
Learning step sizes for unfolded sparse coding

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Sparse coding is typically solved by iterative optimization techniques, such as
2 the Iterative Shrinkage-Thresholding Algorithm (ISTA). Unfolding and learning
3 weights of ISTA using neural networks is a practical way to accelerate estimation.
4 In this paper, we study the selection of adapted step sizes for ISTA. We show
5 that a simple step size strategy can improve the convergence rate of ISTA by
6 leveraging the sparsity of the iterates. However, it is impractical in most large-
7 scale applications. Therefore, we propose a network architecture where only the
8 step sizes of ISTA are learned. We demonstrate that for a large class of unfolded
9 algorithms, if the algorithm converges to the solution of the Lasso, its last layers
10 correspond to ISTA with learned step sizes. Experiments show that our method is
11 competitive with state-of-the-art networks when the solutions are sparse enough.

12 1 Introduction

13 The resolution of convex optimization problems by iterative algorithms has become a key part of
14 machine learning and signal processing pipelines. Amongst these problems, special attention has
15 been devoted to the Lasso (Tibshirani, 1996), due to the attractive sparsity properties of its solution
16 (see Hastie et al. 2015 for an extensive review). For a given input $x \in \mathbb{R}^n$, a dictionary $D \in \mathbb{R}^{n \times m}$
17 and a regularization parameter $\lambda > 0$, the Lasso problem is

$$z^*(x) \in \arg \min_{z \in \mathbb{R}^m} F_x(z) \quad \text{with} \quad F_x(z) \triangleq \frac{1}{2} \|x - Dz\|^2 + \lambda \|z\|_1 . \quad (1)$$

18 A variety of algorithms exist to solve Problem (1), *e.g.* proximal coordinate descent
19 (Tseng, 2001; Friedman et al., 2007), Least Angle Regression (Efron et al., 2004) or proximal
20 splitting methods (Combettes and Bauschke, 2011). The focus of this paper is on the Iterative
21 Shrinkage-Thresholding Algorithm (ISTA, Daubechies et al. 2004), which is a proximal-gradient
22 method applied to Problem (1). ISTA starts from $z^{(0)} = 0$ and iterates

$$z^{(t+1)} = \text{ST} \left(z^{(t)} - \frac{1}{L} D^\top (Dz^{(t)} - x), \frac{\lambda}{L} \right) , \quad (2)$$

23 where ST is the soft-thresholding operator defined as $\text{ST}(x, u) \triangleq \text{sign}(x) \max(|x| - u, 0)$, and L
24 is the greatest eigenvalue of $D^\top D$. In the general case, ISTA converges at rate $1/t$, which can be
25 improved to the *optimal* rate $1/t^2$ (Nesterov, 1983). However, this optimality stands in the worst
26 possible case, and linear rates are achievable in practice (Liang et al., 2014).

27 A popular line of research to improve the speed of Lasso solvers is to try to identify the support
28 of z^* , in order to diminish the size of the optimization problem (El Ghaoui et al., 2012; Ndiaye
29 et al., 2017; Johnson and Guestrin, 2015; Massias et al., 2018). Once the support is identified, larger
30 steps can also be taken, leading to improved rates for first order algorithms (Liang et al., 2014; Poon
31 et al., 2018; Sun et al., 2019).

32 However, these techniques only consider the case where a single Lasso problem is solved. When
 33 one wants to solve the Lasso for many samples $\{x^i\}_{i=1}^N$ – e.g. in dictionary learning (Olshausen and
 34 Field, 1997) – it is proposed by Gregor and Le Cun (2010) to learn a T -layers neural network of
 35 parameters $\Theta, \Phi_\Theta : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that $\Phi_\Theta(x) \simeq z^*(x)$. This Learned-ISTA (LISTA) algorithm
 36 yields better solution estimates than ISTA on new samples for the same number of iterations/layers.
 37 This idea has led to a profusion of literature (summarized in Table A.1 in appendix). Recently, it
 38 has been hinted by Zhang and Ghanem (2018); Ito et al. (2018); Liu et al. (2019) that only a few
 39 well-chosen parameters can be learned while retaining the performances of LISTA.

40 In this article, we study strategies for LISTA where only step sizes are learned. In Section 3, we
 41 propose Oracle-ISTA, an analytic strategy to obtain larger step sizes in ISTA. We show that the
 42 proposed algorithm’s convergence rate can be much better than that of ISTA. However, it requires
 43 computing a large number of Lipschitz constants which is a burden in high dimension. This motivates
 44 the introduction of Step-LISTA (SLISTA) networks in Section 4, where only a step size parameter is
 45 learned per layer. As a theoretical justification, we show in Theorem 4.4 that the last layers of any
 46 deep LISTA network converging on the Lasso must correspond to ISTA iterations with learned step
 47 sizes. We validate the soundness of this approach with numerical experiments in Section 5.

48 2 Notation and Framework

49 **Notation** The ℓ_2 norm on \mathbb{R}^n is $\|\cdot\|$. For $p \in [1, \infty]$, $\|\cdot\|_p$ is the ℓ_p norm. The Frobenius matrix
 50 norm is $\|M\|_F$. The identity matrix of size m is Id_m . ST is the soft-thresholding operator. Iterations
 51 are denoted $z^{(t)}$. $\lambda > 0$ is the regularization parameter. The Lasso cost function is F_x . $\psi_\alpha(z, x)$ is
 52 one iteration of ISTA with step α : $\psi_\alpha(z, x) = \text{ST}(z - \alpha D^\top(Dz - x), \alpha\lambda)$. $\phi_\theta(z, x)$ is one iteration
 53 of a LISTA layer with parameters $\theta = (W, \alpha, \beta)$: $\phi_\theta(z, x) = \text{ST}(z - \alpha W^\top(Dz - x), \beta\lambda)$.

54 The set of integers between 1 and m is $\llbracket 1, m \rrbracket$. Given $z \in \mathbb{R}^m$, the support is $\text{supp}(z) = \{j \in$
 55 $\llbracket 1, m \rrbracket : z_j \neq 0\} \subset \llbracket 1, m \rrbracket$. For $S \subset \llbracket 1, m \rrbracket$, $D_S \in \mathbb{R}^{n \times m}$ is the matrix containing the columns of
 56 D indexed by S . We denote L_S , the greatest eigenvalue of $D_S^\top D_S$. The equicorrelation set is $E =$
 57 $\{j \in \llbracket 1, m \rrbracket : |D_j^\top(Dz^* - x)| = \lambda\}$. The equiregularization set is $\mathcal{B}_\infty = \{x \in \mathbb{R}^n : \|D^\top x\|_\infty = 1\}$.
 58 Neural networks parameters are between brackets, e.g. $\Theta = \{\alpha^{(t)}, \beta^{(t)}\}_{t=0}^{T-1}$. The sign function is
 59 $\text{sign}(x) = 1$ if $x > 0$, -1 if $x < 0$ and 0 if $x = 0$.

60 **Framework** This paragraph recalls some properties of the Lasso. Lemma 2.1 gives the first-order
 61 optimality conditions for the Lasso.

62 **Lemma 2.1** (Optimality for the Lasso). *The Karush-Kuhn-Tucker (KKT) conditions read*

$$z^* \in \arg \min F_x \Leftrightarrow \forall j \in \llbracket 1, m \rrbracket, D_j^\top(x - Dz^*) \in \lambda \partial |z_j^*| = \begin{cases} \{\lambda \text{sign } z_j^*\}, & \text{if } z_j^* \neq 0, \\ [-\lambda, \lambda], & \text{if } z_j^* = 0. \end{cases} \quad (3)$$

63 Defining $\lambda_{\max} \triangleq \|D^\top x\|_\infty$, it holds $\arg \min F_x = \{0\} \Leftrightarrow \lambda \geq \lambda_{\max}$. For some results in Section 3,
 64 we will need the following assumption on the dictionary D :

65 **Assumption 2.2** (Uniqueness assumption). *D is such that the solution of Problem (1) is unique for*
 66 *all λ and x i.e. $\arg \min F_x = \{z^*\}$.*

67 Assumption 2.2 may seem stringent since whenever $m > n$, F_x is not strictly convex. However, it
 68 was shown in Tibshirani (2013, Lemma 4) – with earlier results from Rosset et al. 2004 – that if D is
 69 sampled from a continuous distribution, Assumption 2.2 holds for D with probability one.

