

# On Multi-Layer Basis Pursuit, Efficient Algorithms and Convolutional Neural Networks

Jeremias Sulam, *Member, IEEE*, Aviad Aberdam, Amir Beck, Michael Elad, *Fellow, IEEE*

**Abstract**—Parsimonious representations are ubiquitous in modeling and processing information. Motivated by the recent Multi-Layer Convolutional Sparse Coding (ML-CSC) model, we herein generalize the traditional Basis Pursuit problem to a multi-layer setting, introducing similar sparse enforcing penalties at different representation layers in a symbiotic relation between synthesis and analysis sparse priors. We explore different iterative methods to solve this new problem in practice, and we propose a new Multi-Layer Iterative Soft Thresholding Algorithm (ML-ISTA), as well as a fast version (ML-FISTA). We show that these nested first order algorithms converge, in the sense that the function value of near-fixed points can get arbitrarily close to the solution of the original problem. We further show how these algorithms effectively implement particular recurrent convolutional neural networks (CNNs) that generalize feed-forward ones without introducing any parameters. We present and analyze different architectures resulting unfolding the iterations of the proposed pursuit algorithms, including a new Learned ML-ISTA, providing a principled way to construct deep recurrent CNNs. Unlike other similar constructions, these architectures unfold a global pursuit holistically for the entire network. We demonstrate the emerging constructions in a supervised learning setting, consistently improving the performance of classical CNNs while maintaining the number of parameters constant.

**Index Terms**—Multi-Layer Convolutional Sparse Coding, Network Unfolding, Recurrent Neural Networks, Iterative Shrinkage Algorithms.



## 1 INTRODUCTION

Sparsity has been shown to be a driving force in a myriad of applications in computer vision [30], [42], [43], statistics [38], [39] and machine learning [23], [24], [26]. Most often, sparsity is often enforced not on a particular signal but rather on its representation in a transform domain. Formally, a signal  $\mathbf{x} \in \mathbb{R}^n$  admits a sparse representation in terms of a dictionary  $\mathbf{D} \in \mathbb{R}^{n \times m}$  if  $\mathbf{x} = \mathbf{D}\boldsymbol{\gamma}$ , and  $\boldsymbol{\gamma} \in \mathbb{R}^m$  is sparse. In its simplest form, the problem of seeking for a sparse representation for a signal, possibly contaminated with noise  $\mathbf{w}$  as  $\mathbf{y} = \mathbf{x} + \mathbf{w}$ , can be posed in terms of the following pursuit problem:

$$\min_{\boldsymbol{\gamma}} \|\boldsymbol{\gamma}\|_0 \quad \text{s.t.} \quad \|\mathbf{y} - \mathbf{D}\boldsymbol{\gamma}\|_2^2 \leq \varepsilon,$$

where the  $\ell_0$  pseudo-norm counts the number of non-zero elements in  $\boldsymbol{\gamma}$ . The choice of the (typically overcomplete) dictionary  $\mathbf{D}$  is far from trivial, and has motivated the design of different analytic transforms [9], [16], [27] and the development of dictionary learning methods [2], [30], [36]. The above problem, which is NP-hard in general, is often relaxed by employing the  $\ell_1$  penalty as a surrogate for

the non-convex  $\ell_0$  measure, resulting in the celebrated Basis Pursuit De-Noising (BPDN) problem<sup>1</sup>:

$$\min_{\boldsymbol{\gamma}} \lambda \|\boldsymbol{\gamma}\|_1 + \frac{1}{2} \|\mathbf{y} - \mathbf{D}\boldsymbol{\gamma}\|_2^2. \quad (1)$$

The transition from the  $\ell_0$  to the relaxed  $\ell_1$  case is by now well understood, and the solutions to both problems do coincide under sparse assumptions on the underlying representation (in the noiseless case), or have shown to be close enough in more general settings [17], [41].

This traditional model was recently extended to a multi-layer setting [1], [34], where a signal is expressed as  $\mathbf{x} = \mathbf{D}_1\boldsymbol{\gamma}_1$ , for a sparse  $\boldsymbol{\gamma}_1 \in \mathbb{R}^{m_1}$  and (possibly convolutional) matrix  $\mathbf{D}_1$ , while also assuming that this representation satisfies  $\boldsymbol{\gamma}_1 = \mathbf{D}_2\boldsymbol{\gamma}_2$ , for yet another dictionary  $\mathbf{D}_2 \in \mathbb{R}^{m_1 \times m_2}$  and sparse  $\boldsymbol{\gamma}_2 \in \mathbb{R}^{m_2}$ . Such a construction can be cascaded for a number of  $L$  layers<sup>2</sup>. Under this framework, given the measurement  $\mathbf{y}$ , this multi-layer pursuit problem (or Deep Coding Problem, as first coined in [34]), can be expressed as

$$\min_{\{\boldsymbol{\gamma}_i\}} \|\mathbf{y} - \mathbf{D}_1\boldsymbol{\gamma}_1\|_2^2 \quad \text{s.t.} \quad \{\boldsymbol{\gamma}_{i-1} = \mathbf{D}_i\boldsymbol{\gamma}_i, \|\boldsymbol{\gamma}_i\|_0 \leq s_i\}_{i=1}^L, \quad (2)$$

with  $\mathbf{x} = \boldsymbol{\gamma}_0$ . In this manner, one searches for the closest signal to  $\mathbf{y}$  while satisfying the model assumptions. This can

• J. Sulam and A. Aberdam contributed equally to this work.

• J. Sulam is with the Department of Biomedical Engineering, Johns Hopkins University. A. Aberdam is with the Department of Electrical Engineering, Technion – Israel Institute of Technology. A. Beck is with the School of Mathematical Sciences at Tel Aviv University. M. Elad is with the Department of Computer Science, Technion – Israel Institute of Technology.

The research leading to these results has received funding from the European Research Council under European Unions Seventh Framework Programme, ERC Grant agreement no. 320649. The work of Amir Beck was partially supported by the Israel Science Foundation 1821/16.

1. This problem is also known in the statistical learning community as Least Absolute Shrinkage and Selection Operator (LASSO) [38], where the matrix  $\mathbf{D}$  is given by a set of measurements or descriptors, in the context of a sparse regression problem.

2. In the convolutional setting [35], [37], the notion of sparsity is better characterized by the  $\ell_{0,\infty}$  pseudo-norm, which quantifies the density of non-zeros in the convolutional representations in a local sense. Importantly, however, the BPDN formulation (i.e., employing an  $\ell_1$  penalty), still serves as a proxy for this  $\ell_{0,\infty}$  norm. We refer the reader to [35] for a thorough analysis of convolutional sparse representations.

be understood and analyzed as a projection problem [37], providing an estimate such that  $\hat{\mathbf{x}} = \mathbf{D}_1 \hat{\gamma}_1 = \mathbf{D}_1 \mathbf{D}_2 \hat{\gamma}_2 = \dots = \mathbf{D}_{(1,L)} \hat{\gamma}_L$ , while forcing all intermediate representations to be sparse. Note the notation  $\mathbf{D}_{(i,L)} = \mathbf{D}_i \dots \mathbf{D}_L$  for brevity. Remarkably, the forward pass of neural networks (whose weights at each layer,  $\mathbf{W}_i$ , are set as the transpose of each dictionary  $\mathbf{D}_i$ ) yields stable estimations for the intermediate features or representations  $\hat{\gamma}_i$  provided these are sparse enough [34]. Other generative models have also been recently proposed (as the closely related probabilistic framework in [?], [?]), but the multi-layer sparse model provides a convenient way to study deep learning architectures in terms of pursuit algorithms [?].

As an alternative to the forward pass, one can address the problem in Equation (1) by adopting a projection interpretation and develop an algorithm based on a global pursuit, as in [37]. More recently, the work in [1] showed that this problem can be cast as imposing an analysis prior on the signal's deepest sparse representation. Indeed, the problem in (1) can be written concisely as:

$$\begin{aligned} \min_{\{\gamma_i\}} \quad & \|\mathbf{y} - \mathbf{D}_{(1,L)} \gamma_L\|_2^2 & (3) \\ \text{s.t.} \quad & \|\gamma_L\|_0 \leq s_L, \{ \|\mathbf{D}_{(i,L)} \gamma_L\|_0 \leq s_{i-1} \}_{i=1}^L. \end{aligned}$$

This formulation explicitly shows that the intermediate dictionaries  $\mathbf{D}_{(i,L)}$  play the role of analysis operators, resulting in a representation  $\gamma_L$  which should be orthogonal to as many rows from  $\mathbf{D}_{(i,L)}$  as possible – so as to produce zeros in  $\gamma_i$ . Interestingly, this analysis also allows for *less sparse* representations in shallower layers while still being consistent with the multi-layer sparse model. While a pursuit algorithm addressing (1) was presented in [1], it is greedy in nature and does not scale to high dimensional signals. In other words, there are currently no efficient pursuit algorithms for signals in this multi-layer model that leverage this symbiotic analysis-synthesis priors. More importantly, it is still unclear how the dictionaries could be trained from real data under this scheme. These questions are fundamental if one is to bridge the theoretical benefits of the multi-layer sparse model with practical deep learning algorithms.

In this work we propose a relaxation of the problem in Equation (1), turning this seemingly complex pursuit into a convex multi-layer generalization of the Basis Pursuit (BP) problem<sup>3</sup>. Such a formulation, to the best of our knowledge, has never before been proposed nor studied, though we will comment on a few particular and related cases that have been of interest to the image processing and compressed sensing communities. We explore different algorithms to solve this multi-layer problem, such as variable splitting and the Alternating Directions Method of Multipliers (ADMM) [6], [8] and the Smooth-FISTA from [5], and we will present and analyze two new generalizations of Iterative Soft Thresholding Algorithms (ISTA). We will further show that these algorithms generalize feed-forward neural networks (NNs), both fully-connected and convolutional (CNNs), in a natural way. More precisely: the first iteration of such algorithms implements a traditional CNN, while a new recurrent architecture emerges

with subsequent iterations. In this manner, the proposed algorithms provide a principled framework for the design of recurrent architectures. While other works have indeed explored the unrolling of iterative algorithms in terms of CNNs (e.g. [33], [44]), we are not aware of any work that has attempted nor studied the unrolling of a global pursuit with convergence guarantees. Lastly, we demonstrate the performance of these networks in practice by training our models for image classification, consistently improving on the classical feed-forward architectures without introducing filters nor any other extra parameters in the model.

## 2 MULTI-LAYER BASIS PURSUIT

In this work we propose a convex relaxation of the problem in Equation (1), resulting in a multi-layer BP problem. For the sake of clarity, we will limit our formulations to two layers, but these can be naturally extended to multiple layers – as we will effectively do in the experimental section. This work is centered around the following problem:

$$(P) : \min_{\gamma} \frac{1}{2} \|\mathbf{y} - \mathbf{D}_1 \mathbf{D}_2 \gamma\|_2^2 + \lambda_1 \|\mathbf{D}_2 \gamma\|_1 + \lambda_2 \|\gamma\|_1. \quad (4)$$

This formulation imposes a particular mixture of synthesis and analysis priors. Indeed, if  $\lambda_2 > 0$  and  $\lambda_1 = 0$ , one recovers a traditional Basis Pursuit formulation with a factorized global dictionary. If  $\lambda_1 > 0$ , however, an analysis prior is enforced on the representation  $\gamma$  by means of  $\mathbf{D}_2$ , resulting in a more regularized solution. Note that if  $\lambda_2 = 0$ ,  $\lambda_1 > 0$  and  $\ker \mathbf{D}_2$  is not empty, the problem above becomes ill-posed without a unique solution<sup>4</sup> since  $\ker \mathbf{D}_{(1,2)} \cap \ker \mathbf{D}_2 \neq \{\mathbf{0}\}$ . In addition, unlike previous interpretations of the multi-layer sparse model ([1], [34], [37]), our formulation stresses the fact that there is *one unknown* variable,  $\gamma$ , with different priors enforced on it. Clearly, one may also define and introduce  $\gamma_1 = \mathbf{D}_2 \gamma$ , but this should be interpreted merely as the introduction of auxiliary variables to aid the derivation and interpretation of the respective algorithms. We will expand on this point in later sections.

