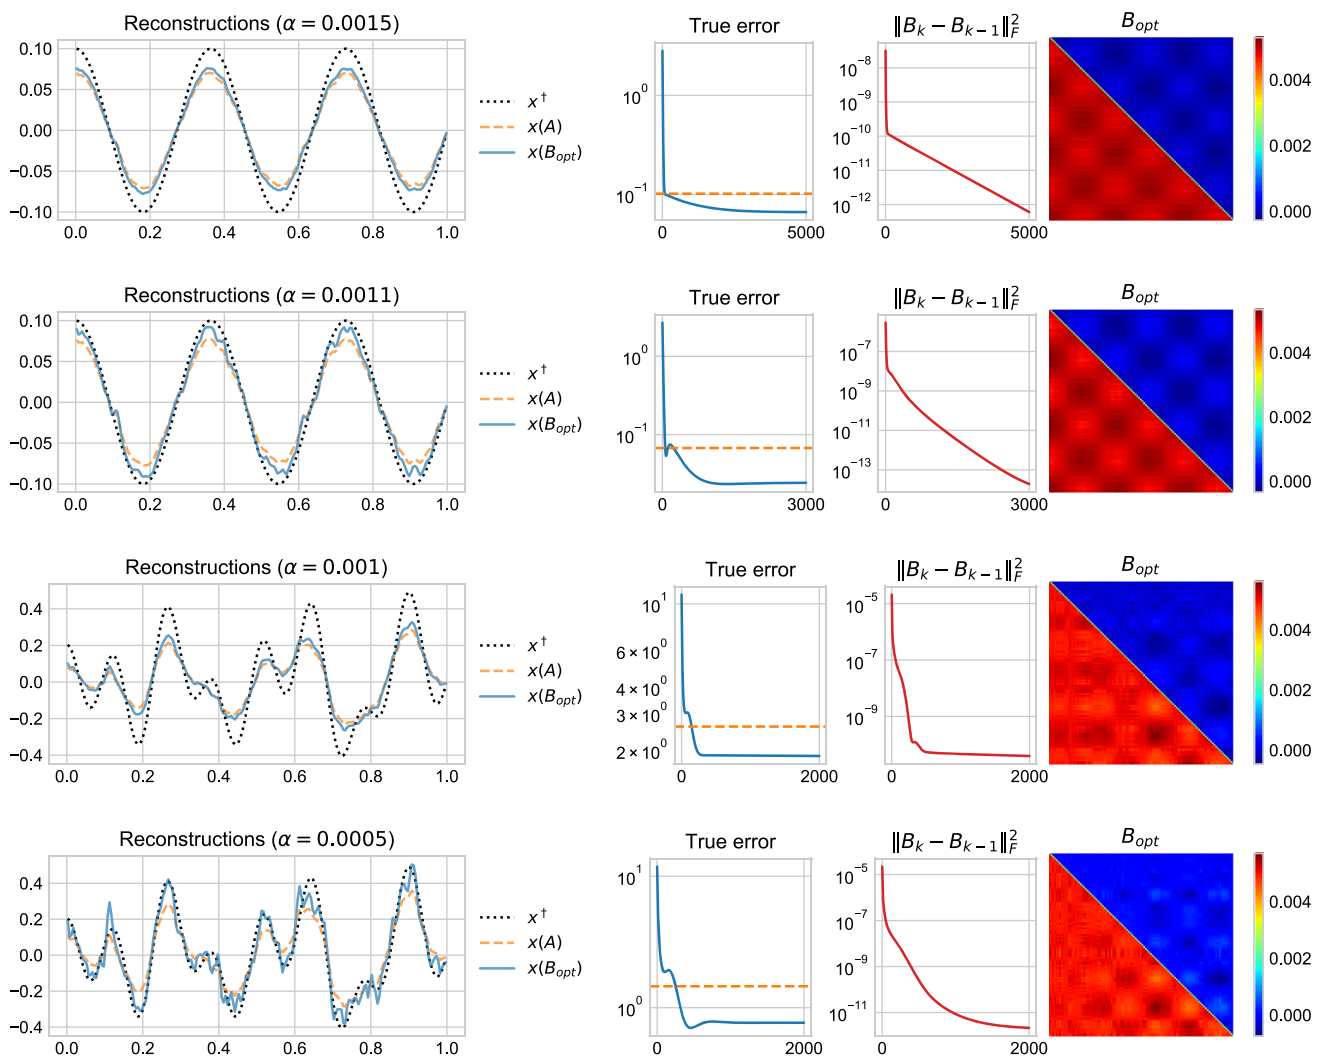


Fig. 8 Reconstructions corresponding to y^δ as in Fig. 6 (first and second row) and Fig. 7 (third and fourth row) for different values of α . The broken line in the second plot of each row indicates the true error of

the standard Tikhonov solution $x(A)$. The horizontal axis in the second and third plots indicates the number of weights updates (Color figure online)

The results confirmed our theoretical findings and showed promising results.

There is obviously, like in deep learning in general, much work to be done in order to have a good understanding of deep inverse priors, but we see much potential in the idea of using deep architectures to regularize inverse problems; especially since an enormous part of the deep learning community is already concerned with the understanding of deep architectures.

Acknowledgements S. Dittmer and T. Kluth acknowledge the support by the Deutsche Forschungsgemeinschaft (DFG) within the framework of GRK 2224/1 “Pi 3 : Parameter Identification—Analysis, Algorithms, Applications”. D. Otero Bager acknowledges the financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project Number 276397488—SFB 1232, sub-project “P02-Heuristic, Statistical and Analytical Experimental Design”. Peter Maass acknowledges funding by the European Union’s Horizon 2020

research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 765374, sub-project ‘Data driven model adaptations of coil sensitivities in MR systems’. The authors would also like to acknowledge the anonymous reviewers for their helpful comments and suggestions.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix 1: A Reminder on Minimization of Tikhonov Functionals and the LISTA Approach

In this section, we consider only linear operators A and we review the well-known theory for the Iterative Soft Shrinkage

