

PAPER

Solving ill-posed inverse problems using iterative deep neural networks

To cite this article: Jonas Adler and Ozan Öktem 2017 *Inverse Problems* **33** 124007

View the [article online](#) for updates and enhancements.

You may also like

- [Public key image encryption algorithm based on pixel information and random number insertion](#)
Huishan Wu, Hongyang Zhu and Guodong Ye
- [Equivariant neural networks for inverse problems](#)
Elena Celledoni, Matthias J Ehrhardt, Christian Etmann et al.
- [Higher-order total variation approaches and generalisations](#)
Kristian Bredies and Martin Holler



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Solving ill-posed inverse problems using iterative deep neural networks

Jonas Adler^{1,2}  and Ozan Öktem¹ 

¹ Department of Mathematics, KTH—Royal Institute of Technology,
100 44 Stockholm, Sweden

² Elekta AB, Box 7593, SE-103 93 Stockholm, Sweden

E-mail: ozan@kth.se

Received 22 May 2017, revised 22 September 2017

Accepted for publication 24 October 2017

Published 22 November 2017



CrossMark

Abstract

We propose a partially learned approach for the solution of ill-posed inverse problems with not necessarily linear forward operators. The method builds on ideas from classical regularisation theory and recent advances in deep learning to perform learning while making use of prior information about the inverse problem encoded in the forward operator, noise model and a regularising functional. The method results in a gradient-like iterative scheme, where the ‘gradient’ component is learned using a convolutional network that includes the gradients of the data discrepancy and regulariser as input in each iteration.

We present results of such a partially learned gradient scheme on a non-linear tomographic inversion problem with simulated data from both the Sheep-Logan phantom as well as a head CT. The outcome is compared against filtered backprojection and total variation reconstruction and the proposed method provides a 5.4 dB PSNR improvement over the total variation reconstruction while being significantly faster, giving reconstructions of 512×512 pixel images in about 0.4 s using a single graphics processing unit (GPU).

Keywords: tomography, deep learning, gradient descent, regularization

(Some figures may appear in colour only in the online journal)

1. Introduction

Inverse problems refer to problems where one seeks to reconstruct parameters characterising the system under investigation from indirect observations. Such problems arise in several areas of science and engineering. Mathematically, an inverse problem can be formulated as reconstructing (estimating) a signal $f_{\text{true}} \in X$ from data $g \in Y$ where

$$g = \mathcal{T}(f_{\text{true}}) + \delta g. \quad (1)$$

In the above, X and Y are topological vector spaces, $\mathcal{T}: X \rightarrow Y$ (forward operator) models how a given signal gives rise to data in absence of noise, and $\delta g \in Y$ is a single sample of a Y -valued random variable that represents the noise component of data.

Many inverse problems, such as those arising in imaging, are naturally formulated when both signal and data are functions. In such case, X is some Banach/Hilbert space of functions defined on a fixed image domain $\Omega \subset \mathbb{R}^d$ and Y is likewise a Banach/Hilbert space of functions defined on a fixed data manifold \mathbb{M} . An important remark here relates to the nature of the data manifold. It can be a subset of Euclidean space, but this is not necessarily the case. In fact, x-ray tomographic imaging leads to inverse problems where elements in the data manifold \mathbb{M} represent lines in \mathbb{R}^d .

1.1. Classical regularisation

A common approach in solving an inverse problem of the form in (1) is to minimise the miss-fit against data. For example by minimising

$$f \rightarrow \mathcal{L}(\mathcal{T}(f), g) \quad (2)$$

where $\mathcal{L}: Y \times Y \rightarrow \mathbb{R}$ is a suitable affine transformation of the data log-likelihood [6]. Then, one may interpret minimising the above as finding a maximum likelihood solution to (1).

This minimisation is a large scale optimisation problem that for typical choices of \mathcal{T} is ill-posed, that is, a solution (if it exists) is unstable with respect to the data g in the sense that small changes to data results in large changes to a reconstruction. Hence, finding a maximum likelihood solution (there may be several) typically leads to over-fitting against data.

Within classical regularisation theory, there are currently three strategies for avoiding over-fitting when solving (1). One is approximate inverse that is applicable to cases when X has a mollifier. The idea is to construct a pseudo-inverse to \mathcal{T} using the mollifier [36]. Another approach is iterative regularisation, which starts out by considering a fixed point iteration scheme for minimising (2). Over-fitting is avoided by stopping the iterates early, which is a feasible strategy if the iterates are semi-convergent [6, 16, 21, 23]. Finally, we have variational regularisation where over-fitting is avoided by introducing a functional $\mathcal{S}: X \rightarrow \mathbb{R}$ (regularisation functional) that encodes *a priori* information about f_{true} and penalises unfeasible solutions [16, 35]. Hence, instead of minimising only the data discrepancy functional, one now seeks to minimise the regularised objective functional by solving

$$\min_{f \in X} [\mathcal{L}(\mathcal{T}(f), g) + \lambda \mathcal{S}(f)] \quad \text{for a fixed } \lambda \geq 0. \quad (3)$$

In the above, λ (regularisation parameter) governs the influence of the *a priori* knowledge encoded by the regularisation functional against the need to fit data.

Formulation in (3) suggests a ‘plug-and-play’ structure where the forward operator and data discrepancy functional can be adapted to the specific application, and the regularisation functional is chosen to capture *a priori* information. This flexibility makes variational methods a powerful framework for image reconstruction, and especially so in cases when the main concern is reconstruction ‘quality’. A typical example in imaging is total variation (TV) regularisation, which is suitable for signals represented by scalar functions of bounded variation with sparse gradient. The corresponding regularisation functional is then given as $\mathcal{S}(f) := \|\nabla f\|_1$.

There are however some drawbacks that come with using variational methods. One is that they are inherently computationally demanding, which is an issue in many applications.

Another is that the prior information needs to be encoded as an explicit functional, which limits the type of *a priori* information that can be accounted for. Finally, how to appropriately choose the regularisation parameter(s) is non-obvious for many applications.

1.2. Machine learning approaches to inverse problems

Machine learning is commonly used for non-linear function approximation under weak assumptions. Applied to the inverse problem in (1), it can be phrased as the problem of finding a (non-linear) mapping $\mathcal{T}_{\Theta}^{\dagger}: Y \rightarrow X$ satisfying the following *pseudo-inverse* property:

$$\mathcal{T}_{\Theta}^{\dagger}(g) \approx f_{\text{true}} \quad \text{whenever data } g \text{ is related to } f_{\text{true}} \text{ as in (1)}.$$

A key element is to parametrise the set of such pseudo-inverse operators by $\Theta \in Z$, where Z is a suitable parameter space. The ‘learning’ part refers to choosing an ‘optimal’ Θ given *training data*, where the concept of optimality is typically quantified through a loss functional that measures the ‘quality’ of a learned pseudo-inverse $\mathcal{T}_{\Theta}^{\dagger}$.

The manner in which the loss functional is specified depends on the type of training data, and here we separate between supervised and unsupervised learning. These two approaches are *fundamentally* different and this article focuses on the supervised learning case since it is the problem with the most structure and we expect learning to give larger improvements over traditional methods.

1.2.1. Supervised learning. In supervised learning, training data are independent identically distributed realisations of a $(Y \times X)$ -valued random variable (g, f) with a known probability distribution μ . Estimating $\Theta \in Z$ from training data can be formulated as minimising a loss functional $\Theta \mapsto L(\Theta)$ that frequently has the following structure:

$$L(\Theta) := F_{\mu} \left[d(\mathcal{T}_{\Theta}^{\dagger}(g), f) \right]. \quad (4)$$

In the above, $\mathcal{T}_{\Theta}^{\dagger}: Y \rightarrow X$ is the pseudo-inverse that is given by $\Theta \in Z$, $d: X \times X \rightarrow \mathbb{R}$ is a ‘distance’ function quantifying the quality of a specific reconstruction, and F_{μ} maps real-valued random variables on $Y \times X$ to real numbers.