Other optimization problems similar to (P) have indeed been proposed, such as the Analysis-LASSO [10], [28], though their observation matrix and the analysis operator ( $\mathbf{D}_{(1,2)}$  and  $\mathbf{D}_2$ , in our case) must be independent, and the latter is further required to be a tight frame [10], [28]. The Generalized Lasso problem [40] is also related to our multi-layer BP formulation, as we will see in the following section. On the other hand, and in the context of image restoration, the work in [7] imposes a Total Variation and a sparse prior on the unknown image, as does the work in [21], thus being closely related to the general expression in (2).

### 2.1 Algorithms

From an optimization perspective, our multi-layer BP problem can be expressed more generally as

$$\min_{\gamma} F(\gamma) = f(\mathbf{D}_2 \gamma) + g_1(\mathbf{D}_2 \gamma) + g_2(\gamma), \quad (5)$$

3. In an abuse of terminology, and for the sake of simplicity, we will refer to the BPDN problem in Equation (1) as BP.

4. It is true that also in Basis Pursuit one can potentially obtain infinite solutions, as the problem is not strongly convex.

**Algorithm 1** ADMM algorithm for a two-layer ML-CSC model.

Input: signal  $\mathbf{y}$ , dictionaries  $\mathbf{D}_i$  and parameters  $\lambda_i$

- 1: **while** not converged **do**
- 2:    $\gamma_2 \leftarrow \arg \min_{\gamma_2} \|\mathbf{y} - \mathbf{D}_1 \mathbf{D}_2 \gamma_2\|_2^2$   
        $+ \frac{\rho}{2} \|\gamma_1 - \mathbf{D}_2 \gamma_2 + \mathbf{u}\|_2^2 + \lambda_2 \|\gamma_2\|_1$
- 3:    $\gamma_1 \leftarrow \arg \min_{\gamma_1} \frac{\rho}{2} \|\gamma_1 - \mathbf{D}_2 \gamma_2 + \mathbf{u}\|_2^2 + \lambda_1 \|\gamma_1\|_1$
- 4:    $\mathbf{u} \leftarrow \mathbf{u} + \rho(\gamma_1 - \mathbf{D}_2 \gamma_2)$

where  $f$  is convex and smooth, and  $g_1$  and  $g_2$  are convex but non-smooth. For the specific problem in (2),  $f(\mathbf{z}) = \frac{1}{2} \|\mathbf{y} - \mathbf{D}_1 \mathbf{z}\|_2^2$ ,  $g_1(\mathbf{z}) = \lambda_1 \|\mathbf{z}\|_1$  and  $g_2(\mathbf{z}) = \lambda_2 \|\mathbf{z}\|_1$ . Since this problem is convex, the choice of available algorithms is extensive. We are interested in high-dimensional settings, however, where interior-point methods and other solvers depending on second-order information might have a prohibitive computational complexity. In this context, the Iterative Soft Thresholding Algorithm (ISTA), and its Fast version (FISTA), are appealing as they only require matrix-vector multiplications and entry-wise operations. The former, originally introduced in [15], provides convergence (in function value) of order  $\mathcal{O}(1/k)$ , while the latter provides an improved convergence rate with order of  $\mathcal{O}(1/k^2)$  [4].

Iterative shrinkage algorithms decompose the total loss into two terms:  $f(\gamma)$ , convex and smooth (with Lipschitz constant  $L$ ), and  $g(\gamma)$ , convex and possibly non smooth. The central idea of ISTA, as a proximal gradient method for finding a minimizer of  $f + g$ , is to iterate the updates given by the proximal operator of  $g(\cdot)$  at the forward-step point:

$$\gamma^{k+1} = \text{prox}_{\frac{1}{L}g} \left( \gamma^k - \frac{1}{L} \nabla f(\gamma^k) \right). \quad (6)$$

Clearly, the appeal of ISTA depends on how effectively the proximal operator can be computed. When  $g(\gamma) = \lambda \|\gamma\|_1$  (as in the original BP formulation), such a proximal mapping becomes separable, resulting in the element-wise shrinkage or soft-thresholding operator. However, this family of methods cannot be readily applied to (P) where  $g(\gamma) = \lambda_1 \|\mathbf{D}_2 \gamma\|_1 + \lambda_2 \|\gamma\|_1$ . Indeed, computing  $\text{prox}_g(\gamma)$  when  $g(\cdot)$  is a sum of  $\ell_1$  composite terms is no longer directly separable, and one must resort to iterative approaches, making ISTA lose its appeal.

The problem (P) is also related to the Generalized Lasso formulation [40] in the compressed sensing community, which reads

$$\min_{\gamma} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\gamma\|_2^2 + \nu \|\mathbf{A}\gamma\|_1. \quad (7)$$

Certainly, the multi-layer BP problem we study can be seen as a particular case of this formulation<sup>5</sup>. With this insight, one might consider solving (P) through the solution of the generalized Lasso [40]. However, such an approach also becomes computationally demanding as it boils down to an iterative algorithm that includes the inversion of linear operators. Other possible solvers might rely on re-weighted  $\ell_2$  approaches [11], but these also require iterative matrix inversions.

5. One can rewrite problem (2) as in (2.1) by making  $\mathbf{X} = \mathbf{D}_{(1,L)}$  and  $\mathbf{A} = [\lambda_1/\nu \mathbf{D}_2^T, \lambda_2/\nu \mathbf{I}]^T$ .

A simple way of tackling problem (P) is the popular Alternating Directions Method of Multipliers (ADMM), which provides a natural way to address these kind of problems through variable splitting and auxiliary variables. For a two layer model, one can rewrite the multi-layer BP as a constrained minimization problem:

$$\begin{aligned} \min_{\gamma_1, \gamma_2} \quad & \frac{1}{2} \|\mathbf{y} - \mathbf{D}_1 \mathbf{D}_2 \gamma_2\|_2^2 + \lambda_1 \|\gamma_1\|_1 + \lambda_2 \|\gamma_2\|_1 \\ \text{s.t.} \quad & \gamma_1 = \mathbf{D}_2 \gamma_2. \end{aligned}$$

ADMM minimizes this constrained loss by constructing an augmented Lagrangian (in normalized form) as

$$\begin{aligned} \min_{\gamma_1, \gamma_2, \mathbf{u}} \quad & \frac{1}{2} \|\mathbf{y} - \mathbf{D}_1 \mathbf{D}_2 \gamma_2\|_2^2 + \lambda_1 \|\gamma_1\|_1 \\ & + \lambda_2 \|\gamma_2\|_1 + \frac{\rho}{2} \|\gamma_1 - \mathbf{D}_2 \gamma_2 + \mathbf{u}\|_2^2, \end{aligned}$$

which can be minimized iteratively by repeating the updates in Algorithm 1. This way, and after merging both  $\ell_2$  terms, the pursuit of the inner-most representation ( $\gamma_2$  in this case) is carried out in terms of a regular BP formulation that can be tackled with a variety of convex methods, including ISTA or FISTA. The algorithm then updates the intermediate representations ( $\gamma_1$ ) by a simple shrinkage operation, followed by the update of the dual variable,  $\mathbf{u}$ . Note that this algorithm is guaranteed to converge (at least in the sequence sense) to a global optimum of (P) due to the convexity of the function being minimized [6], [8].

A third alternative, which does not incur in an additional inner iteration nor inversions, is the Smooth-FISTA approach from [5]. S-FISTA addresses cost functions of the same form as problem (P) by replacing one of the non-smooth functions,  $g_1(\gamma)$  in Eq. (2.1), by a smoothed version in terms of its Moreau envelope. In this way, S-FISTA converges with order  $\mathcal{O}(1/\varepsilon)$  to an estimate that is  $\varepsilon$ -away from the solution of the original problem in terms of function value. We will revisit this method further in the following sections.

Before moving on, we make a short intermission to note here that other Basis Pursuit schemes have been proposed in the context of multi-layer sparse models and deep learning. Already in [34] the authors proposed the Layered Basis Pursuit, which addresses the sequence of pursuits given by

$$\hat{\gamma}_i \leftarrow \arg \min_{\gamma_i} \|\hat{\gamma}_{i-1} - \mathbf{D}_i \gamma_i\|_2^2 + \lambda_i \|\gamma_i\|_1,$$

from  $i = 1$  to  $L$ , where  $\gamma_0 = \mathbf{y}$ . Clearly, each of these can be solved with any BP solver just as well. A related idea was also recently proposed in [?], showing that cascading basis pursuit problems can lead to competitive deep learning constructions. However, the Layered Basis Pursuit formulation, or other similar variations that attempt to *unfold* neural network architectures [33], [44], do not minimize (P) and thus their solutions only represent sub-optimal and heuristic approximations to the minimizer of the multi-layer BP. More clearly, such a series of steps never provide estimates  $\hat{\gamma}_i$  that can generate a signal according to the multi-layer sparse model. As a result, one cannot reconstruct  $\hat{\mathbf{x}} = \mathbf{D}_{(1,L)} \hat{\gamma}_L$ , because each representation is required to explain the next layer *only approximately*, so that  $\hat{\gamma}_{i-1} \neq \mathbf{D}_i \hat{\gamma}_i$ .

## 2.2 Towards Multi-Layer ISTA

We now move to derive the proposed approach to efficiently tackle  $(P)$  while relying on the concept of the *gradient mapping* (see for example [3, Chapter 10]), which we briefly review next. Given a function  $F(\gamma) = f(\gamma) + g(\gamma)$ , where  $f$  is convex and smooth with Lipschitz constant  $L$  and  $g$  is convex, the gradient mapping is the operator given by

$$G_L^{f,g}(\gamma) = L \left[ \gamma - \text{prox}_{\frac{1}{L}g} \left( \gamma - \frac{1}{L} \nabla f(\gamma) \right) \right].$$

Naturally, the ISTA update step in Equation (2.1) can be seen as a “gradient-mapping descent” step, since it can be rewritten as  $\gamma^{k+1} = \gamma^k - \frac{1}{L} G_L^{f,g}(\gamma)$ . Moreover,  $G_L^{f,g}(\gamma)$  provides a sort of generalization of the gradient of  $F(\gamma)$ , since

- 1)  $G_L^{f,g}(\gamma) = \nabla F(\gamma) = \nabla f(\gamma)$  if  $g(\gamma) \equiv 0$ ,
- 2)  $G_L^{f,g}(\gamma) = 0$  if and only if  $\gamma$  is a minimizer of  $F(\gamma)$ .

We refer the reader to [3, Chapter 10] for further details on gradient mapping operators.