Algorithm (ISTA) as well as the slightly more general Proximal Gradient (PG) [11,26] method for minimizing Tikhonov functionals of the type

$$J(x) = \frac{1}{2} \|Ax - y^\delta\|^2 + \alpha R(x). \tag{5.1}$$

We recapitulate the main steps in deriving ISTA and PG, as far as we need it for our motivation. The necessary first-order condition for a minimizer is given by

$$0 \in A^*(Ax - y^\delta) + \alpha \partial R(x). \tag{5.2}$$

Multiplying with an arbitrary real positive number λ and adding x plus rearranging yields

$$x - \lambda A^*(Ax - y^\delta) \in x + \lambda \alpha \partial R(x). \tag{5.3}$$

For convex R , the term of the right hand side is inverted by the (single valued) proximal mapping of $\lambda \alpha R$, which yields

$$\text{Prox}_{\lambda \alpha R}(x - \lambda A^*(Ax - y^\delta)) = x. \tag{5.4}$$

Hence, this is a fixed point condition, which is a necessary condition for all minimizers of J . Turning the fixed point condition into an iteration scheme yields the PG method

$$x^{k+1} = \text{Prox}_{\lambda \alpha R}(x^k - \lambda A^*(Ax^k - y^\delta)) \tag{5.5}$$

$$= \text{Prox}_{\lambda \alpha R}((I - \lambda A^*A)x^k + \lambda A^*y^\delta). \tag{5.6}$$

This structure is also the motivation for LISTA [16] approaches where fully connected networks with L internal layers of identical size are used. Moreover, in some versions of LISTA, the affine maps between the layers are assumed to be identical. The values at the k th layer are denoted by x^k , hence,

$$x^{k+1} = \phi(Wx^k + b). \tag{5.7}$$

LISTA then trains (W, b) on some given training data. More precisely, it trains two matrices $W = I - \lambda A^*A$ and $S = \lambda A^*$ such that

$$x^{k+1} = \phi(Wx^k + Sy^\delta). \tag{5.8}$$

This derivation can be rephrased as follows.