70 **Definition 2.3** (Equicorrelation set). *The KKT conditions motivate the introduction of the equicorre-*
 71 *lation set $E \triangleq \{j \in \llbracket 1, m \rrbracket : |D_j^\top(Dz^* - x)| = \lambda\}$, since $j \notin E \implies z_j^* = 0$, i.e. E contains the*
 72 *support of any solution z^* .*

73 *When Assumption 2.2 holds, we have $E = \text{supp}(z^*)$ (Tibshirani, 2013, Lemma 16).*

74 We consider samples x in the equiregularization set

$$\mathcal{B}_\infty \triangleq \{x \in \mathbb{R}^n : \|D^\top x\|_\infty = 1\}, \quad (4)$$

75 which is the set of x such that $\lambda_{\max}(x) = 1$. Therefore, when $\lambda \geq 1$, the solution is $z^*(x) = 0$ for
 76 all $x \in \mathcal{B}_\infty$, and when $\lambda < 1$, $z^*(x) \neq 0$ for all $x \in \mathcal{B}_\infty$. For this reason, we assume $0 < \lambda < 1$ in
 77 the following.

78 3 Better step sizes for ISTA

79 The Lasso objective is the sum of a L -smooth function, $\frac{1}{2}\|x - D \cdot\|^2$, and a function with an explicit
 80 proximal operator, $\lambda\|\cdot\|_1$. Proximal gradient descent for this problem, with the sequence of step
 81 sizes $(\alpha^{(t)})$ consists in iterating

$$z^{(t+1)} = \text{ST} \left(z^{(t)} - \alpha^{(t)} D^\top (Dz^{(t)} - x), \lambda \alpha^{(t)} \right) . \quad (5)$$

82 ISTA follows these iterations with a constant step size $\alpha^{(t)} = 1/L$. In the following, denote
 83 $\psi_\alpha(z, x) \triangleq \text{ST}(z - \alpha D^\top (Dz^{(t)} - x), \alpha \lambda)$. One iteration of ISTA can be cast as a majorization-
 84 minimization step (Beck and Teboulle, 2009). Indeed, for all $z \in \mathbb{R}^m$,

$$F_x(z) = \frac{1}{2}\|x - Dz^{(t)}\|^2 + (z - z^{(t)})^\top D^\top (Dz^{(t)} - x) + \frac{1}{2}\|D(z - z^{(t)})\|^2 + \lambda\|z\|_1 \quad (6)$$

$$\leq \underbrace{\frac{1}{2}\|x - Dz^{(t)}\|^2 + (z - z^{(t)})^\top D^\top (Dz^{(t)} - x) + \frac{L}{2}\|z - z^{(t)}\|^2 + \lambda\|z\|_1}_{\triangleq Q_{x,L}(z, z^{(t)})} , \quad (7)$$

85 where we have used the inequality $(z - z^{(t)})^\top D^\top (Dz^{(t)} - x) \leq L\|z - z^{(t)}\|^2$. The minimizer of
 86 $Q_{x,L}(\cdot, z^{(t)})$ is $\psi_{1/L}(z^{(t)}, x)$, which is the next ISTA step.

87 **Oracle-ISTA: an accelerated ISTA with larger step sizes** Since the iterates are sparse, this
 88 approach can be refined. For $S \subset \llbracket 1, m \rrbracket$, let us define the S -smoothness of D as

$$L_S \triangleq \max_z z^\top D^\top D z, \text{ s.t. } \|z\| = 1 \text{ and } \text{supp}(z) \subset S , \quad (8)$$

89 with the convention $L_\emptyset = L$. Note that L_S is the greatest eigenvalue of $D_S^\top D_S$ where $D_S \in \mathbb{R}^{n \times |S|}$
 90 is the columns of D indexed by S . For all S , $L_S \leq L$, since L is the solution of Equation (8)
 91 without support constraint. Assume $\text{supp}(z^{(t)}) \subset S$. Combining Equations (6) and (8), we have

$$\forall z \text{ s.t. } \text{supp}(z) \subset S, F_x(z) \leq Q_{x,L_S}(z, z^{(t)}) . \quad (9)$$

92 The minimizer of the r.h.s is $z = \psi_{1/L_S}(z^{(t)}, x)$. Furthermore, the r.h.s. is a tighter upper bound than
 93 the one given in Equation (7) (see illustration in Figure 1). Therefore, using $z^{(t+1)} = \psi_{1/L_S}(z^{(t)}, x)$
 94 minimizes a tighter upper bound, provided that the following condition holds

$$\text{supp}(z^{(t+1)}) \subset S . \quad (\star)$$

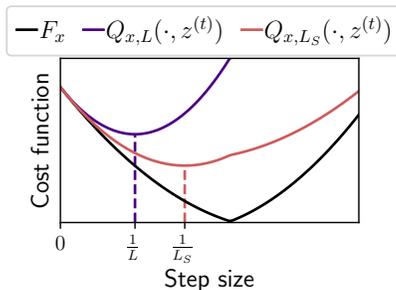


Figure 1: Majorization illustration. If $z^{(t)}$ has support S , $Q_{x,L_S}(\cdot, z^{(t)})$ is a tighter upper bound of F_x than $Q_{x,L}(\cdot, z^{(t)})$ on the set of points of support S .

95 Oracle-ISTA (OISTA) is an accelerated version of ISTA which leverages the sparsity of the iterates
 96 in order to use larger step sizes. The method is summarized in Algorithm 1. OISTA computes
 97 $y^{(t+1)} = \psi_{1/L_S}(z^{(t)}, x)$, using the larger step size $1/L_S$, and checks if it satisfies the support
 98 Condition \star . When the condition is satisfied, the step can be safely accepted. In particular Equation (9)
 99 yields $F_x(y^{(t+1)}) \leq F_x(z^{(t)})$. Otherwise, the algorithm falls back to the regular ISTA iteration
 100 with the smaller step size. Hence, each iteration of the algorithm is guaranteed to decrease F_x . The
 101 following proposition shows that OISTA converges in iterates, achieves finite support identification,
 102 and eventually reaches a safe regime where Condition \star is always true.

Algorithm 1: Oracle-ISTA (OISTA) with larger step sizes

Input: Dictionary D , target x , number of iterations T

$z^{(0)} = 0$

for $t = 0, \dots, T - 1$ **do**

 Compute $S = \text{supp}(z^{(t)})$ and L_S using an oracle ;

 Set $y^{(t+1)} = \psi_{1/L_S}(z^{(t)}, x)$;

if *Condition \star* : $\text{supp}(y^{(t+1)}) \subset S$ **then** Set $z^{(t+1)} = y^{(t+1)}$;

else Set $z^{(t+1)} = \psi_{1/L}(z^{(t)}, x)$;

Output: Sparse code $z^{(T)}$

103 **Proposition 3.1** (Convergence, finite-time support identification and safe regime). *When Assump-*
104 *tion 2.2 holds, the sequence $(z^{(t)})$ generated by the algorithm converges to $z^* = \arg \min F_x$.*

105 *Further, there exists an iteration T^* such that for $t \geq T^*$, $\text{supp}(z^{(t)}) = \text{supp}(z^*) \triangleq S^*$ and*
106 *Condition \star is always satisfied.*

107 *Sketch of proof (full proof in Subsection B.1).* Using Zangwill's global convergence theorem (Zang-
108 *will, 1969)*, we show that all accumulation points of $(z^{(t)})$ are solutions of Lasso. Since the solution
109 *is assumed unique, $(z^{(t)})$ converges to z^* . Then, we show that the algorithm achieves finite-support*
110 *identification with a technique inspired by Hale et al. (2008). The algorithm gets arbitrary close*
111 *to z^* , eventually with the same support. We finally show that in a neighborhood of z^* , the set of*
112 *points of support S^* is stable by $\psi_{1/L_S}(\cdot, x)$. The algorithm eventually reaches this region, and then*
113 *Condition \star is true. \square*

114 It follows that the algorithm enjoys the usual ISTA convergence results replacing L with L_{S^*} .

115 **Proposition 3.2** (Rates of convergence). *For $t > T^*$, $F_x(z^{(t)}) - F_x(z^*) \leq L_{S^*} \frac{\|z^* - z^{(T^*)}\|^2}{2(t - T^*)}$.*

116 *If additionally $\inf_{\|z\|=1} \|D_{S^*} z\|^2 = \mu^* > 0$, then the convergence rate for $t \geq T^*$ is*

117 $F_x(z^{(t)}) - F_x(z^*) \leq (1 - \frac{\mu^*}{L_{S^*}})^{t - T^*} (F_x(z^{(T^*)}) - F_x(z^*))$.