Returning to the problem in (2.1), our first attempt to minimize  $F(\gamma_2) = f(\mathbf{D}_2 \gamma_2) + g_1(\mathbf{D}_2 \gamma_2) + g_2(\gamma_2)$  is a *proximal gradient-mapping* method, and it takes an update of the following form:

$$\gamma_2^{k+1} = \text{prox}_{t_{g_2}} \left( \gamma_2^k - t G_{1/\mu}^{f(\cdot), g_1(\mathbf{D}_2 \cdot)}(\gamma_2^k) \right),$$

for constants  $\mu > 0$  and  $t > 0$  that will be specified shortly. This expression, however, requires the computation of  $\text{prox}_{g_1(\mathbf{D}_2 \cdot)}(\cdot)$ , which is problematic as it involves a composite term<sup>6</sup>. To circumvent this difficulty, we propose the following approximation in terms of  $\gamma_1 = \mathbf{D}_2 \gamma_2$ . In the spirit of the chain rule<sup>7</sup>, we modify the previous update to

$$\gamma_2^{k+1} = \text{prox}_{t_{g_2}} \left( \gamma_2^k - t \mathbf{D}_2^T G_{1/\mu}^{f,g_1}(\gamma_1^k) \right). \quad (8)$$

Importantly, the above update step now involves the  $\text{prox}$  of  $g_1(\cdot)$  as opposed to that of  $g_1(\mathbf{D}_2 \cdot)$ . This way, in the case of  $g_1(\cdot) = \lambda_1 \|\cdot\|_1$ , the proximal mapping of  $g_1$  becomes the soft-thresholding operator with parameter  $\lambda_1$ , i.e.  $\text{prox}_{g_1}(\gamma_1) = \mathcal{T}_{\lambda_1}(\gamma_1)$ . An analogous operator is obtained for  $\text{prox}_{g_2}$  just as well. Therefore, the proposed Multi-Layer ISTA update can be concisely written as

$$\gamma_2^{k+1} = \mathcal{T}_{t\lambda_2} \left( \gamma_2^k - \frac{t}{\mu} \mathbf{D}_2^T \left( \gamma_1^k - \mathcal{T}_{\mu\lambda_1}(\gamma_1^k - \mu \mathbf{D}_1^T (\mathbf{D}_1 \gamma_1^k - \mathbf{y})) \right) \right). \quad (9)$$

A few comments are in place. First, this algorithm results in a nested series of shrinkage operators, involving only matrix-vector multiplications and entry-wise non linear operations. Note that if  $\lambda_1 = 0$ , i.e. in the case of a traditional Basis Pursuit problem, the update above reduces to the update of ISTA. Second, though seemingly complicated at first sight, the resulting operator in (2.2) can be decomposed into simple recursive layer-wise operations, as presented in Algorithm 2. Lastly, because the above update provides a multi-layer extension to ISTA, one can naturally suggest a

6. The proximal of a composition with an affine map is only available for unitary linear transformations. See [3, Chapter 10] and [13] for further details.

7. The step taken to arrive at Equation (2.2) is not actually the chain rule, as the gradient mapping  $G_L^{f,g}$  is not necessarily a gradient of a smooth function.

“fast version” of it by including a momentum term, just as done by FISTA. In other words, ML-FISTA will be given by the iterations

$$\begin{aligned} \gamma_2^{k+1} &= \text{prox}_{t_{g_2}} \left( \mathbf{z}^k - t \mathbf{D}_2^T G_{1/\mu}^{f,g_1}(\mathbf{D}_2 \mathbf{z}^k) \right), \\ \mathbf{z}^{k+1} &= \gamma_2^{k+1} + \rho^k (\gamma_2^{k+1} - \gamma_2^k), \end{aligned} \quad (10)$$

where  $\rho^k = \frac{t_k - 1}{t_{k+1}}$ , and the  $t_k$  parameter is updated according to  $t_{k+1} = \frac{1 + \sqrt{1 + 4t_k^2}}{2}$ . Clearly, one can also write this algorithm in terms of layer-wise operations, as described in Algorithm 3.

---

### Algorithm 2 Multi-Layer ISTA.

---

Input: signal  $\mathbf{y}$ , dictionaries  $\mathbf{D}_i$  and parameters  $\lambda_i$ .

Init: Set  $\gamma_0^k = \mathbf{y} \forall k$  and  $\gamma_L^1 = 0$ .

- 1: **for**  $k = 1 : K$  **do** % for each iteration
  - 2:  $\hat{\gamma}_i \leftarrow \mathbf{D}_{(i,L)} \gamma_L^k \quad \forall i \in [0, L - 1]$
  - 3: **for**  $i = 1 : L$  **do** % for each layer
  - 4:  $\gamma_i^{k+1} \leftarrow \mathcal{T}_{\mu_i \lambda_i} (\hat{\gamma}_i - \mu_i \mathbf{D}_i^T (\mathbf{D}_i \hat{\gamma}_i - \gamma_{i-1}^{k+1}))$
- 

---

### Algorithm 3 Multi-Layer FISTA.

---

Input: signal  $\mathbf{y}$ , dictionaries  $\mathbf{D}_i$  and parameters  $\lambda_i$ .

Set  $\gamma_0^k = \mathbf{y} \forall k$  and  $\mathbf{z} = 0$ .

- 1: **for**  $k = 1 : K$  **do** % for each iteration
  - 2:  $\hat{\gamma}_i \leftarrow \mathbf{D}_{(i,L)} \mathbf{z} \quad \forall i \in [0, L - 1]$
  - 3: **for**  $i = 1 : L$  **do** % for each layer
  - 4:  $\gamma_i^{k+1} \leftarrow \mathcal{T}_{\mu_i \lambda_i} (\hat{\gamma}_i - \mu_i \mathbf{D}_i^T (\mathbf{D}_i \hat{\gamma}_i - \gamma_{i-1}^{k+1}))$
  - 5:  $t_{k+1} \leftarrow \frac{1 + \sqrt{1 + 4t_k^2}}{2}$
  - 6:  $\mathbf{z} \leftarrow \gamma_L^{k+1} + \frac{t_k - 1}{t_{k+1}} (\gamma_L^{k+1} - \gamma_L^k)$
- 

## 2.3 Convergence Analysis of ML-ISTA

One can then inquire – does the update in Equation (2.2) provide convergent algorithm? Does the successive iterations minimize the original loss function? Though these questions have been extensively studied for proximal gradient methods [4], [13], algorithms based on a proximal *gradient mapping* have never been proposed – let alone analyzed. Herein we intend to provide a first theoretical analysis of the resulting multi-layer thresholding approaches.

Let us formalize the problem assumptions, recalling that we are interested in

$$(P) : \quad \min_{\gamma} F(\gamma) = f(\mathbf{D}_2 \gamma) + g_1(\mathbf{D}_2 \gamma) + g_2(\gamma),$$

where  $f : \mathbb{R}^{m_1} \rightarrow \mathbb{R}$  is a quadratic convex function,  $g_1 : \mathbb{R}^{m_1} \rightarrow \mathbb{R}$  is a convex and Lipschitz continuous function with constant  $\ell_{g_1}$  and  $g_2 : \mathbb{R}^{m_2} \rightarrow (-\infty, \infty]$  is a proper closed and convex function that is  $\ell_{g_2}$ -Lipschitz continuous over its domain. Naturally, we will assume that both  $g_1$  and  $g_2$  are *proximable*, in the sense that  $\text{prox}_{\alpha g_1}(\gamma)$  and  $\text{prox}_{\alpha g_2}(\gamma)$  can be efficiently computed for any  $\gamma$  and  $\alpha > 0$ . We will further require that  $g_2$  has a bounded domain<sup>8</sup>.

8. This can be easily accommodated by adding to  $g_2$  a norm bound constraint in the form of an indicator function  $\delta_{B[0,R]}$ , for some large enough  $R > 0$ .

More precisely,  $\text{dom}(g_2) \subseteq \{\gamma_2 : \|\gamma_2\|_2 \leq R\}$ . We denote  $R_1 = \|\mathbf{D}_2\|_2 R$ , so that  $\gamma_2 \in \text{dom}(g_2) \implies \gamma_1 = \mathbf{D}_2 \gamma_2$  satisfies  $\|\gamma_1\|_2 \leq R_1$ .

Note also that the convex and smooth function  $f$  can be expressed as  $f(\gamma_1) = \frac{1}{2} \gamma_1^T \mathbf{Q} \gamma_1 + \mathbf{b}^T \gamma_1 + c$ , for a positive semi-definite matrix  $\mathbf{Q} = \mathbf{D}_1^T \mathbf{D}_1$ . The gradient of  $f$  can then easily be bounded by

$$\|\nabla f(\gamma_1)\|_2 \leq M \equiv \|\mathbf{Q}\|_2 R_1 + \|\mathbf{b}\|_2, \quad (11)$$

for any  $\gamma_1 = \mathbf{D}_2 \gamma_2$  with  $\gamma_2 \in \text{dom}(g_2)$ .

It is easy to see that the algorithm in Equation (2.2) does not converge to the minimizer of the problem (P) by studying its fixed point. The point  $\gamma_2^*$  is a fixed point of ML-ISTA if

$$\begin{aligned} \exists \mathbf{w}_2 \in \partial g_2(\gamma_2^*), \mathbf{w}_1 \in \partial g_1(\hat{\gamma}_1) \\ \text{such that } \mathbf{D}_2^T \nabla f(\mathbf{D}_2 \gamma_2^*) + \mathbf{D}_2^T \mathbf{w}_1 + \mathbf{w}_2 = 0, \end{aligned}$$

where  $\hat{\gamma}_1 = \text{prox}_{g_1 \mu}(\mathbf{D}_2 \gamma_2^* - \mu \nabla f(\mathbf{D}_2 \gamma_2^*))$ . We extend on the derivation of this condition in Section 5.1. This is clearly different from the optimality conditions for problem (P), which is

$$\begin{aligned} \exists \mathbf{w}_2 \in \partial g_2(\gamma_2^*), \mathbf{w}_1 \in \partial g_1(\mathbf{D}_2 \gamma_2^*) \\ \text{such that } \mathbf{D}_2^T \nabla f(\mathbf{D}_2 \gamma_2^*) + \mathbf{D}_2^T \mathbf{w}_1 + \mathbf{w}_2 = 0. \end{aligned}$$

Nonetheless, we will show that the parameter  $\mu$  controls the proximity of “near”-fixed points to the optimal solution in the following sense: as  $\mu$  gets smaller, the objective function of a fixed-point of the ML-ISTA method gets closer to  $F_{\text{opt}}$ , the minimal value of  $F(\gamma)$ . In addition, for points that satisfy the fixed point equation up to some tolerance  $\varepsilon > 0$ , the distance to optimality in terms of the objective function is controlled by both  $\mu$  and  $\varepsilon$ .

We first present the following lemma, stating that the norm of the gradient mapping operator is bounded. For clarity, we defer the proof of this lemma, as well as that of the following theorem, to Section 5.

**Lemma 2.1.** For any  $\mu > 0$  and  $\gamma_1 \in \mathbb{R}^{m_1}$ ,

$$\|G_{1/\mu}^{f, g_1}(\gamma_1)\|_2 \leq M + \ell_{g_1}.$$

We now move to our main convergence result. In a nutshell, it states that the distance from optimality in terms of the objective function value of  $\varepsilon$ -fixed points is bounded by constants multiplying  $\varepsilon$  and  $\mu$ .

**Theorem 2.2.** Let<sup>9</sup>  $\mu \in \left(0, \frac{1}{\|\mathbf{Q}\|_2}\right)$ ,  $t \in \left(0, \frac{4\mu}{3\|\mathbf{D}_2\|_2}\right)$ , and assume  $\tilde{\gamma}_2 \in \text{dom}(g_2)$  and  $\tilde{\gamma}_1 = \mathbf{D}_2 \tilde{\gamma}_2$ . If

$$\frac{1}{t} \left\| \tilde{\gamma}_2 - \text{prox}_{t g_2} \left( \tilde{\gamma}_2 - t \mathbf{D}_2^T G_{1/\mu}^{f, g_1}(\tilde{\gamma}_1) \right) \right\|_2 \leq \varepsilon,$$

then

$$F(\boldsymbol{\alpha}) - F_{\text{opt}} \leq \eta \varepsilon + (\beta + \kappa t) \mu,$$

where  $F_{\text{opt}} \equiv \min_{\gamma_2} F(\gamma_2)$ ,

$$\boldsymbol{\alpha} = \text{prox}_{t g_2} \left( \tilde{\gamma}_2 - t \mathbf{D}_2^T G_{1/\mu}^{f, g_1}(\tilde{\gamma}_1) \right),$$

9. For a matrix  $\mathbf{A}$ ,  $\|\mathbf{A}\|_2$  denotes the spectral norm of  $\mathbf{A}$ :  $\|\mathbf{A}\|_2 = \sqrt{\lambda_{\max}(\mathbf{A}^T \mathbf{A})}$ , where  $\lambda_{\max}(\cdot)$  stands for the maximal eigenvalue of its argument.

and

$$\begin{aligned} \eta &= 2R, \\ \beta &= 2R \|\mathbf{D}_2\|_2 \|\mathbf{Q}\|_2 (M + \ell_{g_1}) + \|\mathbf{Q}\|_2^2 R_1^2 \\ &\quad + 2 \|\mathbf{b}\|_2 \|\mathbf{Q}\|_2 R_1 + \ell_{g_1}^2 + 2 \ell_{g_1} M, \\ \kappa &= \|\mathbf{D}_2\|_2 (\|\mathbf{D}_2\|_2 (M + \ell_{g_1}) + \ell_{g_2}) \|\mathbf{Q}\|_2 (M + \ell_{g_1}). \end{aligned} \quad (12)$$

A consequence of this result is the following.