Lemma 5.1 *Let φ_Θ , $\Theta = (W, b)$, denote a fully connected network with input x^0 and L -internal layers. Further assume that the activation function is identical to a proximal mapping for a convex functional $\lambda \alpha R : X \rightarrow \mathbb{R}$. Assume W*

is restricted, such that $I - W$ is positive definite, i.e., there exists a matrix B such that

$$I - W = \lambda B^*B. \tag{5.9}$$

*Furthermore, we assume that the bias term is fixed as $b = \lambda B^*y^\delta$. Then, $\varphi_\Theta(z)$ is the L th iterate of an ISTA scheme with starting value $x^0 = z$ for minimizing*

$$J_B(x) = \frac{1}{2} \|Bx - y^\delta\|^2 + \alpha R(x). \tag{5.10}$$

Proof Follows directly from Eq. (5.5). □

Appendix 2: Proofs

Proof of Lemma 4.1

F is a functional which maps operators B to real numbers, hence, its derivative is given by

$$\partial F(B) = [\partial x(B)^*] A^*(Ax(B) - y^\delta),$$

which follows from classical variational calculus, see, e.g., [13]. The derivative of $x(B)$ with respect to B can be computed using the fix point condition for a minimizer of J_B , namely

$$\text{Prox}_{\lambda \alpha R}(x(B) - \lambda B^*(Bx - y^\delta)) - x(B) = 0,$$

which is equivalent to

$$\psi(x(B), B) = 0.$$

We apply the implicit function theorem and obtain the derivative

$$\partial x(B) = -\psi_x(x(B), B)^{-1} \psi_B(x(B), B). \tag{5.11}$$

Combining $\partial F(B)$ with $\partial x(B)$ yields the required result. □

Proof of Lemma 4.2

We start with the explicit description of the iteration

$$B^{\ell+1} = B^\ell - \eta \partial F(B^\ell) \tag{5.12}$$

with

$$\partial F(B) = \partial x(B)^* A^*(Ax(B) - y^\delta). \tag{5.13}$$

The derivative of $x(B)$ with respect to B is a linear map $\partial x(B) : \mathcal{L}(X, Y) \rightarrow X$. For $\delta B \in \mathcal{L}(X, Y)$ we obtain

$$\begin{aligned} \partial x(B)(\delta B) &= -(B^*B + \alpha I)^{-2} (\delta B^*B + B^*\delta B) B^*y^\delta \\ &\quad + (B^*B + \alpha I)^{-1} \delta B^*y^\delta. \end{aligned} \tag{5.14}$$

The adjoint operator is a mapping from X to $\mathcal{L}(X, Y)$, which can be derived from the defining relation

$$\langle \partial x(B)(\delta B), z \rangle_X = \langle \delta B, [\partial x(B)]^* z \rangle_{\mathcal{L}(X, Y)}. \tag{5.15}$$

Hence,

$$\begin{aligned} [\partial x(B)]^* z &= -BB^*y^\delta z^*(B^*B + \alpha I)^{-2} \\ &\quad - B(B^*B + \alpha I)^{-2} z(y^\delta)^*B \\ &\quad + y^\delta z^*(B^*B + \alpha I)^{-1}. \end{aligned} \tag{5.16}$$

Here, $y^\delta z^* \in \mathcal{L}(X, Y)$ denotes a linear map, which maps an $x \in X$ to $\langle z, x \rangle_X y^\delta$.

First of all, we now aim at determining explicitly $\partial F(B)$ at the starting point of our iteration, i.e., with $B^0 = A$.