118 *Sketch of proof (full proof in Subsection B.2).* After iteration T^* , OISTA is equivalent to ISTA ap-
119 *plied on $F_x(z)$ restricted to $z \in S^*$. This function is L_{S^*} -smooth, and μ^* -strongly convex if $\mu^* > 0$.*
120 *Therefore, the classical ISTA rates apply with improved condition number. \square*

121 These two rates are tighter than the usual ISTA rates – in the convex case $L \frac{\|z^*\|^2}{2t}$ and in the μ -strongly
122 *convex case $(1 - \frac{\mu^*}{L})^t (F_x(0) - F_x(z^*))$ (Beck and Teboulle, 2009). Finally, the same way ISTA*
123 *converges in one iteration when D is orthogonal ($D^\top D = \text{Id}_m$), OISTA converges in one iteration if*
124 *S^* is identified and D_{S^*} is orthogonal.*

125 **Proposition 3.3.** *Assume $D_{S^*}^\top D_{S^*} = L_{S^*} \text{Id}_{|S^*|}$. Then, $z^{(T^*+1)} = z^*$.*

126 *Proof.* For z s.t. $\text{supp}(z) = S^*$, $F_x(z) = Q_{x, L_S}(z, z^{(T^*)})$. Hence, the OISTA step minimizes
127 F_x . \square

128 **Quantification of the rates improvement in a Gaussian setting** The following proposition gives
129 an asymptotic value for $\frac{L_S}{L}$ in a simple setting.

130 **Proposition 3.4.** *Assume that the entries of $D \in \mathbb{R}^{n \times m}$ are i.i.d centered Gaussian variables with*
131 *variance 1. Assume that S consists of k integers chosen uniformly at random in $\llbracket 1, m \rrbracket$. Assume that*
132 *$k, m, n \rightarrow +\infty$ with linear ratios $m/n \rightarrow \gamma$, $k/m \rightarrow \zeta$. Then*

$$\frac{L_S}{L} \rightarrow \left(\frac{1 + \sqrt{\zeta \gamma}}{1 + \sqrt{\gamma}} \right)^2. \quad (10)$$

133 This is a direct application of the Marchenko-Pastur law (Marchenko and Pastur, 1967). The law is
 134 illustrated on a toy dataset in Figure D.1. In Proposition 3.4, γ is the ratio between the number of
 135 atoms and number of dimensions, and the average size of S is described by $\zeta \leq 1$. In an overcomplete
 136 setting, we have $\gamma \gg 1$, yielding the approximation of Equation (10): $L_S \simeq \zeta L$. Therefore, if z^* is
 137 very sparse ($\zeta \ll 1$), the convergence rates of Proposition 3.2 are much better than those of ISTA.

138 **Example** Figure 2 compares the OISTA, ISTA, and FISTA on a toy problem. The improved rate of
 139 convergence of OISTA is illustrated. Further comparisons are displayed in Figure D.2 for different
 140 regularization parameters λ . While this demonstrates a much faster rate of convergence, it requires
 141 computing several Lipschitz constants L_S , which is cumbersome in high dimension. This motivates
 142 the next section, where we propose to *learn* those steps.

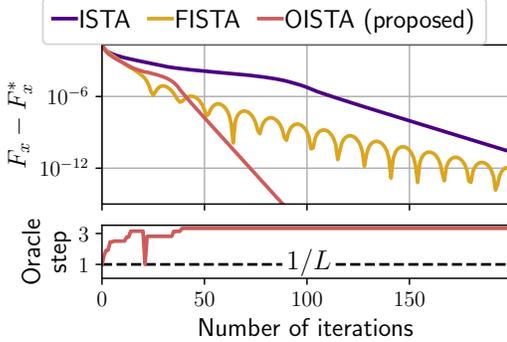


Figure 2: Convergence curves of OISTA, ISTA, and FISTA on a toy problem with $n = 10$, $m = 50$, $\lambda = 0.5$. The bottom figure displays the (normalized) steps taken by OISTA at each iteration. Full experimental setup described in Appendix D.

143 4 Learning unfolded algorithms

144 **Network architectures** At each step, ISTA performs a linear operation to compute an update
 145 in the direction of the gradient $D^\top(Dz^{(t)} - x)$ and then an element-wise non linearity with the
 146 soft-thresholding operator ST . The whole algorithm can be summarized as a recurrent neural network
 147 (RNN), presented in Figure 3a. Gregor and Le Cun (2010) introduced Learned-ISTA (LISTA), a
 148 neural network constructed by unfolding this RNN T times and learning the weights associated to each
 149 layer. The unfolded network, presented in Figure 3b, iterates $z^{(t+1)} = \text{ST}(W_x^{(t)}x + W_z^{(t)}z^{(t)}, \lambda\beta^{(t)})$.
 150 It outputs exactly the same vector as T iterations of ISTA when $W_x^{(t)} = \frac{D^\top}{L}$, $W_z^{(t)} = \text{Id}_m - \frac{D^\top D}{L}$
 151 and $\beta^{(t)} = \frac{1}{L}$. Empirically, this network is able to output a better estimate of the sparse code solution
 152 with fewer operations.

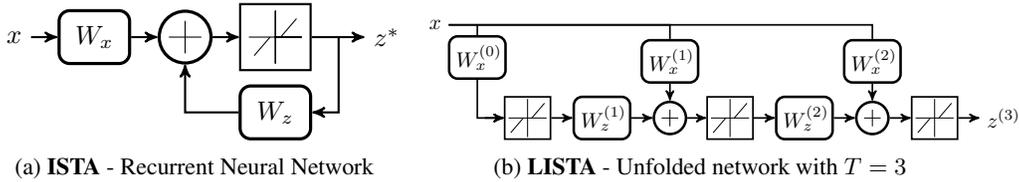


Figure 3: Network architecture for ISTA (left) and LISTA (right).

153 Due to the expression of the gradient, Chen et al. (2018) proposed to consider only a subclass of
 154 the previous networks, where the weights W_x and W_z are coupled via $W_z = \text{Id}_m - W_x^\top D$. This is
 155 the architecture we consider in the following. A layer of LISTA is a function $\phi_\theta : \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}^m$
 156 parametrized by $\theta = (W, \alpha, \beta) \in \mathbb{R}^{n \times m} \times \mathbb{R}_*^+ \times \mathbb{R}_*^+$ such that

$$\phi_\theta(z, x) = \text{ST}(z - \alpha W^\top(Dz - x), \beta\lambda) . \quad (11)$$

157 Given a set of T layer parameters $\Theta^{(T)} = \{\theta^{(t)}\}_{t=0}^{T-1}$, the LISTA network $\Phi_{\Theta^{(T)}} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is
 158 $\Phi_{\Theta^{(T)}}(x) = z^{(T)}(x)$ where $z^{(t)}(x)$ is defined by recursion

$$z^{(0)}(x) = 0, \text{ and } z^{(t+1)}(x) = \phi_{\theta^{(t)}}(z^{(t)}(x), x) \text{ for } t \in \llbracket 0, T-1 \rrbracket . \quad (12)$$

159 Taking $W = D$, $\alpha = \beta = \frac{1}{L}$ yields the same outputs as T iterations of ISTA.

160 To alleviate the need to learn the large matrices $W^{(t)}$, Liu et al. (2019) proposed to use a shared
 161 analytic matrix W_{ALISTA} for all layers. The matrix is computed in a preprocessing stage by

$$W_{\text{ALISTA}} = \arg \min_W \|W^\top D\|_F^2 \quad \text{s.t.} \quad \text{diag}(W^\top D) = \mathbf{1}_m . \quad (13)$$

162 Then, only the parameters $(\alpha^{(t)}, \beta^{(t)})$ are learned. This effectively reduces the number of parameters
 163 from $(nm + 2) \times T$ to $2 \times T$. However, we will see that ALISTA fails in our setup.

164 **Step-LISTA** With regards to the study on step sizes for ISTA in Section 3, we propose to *learn*
 165 approximation of ISTA step sizes for the input distribution using the LISTA framework. The resulting
 166 network, dubbed Step-LISTA (SLISTA), has T parameters $\Theta_{\text{SLISTA}} = \{\alpha^{(t)}\}_{t=0}^{T-1}$, and follows the
 167 iterations:

$$z^{(t+1)}(x) = \text{ST}(z^{(t)}(x) - \alpha^{(t)} D^\top (Dz^{(t)}(x) - x), \alpha^{(t)} \lambda) . \quad (14)$$

168 This is equivalent to a coupling in the LISTA parameters: a LISTA layer $\theta = (W, \alpha, \beta)$ corresponds to
 169 a SLISTA layer if and only if $\frac{\alpha}{\beta} W = D$. This network aims at learning good step sizes, like the ones
 170 used in OISTA, without the computational burden of computing Lipschitz constants. The number of
 171 parameters compared to the classical LISTA architecture Θ_{LISTA} is greatly diminished, making the
 172 network easier to train. Learning curves are shown in Figure ?? in appendix. Figure 4 displays the
 173 learned steps of a SLISTA network on a toy example. The network learns larger step-sizes as the
 174 $1/L_S$'s increase.