**Corollary 2.2.1.** Suppose  $\{\gamma_2^k\}$  is the sequence generated by ML-ISTA with  $\mu \in \left(0, \frac{1}{\|\mathbf{D}_1\|_2^2}\right)$  and  $t \in \left(0, \frac{4\mu}{3\|\mathbf{D}_2\|_2}\right)$ . If  $\|\gamma_2^{k+1} - \gamma_2^k\|_2 \leq t\varepsilon$ , then

$$F(\gamma_2^{k+1}) - F_{\text{opt}} \leq \eta \varepsilon + (\beta + \kappa t) \mu,$$

where  $\eta$ ,  $\beta$  and  $\kappa$  are those given in (2.2).

An additional consequence of Theorem 2.2 is an analogous result for ML-FISTA. Recall that ML-FISTA introduces a momentum term to the update provided by ML-ISTA, and can be written as

$$\begin{aligned} \gamma_2^{k+1} &= \text{prox}_{t g_2} \left( \mathbf{z}^k - t \mathbf{D}_2^T G_{1/\mu}^{f, g_1}(\mathbf{D}_2 \mathbf{z}^k) \right), \\ \mathbf{z}^{k+1} &= \gamma_2^{k+1} + \rho^k (\gamma_2^{k+1} - \gamma_2^k). \end{aligned} \quad (13)$$

We have the following result.

**Corollary 2.2.2.** Let  $\mu \in \left(0, \frac{1}{\|\mathbf{Q}\|_2}\right)$  and  $t \in \left(0, \frac{4\mu}{3\|\mathbf{D}_2\|_2}\right)$ , and assume that  $\{\gamma_2^k\}$  and  $\{\mathbf{z}^k\}$  are the sequences generated by ML-FISTA according to Equation (2.3). If

$$\|\mathbf{z}^k - \gamma_2^{k+1}\|_2 \leq t\varepsilon,$$

then

$$F(\gamma_2^{k+1}) - F_{\text{opt}} \leq \eta \varepsilon + (\beta + \kappa t) \mu,$$

where the constants  $\eta$ ,  $\beta$  and  $\kappa$  are defined in (2.2).

Before moving on, let us comment on the significance of these results. On the one hand, we are unaware of any results for proximal gradient-mapping algorithms, and in this sense, the analysis above presents a first result of its kind. On the other hand, the analysis does not provide a convergence rate, and so they do not reflect any benefits of ML-FISTA over ML-ISTA. As we will see shortly, the empirical convergence of these methods significantly differ in practice.

## 2.4 Synthetic Experiments

We now carry a series of synthetic experiments to demonstrate the effectiveness of the multi-layer Basis Pursuit problem, as well as the proposed iterative shrinkage methods.

First, we would like to illustrate the benefit of the proposed multi-layer BP formulation when compared to the traditional sparse regression problem. In other words, exploring the benefit of having  $\lambda_1 > 0$ . To this end, we construct a two layer model with Gaussian matrices  $\mathbf{D}_1 \in \mathbb{R}^{n \times m_1}$  and  $\mathbf{D}_2 \in \mathbb{R}^{m_1 \times m_2}$ , ( $n = 50, m_1 = 70, m_2 = 60$ ). We construct our signals by obtaining representation  $\gamma_2$  with  $\|\gamma_2\|_0 = 30$  and  $\|\gamma_1\|_0 = 42$ , following the procedure described in [1], so as to provide representations consistent with the multi-layer sparse model. Lastly, we contaminate the signals with Gaussian i.i.d. noise creating

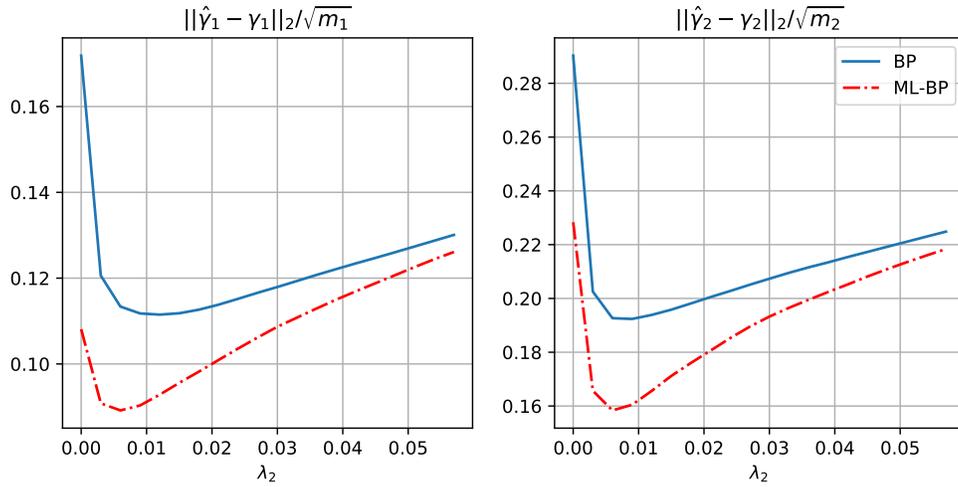


Figure 1. Recovery error for  $\gamma_1$  and  $\gamma_2$  employing BP ( $\lambda_1 = 0$ ) and Multi-Layer BP ( $\lambda_1 > 0$ ).

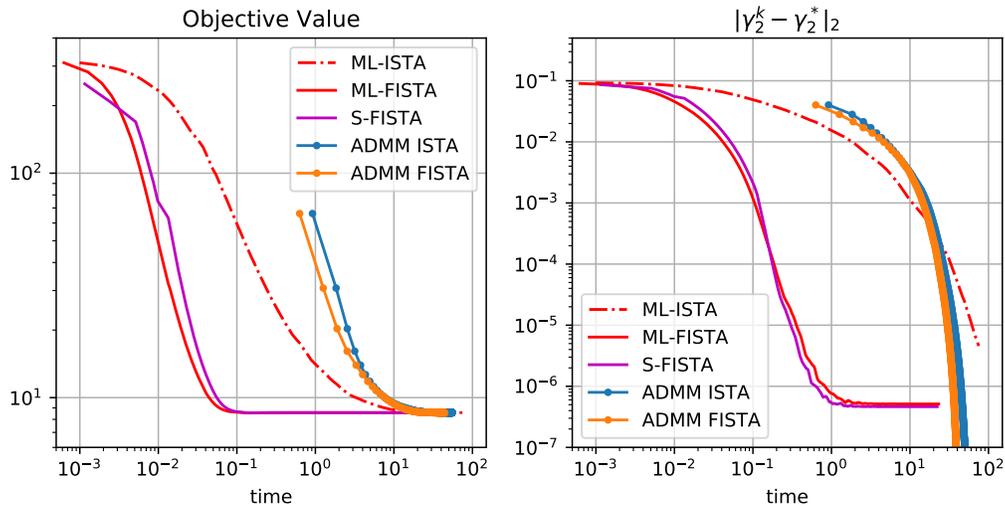


Figure 2. Comparison of different solvers for Multi-Layer Basis Pursuit in terms of objective value (left) and distance to optimal solution (right).

the measurements  $\mathbf{y} = \mathbf{x} + \mathbf{w}$  with SNR=10. We compare minimizing  $(P)$  with  $\lambda_1 = 0$  (which accounts to solving a classical BP problem) with the case when  $\lambda_1 > 0$ , as a function of  $\lambda_2$  when solved with ISTA and ML-ISTA. As can be seen from the results in Figure 1, enforcing the additional analysis penalty on the intermediate representation can indeed prove advantageous and provide a lower recovery error in both  $\gamma_2$  and  $\gamma_1$ . Clearly, this is not true for any value of  $\lambda_1$ , as the larger this parameter becomes, the larger the bias will be in the resulting estimate. For the sake of this demonstration we have set  $\lambda_1$  as the optimal value for each  $\lambda_2$  (with grid search). The theoretical study of the conditions (in terms of the model parameters) under which  $\lambda_1 > 0$  provides a better recovery, and how to determine this parameter in practice, are interesting questions that we defer to future work.

We also employ this synthetic setup to illustrate the convergence properties of the main algorithms presented above: ADMM (employing either ISTA or FISTA for the

inner BP problem), the Smooth-FISTA [5] and the proposed Multi-Layer ISTA and Multi-Layer FISTA. Once again, we illustrate these algorithms for the optimal choice of  $\lambda_1$  and  $\lambda_2$  from the previous experiment, and present the results in Figure 2. We measure the convergence in terms of function value for all methods, as well as the convergence to the solution found with ADMM run until convergence.

ADMM converges in relatively few iterations, though these take a relatively long time due to the inner BP solver. This time is reduced when using FISTA rather than ISTA (as the inner solver converges faster), but it is still significantly slower than any of the other alternatives - even while using warm-start at every iteration, which we employ in these experiments.

Both S-FISTA and our multi layer solvers depend on a parameter that controls the accuracy of their solution and affect their convergence speed – the  $\epsilon$  for the former, and the  $\mu$  for the latter approaches. We set these parameters so as to obtain roughly the same accuracy in terms of recovery

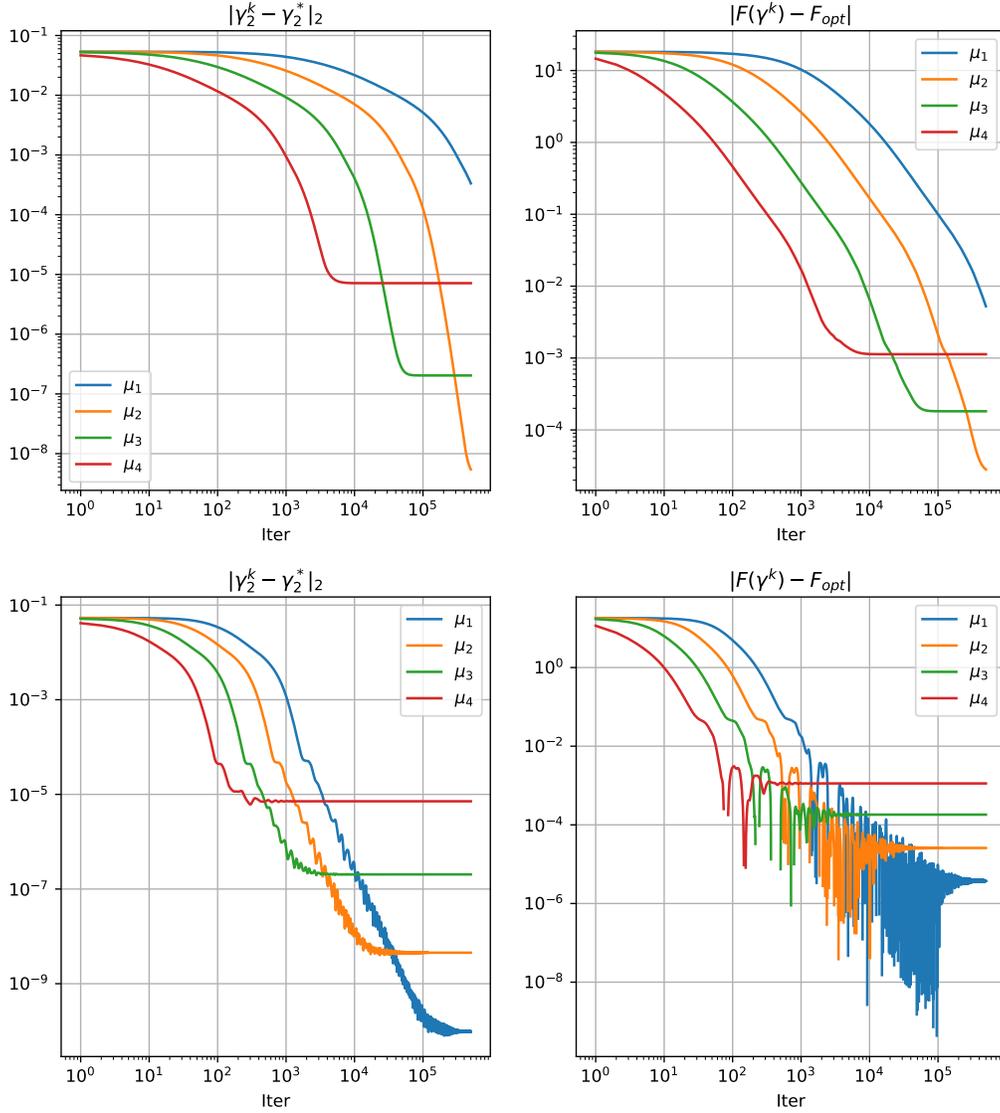


Figure 3. ML-ISTA (top) and ML-FISTA (bottom) evaluation for different values of the parameter  $\mu$ .