A common choice is to use the expected loss w.r.t. the squared distance:

$$L(\Theta) := \mathbb{E}_{\mu} \left[\|\mathcal{T}_{\Theta}^{\dagger}(g) - f\|_X^2 \right]. \quad (5)$$

One may also consider other loss functionals and the method we suggest can easily be adapted to these. As an example, a very conservative reconstruction method would use a loss functional given by the supremum of the ∞ -norm:

$$L(\Theta) := \text{ess sup}_{\mu} \left[\|\mathcal{T}_{\Theta}^{\dagger}(g) - f\|_{\infty} \right].$$

In practical applications, one needs to choose a distribution μ whose samples are representative for the application. As an example, if the application in question is computed tomography (CT) imaging of human heads, then samples from μ should be interpretable as human heads with associated noisy CT data.

Choosing μ such that the aforementioned requirement holds is naturally realised in three different ways. The first is to analytically specify μ , this is, e.g. the case for the randomly generated ellipses in section 3. The second is to choose μ as the empirical distribution derived from available measurements. Finally, these two approaches can be mixed, e.g. by expressing

μ as an analytically known conditional density of \mathbf{g} given \mathbf{f} multiplied with the empirical density of \mathbf{f} . This is the case for the head CT reconstructions in section 3 where we use a finite number of heads and simulate CT data according to a physical noise model.

1.2.2. Unsupervised learning. In unsupervised learning there is no access to input-output pairs as in the supervised learning setting. Instead, here training data are given as elements in Y . The corresponding mathematical setting is to consider these as independent identically distributed realisations of a Y -valued random variable \mathbf{g} with a known probability density μ . A natural choice for a loss function is now to quantify how well the learned reconstruction matches the regularised data discrepancy, i.e.

$$L(\Theta) := \mathbb{E}_\mu \left[\mathcal{L}(\mathcal{T}(\mathcal{T}_\Theta^\dagger(\mathbf{g})), \mathbf{g}) + \mathcal{S}(\mathcal{T}_\Theta^\dagger(\mathbf{g})) \right].$$

This can be interpreted as learning an optimiser for the variational regularisation in (3).

1.3. Survey of the field

Data driven approaches, and in particular deep learning using convolutional neural networks, have shown dramatic improvements over the state-of-the-art in several applications [26]. Likewise, knowledge-driven approaches outlined in section 1.1, and in particular variational regularisation, have provided a flexible toolbox for solving inverse problems.

As already mentioned, usage of variational regularisation is however associated with several challenges: computational feasibility, flexibility in the prior information that can be accounted for, and choice of regularisation parameter(s). A natural idea is thus to combine elements of deep learning and variational regularisation in order to address these challenges. This has been attempted by several authors as we now outline.

1.3.1. Fully learned reconstruction. Approaching the inverse problem directly with machine learning amounts to learning $\mathcal{T}_\Theta^\dagger: Y \rightarrow X$ from data such that it approximates an inverse of \mathcal{T} in (1). An example of such an approach for solving very small scale tomographic reconstruction problems is given in [4, 31].

An obvious disadvantage with such fully learned approaches is that the result is likely to depend on the data manifold, so training data needs to be rich enough to account for all various data manifolds that one is likely to encounter. Furthermore, training data also needs to be rich enough to allow the learning scheme to learn the structure in \mathcal{T} , which is given by the physics laws governing the formation of data from a signal. Finally, in many applications the adequate digitalisations of the signal and data often requires very high dimensional arrays.

The above considerations imply that the parameter space Z used for parametrising possible inverse operators has to be very high dimensional in a fully learned approach. Therefore, the idea of learning $\mathcal{T}_\Theta^\dagger$ from data without using any knowledge of the physics or the data manifold quickly becomes in-feasible due to lack of training data.

1.3.2. Sequential data and knowledge driven reconstruction. The idea here is to separate the learned components from a part that encodes some knowledge about the structure of \mathcal{T} and the data manifold. Formalising this, we assume

$$\mathcal{T}_\Theta^\dagger = \mathcal{B}_\Theta \circ \mathcal{A} \circ \mathcal{C}_\Theta \tag{6}$$

where $\mathcal{A}: Y \rightarrow X$ is a known component that encodes knowledge about the structure of the forward operator \mathcal{T} whereas the operators $\mathcal{B}_\Theta: X \rightarrow X$ and $\mathcal{C}_\Theta: Y \rightarrow Y$ are the learned components.

An important special case is when \mathcal{C}_Θ is not learnt, which entirely separates the computation of \mathcal{A} from the learning of \mathcal{B}_Θ , i.e. the original inverse problem in (1) is recast as learning $\mathcal{B}_\Theta: X \rightarrow X$. This significantly simplifies the implementation since the often demanding task of computing $g \mapsto \mathcal{A}(g)$ can be separated from the learning software. It also ensures the data manifold is not explicitly part of the learning. One may furthermore exploit additional structure, like locality and/or invariance of the operator $\mathcal{A} \circ \mathcal{T}$. A key step for such a sequential data and knowledge driven reconstruction scheme is to have candidates for \mathcal{A} and one natural option is to let it be some pseudo-inverse. As an example, in tomographic applications it can be given by the back-projection or the filtered back-projection operator. Next, when learning \mathcal{B}_Θ from data, one may use approaches that build on the corpus of knowledge that exists for denoising signals in X . An example demonstrating this approach for tomographic reconstruction is given in [20, 22].

On the other hand, for ill-posed inverse problems some information is irreversibly lost when making the assumption in (6) since the learned operators \mathcal{C}_Θ and \mathcal{B}_Θ cannot recover information that is lost by using a pseudo-inverse \mathcal{A} . To alleviate this problem, for linear forward operators we can consider choosing \mathcal{A} as the adjoint of the forward operator, i.e. $\mathcal{A} := \mathcal{T}^*: Y \rightarrow X$. This can be seen as learning to solve the normal equations for (1) since

$$f_{\text{true}} \approx \mathcal{B}_\Theta(\mathcal{T}^*(g)) \implies f_{\text{true}} \approx \mathcal{B}_\Theta(\mathcal{T}^*(\mathcal{T}(f_{\text{true}}))) \implies \mathcal{B}_\Theta \approx (\mathcal{T}^* \circ \mathcal{T})^{-1}.$$

Nonetheless, solving the normal equations for ill-posed problems is often more ill-posed than the original inverse problem, so such a learning procedure would need to include some kind of regularisation.

We conclude by pointing to examples where the operator $\mathcal{C}_\Theta: Y \rightarrow Y$ is learned. One such case is [41] where tomographic reconstruction is performed by learning \mathcal{C}_Θ . Another similar, but more advanced, example is [32] where the operator \mathcal{A} is given by several filtered back-projection (FBP) operators and the learned operator \mathcal{C}_Θ is given by the filter coefficients.

While such approaches can in principle address all of the aforementioned issues with variation methods, they are ultimately limited by what $\mathcal{A}: Y \rightarrow X$ manages to capture about the ‘inverse’ of $\mathcal{T}: X \rightarrow Y$. In conclusion the final reconstruction cannot contain information that is not already present in $\mathcal{A}(g)$.

1.3.3. Learning for variational reconstruction. In these methods the aim is to *solve a variational problem* by using techniques from machine learning.

One option is to use the latter in order to select the regularisation parameter(s). This procedure can be re-formulated as a bi-level optimisation scheme whose mathematical properties (like existence) can be analysed in a functional analytic setting [8, 9, 12–14, 25]. On the other hand any implementation of an iterative scheme for solving the aforementioned parametrised optimisation problem will terminate after a finite number of iterates. Hence, the outcome will not only depend on the formulation of the parametrised optimisation problem, but also on the solution scheme one chooses to use. In conclusion, a bi-level optimisation scheme of the above type does not by itself uniquely determine a reconstruction operator.