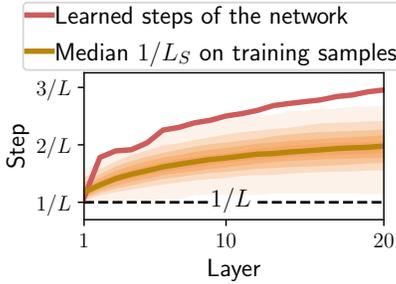


Figure 4: Steps learned with a 20 layers SLISTA network on a 10×20 problem. For each layer t and each training sample x , we compute the support $S(x, t)$ of $z^{(t)}(x)$. The brown curves display the quantiles of the distribution of $1/L_{S(x,t)}$ for each layer t . Full experimental setup described in Appendix D.

175 **Training the network** We consider the framework where the network learns to solve the Lasso on
 176 \mathcal{B}_∞ in an *unsupervised* way. Given a distribution p on \mathcal{B}_∞ , the network is trained by solving

$$\tilde{\Theta}^{(T)} \in \arg \min_{\Theta^{(T)}} \mathcal{L}(\Theta^{(T)}) \triangleq \mathbb{E}_{x \sim p}[F_x(\Phi_{\Theta^{(T)}}(x))] . \quad (15)$$

177 Most of the literature on learned optimization train the network with a different *supervised* objective
 178 (Gregor and Le Cun, 2010; Xin et al., 2016; Chen et al., 2018; Liu et al., 2019). Given a set of pairs
 179 (x^i, z^i) , the supervised approach tries to learn the parameters of the network such that $\Phi_\Theta(x^i) \simeq z^i$
 180 e.g. by minimizing $\|\Phi_\Theta(x^i) - z^i\|^2$. This training procedure differs critically from ours. For instance,
 181 ISTA does not converge for the supervised problem in general while it does for the unsupervised
 182 one. As Proposition 4.1 shows, the unsupervised approach allows to *learn to minimize* the Lasso cost
 183 function F_x .

184 **Proposition 4.1** (Pointwise convergence). *Let $\tilde{\Theta}^{(T)}$ found by solving Problem (15).*
 185 *For $x \in \mathcal{B}_\infty$ such that $p(x) > 0$, $F_x(\Phi_{\tilde{\Theta}^{(T)}}(x)) \xrightarrow{T \rightarrow +\infty} F_x^*$ almost everywhere.*

186 *Proof.* Let $\Theta_{\text{ISTA}}^{(T)}$ the parameters corresponding to ISTA i.e. $\theta_{\text{ISTA}}^{(t)} = (D, 1/L, 1/L)$. For all
 187 T , we have $\mathbb{E}_{x \sim p}[F_x^*] \leq \mathbb{E}_{x \sim p}[F_x(\Phi_{\tilde{\Theta}^{(T)}}(x))] \leq \mathbb{E}_{x \sim p}[F_x(\Phi_{\Theta_{\text{ISTA}}^{(T)}}(x))]$. Since ISTA converges
 188 uniformly on any compact, the right hand term goes to $\mathbb{E}_{x \sim p}[F_x^*]$. Therefore, by the squeeze theorem,
 189 $\mathbb{E}_{x \sim p}[F_x(\Phi_{\tilde{\Theta}^{(T)}}(x)) - F_x^*] \rightarrow 0$. This implies almost sure convergence of $F_x(\Phi_{\tilde{\Theta}^{(T)}}(x)) - F_x^*$ to 0
 190 since it is non-negative. \square

191 **Asymptotical weight coupling theorem** In this paragraph, we show the main result of this paper:
 192 any LISTA network minimizing F_x on \mathcal{B}_∞ reduces to SLISTA in its deep layers ([Theorem 4.4](#)). It
 193 relies on the following Lemmas.

194 **Lemma 4.2** (Stability of solutions around D_j). *Let $D \in \mathbb{R}^{n \times m}$ be a dictionary with non-duplicated
 195 unit-normed columns. Let $c \triangleq \max_{i \neq j} |D_i^\top D_j| < 1$. Then for all $j \in \llbracket 1, m \rrbracket$ and $\varepsilon \in \mathbb{R}^m$ such that
 196 $\|\varepsilon\| < \lambda(1 - c)$ and $D_j^\top \varepsilon = 0$, the vector $(1 - \lambda)e_j$ minimizes F_x for $x = D_j + \varepsilon$.*

197 It can be proven by verifying the KKT conditions (3) for $(1 - \lambda)e_j$, detailed in [Subsection C.1](#).

198 **Lemma 4.3** (Weight coupling). *Let $D \in \mathbb{R}^{n \times m}$ be a dictionary with non-duplicated unit-normed
 199 columns. Let $\theta = (W, \alpha, \beta)$ a set of parameters. Assume that all the couples $(z^*(x), x) \in \mathbb{R}^m \times \mathcal{B}_\infty$
 200 such that $z^*(x) \in \arg \min F_x(z)$ verify $\phi_\theta(z^*(x), x) = z^*(x)$. Then, $\frac{\alpha}{\beta}W = D$.*

201 *Sketch of proof (full proof in [Subsection C.2](#)).* For $j \in \llbracket 1, m \rrbracket$, consider $x = D_j + \varepsilon$, with
 202 $\varepsilon^\top D_j = 0$. For $\|\varepsilon\|$ small enough, $x \in \mathcal{B}_\infty$ and ε verifies the hypothesis of [Lemma 4.2](#),
 203 therefore $z^* = (1 - \lambda)e_j \in \arg \min F_x$. Writing $\phi_\theta(z^*, x) = z^*$ for the j -th coordinate yields
 204 $\alpha W_j^\top (\lambda D_j + \varepsilon) = \lambda \beta$. We can then verify that $(\alpha W_j^\top - \beta D_j^\top)(\lambda D_j + \varepsilon) = 0$. This stands for
 205 any ε orthogonal to D_j and of norm small enough. Simple linear algebra shows that this implies
 206 $\alpha W_j - \beta D_j = 0$. \square

207 [Lemma 4.3](#) states that the Lasso solutions are fixed points of a LISTA layer only if this layer
 208 corresponds to a step size for ISTA. The following theorem extends the lemma by continuity, and
 209 shows that the deep layers of any converging LISTA network must tend toward a SLISTA layer.

210 **Theorem 4.4.** *Let $D \in \mathbb{R}^{n \times m}$ be a dictionary with non-duplicated unit-normed columns. Let
 211 $\Theta^{(T)} = \{\theta^{(t)}\}_{t=0}^T$ be the parameters of a sequence of LISTA networks such that the transfer function
 212 of the layer t is $z^{(t+1)} = \phi_{\theta^{(t)}}(z^{(t)}, x)$. Assume that*

213 (i) *the sequence of parameters converges i.e. $\theta^{(t)} \xrightarrow{t \rightarrow \infty} \theta^* = (W^*, \alpha^*, \beta^*)$,*

214 (ii) *the output of the network converges toward a solution $z^*(x)$ of the Lasso (1) uniformly over
 215 the equiregularization set \mathcal{B}_∞ , i.e. $\sup_{x \in \mathcal{B}_\infty} \|\Phi_{\Theta^{(T)}}(x) - z^*(x)\| \xrightarrow{T \rightarrow \infty} 0$.*

216 Then $\frac{\alpha^*}{\beta^*}W^* = D$.

217 *Sketch of proof (full proof in [Subsection C.3](#)).* Let $\varepsilon > 0$, and $x \in \mathcal{B}_\infty$. Using the triangular in-
 218 equality, we have

$$\|\phi_{\theta^*}(z^*, x) - z^*\| \leq \|\phi_{\theta^*}(z^*, x) - \phi_{\theta^{(t)}}(z^{(t)}, x)\| + \|\phi_{\theta^{(t)}}(z^{(t)}, x) - z^*\| \quad (16)$$

219 Since the $z^{(t)}$ and $\theta^{(t)}$ converge, they are valued over a compact set K . The function $f : (z, x, \theta) \mapsto$
 220 $\phi_\theta(z, x)$ is continuous, piecewise-linear. It is therefore Lipschitz on K . Hence, we have $\|\phi_{\theta^*}(z^*, x) -$
 221 $\phi_{\theta^{(t)}}(z^{(t)}, x)\| \leq \varepsilon$ for t large enough. Since $\phi_{\theta^{(t)}}(z^{(t)}, x) = z^{(t+1)}$ and $z^{(t)} \rightarrow z^*$, $\|\phi_{\theta^{(t)}}(z^{(t)}, x) -$
 222 $z^*\| \leq \varepsilon$ for t large enough. Finally, $\phi_{\theta^*}(z^*, x) = z^*$. [Lemma 4.3](#) allows to conclude. \square

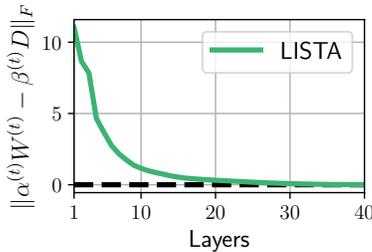


Figure 5: Illustration of [Theorem 4.4](#): for deep layers of LISTA, we have $\|\alpha^{(t)}W^{(t)} - \beta^{(t)}D\|_F \rightarrow 0$, indicating that the network ultimately only learns a step size. Full experimental setup described in [Appendix D](#).