error, and compare their convergence behavior. We can see that ML-FISTA is clearly faster than ML-ISTA, and slightly slightly faster than S-FISTA. Lastly, in order to demonstrate the effect of  $\mu$  in ML-ISTA and ML-FISTA, we run the same algorithms for different values of this parameter (in decreasing order and equispaced in logarithmic scale between  $10^{-2.5}$  and 1) for the same setting, and present the results in Figure 3 for ML-ISTA (top) and ML-FISTA (bottom). These numerical results illustrate the theoretical analysis provided by Theorem 2.2 in that the smaller  $\mu$ , the more accurate the solution becomes, albeit requiring more iterations to converge. These results also reflect the limitation of our current theoretical analysis, which is incapable of providing insights into the convergence rate.

### 3 PRINCIPLED RECURRENT NEURAL NETWORKS

As seen above, the ML-ISTA and ML-FISTA schemes provide efficient solvers for problem (P). Interestingly, if one considers the first iteration of either of the algorithms (with

$\gamma_L^0 = \mathbf{0}$ ), the update of the inner most representation results in

$$\gamma_2 \leftarrow \frac{t}{\mu} \mathcal{T}_{t\lambda_2} \left( \mathbf{D}_2^T \mathcal{T}_{\mu\lambda_1} (\mu \mathbf{D}_1^T (\mathbf{y})) \right), \quad (14)$$

for a two-layer model, for instance. If one further imposes a non-negativity assumption on the representation coefficients, the thresholding operators  $\mathcal{T}_\lambda$  become non-negative projections shifted by a bias of  $\lambda$ . Therefore, the above soft-thresholding operation can be equivalently written as

$$\gamma_2 \leftarrow \text{ReLU} \left( \mathbf{D}_2^T \text{ReLU} (\mathbf{D}_1^T \mathbf{y} + \mathbf{b}_1) + \mathbf{b}_2 \right)$$

where the biases vectors  $\mathbf{b}_1$  and  $\mathbf{b}_2$  account for the corresponding thresholds<sup>10</sup>. Just as pointed out in [34], this is simply the forward pass in a neural network. Moreover, all the analysis presented above holds also in the case of convo-

<sup>10</sup>. Note that this expression is more general in that it allows for different thresholds per atom, as opposed the expression in (3). The latter can be recovered by setting every entry in the bias vector to be  $\lambda_i$ .

lutional dictionaries, where the dictionary atoms are nothing but convolutional filters (transposed) in a convolutional neural network. Could we benefit from this observation to improve on the performance of CNNs?

In this section, we intend to demonstrate how, by interpreting neural networks as approximation algorithms of the solution to a Multi-Layer BP problem, one can boost the performance of typical CNNs *without introducing any parameters in the model*. To this end, we will first impose a generative model on the features  $\gamma_L$  in terms of multi-layer sparse convolutional representations; i.e., we assume that  $\mathbf{y} \approx \mathbf{D}_{(1,L)}\gamma_L$ , for convolutional dictionaries  $\mathbf{D}_i$ . Furthermore, we will adopt a supervised learning setting in which we attempt to minimize an empirical risk over  $N$  training samples of signals  $\mathbf{y}_i$  with labels  $h_i$ . A classifier  $\zeta_\theta(\gamma^*)$ , with parameters  $\theta$ , will be trained on estimates of said features  $\gamma^*(\mathbf{y})$  obtained as the solution of the ML-BP problem; i.e.

$$\min_{\theta, \{\mathbf{D}_i, \lambda_i\}} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(h_i, \zeta_\theta(\gamma^*)) \quad \text{s.t.}$$

$$\gamma^* = \arg \min_{\gamma} \|\mathbf{y} - \mathbf{D}_{(1,L)}\gamma\|_2^2 + \sum_{i=1}^{L-1} \lambda_i \|\mathbf{D}_{(i+1,L)}\gamma\|_1 + \lambda_L \|\gamma\|_1. \quad (15)$$

The function  $\mathcal{L}$  is a loss or cost function to be minimized during training, such as the cross entropy which we employ for the classification case. Our approach to address this bi-level optimization problem is to approximate the solution of the lower-level problem by  $k$  iterations of the ML-ISTA approaches – effectively implemented as  $k$  layers of unfolded recurrent neural networks. This way,  $\gamma^*$  becomes a straight-forward function of  $\mathbf{y}$  and the model parameters ( $\mathbf{D}_i$  and  $\lambda_i$ ), which can be plugged into the loss function  $\mathcal{L}$ . A similar approach is employed by Task Driven Dictionary Learning [29] in which the constraint is a single layer BP (i.e.  $L = 1$ ) that is solved with LARS [18] until convergence, resulting in a more involved algorithm.

Importantly, if only one iteration is employed for the ML-ISTA, and a linear classifier<sup>11</sup> is chosen for  $\zeta_\theta(\gamma^*)$ , the problem in (3) boils down exactly to training a CNN to minimize the classification loss  $\mathcal{L}$ . Naturally, when considering further iterations of the multi-layer pursuit, one is effectively implementing a recurrent neural network with “skip connections”, as depicted in Figure 4 for a two-layer model. These extended networks, which can become very deep, have exactly as many parameters as their traditional *forward-pass* counterparts – namely, the dictionaries  $\mathbf{D}_i$ , biases  $\lambda_i$  and classifier parameters  $\theta$ . Notably, and unlike other popular constructions in the deep learning community (e.g., Residual Neural Networks [22], DenseNet [25], and other similar constructions), these recurrent components and connections follow a precise optimization justification.

The concept of unfolding an iterative sparse coding algorithm is clearly not new. The first instance of such an idea was formalized by the Learned ISTA (LISTA) approach [20]. LISTA decomposes the linear operator of ISTA in terms of 2 matrices, replacing the computation of

$$\mathcal{T}_\lambda(\gamma - \eta \mathbf{D}^T(\mathbf{D}\gamma - \mathbf{y})) \text{ by}$$

$$\mathcal{T}_\lambda(\mathbf{W}\gamma + \mathbf{B}\mathbf{y}), \quad (16)$$

following the equivalences  $\mathbf{W} = \mathbf{I} - \eta \mathbf{D}^T \mathbf{D}$  and  $\mathbf{B} = \mathbf{D}^T$ . Then, it adaptively learns these new operators instead of the initial dictionary  $\mathbf{D}$  in order to provide estimates  $\hat{\gamma}$  that approximate the solution of ISTA. Interestingly, such a decomposition allows for the acceleration of ISTA [32], providing an accurate estimate in very few iterations. A natural question is, then, could we propose an analogous multi-layer Learned ISTA?

There are two main issues that need to be resolved if one is to propose a LISTA-like decomposition in the framework of our multi-layer pursuits. The first one is that the decomposition in (3) has been proposed and analyzed for general matrices (i.e., fully-connected layers in a CNN context), but not for convolutional dictionaries. If one was to naively propose to learn such an (unconstrained) operator  $\mathbf{W}$ , this would result in an enormous amount of added parameters. To resolve this point, in the case where  $\mathbf{D}$  is a convolutional dictionary (as in CNNs) we propose a decomposition of the form

$$\mathcal{T}_\lambda((\mathbf{I} - \mathbf{W}^T \mathbf{W})\gamma + \mathbf{B}\mathbf{y}),$$

where  $\mathbf{W}$  is also constrained to be convolutional, thus controlling the number of parameters<sup>12</sup>. In fact, the number of parameters in a layer of this ML-LISTA is simply twice as many parameters as the conventional case, since the number of convolutional filters in  $\mathbf{W}$  and  $\mathbf{B}$  (and their dimensions) are equal to those in  $\mathbf{D}$ .

The second issue is concerned with the fact that LISTA was proposed as a relaxation of ISTA – a pursuit tackling a *single layer* pursuit problem. To accommodate a similar decomposition in our multi-layer setting, we naturally extend the update to:

$$\hat{\gamma}_1 \leftarrow \mathcal{T}_{\lambda_1}((\mathbf{I} - \mathbf{W}_1^T \mathbf{W}_1)\gamma_1^k + \mathbf{B}_1 \mathbf{y}), \quad (17)$$

$$\gamma_2^{k+1} \leftarrow \mathcal{T}_{\lambda_2}((\mathbf{I} - \mathbf{W}_2^T \mathbf{W}_2)\gamma_2^k + \mathbf{B}_2 \hat{\gamma}_1),$$

for a two-layer model for simplicity. In the context of the supervised classification setting, the learning of the dictionaries  $\mathbf{D}_i$  is replaced by learning the operators  $\mathbf{W}_i$  and  $\mathbf{B}_i$ . Note that this decomposition prevents us from obtaining the dictionaries  $\mathbf{D}_i$ , and so we use<sup>13</sup>  $\gamma_1^k = \mathbf{B}_2^T \gamma_2^k$  in Equation (3).

## 4 EXPERIMENTS

In this final section, we show how the presented algorithms can be used for image classification on three common datasets: MNIST, SVHN and CIFAR10, while improving the performance of CNNs *without introducing any extra parameters in the model*. Recalling the learning formulation in Equation (3), we will compare different architectures resulting from different solvers for the features  $\gamma^*$ . As employing only one iteration of the proposed algorithms

12. For completeness, we have also tested the *traditional* decomposition proposed in Equation (3), resulting in worse performance than that of ML-ISTA – likely due to the significant increase in the number of parameters discussed above.

13. An alternative is to employ  $\gamma_1^k = \mathbf{W}_2 \gamma_2^k$ , but this choice was shown to perform slightly worse in practice.

