



HAL
open science

Dictionary and prior learning with unrolled algorithms for unsupervised inverse problems

Benoît Malézieux, Thomas Moreau, Matthieu Kowalski

► **To cite this version:**

Benoît Malézieux, Thomas Moreau, Matthieu Kowalski. Dictionary and prior learning with unrolled algorithms for unsupervised inverse problems. 2021. hal-03258196v2

HAL Id: hal-03258196

<https://inria.hal.science/hal-03258196v2>

Preprint submitted on 11 Jun 2021 (v2), last revised 8 Feb 2022 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dictionary and prior learning with unrolled algorithms for unsupervised inverse problems

Benoît Malézieux

Université Paris-Saclay, Inria, CEA
Palaiseau, 91120, France
benoit.malezieux@inria.fr

Thomas Moreau

Université Paris-Saclay, Inria, CEA
Palaiseau, 91120, France
thomas.moreau@inria.fr

Matthieu Kowalski

L2S, Université Paris-Saclay–CNRS–CentraleSupélec
Gif-sur-Yvette, 91190, France
matthieu.kowalski@universite-paris-saclay.fr

Abstract

Inverse problems consist in recovering a signal given noisy observations. One classical resolution approach is to leverage sparsity and integrate prior knowledge of the signal to the reconstruction algorithm to get a plausible solution. Still, this prior might not be sufficiently adapted to the data. In this work, we study Dictionary and Prior learning from degraded measurements as a bi-level problem, and we take advantage of unrolled algorithms to solve approximate formulations of Synthesis and Analysis. We provide an empirical and theoretical analysis of automatic differentiation for Dictionary Learning to understand better the pros and cons of unrolling in this context. We find that unrolled algorithms speed up the recovery process for a small number of iterations by improving the gradient estimation. Then we compare Analysis and Synthesis by evaluating the performance of unrolled algorithms for inverse problems, without access to any ground truth data for several classes of dictionaries and priors. While Analysis can achieve good results, Synthesis is more robust and performs better. Finally, we illustrate our method on pattern and structure learning tasks from degraded measurements.

1 Introduction

Linear inverse problems are ubiquitous in observational science such as imaging [Ribes and Schmitt, 2008], neurosciences [Gramfort et al., 2012] or astrophysics [Starck, 2016]. They consist in reconstructing a signal $x \in \mathbb{R}^n$ from remote and noisy measurements $y \in \mathbb{R}^m$ – also called observed signal – which is obtained as a linear transformation $A \in \mathbb{R}^{m \times n}$ of x , corrupted with noise $b \in \mathbb{R}^m$: $y = Ax + b$. As the dimension m of the measurements is usually much smaller than the dimension n of the signal, these problems are ill-posed, meaning that several solutions could be correct given a set of observations. The uncertainty on the measurements, which are generally corrupted by noise, increases the number of potential solutions. Therefore, practitioners rely on prior knowledge of the data to select the most plausible solution among all possible ones.

Efficient and analytic transforms are available and produce satisfying results on specific data, such as wavelets for images or Gaborlets for audio signals [Mallat, 2008]. However, the complexity and the variability of the signals often make it hard to rely on *ad hoc* priors or dictionaries. In the latter case, the prior can be learned from the data. In particular, methods trying to summarize the structure of the observations efficiently by leveraging sparsity have been extensively studied [Elad, 2010]. A classical approach considers that the signal to reconstruct x can be represented as a sparse vector in

an unknown space. A first way to impose such prior is to assume that there exists a forward transform $\Gamma \in \mathbb{R}^{L \times n}$ which sparsify the signal. This formulation is called Analysis [Peyré and Fadili, 2011, Yaghoobi et al., 2013]. Γ is learned in a set of constraints \mathcal{C}_A by solving

$$\min_{x \in \mathbb{R}^n, \Gamma \in \mathcal{C}_A} F_A(x, \Gamma) \triangleq \frac{1}{2} \|Ax - y\|_2^2 + \lambda \|\Gamma^\top x\|_1 . \quad (1)$$

Another approach is to consider that the signal is a linear combination of a small number of atoms, which are the columns of a dictionary $D \in \mathbb{R}^{n \times L}$. This formulation is called Synthesis, as the signal is synthesized from the dictionary. It is often referred to as Dictionary Learning [Olshausen and Field, 1997, Aharon et al., 2006, Mairal et al., 2009]. Even though priors are mostly related to Bayesian frameworks, we will – with a slight abuse of terminology and for the sake of simplicity – refer to Analysis transforms and dictionaries as priors in this work. D is learned in a set of constraints \mathcal{C}_S by solving

$$\min_{z \in \mathbb{R}^L, D \in \mathcal{C}_S} F_S(z, D) \triangleq \frac{1}{2} \|ADz - y\|_2^2 + \lambda \|z\|_1 . \quad (2)$$