223 [Theorem 4.4](#) means that the deep layers of any LISTA network that converges to solutions of the
 224 Lasso correspond to SLISTA iterations: $W^{(t)}$ aligns with D , and $\alpha^{(t)}, \beta^{(t)}$ get coupled. This is
 225 illustrated in [Figure 5](#), where a 40-layers LISTA network is trained on a 10×20 problem with

226 $\lambda = 0.1$. As predicted by the theorem, $\frac{\alpha^{(t)}}{\beta^{(t)}} W^{(t)} \rightarrow D$. The last layers only learn a step size. This
 227 is consistent with the observation of Moreau and Bruna (2017) which shows that the deep layers
 228 of LISTA stay close to ISTA. Further, Theorem 4.4 also shows that it is hopeless to optimize the
 229 unsupervised objective (15) with W_{ALISTA} (13), since this matrix is not aligned with D .

230 5 Numerical Experiments

231 This section provides numerical arguments to compare SLISTA to LISTA and ISTA. All the experi-
 232 ments were run using Python (Python Software Foundation, 2017) and pytorch (Paszke et al., 2017).
 233 The code to reproduce the figures is available online¹.

234 **Network comparisons** We compare the proposed approach SLISTA to state-of-the-art learned
 235 methods LISTA (Chen et al., 2018) and ALISTA (Liu et al., 2019) on synthetic and semi-real cases.

236 In the synthetic case, a dictionary $D \in \mathbb{R}^{n \times m}$ of Gaussian i.i.d. entries is generated. Each column is
 237 then normalized to one. A set of Gaussian i.i.d. samples $(\tilde{x}^i)_{i=1}^N \in \mathbb{R}^n$ is drawn. The input samples
 238 are obtained as $x^i = \tilde{x}^i / \|D^\top \tilde{x}^i\|_\infty \in \mathcal{B}_\infty$, so that for all i , $x^i \in \mathcal{B}_\infty$. We set $m = 256$ and
 239 $n = 64$.

240 For the semi-real case, we used the digits dataset from scikit-learn (Pedregosa et al., 2011) which
 241 consists of 8×8 images of handwritten digits from 0 to 9. We sample $m = 256$ samples at random
 242 from this dataset and normalize it to generate our dictionary D . Compared to the simulated Gaussian
 243 dictionary, this dictionary has a much richer correlation structure, which is known to imper the
 244 performances of learned algorithms (Moreau and Bruna, 2017). The input distribution is generated as
 245 in the simulated case.

246 The networks are trained by minimizing the empirical loss \mathcal{L} (15) on a training set of size $N_{\text{train}} =$
 247 10,000 and we report the loss on a test set of size $N_{\text{test}} = 10,000$. Further details on training are in
 248 Appendix D.

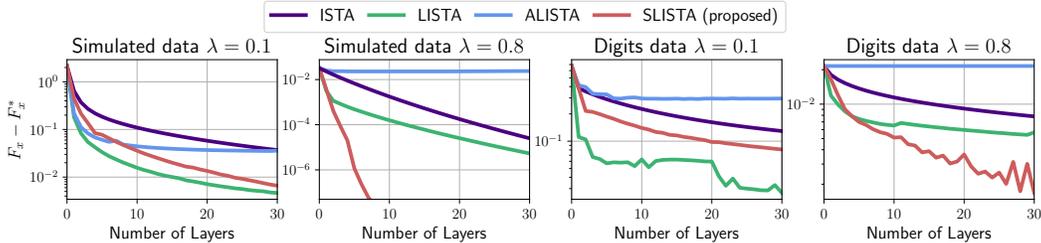


Figure 6: Test loss of ISTA, ALISTA, LISTA and SLISTA on simulated and semi-real data for different regularization parameters.

249 Figure 6 shows the test curves for different levels of regularization $\lambda = 0.1$ and 0.8 . SLISTA performs
 250 best for high λ , even for challenging semi-real dictionary D . In a low regularization setting, LISTA
 251 performs best as SLISTA cannot learn larger steps due to the low sparsity of the solution. In this
 252 unsupervised setting, ALISTA does not converge in accordance with Theorem 4.4.

253 6 Conclusion

254 We showed that using larger step sizes is an efficient strategy to accelerate ISTA for sparse solution
 255 of the Lasso. In order to make this approach practical, we proposed SLISTA, a neural network
 256 architecture which learns such step sizes. Theorem 4.4 shows that the deepest layers of any converging
 257 LISTA architecture must converge to a SLISTA layer. Numerical experiments show that SLISTA
 258 outperforms LISTA in a high sparsity setting. A major benefit of our approach is that it preserves
 259 the dictionary. We plan on leveraging this property to apply SLISTA in convolutional or wavelet
 260 cases, where the structure of the dictionary allows for fast multiplications.

¹ The code can be found in supplementary materials.

261 References

- 262 Jonas Adler, Axel Ringh, Ozan Öktem, and Johan Karlsson. Learning to solve inverse problems
263 using Wasserstein loss. *preprint ArXiv*, 1710.10898, 2017.
- 264 Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse
265 problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- 266 Mark Borgerding, Philip Schniter, and Sundeep Rangan. AMP-inspired deep networks for sparse
267 linear inverse problems. *IEEE Transactions on Signal Processing*, 65(16):4293–4308, 2017.
- 268 Xiaohan Chen, Jialin Liu, Zhangyang Wang, and Wotao Yin. Theoretical linear convergence of
269 unfolded ISTA and its practical weights and thresholds. In *Advances in Neural Information
270 Processing Systems (NIPS)*, pages 9061–9071, 2018.
- 271 Patrick L Combettes and Heinz H. Bauschke. *Convex Analysis and Monotone Operator Theory in
272 Hilbert Spaces*. Springer, 2011. ISBN 9788578110796. doi: 10.1017/CBO9781107415324.004.
- 273 Ingrid Daubechies, Michel Defrise, and Christine De Mol. An iterative thresholding algorithm
274 for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied
275 Mathematics*, 57(11):1413–1457, 2004.
- 276 Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *Ann.
277 Statist.*, 32(2):407–499, 2004.
- 278 Laurent El Ghaoui, Vivian Viallon, and Tarek Rabbani. Safe feature elimination in sparse supervised
279 learning. *J. Pacific Optim.*, 8(4):667–698, 2012.
- 280 Jerome Friedman, Trevor Hastie, Holger Höfling, and Robert Tibshirani. Pathwise coordinate
281 optimization. *The Annals of Applied Statistics*, 1(2):302–332, 2007.
- 282 Raja Giryes, Yonina C. Eldar, Alex M. Bronstein, and Guillermo Sapiro. Tradeoffs between con-
283 vergence speed and reconstruction accuracy in inverse problems. *IEEE Transaction on Signal
284 Processing*, 66(7):1676–1690, 2018.
- 285 Karol Gregor and Yann Le Cun. Learning Fast Approximations of Sparse Coding. In *International
286 Conference on Machine Learning (ICML)*, pages 399–406, 2010.
- 287 Elaine Hale, Wotao Yin, and Yin Zhang. Fixed-point continuation for ℓ_1 -minimization: Methodology
288 and convergence. *SIAM J. Optim.*, 19(3):1107–1130, 2008.
- 289 Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical Learning with Sparsity: The
290 Lasso and Generalizations*. CRC Press, 2015.
- 291 John R. Hershey, Jonathan Le Roux, and Felix Weninger. Deep unfolding: Model-based inspiration
292 of novel deep architectures. *preprint ArXiv*, 1409.2574, 2014.
- 293 Daisuke Ito, Satoshi Takabe, and Tadashi Wadayama. Trainable ISTA for sparse signal recovery. In
294 *IEEE International Conference on Communications Workshops*, pages 1–6, 2018.
- 295 Tyler Johnson and Carlos Guestrin. Blitz: A principled meta-algorithm for scaling sparse optimization.
296 In *International Conference on Machine Learning (ICML)*, pages 1171–1179, 2015.
- 297 Jingwei Liang, Jalal Fadili, and Gabriel Peyré. Local linear convergence of forward–backward under
298 partial smoothness. In *Advances in Neural Information Processing Systems*, pages 1970–1978,
299 2014.
- 300 Jialin Liu, Xiaohan Chen, Zhangyang Wang, and Wotao Yin. ALISTA: Analytic weights are as good
301 as learned weights in LISTA. In *International Conference on Learning Representation (ICLR)*,
302 2019.
- 303 Vladimir A Marchenko and Leonid Andreevich Pastur. Distribution of eigenvalues for some sets of
304 random matrices. *Mathematics of the USSR-Sbornik*, 1(4):457, 1967.
- 305 Mathurin Massias, Alexandre Gramfort, and Joseph Salmon. Celer: a Fast Solver for the Lasso with
306 Dual Extrapolation. In *International Conference on Machine Learning (ICML)*, 2018.
- 307 Thomas Moreau and Joan Bruna. Understanding neural sparse coding with matrix factorization. In
308 *International Conference on Learning Representation (ICLR)*, 2017.
- 309 Eugene Ndiaye, Olivier Fercoq, Alexandre Gramfort, and Joseph Salmon. Gap safe screening rules
310 for sparsity enforcing penalties. *J. Mach. Learn. Res.*, 18(128):1–33, 2017.