Another approach replaces the learning of the optimisation *problem* with learning an optimiser adapted to a given class of optimisation problems. In [3] a stochastic gradient descent method is learned from training data consisting of optimisation problems, each associated with a deep learning problem. The output is thus a trained stochastic gradient descent

method that can be used to train other deep neural networks. This overall ‘learning to learn’ approach can in principle be extended to other use cases, such as solving inverse problems, by merely changing the underlying class of training data. While this approach may to some extent address the performance issues associated with variational methods, it is in the same way limited in the types of priors that can be represented.

1.3.4. Learned iterative reconstruction. The idea here is to refrain from formulating an optimisation problem without compromising upon the ability to account for knowledge about the inverse problem, such as forward operator, data discrepancy, regulariser, etc.

An example of such a scheme for solving inverse problems is [42], which learns an ADMM-like scheme for Fourier inversion. Another is [11], which learns a ‘proximal’ in an ADMM-like scheme for various image processing problems. Finally, [33] considers solving finite dimensional linear inverse problems typically arising in image restoration. The idea is to learn over a broader class of schemes instead of restrict attention to a specific type of scheme, like ADMM above.

Similar to the sequential data and knowledge driven reconstruction methods, the learned iterative reconstruction schemes can address issues related computational feasibility, flexibility in the prior information that can be accounted for, and choice of regularisation parameter(s). On the other hand, they utilise the same *a priori* information as in variational regularisation.

1.4. Contribution and overview of paper

This paper develops a framework for learned iterative reconstruction, generalising the ideas in [33] along several directions. First, we consider solving inverse problems with (possibly) non-linear forward operators. Next, our framework is formulated in a coordinate free functional analytic setting. Furthermore, we consider the issue of proper initialisation and allow for further prior knowledge to be included in the form of a regulariser. Finally, we provide a generic and scalable open source implementation³ of our method based on operator discretization library (ODL) [2] that can be applied to a wide range of realistic inverse problems. This also includes the trained parameter Θ used for generating the results shown in this paper. To show that the approach can handle (non-linear) forward operators in large scale inverse problems, we consider tomographic reconstruction with a non-linear ray transform inversion.

Section 2 derives a partially learned gradient decent scheme for solving (1) in the functional analytic setting. This section also introduces the deep convolutional network that is used later for tomographic reconstruction. Section 3.1 describes the implementation of the partially learned gradient decent scheme in section 2 and software components used for computations. Section 3 tests the performance of the partially learned gradient decent scheme on tomographic inverse problems. The paper concludes with a discussion in section 4 and a summary of future work and conclusions is given in section 5.

2. Solving inverse problems by learned gradient descent

We begin by proving a heuristic motivation that comes from comparing two natural considerations involving gradient descent schemes associated with solving (1). This results in an initial scheme for partially learned gradient descent given in algorithm 1, which is then extended by adding persistent memory and resulting in the scheme in algorithm 2. Next is a description of

³ https://github.com/adler-j/learned_gradient_tomography

how to integrate deep learning with algorithm 2, resulting in the final scheme given in algorithm 3.

2.1. Motivation

The starting point in learning an iterative scheme that combines elements from (deep) machine learning and classical regularisation theory is to consider the *error functional* $E: X \rightarrow \mathbb{R}$ defined as

$$E(f) := \mathbf{d}(f, f_{\text{true}})$$

where $\mathbf{d}: X \times X \rightarrow \mathbb{R}$ is the distance functional that appears in the definition of the loss functional (4) used for training. It measures how good well f approximates f_{true} , and one natural error functional corresponding to (5) is

$$E(f) = \|f - f_{\text{true}}\|_X^2.$$

Ideally, solving (1) would be based on minimising the error functional, which for obvious reasons is not possible. Hence, we need to use a substitute. In variational regularisation theory, such a substitute is given by the regularised objective functional in (3). Much of regularisation theory aims at choosing the objective functional in (3) so that the regularised solution approximates the true signal to be recovered:

$$\arg \min_{f \in X} E(f) \approx \arg \min_{f \in X} [\mathcal{L}(\mathcal{T}(f), g) + \lambda \mathcal{S}(f)]. \quad (7)$$

Assume next that the objective functional in the right hand side of (7) is Fréchet differentiable and (strictly) convex. Then, a simple gradient descent scheme could be used to find a minimum:

$$f_i := f_{i-1} - \sigma \left(\nabla [\mathcal{L}(\mathcal{T}(\cdot), g)](f_{i-1}) + \lambda [\nabla \mathcal{S}](f_{i-1}) \right) \quad (8)$$

where, assuming a differentiable likelihood and forward operator, we note that

$$\nabla [\mathcal{L}(\mathcal{T}(\cdot), g)](f) = [\partial \mathcal{T}](f)^* \left([\nabla \mathcal{L}(\cdot, g)](\mathcal{T}(f)) \right) \quad \text{for any } f \in X.$$

Likewise, considering the left hand side in the same way, a differentiable convex error functional would allow one to use a corresponding scheme for finding a minimum:

$$f_{j+1} := f_j - \sigma [\nabla E](f_j). \quad (9)$$

Since the gradient mapping $\nabla E: X \rightarrow X$ in (9) requires knowledge about the true signal, it is natural to try to learn it from training data while utilising knowledge about the gradient mappings $\nabla [\mathcal{L}(\mathcal{T}(\cdot), g)]$, $\nabla \mathcal{S}: X \rightarrow X$. For this purpose, we introduce the (*learned*) *updating operator* $\Lambda_{\Theta}: X \times X \times X \rightarrow X$ that, given an appropriately selected (learned) parameter $\Theta \in Z$, should satisfy

$$\Lambda_{\Theta} \left(f, \nabla [\mathcal{L}(\mathcal{T}(\cdot), g)](f), \lambda \nabla \mathcal{S}(f) \right) \approx \nabla E(f).$$

These considerations suggests a partially learned gradient descent scheme specified as in algorithm 1.

Algorithm 1. Initial partially learned gradient descent.

-
- 1: Select an initial guess f_0
 - 2: **for** $i = 1, \dots, I$ **do**
 - 3: $\Delta f_i \leftarrow -\sigma \Lambda_{\Theta}(f_{i-1}, \nabla[\mathcal{L}(\mathcal{T}(\cdot), g)](f_{i-1}), \lambda \nabla S(f_{i-1}))$
 - 4: $f_i \leftarrow f_{i-1} + \Delta f_i$
 - 5: $\mathcal{T}_{\Theta}^{\dagger}(g) \leftarrow f_I$
-

Algorithm 1 suffers from several unnecessary shortcomings that are easily addressed. The regularisation parameter λ and the step length σ have to be explicitly chosen, a task that is known to be troublesome in practical applications (section 4.2). One may instead make these part of Θ and thereby learn them from training data. Next, the convergence rate of gradient descent schemes can be accelerated by using information from previous iterates (memory) as in quasi-Newton schemes [27]. For this purpose we introduce persistent memory $s \in X^M$ that allows algorithm 1 to use information from earlier iterates. The learned updating operator now becomes a mapping

$$\Lambda_{\Theta}: X^M \times X \times X \times X \rightarrow X^M \times X. \quad (10)$$

Finally, one often also has the possibility to select the initial iterate f_0 using some suitable pseudo-inverse $\mathcal{T}^{\dagger}: Y \rightarrow X$. Considering these modifications results in the partially learned gradient descent scheme listed in algorithm 2.

Algorithm 2. Partially learned gradient descent.