These two problems are equivalent when the prior is invertible, but this is not true in general. A detailed comparison between Analysis and Synthesis is available in [Elad et al., 2007] in the case of a known prior. In Dictionary Learning, studies have also been provided on identifiability [Gribonval et al., 2015] and minima analysis [Haeffele and Vidal, 2015, Agarwal et al., 2016, Sun et al., 2016].

Analysis and Synthesis can be written as bi-level optimization problems to minimize the cost function with respect to the prior only, as mentioned by Mairal et al. [2009] for Dictionary Learning.

$$\min_{\Gamma \in \mathcal{C}_A} G_A(\Gamma) \triangleq F_A(x^*(\Gamma), \Gamma) \quad \text{with} \quad x^*(\Gamma) = \operatorname{argmin}_{x \in \mathbb{R}^n} F_A(x, \Gamma) , \quad (3)$$

$$\min_{D \in \mathcal{C}_S} G_S(D) \triangleq F_S(z^*(D), D) \quad \text{with} \quad z^*(D) = \operatorname{argmin}_{z \in \mathbb{R}^L} F_S(z, D) . \quad (4)$$

Computing the data representation $x^*(\Gamma)$ or $z^*(D)$ is often referred to as the inner problem, while the global minimization is called the outer problem. The most classical constraint set chosen for \mathcal{C}_S , or \mathcal{C}_A , is the Unit Norm (UN), where each atom is normalized. In [Yaghoobi et al., 2013], the authors notice that UN is not sufficient to guarantee that the learned prior is nontrivial with Analysis. Therefore, they introduce the Unit Norm Tight Frame constraint (UNTF) : $\{U \text{ s.t. } UU^\top = I, \forall i \in I \|U_i\|_2 = 1\}$ which is a subset of the Stiefel manifold corresponding to orthonormal k-frames with normalized atoms. Finally, \mathcal{C}_S and \mathcal{C}_A can also be chosen as normalized convolutional kernels. It corresponds to Convolutional Dictionary Learning [Grosse et al., 2007] for Synthesis and to learning finite difference schemes for Analysis – as done in [Chambolle and Pock, 2020, Kobler et al., 2020] for Total Variation. One can notice that TV belongs to this set of constraints.

A critical step in the design and selection of such priors is evaluating the reconstructed signal x . When ground truth data are available, the prior can be learned in a supervised setting. For instance, supervised Dictionary Learning has been used in the context of image processing [Mairal et al., 2009]. Other frameworks based on Plug-and-Play [Brifman et al., 2016] and Deep Learning [Chan et al., 2016, Romano et al., 2017, Rick Chang et al., 2017] propose to integrate a pre-trained denoiser in an iterative algorithm to solve the problem. However, ground truth data is rarely accessible for inverse problems outside audio and imaging. The prior itself also provides information on the structure of the signal. A typical example is the study of magnetoencephalography recordings, where one aims to analyze the electrical activity in the brain from measurements of the magnetic field around the scalp of the patient. In [Dupré la Tour et al., 2018], the authors learn a dictionary directly on the output of the sensors and link the learned patterns to known physiological phenomena.

Classical prior learning methods solve problems (3) and (4) through *Alternating Minimization* (AM) [Mairal et al., 2009, Peyré and Fadili, 2011]. It consists in minimizing the cost function over z or x with a fixed prior D or Γ and then performing gradient descent to optimize the prior with a fixed z or x . While AM provides a simple strategy to perform prior learning, it can be inefficient on large-scale data sets due to the need to precisely resolve the inner problems. Over the past years, many studies have proposed to use algorithm unrolling, either for Analysis [Chambolle and Pock, 2020, Lecouat et al., 2020] or Synthesis [Scetbon et al., 2019, Tolooshams et al., 2020], to overcome that issue. The core idea consists of unrolling the algorithm used to solve the inner problem and then computing the gradient with respect to the prior by back-propagating through the iterates of