11. Or, in fact, any other neural-network-based classifier acting on the obtained features.

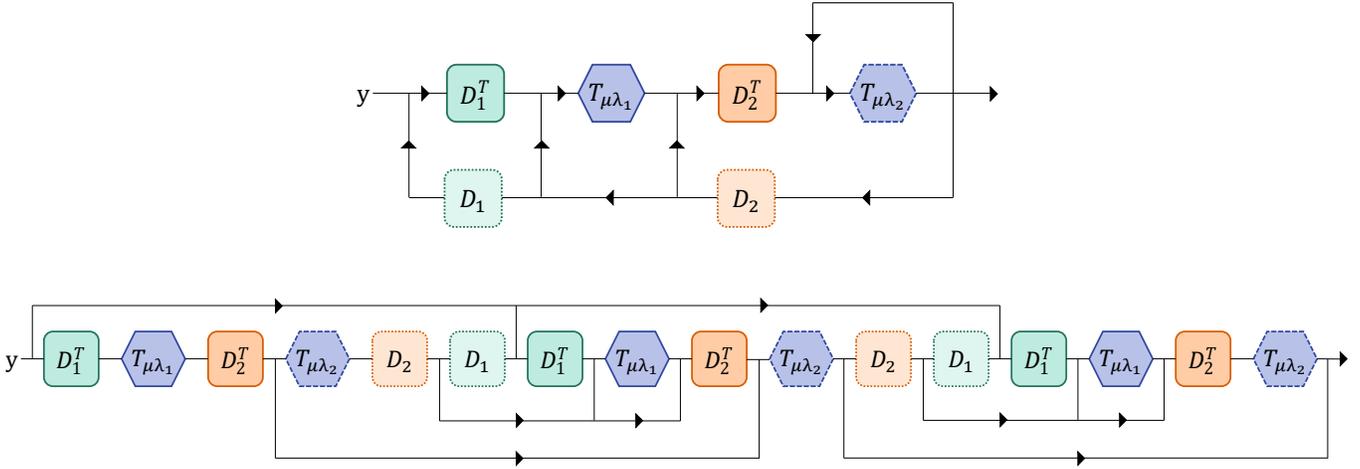


Figure 4. ML-ISTA graph interpretation for a two layer model as a recurrent neural network (top), and its unfolded version for 2 iterations (bottom).

recovers a traditional feed-forward network, we will employ such a basic architecture as our baseline and compare it with the Multi Layer ISTA and FISTA, for different number of iterations or unfoldings. Also for this reason, we deliberately avoid using training “tricks” popular in the deep learning community, such as batch normalization, drop-out, etc., so as to provide clear experimental setups that facilitate the understanding and demonstration of the presented ideas.

For the MNIST case, we construct a standard (LeNet-style) CNN with 3 convolutional layers (i.e., dictionaries) with 32, 64 and 512 filters, respectively<sup>14</sup>, and a final fully-connected layer as the classifier  $\zeta(\gamma^*)$ . We also enforce non-negativity constraints on the representations, resulting in the application of ReLUs and biases as shrinkage operators. For SVHN we use an analogous model, though with three input channels and slightly larger filters to accommodate the larger input size. For CIFAR, we define a ML-CSC model with 3 convolutional layers, and the classifier function  $\zeta(\gamma^*)$  as a three-layer CNN. This effectively results in a 6 layers architecture, out of which the first three are unfolded in the context of the multi-layer pursuits. All models are trained with SGD with momentum, decreasing the learning rate every so many iterations. In particular, we make use of a PyTorch implementation, and training code is made available<sup>15</sup> online.

In order to demonstrate the effect of the ML-ISTA iterations (or unfoldings), we first depict the test error as a function of the training epochs for different number of such iterations in Figure 5. Recall that the case of 0 unfoldings corresponds to the typical feed-forward CNN, while the case with 6 unfoldings effectively implements a 18-layers-deep architecture, alas having the same number of parameters. As can be seen, further unfoldings improve on the resulting performance.

Moving to a more complete comparison, we demonstrate the ML-ISTA and ML-FISTA architectures when compared to some of the models mentioned above; namely:

14. Kernel sizes of  $6 \times 6$ ,  $6 \times 6$  and  $4 \times 4$ , respectively, with stride of 2 in the first two layers.

15. Available through the first author’s website.

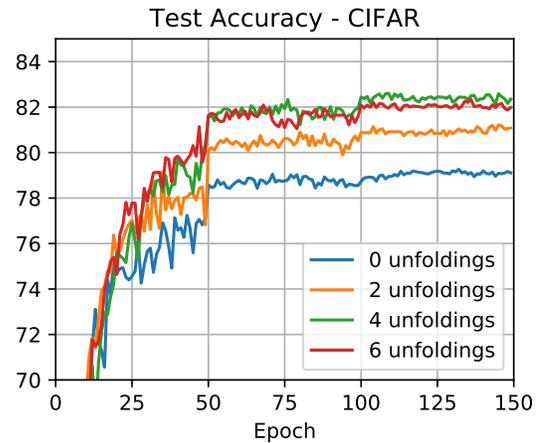


Figure 5. Training ML-ISTA for different number of unfoldings, on CIFAR10. The case of 0 unfoldings corresponds to the traditional feed-forward convolutional network. All networks have the same number of parameters.

- *ML-LISTA*: replacing the learning of the convolutional dictionaries (or filters) by the learning of the (convolutional) factors  $\mathbf{W}_i$  and  $\mathbf{B}_i$ , as indicated in Equation (3).
- *Layered Basis Pursuit*: the approach proposed in [34], which unrolls the iteration of ISTA for a single-layer BP problem *at each layer*. In contrast, the proposed ML-ISTA/FISTA unrolls the iterations of the entire Multi-Layer BP problem.
- An “*All-Free*” model: What if one ignores the generative model (and the corresponding pursuit interpretation) and simply frees all the filters to be adaptively learned? In order to study this question, we train a model with the same depth and an analogous recurrent architecture as the unfolded ML-ISTA/FISTA networks, but where all the filters of the different layers are free to be learned and to provide the best possible performance.

It is worth stressing that the ML-ISTA, ML-FISTA and Lay-

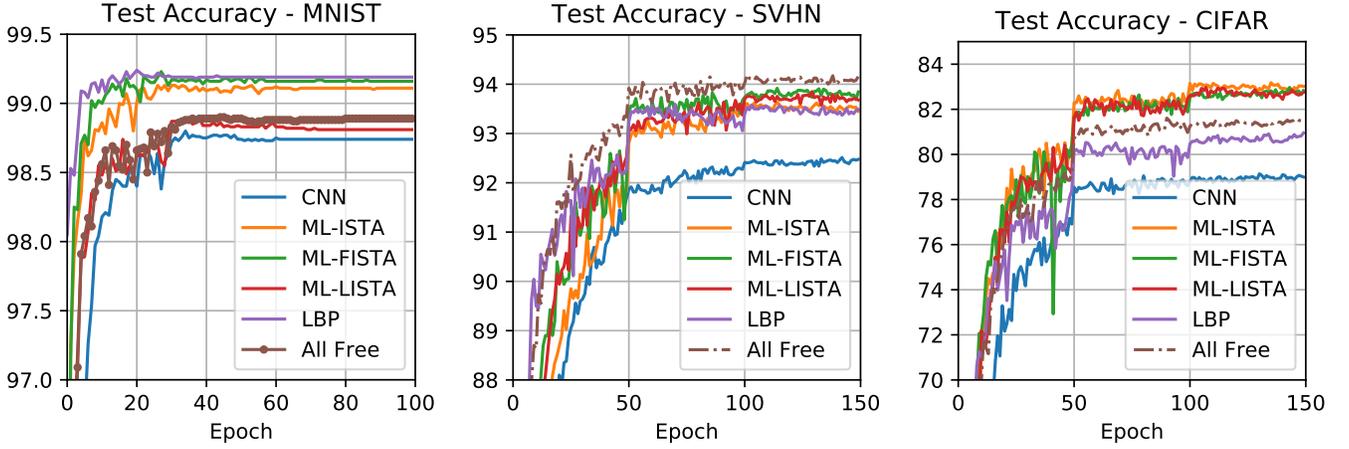


Figure 6. Comparison of different architectures on the SVHN dataset, with a feed-forward network as baseline. All networks have the same number of parameters.

ered BP have all the same number of parameters as the feed-forward CNN. The ML-LISTA version has twice as many parameters, while the All-Free version has order  $\mathcal{O}(LK)$  more parameters, where  $L$  is the number of layers and  $K$  is the number of unfoldings.

The accuracy as a function of the iterations for all models are presented in Figure 6, and the final results are detailed in Table 1. A first observation is that most “unrolled” networks provide an improvement over the baseline feed-forward architecture. Second, while the Layered BP performs very well on MNIST, it falls behind on the other two more challenging datasets. Recall that while this approach unfolds the iterations of a pursuit, it does so one layer at a time, and does not address a global pursuit problem as the one we explore in this work.

Third, the performances of the ML-ISTA, ML-FISTA and ML-LISTA are comparable. This is interesting, as the LISTA-type decomposition does not seem to provide an importance advantage over the unrolled multi-layer pursuits. Forth, and most important of all, freeing all the parameters in the architecture does not provide important improvements over the ML-ISTA networks. Limited training data is not likely to be the cause, as ML-ISTA/FISTA outperforms the larger model even for CIFAR, which enjoys a rich variability in the data and while using data-augmentation. This is noteworthy, and this result seems to indicate that the consideration of the multi-layer sparse model, and the resulting pursuit, does indeed provide an (approximate) solution to the problem behind CNNs.

## 5 PROOFS OF MAIN THEOREMS

### 5.1 Fixed Point Analysis

A vector  $\gamma_2^*$  is a fixed point of the ML-ISTA update from Equation (2.2) iff

$$\gamma_2^* = \text{prox}_{t g_2} \left( \gamma_2^* - t \mathbf{D}_2^T G_{1/\mu}^{f, g_1}(\mathbf{D}_2 \gamma_2^*) \right).$$

By the second prox theorem [3, Theorem 6.39], we have that

$$-t \mathbf{D}_2^T G_{1/\mu}^{f, g_1}(\mathbf{D}_2 \gamma_2^*) \in t \partial g_2(\gamma_2^*),$$

Model	MNIST	SVHN	CIFAR 10
Feed-Forward	98.78 %	92.44 %	79.00 %
Layered BP	99.19 %	93.42 %	80.73 %
ML-ISTA	99.10 %	93.52 %	82.93 %
ML-FISTA	99.16 %	93.79 %	82.79 %
ML-LISTA	98.81 %	93.71 %	82.68 %
All-Free	98.89 %	94.06 %	81.48 %

Table 1  
Classification results for different architectures for MNIST, SVHN and CIFAR10.

or, equivalently, there exists  $\mathbf{w}_2 \in \partial g_2(\gamma_2^*)$  so that

$$\mathbf{D}_2^T G_{1/\mu}^{f, g_1}(\mathbf{D}_2 \gamma_2^*) + \mathbf{w}_2 = 0.$$

Employing the definition of  $G_{1/\mu}^{f, g_1}(\mathbf{D}_2 \gamma_2^*)$ ,

$$\mathbf{D}_2^T \frac{1}{\mu} \left( \mathbf{D}_2 \gamma_2^* - \text{prox}_{\mu g_1}(\mathbf{D}_2 \gamma_2^* - \mu \nabla f(\mathbf{D}_2 \gamma_2^*)) \right) + \mathbf{w}_2 = 0. \quad (18)$$

Next, denote

$$\hat{\gamma}_1 = \text{prox}_{\mu g_1}(\mathbf{D}_2 \gamma_2^* - \mu \nabla f(\mathbf{D}_2 \gamma_2^*)). \quad (19)$$

Employing the second prox theorem on (5.1), we have that the above is equivalent to the existence of  $\mathbf{w}_1 \in \partial g_1(\hat{\gamma}_1)$  for which  $\mathbf{D}_2 \gamma_2^* - \mu \nabla f(\mathbf{D}_2 \gamma_2^*) - \hat{\gamma}_1 = \mu \mathbf{w}_1$ . Thus, (5.1) amounts to

$$\mathbf{D}_2^T \frac{1}{\mu} (\mathbf{D}_2 \gamma_2^* - \mathbf{D}_2 \gamma_2^* + \mu \nabla f(\mathbf{D}_2 \gamma_2^*) + \mu \mathbf{w}_1) + \mathbf{w}_2 = 0,$$

for some  $\mathbf{w}_2 \in \partial g_2(\gamma_2^*)$  and  $\mathbf{w}_1 \in \partial g_1(\hat{\gamma}_1)$ . Simplifying the expression above we arrive at the fixed-point condition of ML-ISTA, which is

$$\begin{aligned} \exists \mathbf{w}_2 \in \partial g_2(\gamma_2^*), \mathbf{w}_1 \in \partial g_1(\hat{\gamma}_1) \\ \text{so that } \mathbf{D}_2^T \nabla f(\mathbf{D}_2 \gamma_2^*) + \mathbf{D}_2^T \mathbf{w}_1 + \mathbf{w}_2 = 0. \end{aligned}$$