- 311 Yurii Nesterov. A method for solving a convex programming problem with rate of convergence
312 $O(1/k^2)$. *Soviet Math. Doklady*, 269(3):543–547, 1983.
- 313 Bruno A. Olshausen and David J Field. Sparse coding with an incomplete basis set: a strategy
314 employed by V1, 1997.
- 315 Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito,
316 Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in
317 PyTorch. In *NIPS Autodiff Workshop*, 2017.
- 318 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Pretten-
319 hofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and
320 E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*,
321 12:2825–2830, 2011.
- 322 Clarice Poon, Jingwei Liang, and Carola-Bibiane Schönlieb. Local convergence properties of SAGA
323 and prox-SVRG and acceleration. In *International Conference on Machine Learning (ICML)*,
324 2018.
- 325 Python Software Foundation. Python Language Reference, version 3.6. <http://python.org/>, 2017.
- 326 Saharon Rosset, Ji Zhu, and Trevor Hastie. Boosting as a regularized path to a maximum margin
327 classifier. *J. Mach. Learn. Res.*, 5:941–973, 2004.
- 328 Pablo Sprechmann, Alex M. Bronstein, and Guillermo Sapiro. Learning efficient structured sparse
329 models. In *International Conference on Machine Learning (ICML)*, pages 615–622, 2012.
- 330 Pablo Sprechmann, Roei Litman, and TB Yakar. Efficient supervised sparse analysis and synthesis
331 operators. In *Advances in Neural Information Processing Systems (NIPS)*, pages 908–916, 2013.
- 332 Yifan Sun, Halyun Jeong, Julie Nutini, and Mark Schmidt. Are we there yet? manifold identification
333 of gradient-related proximal methods. In *Proceedings of Machine Learning Research*, volume 89
334 of *Proceedings of Machine Learning Research*, pages 1110–1119. PMLR, 2019.
- 335 Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical
336 Society: Series B (Methodological)*, 58(1):267–288, 1996.
- 337 Ryan Tibshirani. The lasso problem and uniqueness. *Electron. J. Stat.*, 7:1456–1490, 2013.
- 338 Paul Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization.
339 *J. Optim. Theory Appl.*, 109(3):475–494, 2001.
- 340 Zhangyang Wang, Qing Ling, and Thomas S. Huang. Learning deep ℓ_0 encoders. In *AAAI Conference
341 on Artificial Intelligence*, pages 2194–2200, 2015.
- 342 Bo Xin, Yizhou Wang, Wen Gao, and David Wipf. Maximal sparsity with deep networks? In
343 *Advances in Neural Information Processing Systems (NIPS)*, pages 4340–4348, 2016.
- 344 Yan Yang, Jian Sun, Huibin Li, and Zongben Xu. Deep ADMM-Net for compressive sensing MRI.
345 In *Advances in Neural Information Processing Systems (NIPS)*, pages 10–18, 2017.
- 346 Willard I Zangwill. Convergence conditions for nonlinear programming algorithms. *Management
347 Science*, 16(1):1–13, 1969.
- 348 Jian Zhang and Bernard Ghanem. ISTA-Net: Interpretable optimization-inspired deep network for
349 image compressive sensing. In *IEEE Computer Society Conference on Computer Vision and
350 Pattern Recognition*, pages 1828–1837, 2018.

351 A Unfolded optimization algorithms literature summary

352 In Table A.1, we summarize the prolific literature on learned unfolded optimization procedures
 353 for sparse recovery. A particular focus is set on the chosen training loss training which is either
 354 supervised, with a regression of z^i from the input x^i for a given training set (x^i, z^i) , or unsupervised,
 355 where the objective is to minimize the Lasso cost function F_x for each training point x .

Table A.1: Neural network for sparse coding

Reference	Base Algo	Train Loss	Coupled weights	Remarks
Gregor and Le Cun (2010)	ISTA / CD	supervised	×	–
Sprechmann et al. (2012)	Block CD	unsupervised	×	Group ℓ_1
Sprechmann et al. (2013)	ADMM	supervised	N/A	–
Hershey et al. (2014)	NMF	supervised	×	NMF
Wang et al. (2015)	IHT	supervised	×	Hard-thresholding
Xin et al. (2016)	IHT	supervised	×/✓	Hard-thresholding
Giryas et al. (2018)	PGD/IHT	supervised	N/A	Group ℓ_1
Yang et al. (2017)	ADMM	supervised	N/A	–
Adler et al. (2017)	ADMM	supervised	N/A	Wasserstein distance with z^*
Borgerding et al. (2017)	AMP	supervised	×	–
Moreau and Bruna (2017)	ISTA	unsupervised	×	–
Chen et al. (2018)	ISTA	supervised	✓	Linear convergence rate
Ito et al. (2018)	ISTA	supervised	✓	MMSE shrinkage non-linearity
Zhang and Ghanem (2018)	PGD	supervised	✓	Sparsity of Wavelet coefficients
Liu et al. (2019)	ISTA	supervised	✓	Analytic weight W_{ALISTA}
Proposed	ISTA	unsupervised	✓	–

356 B Proofs of Section 3’s results

357 B.1 Proof of Proposition 3.1

358 We consider that the solution of the Lasso is unique, following the result of Tibshirani (2013)[Lemmas
 359 4 and 16] when the entries of D and x come from a continuous distribution.

360 **Proposition 3.1** (Convergence, finite-time support identification and safe regime). *When Assump-*
 361 *tion 2.2 holds, the sequence $(z^{(t)})$ generated by the algorithm converges to $z^* = \arg \min F_x$.*

362 *Further, there exists an iteration T^* such that for $t \geq T^*$, $\text{supp}(z^{(t)}) = \text{supp}(z^*) \triangleq S^*$ and*
 363 *Condition \star is always satisfied.*

364 *Proof.* Let $z^{(t)}$ be the sequence of iterates produced by Algorithm 1. We have a descent function

$$F_x(z^{(t+1)}) - F_x(z^{(t)}) \leq -\frac{\gamma}{2} \|z^{(t+1)} - z^{(t)}\|^2 \leq -\frac{\min \|D_j\|}{2} \|z^{(t+1)} - z^{(t)}\|^2, \quad (17)$$

365 where $\gamma = L_S$ if Condition \star is met, and L otherwise. Additionally, the iterates are bounded because
 366 $F_x(z^{(t)})$ decreases at each iteration and F_x is coercive. Hence we can apply Zangwill’s Global
 367 Convergence Theorem (Zangwill, 1969). Any z^* accumulation point of $(z^{(t)})_{t \in \mathbb{N}}$ is a minimizer of
 368 F_x .

369 Since we only consider the case where the minimizer is unique, the bounded sequence $(z^{(t)})_{t \in \mathbb{N}}$ has
 370 a unique accumulation point, thus converges to z^* .

371 The support identification is a simplification of a result of Hale et al. (2008), we include it here for
 372 completeness.

373 **Lemma B.1** (Approximation of the soft-thresholding). *Let $z \in \mathbb{R}, \nu > 0$. For ϵ small enough, we*
 374 *have*

$$\text{ST}(z + \epsilon, \nu) = \begin{cases} 0, & \text{if } |z| < \nu, \\ \max(0, \epsilon) \text{sign}(z), & \text{if } |z| = \nu, \\ z + \epsilon - \nu \text{sign } z, & \text{if } |z| > \nu. \end{cases} \quad (18)$$

375 Let $\rho > 0$ be such that Equation (18) holds for $\nu = \lambda/L$, every $\epsilon < \rho$, and every $z = z_j^* -$
 376 $\frac{1}{L}D_j^\top(Dz^* - x)$.

377 Let $t \in \mathbb{N}$ such that $z^{(t)} = z^* + \epsilon$, with $\|\epsilon\| \leq \rho$. With $\epsilon' \triangleq (\text{Id} - \frac{1}{L}D^\top D)\epsilon$, we also have $\|\epsilon'\| \leq \rho$.
 378 Let $j \in \llbracket 1, m \rrbracket$.