From this follows the rather lengthy expression

$$\begin{aligned} \partial F(A) &= \partial x(A)^* A^*(Ax(A) - y^\delta) \\ &= \alpha AA^*y^\delta(y^\delta)^*A(A^*A + \alpha I)^{-3} \\ &\quad + \alpha A(A^*A + \alpha I)^{-3} A^*y^\delta(y^\delta)^*A \\ &\quad - \alpha y^\delta(y^\delta)^*A(A^*A + \alpha I)^{-2}. \end{aligned} \tag{5.17} \tag{5.18}$$

This enables us to compute the update

$$B_1 = A - \eta \partial F(A) \tag{5.19}$$

as well as the output of the analytic deep prior approach after one iteration of updating B (assuming a suitably chosen η)

$$x(B_1) = (B_1^*B_1 + \alpha I)^{-1} B_1^*y^\delta. \tag{5.20}$$

Proof of Lemma 4.3

A lengthy computation exploiting $B^0 = A$ and $\beta_0 = \sigma$ shows that the singular value β_ℓ of u in the spectral decomposition of B^ℓ obeys the iteration

$$\beta_{\ell+1} = \beta_\ell - \eta \sigma (\sigma + \delta)^2 (\alpha + \beta_\ell^2 - \sigma \beta_\ell) \frac{\beta_\ell^2 - \alpha}{(\beta_\ell^2 + \alpha)^3}, \tag{5.21}$$

i.e.,

$$B^{\ell+1} = B^\ell - c_\ell v u^* \tag{5.22}$$

with

$$c_\ell = c(\alpha, \delta, \sigma, \eta) = \eta \sigma (\sigma + \delta)^2 (\alpha + \beta_\ell^2 - \sigma \beta_\ell) \frac{\beta_\ell^2 - \alpha}{(\beta_\ell^2 + \alpha)^3}.$$

We will now consider the stability of the fixed points of the sequence $x(B^\ell)$, i.e., we will analyze the fixed points of the iteration described in (5.21), that is,

$$\beta_{\ell+1} = \beta_\ell - c(\beta_\ell), \tag{5.23}$$

where

$$c(\beta) = \eta \sigma (\sigma + \delta)^2 (\alpha + \beta^2 - \sigma \beta) \frac{\beta^2 - \alpha}{(\beta^2 + \alpha)^3}. \tag{5.24}$$

This iteration in-turn gives you via the Tikhonov filter function, the sequence

$$x(\beta_\ell) = \frac{\beta_\ell}{\beta_\ell^2 + \alpha} (\sigma + \delta) u \tag{5.25}$$

of reconstructions. To find the fixed points of the iteration, we analyze the real roots of c , which are

- $\beta^{(1)} = \sqrt{\alpha}$,
- $\beta^{(2)} = -\sqrt{\alpha}$,
- $\beta^{(3)} = \frac{\sigma}{2} + \sqrt{\frac{\sigma^2}{4} - \alpha}$, for $\sigma \geq 2\sqrt{\alpha}$ and
- $\beta^{(4)} = \frac{\sigma}{2} - \sqrt{\frac{\sigma^2}{4} - \alpha}$, for $\sigma \geq 2\sqrt{\alpha}$.

Simple calculations show that

- $\partial_\beta c(\beta^{(1)}) \begin{cases} > 0, & \sigma < 2\sqrt{\alpha} \\ \leq 0, & \text{otherwise.} \end{cases}$
- $\partial_\beta c(\beta^{(2)}) > 0$
- $\partial_\beta c(\beta^{(3)}) > 0$, for $\sigma \geq 2\sqrt{\alpha}$ and
- $\partial_\beta c(\beta^{(4)}) > 0$ for $\sigma \geq 2\sqrt{\alpha}$.