this algorithm. This method has been popularized by [Gregor and LeCun \[2010\]](#), who first proposed to unroll ISTA [[Daubechies et al., 2004](#)] – a proximal gradient descent algorithm designed for the Lasso [[Tibshirani, 1996](#)] – to speed up the computation of $z^*(D)$. The $N + 1$ -th layer of this network – called LISTA – is obtained as $z_{N+1} = ST_{\frac{\lambda}{L}}(W^1 y + W^2 z_N)$, with ST being the soft-thresholding operator. This work has led to many contributions aiming at improving this method and providing theoretical justifications in a supervised [[Chen et al., 2018](#), [Liu and Chen, 2019](#)] or unsupervised [[Moreau and Bruna, 2017](#), [Ablin et al., 2019](#)] setting. For such unrolled algorithms, the weights W^1 and W^2 can be re-parameterized as functions of D – as illustrated in [Figure A.1](#) in the appendix – such that the output $z_N(D)$ of the network matches the result of N iterations of ISTA, *i.e.*

$$W_D^1 = \frac{1}{L}(AD)^\top \quad \text{and} \quad W_D^2 = (I - \frac{1}{L}(AD)^\top AD), \quad \text{where} \quad L = \|AD\|^2 \quad (5)$$

Then, the dictionary can be learned by minimizing the loss $F_S(z_N(D), D)$ over D with back-propagation. We will refer to this approach as *Deep Dictionary Learning* (DDL).

Variants of DDL with different kinds of regularization [[Scetbon et al., 2019](#), [Tolooshams et al., 2020](#), [Lecouat et al., 2020](#)], image processing based on metric learning [[Tang et al., 2020](#)], and classification tasks with scattering [[Zarka et al., 2019](#)] have been proposed in the literature, among others. Networks adapted to learn Analysis priors have also been studied [[Chambolle and Pock, 2020](#), [Jiu and Pustelnik, 2021](#)]. Indeed, the corresponding optimization problem can be solved with a primal-dual algorithm like Condat-Vu [[Condat, 2013](#), [Vũ, 2013](#)], which consists of a primal descent with step size τ and a dual ascent with step size σ . N iterations of Condat-Vu can be unrolled to obtain $x_N(\Gamma)$ on the same principle as LISTA and DDL in Synthesis. This architecture is close to the network designed in [[Jiu and Pustelnik, 2021](#)] in a supervised learning context, and we will denote it *Deep Primal-Dual Prior Learning* (DPDPL). However, most of these works have only considered supervised learning, which is not adapted to inverse problems with no ground truth data.

In this work, we propose to study how unrolled algorithms can be used to learn priors from the point of view of bi-level optimization, with access to degraded observations only. First, we compare Deep Learning to Alternating Minimization in [Section 2](#). We study the convergence rate of the estimated gradients for both methods and show that unrolled algorithms better estimate the gradient for a small number of iterations. We also analyze the convergence of the Jacobian computed with automatic differentiation towards the true Jacobian and find out that its stability is guaranteed on the support of the sparse code only. Then we compare networks based on Analysis and Synthesis in [Section 3](#). We evaluate their performance on several inverse problems without access to any ground truth data for three classes of priors. While Analysis can achieve good results, Synthesis is more robust and performs better. Finally, we apply our method to pattern and structure learning from degraded measurements.

2 Bi-level optimization for Deep Prior Learning

As $x^*(\Gamma)$ and $z^*(D)$ do not have a closed-form expression, neither G_A nor G_S can be computed directly. A solution is to replace the inner problem $z^*(D)$ (resp. $x^*(\Gamma)$) by an approximation $z_N(D)$ (resp. $x_N(\Gamma)$) obtained through N iterations of a numerical optimization algorithm or its unrolled version. This reduces the problem to minimizing $G_{S,N}(D) \triangleq F_S(z_N(D), D)$ for Synthesis or $G_{A,N}(\Gamma) \triangleq F_A(x_N(\Gamma), \Gamma)$ for Analysis. When results apply to both Analysis and Synthesis, we use G_N, G to denote either $G_{A,N}, G_A$ or $G_{S,N}, G_S$. The first question is how sub-optimal global solutions of G_N are compared to the ones of G . [Proposition 2.1](#) shows that the global minima of G_N converge as fast as the numerical approximation z_N or x_N in function value.