379 If $j \notin E$, $|z_j^* - \frac{1}{L}D_j^\top(Dz^* - x)| = |\frac{1}{L}D_j^\top(Dz^* - x)| < \lambda/L$ hence $\text{ST}(z_j^* - \frac{1}{L}D_j^\top(Dz^* - x) +$
 380 $\epsilon'_j, \lambda/L) = 0$.

381 If $j \in E$, $|z_j^* - \frac{1}{L}D_j^\top(Dz^* - x)| = |z_j^* + \frac{\lambda}{L} \text{sign } z_j^*| > \lambda/L$, and $\text{sign } \text{ST}(z_j^* - \frac{1}{L}D_j^\top(Dz^* - x) +$
 382 $\epsilon'_j, \lambda/L) = \text{sign } z_j^*$.

383 The same reasoning can be applied with ρ' such that Equation (18) holds for $\nu = \lambda/L_{S^*}$, every
 384 $\epsilon < \rho'$, and every $z = z_j^* - \frac{1}{L_{S^*}}D_j^\top(Dz^* - x)$. If we introduce $\eta > 0$ such that $\|\epsilon\| \leq \eta \implies$
 385 $\|(\text{Id} - \frac{1}{L_{S^*}}D^\top D)\epsilon\| \leq \rho'$, in the ball of center z^* and radius η , the iteration with step size L_{S^*}
 386 identifies the support.

387 Additionnally, since $\text{Id} - \frac{1}{L_{S^*}}D_{S^*}^\top D_{S^*}$ is non expansive on vectors which support is S^* , the iterations
 388 with the step L_{S^*} never leave this ball once they have entered it.

389 Therefore, once the iterates enter $\mathcal{B}(z^*, \min(\eta, \rho))$, Condition \star is always satisfied.

390

□

391 B.2 Proof of Proposition 3.2

392 **Proposition 3.2** (Rates of convergence). *For $t > T^*$, $F_x(z^{(t)}) - F_x(z^*) \leq L_{S^*} \frac{\|z^* - z^{(T^*)}\|^2}{2(t - T^*)}$.*

393 *If additionally $\inf_{\|z\|=1} \|D_{S^*} z\|^2 = \mu^* > 0$, then the convergence rate for $t \geq T^*$ is*

$$394 F_x(z^{(t)}) - F_x(z^*) \leq (1 - \frac{\mu^*}{L_{S^*}})^{t - T^*} (F_x(z^{(T^*)}) - F_x(z^*)).$$

395 *Proof.* For $t \geq T^*$, the iterates support is S^* and the objective function is L_{S^*} -smooth instead of
 396 L -smooth. It is also μ^* strongly convex if $\mu^* > 0$. The obtained rates are a classical result of the
 397 proximal gradient descent method in these cases. □

398 C Proof of Section 4's Lemmas

399 C.1 Proof of Lemma 4.2

400 **Lemma 4.2** (Stability of solutions around D_j). *Let $D \in \mathbb{R}^{n \times m}$ be a dictionary with non-duplicated*
 401 *unit-normed columns. Let $c \triangleq \max_{l \neq j} |D_l^\top D_j| < 1$. Then for all $j \in \llbracket 1, m \rrbracket$ and $\epsilon \in \mathbb{R}^m$ such that*
 402 *$\|\epsilon\| < \lambda(1 - c)$ and $D_j^\top \epsilon = 0$, the vector $(1 - \lambda)e_j$ minimizes F_x for $x = D_j + \epsilon$.*

403 *Proof.* Let $j \in \llbracket 1, m \rrbracket$ and let $\epsilon \in \mathbb{R}^m \cap D_j^\perp$ be a vector such that $\|\epsilon\| < \lambda(1 - c)$.

404 For notation simplicity, we denote $z^* = z^*(D_j - \epsilon)$.

$$D_j^\top(Dz^* - D_j - \epsilon) = D_j^\top(-\lambda D_j - \epsilon) = -\lambda = -\lambda \text{sign } z_j^*, \quad (19)$$

405 since $1 - \lambda > 0$. For the other coefficients $l \in \llbracket 1, m \rrbracket \setminus \{j\}$, we have

$$|D_l^\top (Dz^* - D_j - \varepsilon)| = |D_l^\top (-\lambda D_j - \varepsilon)|, \quad (20)$$

$$= |\lambda D_l^\top D_j + D_l^\top \varepsilon|, \quad (21)$$

$$\leq \lambda |D_l^\top D_j| + |D_l^\top \varepsilon|, \quad (22)$$

$$\leq \lambda c + \|D_l\| \|\varepsilon\|, \quad (23)$$

$$\leq \lambda c + \|\varepsilon\| < \lambda, \quad (24)$$

$$(25)$$

406 Therefore, $(1 - \lambda)e_j$ verifies the KKT conditions (3) and $z^*(D_j + \varepsilon) = (1 - \lambda)e_j$. \square

407 C.2 Proof of Lemma 4.3

408 **Lemma 4.3** (Weight coupling). *Let $D \in \mathbb{R}^{n \times m}$ be a dictionary with non-duplicated unit-normed*
 409 *columns. Let $\theta = (W, \alpha, \beta)$ a set of parameters. Assume that all the couples $(z^*(x), x) \in \mathbb{R}^m \times \mathcal{B}_\infty$*
 410 *such that $z^*(x) \in \arg \min F_x(z)$ verify $\phi_\theta(z^*(x), x) = z^*(x)$. Then, $\frac{\alpha}{\beta}W = D$.*

411 *Proof.* Let $x \in \mathcal{B}_\infty$ be an input vector and $z^*(x) \in \mathbb{R}^m$ be a solution for the Lasso at level $\lambda > 0$.
 412 Let $j \in \llbracket 1, m \rrbracket$ be such that $z_j^* > 0$. The KKT conditions (3) gives

$$D_j^\top (Dz^*(x) - x) = -\lambda. \quad (26)$$

413 Suppose that $z^*(x)$ is a fixed point of the layer, then we have

$$\text{ST}(z_j^*(x) - \alpha W_j^\top (Dz^*(x) - x), \lambda\beta) = z_j^*(x) > 0. \quad (27)$$

414 By definition, $\text{ST}(a, b) > 0$ implies that $a > b$ and $\text{ST}(a, b) = a - b$. Thus,

$$z_j^*(x) - \alpha W_j^\top (Dz^*(x) - x) - \lambda\beta = z_j^*(x) \quad (28)$$

$$\Leftrightarrow \alpha W_j^\top (Dz^*(x) - x) + \lambda\beta = 0 \quad (29)$$

$$\Leftrightarrow \alpha W_j^\top (Dz^*(x) - x) - \beta D_j^\top (Dz^*(x) - x) = 0 \quad \text{by (26)} \quad (30)$$

$$\Leftrightarrow (\alpha W_j - \beta D_j)^\top (Dz^*(x) - x) = 0. \quad (31)$$

415 As the relation (31) must hold for all $x \in \mathcal{B}_\infty$, it is true for all $D_j + \varepsilon$ for all $\varepsilon \in \mathcal{B}(0, \lambda(1 - c)) \cap D_j^\perp$.
 416 Indeed, in this case, $\|D^\top (D_j + \varepsilon)\|_\infty = 1$. D verifies the conditions of Lemma 4.2, and thus
 417 $z^* = (1 - \lambda)e_j$, i.e.

$$(\alpha W_j - \beta D_j)^\top (D(1 - \lambda)e_j - (D_j + \varepsilon)) = 0 \quad (32)$$

$$(\alpha W_j - \beta D_j)^\top (-\lambda D_j - \varepsilon) = 0 \quad (33)$$

418 Taking $\varepsilon = 0$ yields $(\alpha W_j - \beta D_j)^\top D_j = 0$, and therefore Eq. (33) becomes $(\alpha W_j - \beta D_j)^\top \varepsilon = 0$
 419 for all ε small enough and orthogonal to D_j , which implies $\alpha W_j - \beta D_j = 0$ and concludes our
 420 proof. \square

421 C.3 Proof of Theorem 4.4

422 **Theorem 4.4.** *Let $D \in \mathbb{R}^{n \times m}$ be a dictionary with non-duplicated unit-normed columns. Let*
 423 *$\Theta^{(T)} = \{\theta^{(t)}\}_{t=0}^T$ be the parameters of a sequence of LISTA networks such that the transfer function*
 424 *of the layer t is $z^{(t+1)} = \phi_{\theta^{(t)}}(z^{(t)}, x)$. Assume that*

425 (i) *the sequence of parameters converges i.e. $\theta^{(t)} \xrightarrow{t \rightarrow \infty} \theta^* = (W^*, \alpha^*, \beta^*)$,*

426 (ii) *the output of the network converges toward a solution $z^*(x)$ of the Lasso (1) uniformly over*
 427 *the equiregularization set \mathcal{B}_∞ , i.e. $\sup_{x \in \mathcal{B}_\infty} \|\Phi_{\Theta^{(T)}}(x) - z^*(x)\| \xrightarrow{T \rightarrow \infty} 0$.*

428 Then $\frac{\alpha^*}{\beta^*}W^* = D$.

429 *Proof.* For simplicity of the notation, we will drop the x variable whenever possible, i.e. $z^* = z^*(x)$
430 and $\phi_\theta(z) = \phi_\theta(z, x)$. We denote $z^{(t)} = \Phi_{\Theta^{(t)}}(x)$ the output of the network with t layers.