This leads to the single attractive fixed point $\beta^{(1)}$ for $\sigma < 2\sqrt{\alpha}$ and the two attractive fixed points $\beta^{(3)}$ and $\beta^{(4)}$ otherwise. Since,

$$x(\beta^{(3)}) = x(\beta^{(4)}), \tag{5.26}$$

we therefore have a unique reconstruction, namely

$$x = \begin{cases} \frac{1}{2\sqrt{\alpha}} (\sigma + \delta) u, & \sigma < 2\sqrt{\alpha} \\ \frac{1}{\sigma} (\sigma + \delta) u, & \text{otherwise.} \end{cases} \tag{5.27}$$

□

Proof of Theorem 4.1

Let $B = \sum \beta_i v_i u_i^*$. We want to find $\{\beta_i\}$ to minimize

$$F(B) = \|Ax(B, y^\delta) - y^\delta\|^2. \tag{5.28}$$

The Tikhonov solution is given by

$$x(B) = \sum \frac{\beta_i}{\beta_i^2 + \alpha} \langle y^\delta, v_i \rangle u_i, \tag{5.29}$$

the result of applying the operator A to $x(B)$ is

$$Ax(B) = \sum \frac{\sigma_i \beta_i}{\beta_i^2 + \alpha} \langle y^\delta, v_i \rangle v_i \tag{5.30}$$

and

$$y^\delta = \sum \langle y^\delta, v_i \rangle v_i. \tag{5.31}$$

Inserting (5.30) and (5.31) in (5.28) yields

$$F(B) = \sum \left| \left(\frac{\sigma_i \beta_i}{\beta_i^2 + \alpha} - 1 \right) \langle y^\delta, v_i \rangle \right|^2. \tag{5.32}$$

In order to minimize $F(B)$, we should set $\frac{\sigma_i \beta_i}{\beta_i^2 + \alpha} = 1$ which implies $\beta_i^2 - \sigma_i \beta_i + \alpha = 0$. The roots of the previous equation are $\beta_i = \frac{\sigma_i}{2} \pm \sqrt{\frac{\sigma_i^2}{4} - \alpha}$ and they are real only if $\frac{\sigma_i^2}{4} \geq \alpha$. If it does not hold then $\frac{\alpha \beta_i}{\beta_i^2 + \alpha} < 1$ and the optimal choice is to find its maximum value which is attained at $\beta_i = \sqrt{\alpha}$.

Therefore, we set

$$\beta_i = \begin{cases} \frac{\sigma_i}{2} + \sqrt{\frac{\sigma_i^2}{4} - \alpha} & \sigma \geq 2\sqrt{\alpha} \\ \sqrt{\alpha} & \sigma < 2\sqrt{\alpha} \end{cases} \tag{5.33}$$

and we minimize every term in the sum (5.32), which means we have found singular values $\{\beta_i\}$ that minimize $F(B)$. □

Proof of Theorem 4.2

In order to prove that K_α is a proper order optimal regularization method, we need to check if the corresponding filters F_α from (4.22) satisfy the three conditions of optimality [21,28].

These conditions state that a filter $F_\alpha : \mathbb{R} \rightarrow \mathbb{R}$ is an order optimal regularization filter if $\exists \gamma, c_1, c_2, c_3 > 0$ such that

1. $\sup_\sigma |F_\alpha(\sigma)\sigma^{-1}| \leq c_1 \alpha^{-\gamma}$
2. $\sup_\sigma |1 - F_\alpha(\sigma)| \sigma^\nu < c_2 \alpha^\nu$
3. $\forall \alpha > 0, \sigma > 0 : |F_\alpha(\sigma)| \leq c_3$

In the following, we show that they hold $\forall \nu > 0$ with $\gamma = \frac{1}{2}, c_1 = \frac{1}{2}, c_2 = 2^\nu, c_3 = 1$:

- i. If $\sigma \geq 2\sqrt{\alpha}$
 1. $\sup_\sigma |F_\alpha(\sigma)\sigma^{-1}| = \sup_\sigma |\sigma^{-1}| \leq \frac{1}{2}\alpha^{-\frac{1}{2}}$
 2. $\sup_\sigma |1 - F_\alpha(\sigma)| \sigma^\nu = 0 \leq \alpha^\nu$
 3. $\forall \alpha > 0, \sigma > 0 : |F_\alpha(\sigma)| = 1$
- ii. If $\sigma < 2\sqrt{\alpha}$
 1. $\sup_\sigma |F_\alpha(\sigma)\sigma^{-1}| = \frac{1}{2}\alpha^{-\frac{1}{2}}$
 2. $\sup_\sigma |1 - F_\alpha(\sigma)| \sigma^\nu = \sup_\sigma \left| \frac{2\sqrt{\alpha} - \sigma}{2\sqrt{\alpha}} \right| \sigma^\nu \leq 2^\nu \alpha^{\frac{\nu}{2}}$
 3. $\forall \alpha > 0, \sigma > 0 : |F_\alpha(\sigma)| = \frac{\sigma}{2\sqrt{\alpha}} \leq 1$

□

Appendix 3: Numerical Experiments

In this section, we provide details about the implementation of the analytic deep inverse prior and the academic example. We start by discretizing the integration operator, which yields the matrix $A_n \in \mathbb{R}^{n \times n}$, that has $\frac{h}{2}$ on the main diagonal, h everywhere under the main diagonal and 0 above (here $h = \frac{1}{n}$). In our experiments, we use $n = 200$.

The analytic deep inverse prior network φ_Θ^L is implemented using Python and Tensorflow [1]. Initially, we create the matrix $B \in \mathbb{R}^{n \times n}$ and add L fully connected layers to the network, all having the same parameters $\Theta = (W, b)$, with weight matrix $W = I - \lambda B^T B$, bias $b = \lambda B^T y^\delta$ and activation function given by the ℓ_2 proximal operator. That means the network contains in total 4×10^4 parameters (the number of components in B). For the experiments shown in the paper, the input z is randomly initialized with a small norm and λ is $\frac{1}{\mu}$, where μ is the biggest eigenvalue of $A^T A$.

We follow the DIP approach and minimize (1.1) using gradient descent. To guarantee that $\varphi_\Theta^L(z) = x(B)$ holds, the network should have thousands of layers, because of the slow convergence of the PG method. This is prohibitive from the implementation point of view. Therefore, we consider only a reduced network with a small number of layers, $L = 10$, and at each iteration we set the input of the network to be the network’s output after the previous iteration. This is equivalent to adding L new identical layers after each update of B , with

$$W_i = I - \lambda B_i^T B_i \tag{5.34}$$

and

$$b_i = \lambda B_i^T y^\delta, \tag{5.35}$$




Fig. 9 The implicit network with $(k + 1)L$ layers. Here, $\phi_{\Theta_k}^L$ refers to a block of L identical fully connected layers with weights $\Theta_k = (W_k, b_k)$

where B_i refers to the value of B at the i th iteration. After k iterations, we implicitly create a network that has $(k + 1)L$ layers (Fig. 9), however, each time we update B , we back-propagate only through the last L layers.