Proposition 2.1. *Let $U^* = \operatorname{argmin}_U G(U)$ and $U_N^* = \operatorname{argmin}_U G_N(U)$, where N is the number of unrolled iterations. We denote by $K(U^*)$ a constant depending on U^* , and by $C(N)$ the convergence speed of the algorithm, which approximates the inner problem solution. We have*

$$G_N(U_N^*) - G(U^*) \leq K(U^*)C(N) .$$

The proofs of all theoretical results are deferred to [Appendix B](#). [Proposition 2.1](#) implies that when z_N is computed with FISTA, the function value for global minima of G_N converges with speed $C(N) = \frac{1}{N^2}$ towards the value of the global minima of F_S . Therefore, suitable solutions for (3) and (4) can be obtained by solving the inner problem approximately, given that the optimization

procedure is efficient enough to find a proper minimum of G_N . As the computational cost of z_N or x_N increases with N , the choice of N results in a trade-off between the precision of the solution and the computational efficiency, which has to be considered for large data sets.

2.1 Gradient estimation in Dictionary Learning

Analysis or Synthesis based prior learning is a non-convex problem, meaning that good or poor local minima of G_N may be reached depending on the initialization, the optimization path, and the structure of the problem. Therefore, a gradient descent on G_N is not guaranteed to find an adequate minimizer of G , even though the global optima might be close. While complete theoretical analysis of these problems is arduous, we propose to study the correlation between the gradient obtained with G_N and the actual gradient of G . We restrain our study to Synthesis, as the gradient of G is well defined in this case. We use the notation from [Ablin et al., 2020] for the gradient estimates, which analyzes similar gradient estimates for bi-level optimization with a smooth and differentiable loss.

Once $z^*(D)$ is known, [Danskin, 1967, Thm 1] states that $g^*(D) = \nabla G_S(D)$ is equal to $\nabla_2 F_S(z^*(D), D)$, where ∇_2 indicates that the gradient is computed relatively to the second variable in F_S . Even though the inner problem is non-smooth, this result holds as long as the solution $z^*(D)$ is unique. In the following, we will assume that $D^\top D$ is invertible on the support of $z^*(D)$, which implies the uniqueness of $z^*(D)$. While both AM and DDL try to minimize G_S , they differ in how they estimate its gradient. AM relies on the analytical formula of g^* and uses an approximation z_N of z^* , leading to the approximate gradient $g_N^1(D) = \nabla_2 F_S(z_N(D), D)$. We evaluate how well g_N^1 approximates g^* in Proposition 2.2.

Proposition 2.2. *Let $D \in \mathbb{R}^{n \times L}$. Then, there exists a constant $L > 0$ such that for every number of iterations N*

$$\|g_N^1 - g^*\| \leq L \|z_N(D) - z^*(D)\| . \quad (6)$$

Proposition 2.2 shows that g_N^1 converges as fast as the iterates of ISTA converges. DDL computes the gradient via automatic differentiation through $z_N(D)$. As opposed to AM, this corresponds to a direct minimization of the loss $G_N(D)$. Automatic differentiation yields a gradient $g_N^2(D)$ such that

$$g_N^2(D) \in \nabla_2 F_S(z_N(D), D) + J_N^+ \left(\partial_1 F_S(z_N(D), D) \right) , \quad (7)$$

where $J_N : \mathbb{R}^{n \times L} \rightarrow \mathbb{R}^L$ is the weak Jacobian of $z_N(D)$ with respect to D and J_N^+ denotes its adjoint operator. Here, the product between J_N^+ and $\partial_1 F_S(z_N(D), D)$ is computed via automatic differentiation.

Proposition 2.3. *Let $D \in \mathbb{R}^{n \times L}$. Let S^* be the support of $z^*(D)$, S_N be the support of z_N and $\tilde{S}_N = S_N \cup S^*$. Let $R(J, \tilde{S}) = J^+ (\nabla_{z,z}^2 f(z^*, D) \odot \mathbb{1}_{\tilde{S}}) + \nabla_{D,z}^2 f(z^*, D) \odot \mathbb{1}_{\tilde{S}}$. Then there exists a constant $L > 0$ and a sub-sequence of (F)ISTA iterates $z_{\phi(N)}$ such that for all $N \in \mathbb{N}$:*

$$\begin{aligned} \exists g_{\phi(N)}^2 \in \nabla_D f(z_{\phi(N)}, D) + J_{\phi(N)}^+ \left(\nabla_z f(z_{\phi(N)}, D) + \lambda \partial_{\|\cdot\|_1}(z_{\phi(N)}) \right) \text{ s.t. :} \\ \left\| g_{\phi(N)}^2 - g^* \right\| \leq \left\| R(J_{\phi(N)}, \tilde{S}_{\phi(N)}) \right\| \|z_{\phi(N)} - z^*\| + \frac{L}{2} \|z_{\phi(N)} - z^*\|^2 . \end{aligned} \quad (8)$$