-
- 1: $f_0 \leftarrow \mathcal{T}^{\dagger}(g)$.
 - 2: Initialize ‘memory’ $s_0 \in X^M$.
 - 3: **for** $i = 1, \dots, I$ **do**
 - 4: $(s_i, \Delta f_i) \leftarrow \Lambda_{\Theta}(s_{i-1}, f_{i-1}, \nabla[\mathcal{L}(\mathcal{T}(\cdot), g)](f_{i-1}), \nabla S(f_{i-1}))$
 - 5: $f_i \leftarrow f_{i-1} + \Delta f_i$
 - 6: $\mathcal{T}_{\Theta}^{\dagger}(g) \leftarrow f_I$
-

2.2. Parametrising the learned updating operators

The goal here is to specify the class of learned updating operators that are parametrised by $\Theta \in Z$. Following the paradigm in (deep) neural networks, we start by defining a family of affine operators

$$\mathcal{W}_{w_n, b_n}: X^{c_{n-1}} \rightarrow X^{c_n} \quad \text{for } n = 0, \dots, N, \quad (11)$$

parametrised by linear mappings $w_n: X^{c_{n-1}} \rightarrow X^{c_n}$ (weights) and $b_n \in X^{c_n}$ (biases). Here, N is usually referred to as the depth of the neural network that will eventually define the learned updating operator and c_n is the number of channels in the n :th layer. Next, we introduce a family of non-linear operators

$$\mathcal{A}_n: X^{c_n} \rightarrow X^{c_n} \quad (12)$$

that are given by point-wise application of a fixed non-linear scalar function, henceforth called the response function.

By chaining compositions, we now define a parametrised family of learned updating operators as

$$\Lambda_{\Theta} := (\mathcal{A}_N \circ \mathcal{W}_{w_N, b_N}) \circ \dots \circ (\mathcal{A}_1 \circ \mathcal{W}_{w_1, b_1})$$

with $\Theta := ((w_N, b_N), \dots, (w_1, b_1))$. In order to match the domain and range of the operator in (10), we need to assume that $c_0 = M + 3$ and $c_N = M + 1$.

Such parametrised operators are used in machine learning applications for two primary reasons: computability and descriptive power. In order to learn the parameters Θ from training data, a (stochastic) gradient descent method is typically applied in which case the derivative $\partial\Lambda_{\Theta}/\partial\Theta$ and its adjoint $[\partial\Lambda_{\Theta}/\partial\Theta]^*$ needs to be repeatedly computed and here one may use the chain rule. This becomes particularly easy to perform in a computationally feasible manner for learned updating operators of this form. Furthermore, introducing the non-linear component in (12) allows the learned operator to approximate a large set of non-linear operators [19].

2.2.1. Choice of affine and non-linear operator families. Our next step is to further narrow down the generative models for the operator families \mathcal{W}_{w_i, b_i} and \mathcal{A}_i . We start by writing the affine operator in (11) as

$$\mathcal{W}_{w_n, b_n} = (\mathcal{W}_{w_n, b_n}^1, \dots, \mathcal{W}_{w_n, b_n}^{c_n})$$

where the components

$$\mathcal{W}_{w_n, b_n}^l : X^{c_{n-1}} \rightarrow X \quad \text{for } l = 1, \dots, c_n$$

represent the affine transformation for the l :th channel in the n :th layer. In particular, using linearity these can be represented as

$$\mathcal{W}_{w_n, b_n}^l(f_1, \dots, f_{c_{n-1}}) = b_n^l + \sum_{j=1}^{c_{n-1}} w_n^{j,l}(f_j)$$

where $w_n^{j,l} : X \rightarrow X$ is a channel-wise linear operator.

However, optimizing over the set of all linear operators would result in a very large number of parameters. A solution is to assume that $w_n^{j,l}$ is translation invariant, which indicates that we seek to use a convolutional network architecture.

Convolution operators are useful for representing translation invariant image features at a specific scale that is governed by the size of the support of the kernel. Hence, hierarchically organising such operators with small kernels in layers provides a (deep) convolutional neural network (CNN) with vastly reduced number of network parameters that is efficient to implement, and at the same time capable of representing various image features at different scales.

In the case of tomographic reconstruction, the situation is however different since there are many non-local dependencies (all pixels/voxels on a line contribute to the value of the ray-transform). Thus, applying convolution network architectures to directly learn the reconstruction would be problematic. On the other hand, the scheme outlined in algorithm 2 includes the forward operator that accounts for these global dependencies. For this reason, it is feasible to use convolution network architectures to learn the learned updating operator in (10).

By using convolution operators, we find that the affine operators can be written

$$\mathcal{W}_{w_n, b_n}^l(f_1, \dots, f_{c_{n-1}}) = b_n^l + \sum_{j=1}^{c_{n-1}} w_n^{j,l} * f_j$$

where $b_n^l \in \mathbb{R}$ represents the bias and w_n is given as a ‘matrix’ of convolution kernels $w_n^{j,l} \in X$. Hence, our parameter space becomes

$$Z = (X^{c_N \times c_{N-1}} \times \mathbb{R}^{c_N}) \times \dots \times (X^{c_1 \times c_0} \times \mathbb{R}^{c_1}).$$

Finally, the non-linear response functions \mathcal{A}_i in (12) can be chosen in different ways and we will be using the rectified linear unit (ReLU) [29]

$$\text{relu}(x) = \begin{cases} x & \text{if } x > 0 \\ 0 & \text{else.} \end{cases}.$$

2.3. The partially learned gradient descent algorithm

A number of hyper-parameters needs to be chosen prior to learning. These are the number of layers $N \in \mathbb{N}$, the number of channels $c_1, \dots, c_{N-1} \in \mathbb{N}$ in each layer, the number of iterations I , and the size of the memory M .

In the examples shown in in section 3, we let the weights w_i be represented by 3×3 pixel convolutions and we used $N = 3$ layers. The number of convolutions in each layer was selected as $m_1 = 32$ and $m_2 = 32$. We selected the number of iterations to be $I = 10$ and the amount of memory to be $M = 5$. Such low numbers were selected in order to reduce the space of allowed parameters which in turn should help reduce over-fitting. All parameters were selected by simple trial and error and it is likely that a more sophisticated set-up would give better results or be better suited for a particular application. For further details, see the supplemental source code.

Once the hyper-parameters are chosen, one may now fully specify the partially learned gradient descent, which is done in algorithm 3. The scheme learns the updating operator by learning the scalars $b_n^l \in \mathbb{R}$ and the convolution kernels (functions) $w_n^{j,l}$ from training data. The resulting learned updating operator can then be used to solve the inverse problem in (1).

Algorithm 3. Partially learned gradient descent.

```

1:  $s_0 \leftarrow 0$ 
2:  $f_0 \leftarrow \mathcal{T}^\dagger(g)$ 
3: for  $i = 1, \dots, I$  do
4:    $u_i^1 \leftarrow (f_{i-1}, s_{i-1}, \nabla[\mathcal{L}(\mathcal{T}(\cdot), g)](f_{i-1}), \nabla S(f_{i-1}))$ 
5:    $u_i^2 \leftarrow \text{relu}(\mathcal{W}_{w_1, b_1}(u_i^1))$ 
6:    $u_i^3 \leftarrow \text{relu}(\mathcal{W}_{w_2, b_2}(u_i^2))$ 
7:    $(u_i^4, \Delta f_i) \leftarrow \mathcal{W}_{w_3, b_3}(u_i^3)$ 
8:    $s_i \leftarrow \text{relu}(u_i^4)$ 
9:    $f_i \leftarrow f_{i-1} + \Delta f_i$ 
10:  $\mathcal{T}_{\Theta}^\dagger(g) \leftarrow f_i$ 

```

3. Implementation and evaluation

The algorithm was tested on the two-dimensional computed tomography problem. The signal is in this case real valued functions defined on a domain in \mathbb{R}^2 representing images and X is a suitable vector space of such functions. The corresponding forward operator is expressible in terms of the ray transform $\mathcal{P} : X \rightarrow Y$, which integrates the signal over a set of lines \mathbb{M} given by the acquisition geometry. Hence, elements in Y are functions on lines

$$\mathcal{P}(f)(\ell) = \int_{\ell} f(x) dx \quad \text{for } \ell \in \mathbb{M}.$$

As training data we consider CT simulations from two particular types of phantoms with different forward operators and noise models:

Ellipses: Training data is randomly generated ellipses on a 128×128 pixel domain. The projection geometry was selected as a sparse 30 view parallel beam geometry with 5% additive Gaussian noise added to the projections. In this case, the log-likelihood was selected as the squared L^2 norm $\mathcal{L}(\cdot, g) := \frac{1}{2} \|\cdot - g\|_Y^2$ which implies

$$\nabla[\mathcal{L}(\mathcal{P}(\cdot), g)](f) = \mathcal{P}^*(\mathcal{P}(f) - g)$$

The phantoms were generated ‘on the fly’, giving an effectively infinite dataset.