## 5.2 Proof of Lemma 2.1

*Proof.* Denote  $S_\mu(\gamma_1) = \text{prox}_{\mu g_1}(\gamma_1 - \mu \nabla f(\gamma_1))$ . Then, by the second prox theorem ([3, Theorem 6.39]), we have that

$$\gamma_1 - \mu \nabla f(\gamma_1) - S_\mu(\gamma_1) \in \mu \partial g_1(S_\mu(\gamma_1)).$$

Dividing by  $\mu$  and employing the definition of the gradient mapping, we obtain

$$G_{1/\mu}^{f,g_1}(\gamma_1) \in \nabla f(\gamma_1) + \partial g_1(S_\mu(\gamma_1)).$$

By the  $\ell_{g_1}$ -Lipschitz continuity of  $g_1$  [3, Theorem 3.61], it follows that  $\|\mathbf{z}\|_2 \leq \ell_{g_1}$  for any  $\mathbf{z} \in \partial g_1(S_\mu(\gamma_1))$ . This, combined with the bound  $\|\nabla f(\gamma_1)\|_2 \leq M$  from Equation (2.3), provides the desired claim.  $\square$

## 5.3 Proof of Theorem 2.2

*Proof.* Denote

$$\begin{aligned} \mathbf{a}_1 &= \frac{1}{t} \left[ \tilde{\gamma}_2 - \text{prox}_{t g_2} \left( \tilde{\gamma}_2 - t \mathbf{D}_2^T (\mathbf{I} - \mu \mathbf{Q}) G_{1/\mu}^{f,g_1}(\tilde{\gamma}_1) \right) \right], \\ \mathbf{a}_2 &= \frac{1}{t} \left[ \tilde{\gamma}_2 - \text{prox}_{t g_2} \left( \tilde{\gamma}_2 - t \mathbf{D}_2^T G_{1/\mu}^{f,g_1}(\tilde{\gamma}_1) \right) \right]. \end{aligned}$$

By the triangle inequality,

$$\|\mathbf{a}_1\|_2 \leq \|\mathbf{a}_2\|_2 + \|\mathbf{a}_1 - \mathbf{a}_2\|_2. \quad (20)$$

We will upper-bound the right-hand side of this inequality. First, employing the non-expansiveness property of prox operators ([14, Lemma 2.4]), we can write

$$\begin{aligned} \|\mathbf{a}_1 - \mathbf{a}_2\|_2 &= \frac{1}{t} \left\| \text{prox}_{t g_2} \left( \tilde{\gamma}_2 - t \mathbf{D}_2^T G_{1/\mu}^{f,g_1}(\tilde{\gamma}_1) \right) \right. \\ &\quad \left. - \text{prox}_{t g_2} \left( \tilde{\gamma}_2 - t \mathbf{D}_2^T (\mathbf{I} - \mu \mathbf{Q}) G_{1/\mu}^{f,g_1}(\tilde{\gamma}_1) \right) \right\|_2 \\ &\leq \left\| \mathbf{D}_2^T (\mathbf{I} - \mu \mathbf{Q}) G_{1/\mu}^{f,g_1}(\tilde{\gamma}_1) - \mathbf{D}_2^T G_{1/\mu}^{f,g_1}(\tilde{\gamma}_1) \right\|_2 \\ &= \mu \left\| \mathbf{D}_2^T \mathbf{Q} G_{1/\mu}^{f,g_1}(\tilde{\gamma}_1) \right\|_2 \\ &\leq \mu \|\mathbf{D}_2\|_2 \|\mathbf{Q}\|_2 (M + \ell_{g_1}), \end{aligned} \quad (21)$$

where the last inequality follows from the definition of operator norms and Lemma 2.1. Also,  $\|\mathbf{a}_2\|_2 \leq \varepsilon$  by assumption. Thus, from (5.3),

$$\|\mathbf{a}_1\|_2 \leq \varepsilon + \mu \|\mathbf{D}_2\|_2 \|\mathbf{Q}\|_2 (M + \ell_{g_1}). \quad (22)$$

Consider now the function  $H_\mu : \mathbb{R}^{m_1} \rightarrow \mathbb{R}$  given by

$$\begin{aligned} H_\mu(\gamma) &= \frac{1}{2} \gamma^T (\mathbf{Q} - \mu \mathbf{Q}^2) \gamma + \mathbf{b}^T (\mathbf{I} - \mu \mathbf{Q}) \gamma \\ &\quad + M_{g_1}^\mu ((\mathbf{I} - \mu \mathbf{Q}) \gamma - \mu \mathbf{b}), \end{aligned}$$

where  $M_{g_1}^\mu$  is the Moreau envelope of  $g_1$  with smoothness parameter  $\mu$  [31]. Note that  $H_\mu$  is convex since  $\mu < \frac{1}{\|\mathbf{Q}\|_2}$  implies  $\mathbf{Q} - \mu \mathbf{Q}^2 \succeq 0$  and the Moreau envelope of a convex function is convex. Recall that the gradient of the Moreau envelope is given by  $\nabla M_{g_1}^\mu(\gamma) = \frac{1}{\mu} (\gamma - \text{prox}_{\mu g_1}(\gamma))$  (see e.g. [3, Theorem 6.60]), and so

$$\begin{aligned} \nabla H_\mu(\gamma) &= \frac{1}{\mu} (\mathbf{I} - \mu \mathbf{Q}) \left[ \gamma - \text{prox}_{\mu g_1}(\gamma - \mu (\mathbf{Q} \gamma + \mathbf{b})) \right] \\ &= (\mathbf{I} - \mu \mathbf{Q}) G_{1/\mu}^{f,g_1}(\gamma). \end{aligned}$$

Consider now  $\gamma_1 = \mathbf{D}_2 \gamma_2$  and  $\tilde{H}_\mu(\gamma_2) \equiv H_\mu(\mathbf{D}_2 \gamma_2)$ . Applying the chain rule yields

$$\nabla \tilde{H}_\mu(\gamma_2) = \mathbf{D}_2^T (\mathbf{I} - \mu \mathbf{Q}) G_{1/\mu}^{f,g_1}(\gamma_1).$$

Thus, we can conclude that  $\mathbf{a}_1$  is nothing else than the gradient mapping of  $\tilde{H}_\mu$  and  $g_2$ , and so the inequality in (5.3) can be rewritten as

$$\left\| G_{1/t}^{\tilde{H}_\mu, g_2}(\tilde{\gamma}_2) \right\|_2 \leq \varepsilon + \mu \|\mathbf{D}_2\|_2 \|\mathbf{Q}\|_2 (M + \ell_{g_1}). \quad (23)$$

Gradient mapping operators are firmly non-expansive with constant  $\frac{3\mu}{4}$  ([3, Lemma 10.11]), from which it follows that  $H_\mu$  is  $\frac{4}{3\mu}$ -smooth. Denote  $F_\mu(\gamma_2) = H_\mu(\mathbf{D}_2 \gamma_2) + g_2(\gamma_2)$ , and one of its minimizers by  $\gamma_\mu^* \in \arg \min F_\mu(\gamma_2)$ . Moreover, define

$$\hat{\gamma} = \text{prox}_{t g_2} \left( \tilde{\gamma}_2 - t \mathbf{D}_2^T (\mathbf{I} - \mu \mathbf{Q}) G_{1/\mu}^{f,g_1}(\tilde{\gamma}_1) \right).$$

By the fundamental prox-grad inequality ([3, Theorem 10.16]), and since  $t \in \left(0, \frac{4\mu}{3\|\mathbf{D}_2\|_2}\right)$ , it follows that

$$F_\mu(\gamma_\mu^*) - F_\mu(\hat{\gamma}) \geq \frac{1}{2t} \|\gamma_\mu^* - \hat{\gamma}\|_2^2 - \frac{1}{2t} \|\gamma_\mu^* - \tilde{\gamma}_2\|_2^2.$$

Then, by the three-points lemma (see [12]), we may rewrite

$$\|\gamma_\mu^* - \hat{\gamma}\|_2^2 - \|\gamma_\mu^* - \tilde{\gamma}_2\|_2^2 = 2 \langle \gamma_\mu^* - \hat{\gamma}, \tilde{\gamma}_2 - \hat{\gamma} \rangle - \|\tilde{\gamma}_2 - \hat{\gamma}\|_2^2.$$

Thus,

$$\begin{aligned} F_\mu(\hat{\gamma}) - F_\mu(\gamma_\mu^*) &\leq \frac{1}{t} \langle \hat{\gamma} - \gamma_\mu^*, \tilde{\gamma}_2 - \hat{\gamma} \rangle + \frac{1}{2t} \|\tilde{\gamma}_2 - \hat{\gamma}\|_2^2 \\ &= \frac{1}{t} \langle \hat{\gamma} - \tilde{\gamma}_2, \tilde{\gamma}_2 - \hat{\gamma} \rangle \\ &\quad + \frac{1}{t} \langle \tilde{\gamma}_2 - \gamma_\mu^*, \tilde{\gamma}_2 - \hat{\gamma} \rangle + \frac{1}{2t} \|\tilde{\gamma}_2 - \hat{\gamma}\|_2^2 \\ &= -\frac{1}{2t} \|\tilde{\gamma}_2 - \hat{\gamma}\|_2^2 + \frac{1}{t} \langle \tilde{\gamma}_2 - \gamma_\mu^*, \tilde{\gamma}_2 - \hat{\gamma} \rangle \\ &\leq \langle \tilde{\gamma}_2 - \gamma_\mu^*, G_{1/t}^{\tilde{H}_\mu, g_2}(\tilde{\gamma}_2) \rangle \\ &\leq 2R \|G_{1/t}^{\tilde{H}_\mu, g_2}(\tilde{\gamma}_2)\|_2, \end{aligned}$$

where the last passage uses the Cauchy-Schwarz inequality along with  $\|\tilde{\gamma}_2 - \gamma_\mu^*\|_2 \leq 2R$ . Combining the above with Inequality (5.3) yields

$$F_\mu(\hat{\gamma}) - F_\mu(\gamma_\mu^*) \leq 2R\varepsilon + 2\mu R \|\mathbf{Q}\|_2 \|\mathbf{D}_2\|_2 (M + \ell_{g_1}). \quad (24)$$

Finally, we will connect between  $F_\mu(\gamma_\mu^*)$ ,  $F_\mu(\tilde{\gamma}_2)$  and  $F(\gamma_\mu^*)$ ,  $F(\tilde{\gamma}_2)$ , respectively. Note that for any  $\gamma_2 \in \text{dom}(g_2)$  and  $\gamma_1 = \mathbf{D}_2 \gamma_2$ ,

$$\begin{aligned} \left| H_\mu(\gamma_1) - f(\gamma_1) - g_1(\gamma_1) \right| &= \left| -\mu \left( \frac{1}{2} \|\mathbf{Q} \gamma_1\|_2^2 + \mathbf{b}^T \mathbf{Q} \gamma_1 \right) \right. \\ &\quad \left. + M_{g_1}^\mu ((\mathbf{I} - \mu \mathbf{Q}) \gamma_1 - \mu \mathbf{b}) - g_1(\gamma_1) \right|. \end{aligned}$$

Moreover, this expression can be upper-bounded by

$$\begin{aligned} &\left( \frac{1}{2} \|\mathbf{Q}\|_2^2 R_1^2 + \|\mathbf{b}\|_2 \|\mathbf{Q}\|_2 R_1 \right) \mu \\ &\quad + \left| M_{g_1}^\mu ((\mathbf{I} - \mu \mathbf{Q}) \gamma_1 - \mu \mathbf{b}) - g_1(\gamma_1) \right|. \end{aligned}$$