431 Let $\epsilon > 0$. By hypothesis (i), there exists T_0 such that for all $t \geq T_0$,

$$\|W^{(t)} - W^*\| \leq \epsilon \quad |\alpha^{(t)} - \alpha^*| \leq \epsilon \quad |\beta^{(t)} - \beta^*| \leq \epsilon. \quad (34)$$

432 By hypothesis (ii), there exists T_1 such that for all $t \geq T_1$ and all $x \in \mathcal{B}_\infty$,

$$\|z^{(t)} - z^*\| \leq \epsilon. \quad (35)$$

433 Let $x \in \mathcal{B}_\infty$ be an input vector and $t \geq \max(T_0, T_1)$. Using (35), we have

$$\|z^{(t+1)} - z^{(t)}\| \leq \|z^{(t+1)} - z^*\| + \|z^{(t)} - z^*\| \leq 2\epsilon \quad (36)$$

434 By (i), there exist a compact set $\mathcal{K}_1 \subset \mathbb{R}^{n \times m} \times \mathbb{R}_*^+ \times \mathbb{R}_*^+$ s.t. $\theta^{(t)} \in \mathcal{K}_1$ for all $t \in \mathbb{N}$ and
435 $\theta^* \in \mathcal{K}$. The input x is taken in a compact set \mathcal{B}_∞ and as $z^* = \arg \min_z F_x(z)$, we have
436 $\lambda \|z\|_1 \leq F_x(z^*) \leq F_x(0) = \|x\|$ thus z^* is also in a compact set \mathcal{K}_2 .

437 We consider the function $f(z, x, \theta) = \text{ST}(z - \alpha W^\top(Dz - x), \beta)$ on the compact set $\mathcal{K}_2 \times \mathcal{B}_\infty \times \mathcal{K}_1$.
438 This function is continuous and piece-wise linear on a compact set. It is thus L -Lipschitz and thus

$$\|\phi_{\theta^{(t)}}(z^{(t)}) - \phi_{\theta^{(t)}}(z^*)\| \leq L \|z^{(t)} - z^*\| \leq L\epsilon \quad (37)$$

$$\|\phi_{\theta^*}(z^*) - \phi_{\theta^{(t)}}(z^*)\| \leq L \|\theta^{(t)} - \theta^*\| \leq L\epsilon \quad (38)$$

439 Using these inequalities, we get

$$\begin{aligned} \|\phi_{\theta^*}(z^*, x) - z^*\| &\leq \underbrace{\|\phi_{\theta^*}(z^*) - \phi_{\theta^{(t)}}(z^*)\|}_{< L\epsilon \text{ by (38)}} + \underbrace{\|\phi_{\theta^{(t)}}(z^*) - \phi_{\theta^{(t)}}(z^{(t)})\|}_{< L\epsilon \text{ by (37)}} \\ &\quad + \underbrace{\|\phi_{\theta^{(t)}}(z^{(t)}) - z^{(t)}\|}_{< 2\epsilon \text{ by (36)}} + \underbrace{\|z^{(t)} - z^*\|}_{< \epsilon \text{ by (35)}} \\ &\leq (2L + 3)\epsilon. \end{aligned} \quad (39)$$

440 As this result holds for all $\epsilon > 0$ and all $x \in \mathcal{B}_\infty$, we have $\phi_{\theta^*}(z^*) = z^*$ for all $x \in \mathcal{B}_\infty$. We can
441 apply the Lemma 4.3 to conclude this proof. \square

442 D Experimental setups and supplementary figures

443 **Dictionary generation:** Unless specified otherwise, to generate synthetic dictionaries, we first draw
444 a random i.i.d. Gaussian matrix $\hat{D} \in \mathbb{R}^{n \times m}$. The dictionary is obtained by normalizing the columns:

$$445 D_{ij} = \frac{1}{\|\hat{D}_{:,i}\|} \hat{D}_{ij}.$$

446 **Samples generation:** The samples x are generated as follows: Random i.i.d. Gaussian samples
447 $\hat{x} \in \mathbb{R}^n$ are generated. We then normalize them: $x = \frac{1}{\|D^\top \hat{x}\|_\infty} \hat{x}$, so that $x \in \mathcal{B}_\infty$.

448 **Training the networks** Since the loss function and the network are continuous but non-differentiable,
449 we use sub-gradient descent for training. The sub-gradient of the cost function with respect to the
450 parameters of the network is computed by automatic differentiation. We use full-batch sub-gradient
451 descent with a backtracking procedure to find a suitable learning rate. To verify that we do not overfit
452 the training set, we always check that the test loss and train loss are comparable.

453 Main text figures setup

- 454 • **Figure 2:** We generate a random dictionary of size 10×50 . We take $\lambda = 0.5$, and a random
455 sample $x \in \mathcal{B}_\infty$. F_x^* is computed by iterating ISTA for 10000 iterations.
- 456 • **Figure 4:** We generate a random dictionary of size 10×20 . We take $\lambda = 0.2$. We generate
457 a training set of $N = 1000$ samples $(x^i)_{i=1}^{1000} \in \mathcal{B}_\infty$. A 20 layers SLISTA network is trained
458 by gradient descent on these data. We report the learned step sizes. For each layer t of
459 the network and each training sample x , we compute the support at the output of the t -th
460 layer, $S(x, t) = \text{supp}(z^{(t)}(x))$. For each t , we display the quantiles of the distribution of
461 the $(1/L_{S(x^i, t)})_{i=1}^{1000}$.
- 462 • **Figure 5:** A random 10×20 dictionary is generated. We take 1000 training samples, and
463 $\lambda = 0.05$. A 40 layers LISTA network is trained by gradient descent on those samples. We
464 report the quantity $\|\alpha^{(t)} W^{(t)} - \beta^{(t)} D\|_F$ for each layer t .

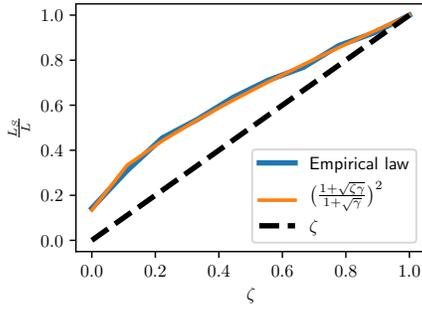


Figure D.1: Illustration of Proposition 3.4. A toy Gaussian dictionary is generated with $n = 200$, $m = 600$ so that $\gamma = 3$. We compute its Lipschitz constant L . For ζ between 0 and 1, we extract $\lfloor \zeta m \rfloor$ columns at random and compute the corresponding Lipschitz constant L_S . The plot shows an almost perfect fit between the empirical law and the theoretical limit (10).

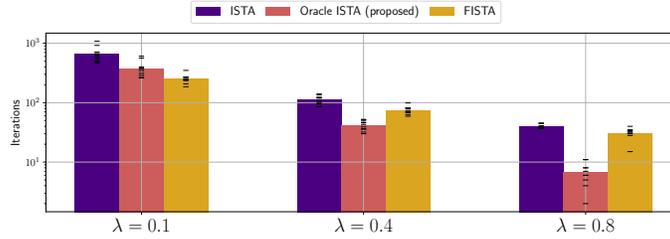


Figure D.2: Comparison between ISTA, FISTA and Oracle-ISTA for different levels of regularization on a Gaussian dictionary, with $n = 100$ and $m = 200$. We report the average number of iterations taken to reach a point z such that $F_x(z) < F_x^* + 10^{-13}$. The experiment is repeated 10 times, starting from random points in \mathcal{B}_∞ . OISTA is always faster than ISTA, and is faster than FISTA for high regularization.

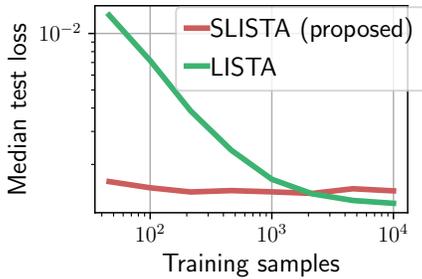


Figure D.3: Learning curves of SLISTA and LISTA. Random Gaussian dictionaries with $n = 10$ and $m = 20$ are generated. We take $\lambda = 0.3$. Networks with 10 layers are fit on those dictionaries, and their test loss is reported for different number of training samples. The process is repeated 100 times; the curves shown display the median of the test-loss.