Heads: The training is simulated projections of 512×512 pixel, 256×256 mm slices of CT scans of human heads as provided by Elekta (Elekta AB, Stockholm, Sweden). The acquisition geometry defining the data manifold was selected as a fan beam geometry with source-axis distance of 500 mm, source-detector distance 1000 mm, 1000 pixel, and 1000 angles.

Here, in order to get a accurate noise model we used a non-linear forward operator given by

$$\mathcal{T}(f)(\ell) = \lambda \exp(-\mu \mathcal{P}(f)(\ell))$$

where $\lambda \in \mathbb{R}^+$ is the mean number of photons per pixel, taken to be 10 000, and $\mu \in \mathbb{R}^+$ is the linear attenuation coefficient which was taken to be that of water ($\approx 0.2 \text{ cm}^{-1}$). Poisson noise was added to the projections, and given 10 000 photons per pixel, which corresponds to a low dose scan. For this type of noise, the log-likelihood is given by the Kullback–Leibler divergence and the data discrepancy becomes

$$\mathcal{L}(\mathcal{T}(f), g) := \int_{\mathbb{M}} \left(\mathcal{T}(f)(\ell) + g(\ell) \log \left(\frac{g(\ell)}{\mathcal{T}(f)(\ell)} \right) \right) d\ell$$

which implies that

$$\nabla[\mathcal{L}(\mathcal{T}(\cdot), g)](f) = [\partial \mathcal{T}(f)]^* \left(1.0 - \frac{g}{\mathcal{T}(f)} \right).$$

In the above, the adjoint of the derivative of the forward operator applied in a perturbation $\delta g \in Y$ is given by

$$[\partial \mathcal{T}(f)]^*(\delta g) = \mathcal{P}^* \left(-\mu \underbrace{\lambda \exp(-\mu \mathcal{P}(f)(\cdot))}_{\mathcal{T}(f)(\cdot)} \delta g(\cdot) \right),$$

which after some simplifications gives the following expression for the gradient:

$$\nabla[\mathcal{L}(\mathcal{T}(\cdot), g)](f) = -\mu \mathcal{P}^*(\mathcal{T}(f) - g) \quad \text{for } f \in X.$$

The training used 500 CT scans with a total of 41 000 slices.

For both cases the regulariser was selected as the Dirichlet energy, e.g.

$$\mathcal{S}(f) := \frac{1}{2} \|\nabla f\|_2^2 \implies \nabla \mathcal{S}(f) = \nabla^*(\nabla f)$$

which acts as a smoothing, there by reducing noise [7]. See figure 1 for examples of the data used for training and validation.

3.1. Implementation

The methods described above were implemented in Python using ODL [2] and Tensorflow [1]. All operator-related components, such as the forward operator \mathcal{T} , were implemented in ODL, and these were then converted into Tensorflow layers using the `as_tensorflow_layer` functionality of ODL. The neural network layers and training were implemented using Tensorflow.

The implementation utilises abstract ODL structures for representing functional analytic notions and is therefore generic, yet easily adaptable to other inverse problems. We used the ODL operator `RayTransform` in order to evaluate the ray transform and its adjoint using the GPU accelerated `'astra_gpu'` backend [40]. The pseudo-inverse \mathcal{T}^\dagger was given by the filtered back-projection algorithm implemented in ODL as `fbp_op` with no additional smoothing filter.

We emphasise that the functional analytic formulation of algorithm 3 is critical to handle problems of this scale. As an example, storing the ray transform used for the heads dataset as a sparse matrix of floating point numbers would require about 1 GB of GPU memory.

3.1.1. Training. We trained the parameters Θ using the `RMSPropOptimizer` optimiser in Tensorflow. We initially used 10^5 batches on the ellipses problem, where each batch contained 20 tomography problems with a learning rate starting at 10^{-3} and decayed according to the inverse of the iteration number down to about 10^{-5} . This training took four days on a workstation with a single Nvidia GTX Titan GPU. These parameters were then used as an initial guess for the heads problem, once again trained according to the same scheme but with the learning rate starting at 10^{-5} and decreased to about 10^{-7} and with each batch containing only one tomographic problems due to memory limitations of the current implementation. This training took four days on the aforementioned hardware.

3.1.2. Comparison. We compare the performance of the partially learned iterative algorithm against the FBP algorithm and TV regularisation. The FBP reconstructions were performed using a Hann filter with bandwidth selected to maximise peak signal to noise ratio (PSNR). The data discrepancy in the TV regularisation matched the one used in the partially learned algorithm and the regularisation parameter was selected to maximise PSNR.

We solve the TV regularised problem without smoothing using the the generic ODL implementation of the non-linear primal dual hybrid gradient (PDHG) optimisation method [39]. This is needed since the forward operator in the heads dataset is non-linear. We used 1000 iterations at which point the objective function was stationary. For the ellipses dataset, the evaluation was performed on the modified Shepp–Logan phantom, while on the heads dataset a slice through the nasal region was used.

3.2. Results

We compare the reconstructions of the partially learned algorithm with the FBP and TV reconstructions for both the ellipse and head datasets and computed the PSNR, runtime and performed a visual comparison. The quantitative results are given in table 1, which visualisations are available in figures 2 and 3. We also display some partial results of the iterative algorithm for reference in figure 4.




Figure 1. Examples from training data. Top row shows the phantoms, middle row the simulated tomographic data, and the bottom row is the initial guess obtained using FBP. The left column is the random ellipses, the middle column is the Shepp–Logan phantom, and the right column is slice of head phantom shown with window $[-200, 200]$ HU.

We note that for the ellipse data, the FBP algorithm performs very poorly under the high noise while the TV and learned methods give comparable results. This is expected given that TV regularisation is very competitive for images of this type, nonetheless the learned method does outperform the TV method by approximately 2 dB and the visual result looks slightly more appealing, with less randomly occurring structures and significantly less stair-casing.

For the head dataset where the noise is lower, the filtered back-projection reconstruction performs much better, and is arguably comparable to the TV regularised reconstruction. The learned reconstruction provides (perhaps too) smooth images, where we note that especially in the boundary regions, e.g. in the air-skin boundaries and around the bone the algorithm performs amiably.

In addition to the visual results, we see that the learned algorithm is significantly better than the TV reconstruction w.r.t the PSNR, giving an >5 dB improvement. Finally, the runtime of

Table 1. Comparison of the learned method with standard methods.

Method	PSNR (dB)		Runtime (ms)	
	Ellipses	Heads	Ellipses	Heads
FBP	19.75	36.12	4	130
Learned	32.02	43.82	58	430
TV	29.83	38.40	11 963	173 845




Figure 2. Reconstructing Shepp–Logan phantom using FBP, TV and the partially learned gradient scheme. Data is simulated from the Shepp–Logan phantom, which attains values between $[0, 1]$. All images are shown using a window set to $[0.1, 0.4]$ for improved contrast.

the algorithm, while being slightly slower than traditional filtered-back-projection, is significantly faster than the TV method as shown in table 1.

3.3. Impact of including gradient mappings

The impact of including the gradient mappings

$$\nabla[\mathcal{L}(\mathcal{T}(\cdot), g)], \nabla S: X \rightarrow X$$




Figure 3. Reconstructing a head phantom using FBP, TV and the partially learned gradient scheme. Data is simulated from a physiological head phantom, which includes some weak streaks (so these are part of the ground truth). All images are shown using a window set to $[-200, 200]$ HU.

in the partially learned gradient scheme can be empirically analysed. We do this by training the network in the exact same manner with and without the gradients added and then performing 100 reconstructions of the Shepp–Logan phantom with the respective methods.