Further, from basic properties of the Moreau envelope (in particular Theorem 10.51 in [3]) and the  $\ell_{g_1}$ -Lipschitz prop-

erty of  $g_1$ , we have that

$$\begin{aligned} & |M_{g_1}^\mu((\mathbf{I} - \mu\mathbf{Q})\gamma_1 - \mu\mathbf{b}) - g_1(\gamma)| \leq \\ & |M_{g_1}^\mu((\mathbf{I} - \mu\mathbf{Q})\gamma_1 - \mu\mathbf{b}) - g_1((\mathbf{I} - \mu\mathbf{Q})\gamma_1 - \mathbf{b})| \\ & \quad + |g_1((\mathbf{I} - \mu\mathbf{Q})\gamma_1 - \mathbf{b}) - g_1(\gamma_1)| \\ & \leq \frac{\ell_{g_1}^2}{2}\mu + \ell_{g_1}M\mu. \end{aligned}$$

We thus obtain that for any  $\gamma_2 \in \text{dom}(g_2)$

$$|H_\mu(\gamma_1) - f(\gamma_1) - g_1(\gamma_1)| = |F_\mu(\gamma_2) - F(\gamma_2)| \leq C\mu, \quad (25)$$

where

$$C = \frac{R_1^2}{2}\|\mathbf{Q}\|_2^2 + \|\mathbf{b}\|_2\|\mathbf{Q}\|_2R_1 + \frac{\ell_{g_1}^2}{2} + \ell_{g_1}M.$$

From this, we have that

$$F(\hat{\gamma}) \leq F_\mu(\hat{\gamma}) + C\mu.$$

Recall now that  $\alpha = \text{prox}_{t g_2}(\tilde{\gamma}_2 - t\mathbf{D}_2^T G_{1/\mu}^{f, g_1}(\tilde{\gamma}_1))$ . Then,

$$\begin{aligned} |F(\alpha) - F(\hat{\gamma})| & \leq |f(\mathbf{D}_2\alpha) - f(\mathbf{D}_2\hat{\gamma})| \\ & \quad + |g_1(\mathbf{D}_2\alpha) - g_1(\mathbf{D}_2\hat{\gamma})| \\ & \quad + |g_2(\alpha) - g_2(\hat{\gamma})| \\ & \leq (\|\mathbf{D}_2\|_2(M + \ell_{g_1}) + \ell_{g_2})\|\alpha - \hat{\gamma}\|_2, \end{aligned}$$

where we have used the Lipschitz continuity of  $g_1$  and  $g_2$ . Next, note that  $\|\alpha - \hat{\gamma}\|_2 = t\|\mathbf{a}_1 - \mathbf{a}_2\|$  from (5.3). This way,

$$F(\alpha) \leq F(\hat{\gamma}) + t\kappa\mu, \quad (26)$$

where  $\kappa = \|\mathbf{D}_2\|_2(\|\mathbf{D}_2\|_2(M + \ell_{g_1}) + \ell_{g_2})\|\mathbf{Q}\|_2(M + \ell_{g_1})$ . Returning to (5.3), and because

$$\min_{\gamma_2} F(\gamma_2) \geq \min_{\gamma_2} F_\mu(\gamma_2) - C\mu,$$

we have that

$$F_{\text{opt}} \geq F_\mu(\gamma_\mu^*) - C\mu. \quad (27)$$

Finally, combining (5.3), (5.3) and (5.3), we obtain

$$F(\alpha) - F_{\text{opt}} \leq 2R\varepsilon + t\mu\kappa + (2R\|\mathbf{D}_2\|_2\|\mathbf{Q}\|_2(M + \ell_{g_1}) + 2C)\mu. \quad \square$$

## 6 CONCLUSION

Motivated by the multi-layer sparse model, we have introduced a multi-layer basis pursuit formulation which enforces an  $\ell_1$  penalty on the intermediate representations of the ML-CSC model. We showed how to solve this problem effectively through multi-layer extensions of iterative thresholding algorithms, building up on a projected gradient mapping approach. We showed that  $\varepsilon$ -fixed points provide approximations that are arbitrarily close, in function value, to the optimal solution. Other theoretical questions, such as those of convergence rates, constitute part of ongoing research.

We further showed how these algorithms generalize feed-forward CNN architectures by principled residual ones, improving on their performance as subsequent iterations are considered. It is intriguing how one could employ the recent results in [19], [32] to the analysis of

our resulting unfolded networks, or to understand why the Learned ML-LISTA does not provide further benefits over ML-ISTA/FISTA in the studied cases. More broadly, we believe that the study and analysis of these problems will likely contribute to the further understanding of deep learning.

## REFERENCES

- [1] Aberdam, A., Sulam, J., Elad, M.: Multi layer sparse coding: the holistic way. *To appear in SIAM Journal on Mathematics of Data Science* (2018)
- [2] Aharon, M., Elad, M., Bruckstein, A.: K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing* **54**(11), 4311–4322 (2006)
- [3] Beck, A.: First-order methods in optimization, *MOS-SIAM Series on Optimization*, vol. 25. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA (2017)
- [4] Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences* **2**(1), 183–202 (2009)
- [5] Beck, A., Teboulle, M.: Smoothing and first order methods: A unified framework. *SIAM Journal on Optimization* **22**(2), 557–580 (2012)
- [6] Bertsekas, D.P.: Constrained optimization and Lagrange multiplier methods. Computer Science and Applied Mathematics. Academic Press, Inc., New York-London (1982)
- [7] Bioucas-Dias, J.M., Figueiredo, M.A.: An iterative algorithm for linear inverse problems with compound regularizers. In: *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*, pp. 685–688. IEEE (2008)
- [8] Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J., et al.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* **3**(1), 1–122 (2011)
- [9] Candes, E.J., Donoho, D.L.: Curvelets: A surprisingly effective nonadaptive representation for objects with edges. Tech. rep., Stanford Univ Ca Dept of Statistics (2000)
- [10] Candes, E.J., Eldar, Y.C., Needell, D., Randall, P.: Compressed sensing with coherent and redundant dictionaries. arXiv preprint arXiv:1005.2613 (2010)
- [11] Chartrand, R., Yin, W.: Iteratively reweighted algorithms for compressive sensing. In: *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pp. 3869–3872. IEEE (2008)
- [12] Chen, G., Teboulle, M.: Convergence analysis of a proximal-like minimization algorithm using Bregman functions. *SIAM J. Optim.* **3**(3), 538–543 (1993). DOI 10.1137/0803026. URL <https://doi.org/10.1137/0803026>
- [13] Combettes, P.L., Pesquet, J.C.: Proximal splitting methods in signal processing. In: *Fixed-point algorithms for inverse problems in science and engineering*, pp. 185–212. Springer (2011)
- [14] Combettes, P.L., Wajs, V.R.: Signal recovery by proximal forward-backward splitting. *Multiscale Model. Simul.* **4**(4), 1168–1200 (2005)
- [15] Daubechies, I., Defrise, M., De Mol, C.: An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on pure and applied mathematics* **57**(11), 1413–1457 (2004)
- [16] Do, M.N., Vetterli, M.: The contourlet transform: an efficient directional multiresolution image representation. *IEEE Transactions on image processing* **14**(12), 2091–2106 (2005)
- [17] Donoho, D.L., Elad, M.: Optimally sparse representation in general (nonorthogonal) dictionaries via  $\ell_1$  minimization. *Proceedings of the National Academy of Sciences* **100**(5), 2197–2202 (2003)
- [18] Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., et al.: Least angle regression. *The Annals of statistics* **32**(2), 407–499 (2004)
- [19] Giryes, R., Eldar, Y.C., Bronstein, A.M., Sapiro, G.: Tradeoffs between convergence speed and reconstruction accuracy in inverse problems. *IEEE Transactions on Signal Processing* **66**(7), 1676–1690 (2018)
- [20] Gregor, K., LeCun, Y.: Learning fast approximations of sparse coding. In: *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pp. 399–406. Omnipress (2010)

- [21] Haeffele, B., Young, E., Vidal, R.: Structured low-rank matrix factorization: Optimality, algorithm, and applications to image processing. In: E.P. Xing, T. Jebara (eds.) Proceedings of the 31st International Conference on Machine Learning, *Proceedings of Machine Learning Research*, vol. 32, pp. 2007–2015. PMLR, Beijing, China (2014)
- [22] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778 (2016)
- [23] Henaff, M., Jarrett, K., Kavukcuoglu, K., LeCun, Y.: Unsupervised learning of sparse features for scalable audio classification. In: ISMIR, vol. 11, p. 2011. Citeseer (2011)
- [24] Huang, F.J., Boureau, Y.L., LeCun, Y., et al.: Unsupervised learning of invariant feature hierarchies with applications to object recognition. In: Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on, pp. 1–8. IEEE (2007)
- [25] Huang, G., Liu, Z., Weinberger, K.Q., van der Maaten, L.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition, vol. 1, p. 3 (2017)
- [26] Kavukcuoglu, K., Ranzato, M., LeCun, Y.: Fast inference in sparse coding algorithms with applications to object recognition. arXiv preprint arXiv:1010.3467 (2010)
- [27] Kutyniok, G., Labate, D.: Shearlets: Multiscale analysis for multivariate data. Springer Science & Business Media (2012)
- [28] Lin, J., Li, S.: Sparse recovery with coherent tight frames via analysis dantzig selector and analysis lasso. *Applied and Computational Harmonic Analysis* **37**(1), 126–139 (2014)
- [29] Mairal, J., Bach, F., Ponce, J.: Task-driven dictionary learning. *IEEE transactions on pattern analysis and machine intelligence* **34**(4), 791–804 (2012)
- [30] Mairal, J., Bach, F., Ponce, J., Sapiro, G.: Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research* **11**(Jan), 19–60 (2010)
- [31] Moreau, J.J.: Proximité et dualité dans un espace hilbertien. *Bull. Soc. Math. France* **93**, 273–299 (1965)
- [32] Moreau, T., Bruna, J.: Understanding trainable sparse coding via matrix factorization. In: Proceedings of the International Conference on Learning Representations (2017)
- [33] Murdock, C., Chang, M.F., Lucey, S.: Deep component analysis via alternating direction neural networks. arXiv preprint arXiv:1803.06407 (2018)
- [34] Pappas, V., Romano, Y., Elad, M.: Convolutional neural networks analyzed via convolutional sparse coding. *The Journal of Machine Learning Research* **18**(1), 2887–2938 (2017)
- [35] Pappas, V., Sulam, J., Elad, M.: Working locally thinking globally: Theoretical guarantees for convolutional sparse coding. *IEEE Transactions on Signal Processing* **65**(21), 5687–5701 (2017)
- [36] Sulam, J., Ophir, B., Zibulevsky, M., Elad, M.: Trainlets: Dictionary learning in high dimensions. *IEEE Transactions on Signal Processing* **64**(12), 3180–3193 (2016)
- [37] Sulam, J., Pappas, V., Romano, Y., Elad, M.: Multilayer convolutional sparse modeling: Pursuit and dictionary learning. *IEEE Transactions on Signal Processing* **66**(15), 4090–4104 (2018)
- [38] Tibshirani, R.: Regression shrinkage and selection via the lasso: a retrospective. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(3), 273–282 (2011)
- [39] Tibshirani, R., Wainwright, M., Hastie, T.: *Statistical learning with sparsity: the lasso and generalizations*. Chapman and Hall/CRC (2015)
- [40] Tibshirani, R.J.: *The solution path of the generalized lasso*. Stanford University (2011)
- [41] Tropp, J.A.: Just relax: Convex programming methods for identifying sparse signals in noise. *IEEE transactions on information theory* **52**(3), 1030–1051 (2006)
- [42] Wright, J., Ma, Y., Mairal, J., Sapiro, G., Huang, T.S., Yan, S.: Sparse representation for computer vision and pattern recognition. *Proceedings of the IEEE* **98**(6), 1031–1044 (2010)
- [43] Zeiler, M.D., Krishnan, D., Taylor, G.W., Fergus, R.: Deconvolutional networks. In: Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on, pp. 2528–2535. IEEE (2010)
- [44] Zhang, J., Ghanem, B.: Ista-net: Interpretable optimization-inspired deep network for image compressive sensing. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1828–1837 (2018)