References


- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: large-scale machine learning on heterogeneous systems (2015). <https://www.tensorflow.org/>. Software available from tensorflow.org
- Adler, J., Öktem, O.: Learned primal-dual reconstruction. *IEEE Trans. Med. Imaging* **37**(6), 1322–1332 (2018)
- Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M.W., Pfau, D., Schaul, T., Shillingford, B., De Freitas, N.: Learning to learn by gradient descent by gradient descent. In: *Advances in Neural Information Processing Systems*, pp. 3981–3989 (2016)
- Anonymous: on the spectral bias of neural networks. In: *Submitted to International Conference on Learning Representations (under review)* (2019). <https://openreview.net/forum?id=r1gR2sC9FX>. Accessed 28 Oct 2019
- Arridge, S., Maass, P., Öktem, O., Schönlieb, C.B.: Solving inverse problems using data-driven models. *Acta Numer.* **28**, 1–174 (2019)
- Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.* **2**(1), 183–202 (2009). <https://doi.org/10.1137/080716542>
- Bora, A., Jalal, A., Price, E., Dimakis, A.G.: Compressed sensing using generative models. In: *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6–11 August 2017*, pp. 537–546 (2017). <http://proceedings.mlr.press/v70/bora17a.html>
- Bruna, J., Mallat, S.: Invariant scattering convolution networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(8), 1872–1886 (2013)
- Chen, X., Liu, J., Wang, Z., Yin, W.: Theoretical linear convergence of unfolded ISTA and its practical weights and thresholds. In: *Advances in Neural Information Processing Systems*, pp. 9061–9071 (2018)
- Cheng, Z., Gadelha, M., Maji, S., Sheldon, D.: A Bayesian perspective on the deep image prior. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
- Combettes, P., Wajs, V.: Signal recovery by proximal forward-backward splitting. *Multiscale Model. Simul.* **4**(4), 1168–1200 (2005). <https://doi.org/10.1137/050626090>
- Daubechies, I., Defrise, M., De Mol, C.: An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Commun. Pure Appl. Math.* **57**(11), 1413–1457 (2004). <https://doi.org/10.1002/cpa.20042>
- Engl, H.W., Hanke, M., Neubauer, A.: *Regularization of Inverse Problems, Mathematics and Its Applications*, vol. 375. Kluwer Academic Publishers Group, Dordrecht (1996)
- Forster, D., Sheikh, A.S., Lücke, J.: Neural simpletrons: learning in the limit of few labels with directed generative networks. *Neural Comput.* **8**(30), 2113–2174 (2018)
- Giryas, R., Eldar, Y.C., Bronstein, A.M., Sapiro, G.: Tradeoffs between convergence speed and reconstruction accuracy in inverse problems. *IEEE Trans. Signal Process.* **66**(7), 1676–1690 (2018)
- Gregor, K., LeCun, Y.: Learning fast approximations of sparse coding. In: *ICML 2010—Proceedings, 27th International Conference on Machine Learning*, pp. 399–406 (2010)
- Hauptmann, A., Lucka, F., Betcke, M., Huynh, N., Adler, J., Cox, B., Beard, P., Ourselin, S., Arridge, S.: Model-based learning for accelerated, limited-view 3-d photoacoustic tomography. *IEEE Trans. Med. Imaging* **37**(6), 1382–1393 (2018)
- Jin, K.H., McCann, M.T., Froustey, E., Unser, M.: Deep convolutional neural network for inverse problems in imaging. *IEEE Trans. Image Process.* **26**(9), 4509–4522 (2017)
- Lempitsky, V., Vedaldi, A., Ulyanov, D.: Deep image prior. In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9446–9454 (2018). <https://doi.org/10.1109/CVPR.2018.00984>
- Liu, J., Chen, X., Wang, Z., Yin, W.: ALISTA: analytic weights are as good as learned weights in LISTA. In: *International Conference on Learning Representations* (2019). <https://openreview.net/forum?id=B1lnzn0ctQ>
- Louis, A.K.: *Inverse und Schlecht Gestellte Probleme*. Vieweg+Teubner Verlag, Wiesbaden (1989)
- Maass, P.: *Deep Learning for Trivial Inverse Problems*, pp. 195–209. Springer, Cham (2019). https://doi.org/10.1007/978-3-319-73074-5_6
- Mataev, G., Elad, M., Milanfar, P.: Deepred: deep image prior powered by red (2019). arXiv preprint [arXiv:1903.10176](https://arxiv.org/abs/1903.10176)
- Meinhardt, T., Möller, M., Hazirbas, C., Cremers, D.: Learning proximal operators: using denoising networks for regularizing inverse imaging problems. In: *IEEE International Conference on Computer Vision*, pp. 1781–1790 (2017)
- Moreau, T., Bruna, J.: Understanding trainable sparse coding via matrix factorization (2016). arXiv preprint [arXiv:1609.00285](https://arxiv.org/abs/1609.00285)
- Nesterov, Y.: *Lectures on Convex Optimization*. Springer Optimization and Its Applications. Springer (2019). <https://books.google.de/books?id=JSyNtQEACAAJ>
- Papayan, V., Romano, Y., Sulam, J., Elad, M.: Theoretical foundations of deep learning via sparse representations: a multilayer sparse model and its connection to convolutional neural networks. *IEEE Signal Process. Mag.* **35**(4), 72–89 (2018)
- Rieder, A.: *Keine Probleme mit inversen Problemen: eine Einführung in ihre stabile Lösung*. Vieweg, Wiesbaden (2003)
- Saxe, A.M., McClelland, J.L., Ganguli, S.: Exact solutions to the nonlinear dynamics of learning in deep linear neural networks (2013). arXiv preprint [arXiv:1312.6120](https://arxiv.org/abs/1312.6120)
- Sprechmann, P., Bronstein, A.M., Sapiro, G.: Learning efficient sparse and low rank models. *IEEE Trans. Pattern Anal. Mach. Intell.* **37**(9), 1821–1833 (2015)
- Sulam, J., Aberdam, A., Beck, A., Elad, M.: On multi-layer basis pursuit, efficient algorithms and convolutional neural networks. *IEEE Trans. Pattern Anal. Mach. Intell.* (2019). arXiv preprint [arXiv:1806.00701](https://arxiv.org/abs/1806.00701)
- Van Veen, D., Jalal, A., Price, E., Vishwanath, S., Dimakis, A.G.: Compressed sensing with deep image prior and learned regularization (2018). arXiv preprint [arXiv:1806.06438](https://arxiv.org/abs/1806.06438)
- Vonesch, C., Unser, M.: A fast iterative thresholding algorithm for wavelet-regularized deconvolution - art. no. 67010d. In: *Wavelets Xii, Pts 1 And 2*, vol. 6701, pp. D7010–D7010 (2007)