Without the gradients, the PSNR was 29.65 dB while it was 30.51 dB with the gradient of the data discrepancy. This can be compared to 32.02 dB with both gradients. Visual inspection also indicates that adding the gradients provides a sharper reconstruction with more detail, where especially in the case of no gradients the small inserts are barely visible. Performance wise, the method took 19 ms without the gradients, 64 ms with the gradient of the data discrepancy and 66 ms with both gradients. See figure 5 for a visual comparison.




Figure 4. Iterates of the partially learned gradient scheme when applied to reconstruct the Shepp–Logan phantom. The initial iterate is given by the FBP method.

4. Discussion

The partially learned gradient scheme differs significantly from the current paradigm for regularisation of inverse problems, so there are several remarks that deserve a closer discussion.

4.1. Theory

The partially learned gradient scheme is presented with a strong emphasis on the algorithmic aspects, its implementation and its performance. There are however several interesting theoretical issues that deserve closer attention.




Figure 5. Comparison of reconstructions using the partially learned gradient scheme with and without the gradient information. Note that gradient of data discrepancy includes the derivative of the forward operator.

4.1.1. Deep neural networks in function spaces. The scheme in algorithm 2 is formulated in a functional analytic setting. However, the theory for deep neural networks is not well established in the infinite dimensional setting and there are several open issues that remain to be answered, such as determining what class of operators can be approximated by a given deep neural network and to what accuracy [19].

Another aspect relates to usage of probabilistic notions in infinite dimensional vector spaces. The classical theory for deep learning deals with finite, or at most countable data where the law of large numbers holds. In the infinite dimensional setting one needs to be more careful regarding which topologies that are used. It is clearly advantageous to work with spaces where one can prove various forms of weak convergence of probability measures and state and prove results corresponding to the law of large numbers, see [37, 38]. Hence, applying deep learning to infinite dimensional spaces is associated with a number of fundamental questions regarding convergence of the learning, and if it converges, in what sense?

4.1.2. Regularising properties. A theoretical topic of interest is to prove that the given reconstruction scheme constitutes a formal regularisation in the sense of [35], that is proving existence, stability and convergence.




Figure 6. Example of partially learned reconstruction when applied to fully sampled and noiseless data while the parameter Θ is trained on sparsely sampled and noisy data.

Existence for $\mathcal{T}_{\Theta}^{\dagger}: Y \rightarrow X$ is a non-issue since this operator is given by a finite number of compositions of well-defined operators. Next, assume that the gradients of the data discrepancy $\mathcal{L}(\mathcal{T}(\cdot), g): X \rightarrow \mathbb{R}$ and the regulariser $\mathcal{S}: X \rightarrow \mathbb{R}$ are Lipschitz continuous. Then the partially learned pseudo-inverse $\mathcal{T}_{\Theta}^{\dagger}$ is also Lipschitz continuous, which in turn implies stability. The final consideration concerns convergence, which is formally defined as

$$\left\| \mathcal{T}_{\Theta}^{\dagger}(\mathcal{T}(f_{\text{true}}) + \delta g) - f^* \right\|_X \rightarrow 0 \quad \text{whenever} \quad \|\delta g\|_Y \rightarrow 0$$

for some parameter choice rule for the hyper-parameters in section 2.2 and the training data, which uniquely define Θ , and where $f^* \in X$ is some minimum norm solution to (1). Clearly, the above convergence criteria can only be satisfied in general if the hyper-parameters and training data used for learning Θ are re-chosen as the data noise tends to zero. To conclude, in figure 6 we compare the result of applying the partially learned reconstruction scheme to noiseless data while training the parameter Θ against noisy data. It is reasonable to expect a significantly lower error if the method was re-trained on noiseless data, but we are currently unable to give a rigorous proof that this would converge to zero.

4.2. Use cases

The framework for partially learned reconstruction was primarily motivated by a number of use cases involving ill-posed inverse problems. A number of challenges naturally arise when classical regularisation is applied to solve the associated inverse problems and below we describe how these challenges can be resolved using a partially learned reconstruction scheme.

4.2.1. Computational feasibility. The forward operator is an important part of a regularisation and the more accurately it models the relation between signal and data, the better the outcome. Usage of more accurate forward models is however almost always computationally more demanding. Likewise, more elaborate regularisation schemes that are better at utilising the available *a priori* knowledge are often also computationally more demanding. As an example, several of the more advanced regularisers in the literature exploit some kind of sparsity using a L_1 -like norm [35]. Such regularisers typically give rise to non-differentiable

objective functional that require using optimisation algorithms from non-smooth analysis for their efficient solution. Finally, there may also be reconstruction parameters that, unlike regularisation parameter(s), do not influence the reconstruction quality. The role of the reconstruction parameters is to ensure the method is as efficient as possible, so these affect the speed of reconstruction.

Computational feasibility becomes especially critical in imaging applications since these involve very large-scale data structures. In such setting, variational and iterative regularisation schemes quickly become infeasible even for applications with moderate time requirements despite usage of state of the art algorithms.

The learned method algorithm 2 improves upon this by having an *a priori* defined run-time which can be tweaked by using more or less iterates or a more complicated updating operator. By learning, we thus learn a optimal reconstruction scheme for a given execution time. Note that the run-time of the method on our examples is significantly faster than the TV regularised method.

4.2.2. Nuisance parameters. Nuisance parameters are additional unknowns that need to be reconstructed alongside the signal. They are not of primary interest, but they nevertheless need to be reconstructed. As an example, in certain tomographic applications the acquisition geometry (sampling of the data manifold \mathbb{M}) is partially unknown, so the nuisance parameters would be those needed for a precise description of said geometry. Another is use of a more accurate forward model, which often introduces nuisance parameters.

A common approach is to adopt an intertwined scheme in which each iterate involves updating the signal by reconstructing it from data using the previous value for the nuisance parameter(s), followed by updating the nuisance parameter(s) by reconstructing them from data and making use of the previous value of recently updated signal.

Our learned reconstruction scheme algorithm 2 can easily be extended to include such intertwined schemes.

4.2.3. Regularisation parameter selection rule. Regularisation parameter(s) govern the balancing between preventing over-fitting against the need to have a solution that generates data, which is consistent with measurements. To have an appropriate parameter choice rule is critical for success.

Unfortunately, there is little theory to guide how to choose the regularisation parameter(s). Mathematical results often study asymptotic behaviour of a parameter choice rule as data noise level tends to zero. Results mainly cover the case when noise in data is additive Gaussian and its magnitude can be reliably estimated [16], even though there are extensions for other noise types as well. Nevertheless, many of these assumptions are often not met in reality.

Some work has been done in selecting an optimal parameter using learning [8]. The proposed method encompasses this since the regularisation parameter (and other optimisation related parameters) are included in the learned updating operator and thus optimally selected from the training data.

4.2.4. Feature reconstruction. Reconstructing the signal is in many applications merely one part of a more elaborate scheme of transforming measured data to knowledge. As an example, in tomographic imaging the reconstructed image serves as input for an image analysis part. The latter often involves complex procedures, like segmentation and object recognition, that currently require involvement of human expertise.

There is a growing trend in including some of these into the inverse problem that is referred to as feature reconstruction. To some extent, compressed sensing can be seen as an example of feature reconstruction where the sparse coding dictionary is the feature extraction part. Other examples are joint image reconstruction and segmentation [5, 28, 34] and shape based reconstruction [17, 18, 30]. Such feature reconstruction methods are however hard to analyse theoretically and current methods are limited in the type of feature extraction capabilities they can include. They also tend to be computationally demanding.

It is natural to perform feature reconstruction by adding a feature extraction network to the learned reconstruction scheme such as in [15]. The proposed framework could in a similar way be extended to feature reconstruction by composing the learned reconstruction operator $\mathcal{T}_{\Theta}^{\dagger}$ with a feature extraction operator $\mathcal{R}: X \rightarrow \mathbb{F}$ where \mathbb{F} is a vector space of features. If the latter is differentiable, which is the case for deep learning based feature extractors, then we can define the loss functional (4) using the composed operator $\mathcal{R} \circ \mathcal{T}_{\Theta}^{\dagger}$. This allows for *truly end-to-end optimisation of task dependent reconstruction schemes for general inverse problems*.

4.3. Stability

A general question often asked when learning is applied to some problem is whether the method generalises to other problems, e.g. if a method that is trained on a specific dataset can be applied to another dataset or to what extent one can change to forward operator without re-training.

Note first that the partially learned gradient scheme does not have an explicit regularisation parameter, instead its regularisation properties are implicitly contained in the training dataset (and to some extent in the hyper-parameters). Hence, a significant change in the training dataset (notably, a change of scaling) would require a re-training. On the other hand, empirical numerical experience suggests that dependence is relatively weak, at least for the tomographic reconstruction problems we considered. Specifically, we were able to successfully pre-train the system using a simplified acquisition geometry, a linearised forward operator, different domain size and significantly simplified phantoms and then successfully use this to train the network for the much more complicated heads dataset.

Finally, numerical experiments also suggests that changing the forward operator requires only a modest fine-tuning where the given parameters Θ can be used as an initial guess.

5. Conclusion and future work

We have presented a partially learned approach for solving ill-posed inverse problems that can integrate prior knowledge about the inverse problem with learning from training data. The presented method works with any non-linear operator and the method could easily be applied to a wide range of problems. Numerical experiments on tomographic data shows that the method gives notably better reconstructions than traditional FBP and TV regularisation. Furthermore, adding prior information improves the reconstruction. In conclusion, using prior knowledge about the forward operator, data acquisition, data noise model and regulariser can significantly improve the performance of deep learning based approaches for solving inverse problems, and especially so when the available training data is much smaller than the size of the parameter space.

An obvious next step is to tackle fully three-dimensional tomographic problems while training on two-dimensional datasets. It would also be of interest to improve upon the choice

of regulariser by adding more regularisers and/or more advanced regularisers such as wavelet based regularisers. Other more elaborate extensions are outlined below.

5.1. Extension to other iterative schemes

The given iterative method is based of the gradient descent scheme, but this scheme is known to be sub-optimal in the case of non-differentiable objective functions. A natural extension of the scheme is thus to instead consider iterative schemes better suited for this use case. One such iterative scheme is the (non-linear) PDHG algorithm [10, 39] for solving problems of the form

$$\min_{f \in X} [\mathcal{F}(\mathcal{K}(f)) + \mathcal{G}(f)]$$

where $\mathcal{K}: X \rightarrow U$ is a (possibly non-linear) operator between Banach spaces X and U . The scheme is given by algorithm 4 and the proximal operators in algorithm 4 are given by

$$\begin{aligned} \text{prox}_{\sigma\mathcal{F}^*}(h) &= \arg \min_{h' \in U} \left[\mathcal{F}^*(h') + \frac{1}{2\sigma} \|h' - h\|_U^2 \right] \\ \text{prox}_{\tau\mathcal{G}}(f) &= \arg \min_{f' \in X} \left[\mathcal{G}(f') + \frac{1}{2\tau} \|f' - f\|_X^2 \right] \end{aligned}$$

where \mathcal{F}^* is the Fenchel conjugate of \mathcal{F} . The special case of TV regularised reconstruction for (1) amounts to selecting

$$\mathcal{K}: X \rightarrow Y \times X^d \quad \text{as} \quad \mathcal{K}(f) := [\mathcal{T}(f), \nabla f]$$

where d is the dimension of the space and

$$\mathcal{F}([y_1, y_2]) := \|y_1 - g\|_2^2 + \|y_2\| \quad \text{and} \quad \mathcal{G}(f) := 0.$$

The resulting algorithm is summarised in algorithm 4.

Algorithm 4. Non-linear PDHG algorithm.

- 1: Given: $\sigma, \tau > 0$ s.t. $\sigma\tau \|\mathcal{K}\|^2 < 1$, $\theta \in [0, 1]$ and $f_0 \in X, h_0 \in U$.
 - 2: **for** $i = 1, \dots, I$ **do**
 - 3: $h^{i+1} = \text{prox}_{\sigma\mathcal{F}^*}(h^i + \sigma\mathcal{K}(\bar{f}^i))$
 - 4: $f^{i+1} = \text{prox}_{\tau\mathcal{G}}(f^i - \tau[\partial\mathcal{K}(f^i)]^*(h^{i+1}))$
 - 5: $\bar{f}^{i+1} = f^{i+1} + \theta(f^{i+1} - f^i)$
-

To introduce a learning component, one may either learn the primal proximal ($\text{prox}_{\tau\mathcal{G}}$) or the dual proximal ($\text{prox}_{\sigma\mathcal{F}^*}$), or both. Some recent papers have approached learning the primal proximal operator [11, 42] in the scope of ADMM, but these do not consider learning the dual. It is likely that learning the dual proximal offers an advantage since this allows the inclusion of various operators into the learning. To illustrate this, one can learn a proximal operator for the directional wavelet coefficients of a signal. This is successfully done for denoising [24], and would likely be useful for reconstruction as well.

Learning the dual proximal also allows one to incorporate memory into the algorithm. This can be done for the above case by extending the operator \mathcal{K} so that it also contains a zero component:

$$\mathcal{K}: X \rightarrow Y \times X^d \times X^M \quad \text{with} \quad \mathcal{K}(f) := [\mathcal{T}(f), \nabla f, 0].$$




Figure 7. Reconstruction with a partially learned dual proximal using the method in algorithm 4.

We intend to further elaborate on this approach in an upcoming paper, at this stage we settle with providing an example reconstruction shown in figure 7.

5.2. Choice of discretisation

The given examples were performed using the simplest discretisation of the space X , a pixel basis, but algorithm 2 also works with other representations such as Fourier, wavelet or shearlet coefficients. This could in many cases be better suited for the inverse problem in question and especially so if the operator \mathcal{T} or the regulariser \mathcal{S} has a simple form in this representation, as with the Fourier transform for MRI.

5.3. Choice of error functional

We simply investigated the squared norm error function, $E(f) = \|f - f_{\text{true}}\|_X^2$, but experience tells us that this is perhaps not the best predictor of human observer performance on a given image. For example, the given algorithm gives a 5 dB improvement over the TV algorithm, but by visual inspection the improvement is not equally drastic. A possible way to improve this and further leverage the power of the learning approach is to use a more sophisticated error functional. Here, performing end-to-end optimisation should be a feasible alternative, instead maximise some type of task based measure.

Acknowledgments

The work of both authors was supported by the Swedish Foundation of Strategic Research grant AM13-0049 and Industrial PhD grant ID14-0055. Jonas Adler's work was also supported by Elekta.