34. Xin, B., Wang, Y., Gao, W., Wipf, D., Wang, B.: Maximal sparsity with deep networks? In: *Advances in Neural Information Processing Systems*, pp. 4340–4348 (2016)
35. Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning requires rethinking generalization (2016). arXiv preprint [arXiv:1611.03530](https://arxiv.org/abs/1611.03530)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Sören Dittmer is a Ph.D. student and member of the Research Training Group π^3 at the Center for Industrial Mathematics (ZeTeM) at the University of Bremen, Germany. His current research interests include deep learning, inverse problems, harmonic analysis and signal processing.



Tobias Kluth is a mathematician by training and he did his Ph.D. in computer science related to neuroscience. He is currently a postdoc at the Center for Industrial Mathematics (ZeTeM) at the University of Bremen. His major research interests include the theoretical fields of nonlinear inverse problems, partial differential equations, (nonlinear) signal processing and applications in magnetic particle imaging, electrical impedance tomography, image processing/encoding and computer/human

vision (neural behavior).



Peter Maass is a Professor for Applied Mathematics and the Director of the Center for Industrial Mathematics (ZeTeM) at University of Bremen, Germany, since 2009. He held positions as Assistant Professor at Tufts University, Medford, MA, USA and Saarland University, Saarbrücken, Germany, before he was appointed as a Full Professor of Numerical Analysis at University of Potsdam, Germany, in 1993. Peter Maass is an Adjunct Professor at Clemson University, SC, USA, since 2010.

He holds several patents in the field of image processing. His current research interests include deep learning, inverse problems and wavelet analysis with an emphasis on applications in medical imaging. Peter Maass was awarded an honorary doctorate by the University of Saarland, Germany, in 2018.



Daniel Otero Baguer is a Ph.D. student at the Center for Industrial Mathematics (ZeTeM) at the University of Bremen. He got a Diploma in Computer Science at the University of Havana in 2015. His research interests include deep learning, inverse problems and applications in medical imaging.