ORCID iDs

Jonas Adler  <https://orcid.org/0000-0001-9928-3407>

Ozan Öktem  <https://orcid.org/0000-0002-1118-6483>

References

- [1] Abadi M *et al* 2015 TensorFlow: large-scale machine learning on heterogeneous systems *OSDI'16 Proc. of the 12th USENIX Conf. on Operating Systems Design and Implementation (Savannah, GA, 2–4 November 2016)*
- [2] Adler J, Kohr H and Öktem O 2017 Operator discretization library (ODL). Software available from <https://github.com/odlgroup/odl>
- [3] Andrychowicz M, Denil M, Colmenarejo S G, Hoffman M W, Pfau D, Schaul T and de Freitas N 2016 Learning to learn by gradient descent by gradient descent *Proc. of the 34th Int. Conf. on Machine Learning (Sydney, PMLR 70, 2017)* (arXiv:1606.04474 [cs.NE])
- [4] Argyrou M, Maintas D, Tsoumpas C and Stiliaris E 2012 Tomographic image reconstruction based on artificial neural network (ANN) techniques *IEEE Nuclear Science Symp. and Medical Imaging Conf.*
- [5] Batenburg K J and Sijbers J 2011 DART: a practical reconstruction algorithm for discrete tomography *IEEE Trans. Image Process.* **20** 2542–53
- [6] Bertero M, Lantéri H and Zanni L 2008 Iterative image reconstruction: a point of view *Proc. of the Interdisciplinary Workshop on Mathematical Methods in Biomedical Imaging and Intensity-Modulated Radiation (IMRT) (Pisa, Italy) ed Y Censor et al* pp 37–63
- [7] Burger M and Osher S 2013 A Guide to the TV Zoo *Level Set and PDE Based Reconstruction Methods in Imaging* vol 2090 eds M Burger and S Osher (Berlin: Springer) pp 1–70
- [8] Calatroni L, Cao C, De Los Reyes J C, Schönlieb C-B and Valkonen T 2015 Bilevel approaches for learning of variational imaging models *Variational Methods In Imaging and Geometric Control* vol 18 (Berlin: de Gruyter & Co) pp 252–90
- [9] Cao C, De Los Reyes J C and Schönlieb C-B 2017 Learning optimal spatially-dependent regularization parameters in total variation image restoration *Inverse Problems* **33** 074005
- [10] Chambolle A and Pock T 2010 A first-order primal-dual algorithm for convex problems with applications to imaging *Technical Report* 00490826, HAL-archives. Report available from <https://hal.archives-ouvertes.fr/hal-00490826>
- [11] Chang J H R, Li C-L, Poczós B, VijayaKumar B V K and Sankaranarayanan A C 2017 One network to solve them all—solving linear inverse problems using deep projection models (arXiv:1703.09912 [cs.CV])
- [12] De Los Reyes J C and Schönlieb C-B 2013 Image denoising: learning noise distribution via PDE-constrained optimisation *Inverse Problems Imaging* **7** 1183–214
- [13] De Los Reyes J C, Schönlieb C-B and Valkonen T 2016 The structure of optimal parameters for image restoration problems *J. Math. Anal. Appl.* **434** 464–500
- [14] De Los Reyes J C, Schönlieb C-B and Valkonen T 2017 Bilevel parameter learning for higher-order total variation regularisation models *J. Math. Imaging Vis.* **57** 1–25
- [15] Diamond S, Sitzmann V, Boyd S, Wetzstein G and Heide F 2017 Dirty pixels: optimizing image classification architectures for raw sensor data (arXiv:1701.06487 [cs.CV])
- [16] Engl H W, Hanke M and Neubauer A 2000 Regularization of inverse problems *Mathematics and its Applications* vol 375 (Dordrecht: Kluwer)
- [17] Fidler T, Grasmair M and Scherzer O 2012 Shape reconstruction with *a priori* knowledge based on integral invariants *SIAM J. Imaging Sci.* **5** 726–45
- [18] Gopinath A, Xu G, Ress D, Öktem O, Subramaniam S and Bajaj C 2012 Shape-based regularization of electron tomographic reconstruction *IEEE Trans. Med. Imaging* **31** 2241–52
- [19] Guss W H 2016 Deep function machines: generalized neural networks for topological layer expression (arXiv:1612.04799 [stat.ML])
- [20] Hammernik K, Würfl T, Pock T and Maier A 2017 A deep learning architecture for limited-angle computed tomography reconstruction *Bildverarbeitung für die Medizin 2017: Algorithmen—Systeme—Anwendungen. Proc. des Workshops (Heidelberg, vom 12–bis 14 März 2017) ed K H Maier-Hein et al* (Berlin: Springer) pp 92–7
- [21] Hansen P-C 1997 *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion (SIAM Monographs on Mathematical Modeling and Computation vol 4)* (Philadelphia: SIAM)
- [22] Jin K H, McCann M T, Froustey E and Unser M 2017 Deep convolutional neural network for inverse problems in imaging *IEEE Trans. Image Process.* **26** 4509–22

- [23] Kaltenbacher B, Neubauer A and Scherzer O 2008 *Iterative Regularization Methods for Nonlinear Ill-posed Problems (Radon Series on Computational and Applied Mathematics vol 6)* (Berlin: Walter de Gruyter)
- [24] Kang E, Min J and Ye J C 2016 Wavenet: a deep convolutional neural network using directional wavelets for low-dose x-ray CT reconstruction *Med. Phys.* **44** 360–75
- [25] Kunisch K and Pock T 2013 A bilevel optimization approach for parameter learning in variational models *SIAM J. Imaging Sci.* **6** 938–83
- [26] LeCun Y, Bengio Y and Hinton G 2015 Deep learning *Nature* **521** 436–44
- [27] Liu D C and Nocedal J 1989 On the limited memory bfgs method for large scale optimization *Math. Program.* **45** 503–28
- [28] Louis A K 2011 Feature reconstruction in inverse problems *Inverse Problems* **27** 065010
- [29] Nair V and Hinton G E 2010 Rectified linear units improve restricted boltzmann machines ed J Fürnkranz and T Joachims *Proc. of the 27th Int. Conf. on Machine Learning* (Omnipress) pp 807–14
- [30] Öktem O, Chen C, Domaniç N O, Ravikumar P and Bajaj C 2017 Shape-based image reconstruction using linearized deformations *Inverse Problems* **33** 035004
- [31] Paschalis P, Giokaris N D, Karabarbounis A, Loudos G K, Maintas D, Papanicolas C N, Spanoudaki V, Tsoumpas Ch and Stiliaris E 2004 Tomographic image reconstruction using artificial neural networks *Nucl. Instrum. Methods Phys. Res. A* **527** 211–5
- [32] Pelt D M and Batenburg K J 2013 Fast tomographic reconstruction from limited data using artificial neural networks *IEEE Trans. Image Process.* **22** 5238–51
- [33] Putzky P and Welling M 2017 Recurrent inference machines for solving inverse problems *Submitted to ICLR 2017 (Toulon, France, 24–6 April 2017)* (report available from <https://openreview.net/pdf?id=HkSOIP9lg>)
- [34] Ramlau R and Ring W 2007 A Mumford–Shah level-set approach for the inversion and segmentation of x-ray tomography data *J. Comput. Phys.* **221** 539–57
- [35] Scherzer O, Grasmair M, Grossauer H, Haltmeier M and Lenzen F 2009 *Variational Methods in Imaging (Applied Mathematical Sciences vol 167)* (New York: Springer)
- [36] Schuster T 2007 *The Method of Approximate Inverse: Theory and Applications (Lecture Notes in Mathematics vol 1906)* (Heidelberg: Springer)
- [37] Stuart A M 2010 Inverse problems: a Bayesian perspective *Acta Numer.* 451–559
- [38] Vakhania N N, Tarieladze V I and Chobanyan S A 1987 *Probability Distributions on Banach Spaces (Mathematics and Its Applications (Soviet Series))* (Dordrecht: Kluwer)
- [39] Valkonen T 2014 A primal-dual hybrid gradient method for nonlinear operators with applications to MRI *Inverse Problems* **30** 055012
- [40] van Aarle W, Palenstijn W J, Cant J, Janssens E, Bleichrodt F, Dabrovolski A, Beenhouwer J, Batenburg K J and Sijbers J 2016 Fast and flexible x-ray tomography using the ASTRA toolbox *Opt. Express* **24** 25129–47
- [41] Würfl T, Ghesu F C, Christlein V and Maier A 2016 Deep learning computed tomography *MICCAI: Medical Image Computing and Computer-Assisted Intervention (Lecture Notes in Computer Science vol 9902)* ed S Ourselin *et al* (New York: Springer) pp 432–40
- [42] Yang Y, Sun J, Li H and Xu Z 2016 Deep ADMM-Net for compressive sensing MRI *Advances in Neural Information Processing Systems vol 29*, ed D D Lee *et al* (Curran Associates) pp 10–8 (report available from <http://papers.nips.cc/paper/6406-deep-admm-net-for-compressive-sensing-mri.pdf>)