Proposition 2.3 shows that g_N^2 may converge faster than g_N^1 once the support is reached. However, its estimation is difficult because of the sub-differential. The convergence behavior of g_N^2 is also driven by $R(J_N, \tilde{S}_N)$ and thus by the weak Jacobian computed via back-propagation. For the sake of clarity, we only carry out our analysis in the case of denoising with $A = I$. For linear inverse problems, one can consider AD instead of D and use the chain rule. We first compute a closed-form expression of the weak Jacobian of $z^*(D)$ and $z_N(D)$, in the case of denoising. We then show that $R(J_N, \tilde{S}_N) \leq L \|J_N - J^*\|$ and we analyze the convergence of J_N towards J^* .

Study of the Jacobian. The computation of the Jacobian can be done by differentiating through ISTA. In Theorem 2.4, we show that J_{N+1} depends on J_N and the past iterate z_N , and converges towards a fixed point. This formula can be used to compute the Jacobian during the forward pass, avoiding the computational cost of back-propagation and saving memory.




Figure 1: Average convergence of J_l^N towards J_l^* for two random synthetic samples. $\|J_l^* - J_l^N\|$ converges linearly on the support in both cases. However, for sample 2, full back-propagation makes the convergence unstable. Truncated back-propagation improves the convergence behavior.

Theorem 2.4. At iteration $N + 1$ of ISTA, the weak Jacobian of z_{N+1} relatively to D_l , where D_l is the l -th row of D , is given by induction:

$$\frac{\partial(z_{N+1})}{\partial D_l} = \mathbb{1}_{|z_{N+1}|>0} \odot \left(\frac{\partial(z_N)}{\partial D_l} - \frac{1}{L} \left(D_l z_N^\top + (D_l^\top z_N - y_l) Id_n + D^\top D \frac{\partial(z_N)}{\partial D_l} \right) \right) . \quad (9)$$

$\frac{\partial(z_N)}{\partial D_l}$ will be denoted by J_l^N . It converges towards the weak Jacobian J_l^* of $z^*(D)$, whose values are

$$J_l^*_{S^*} = -(D_{:,S^*}^\top D_{:,S^*})^{-1} (D_l z^*{}^\top + (D_l^\top z^* - y_l) Id_n)_{S^*} , \quad (10)$$

on the support S^* of z^* , and 0 elsewhere. Moreover, $R(J^*, S^*) = 0$.

This result is similar to Bertrand et al. [2020] where the Jacobian of z is computed over λ to perform hyper-parameter optimization in Lasso-type models. Note that the weak Jacobian with a sensing matrix A is also derived in Corollary B.2. Using $R(J^*, S^*) = 0$, we can write

$$\left\| R(J_N, \tilde{S}_N) \right\| \leq \left\| R(J_N, \tilde{S}_N) - R(J^*, S^*) \right\| \leq L \|J_N - J^*\| , \quad (11)$$

as $\|\nabla_{z,z}^2 f(z^*, D)\|_2 = L$. If the back-propagation were to output an accurate estimate J_N of the weak Jacobian J^* , $\left\| R(J_N, \tilde{S}_N) \right\|$ would be 0, and the convergence rate of g_N^2 could be twice as fast as the one of g_N^1 . To quantify this, we now analyze the convergence of J_N towards J^* . In Proposition 2.5, we compute an upper bound of $\|J_l^N - J_l^*\|$ with possible usage of truncated back-propagation [Shaban et al., 2019]. Truncated back-propagation of depth K corresponds to an initial estimate of the Jacobian $J^{N-K} = 0$ and iterating the induction (9).

Proposition 2.5. Let N be the number of iterations and K be the back-propagation depth. We assume that $\forall n \geq N - K$, $S^* \subset S_n$. Let $\bar{E}_N = S_n \setminus S^*$, let L be the largest eigenvalue of $D_{:,S^*}^\top D_{:,S^*}$, and let μ_n be the smallest eigenvalue of $D_{:,S_n}^\top D_{:,S_{n-1}}$. Let $B_n = \left\| P_{\bar{E}_n} - D_{:, \bar{E}_n}^\top D_{:, S^*}^\dagger P_{S^*} \right\|$, where P_S is the projection on \mathbb{R}^S and D^\dagger is the pseudo-inverse of D . We have

$$\|J_l^N - J_l^*\| \leq \prod_{k=1}^K \left(1 - \frac{\mu_{N-k}}{L}\right) \|J_l^*\| + \frac{2}{L} \|D_l\| \sum_{k=0}^{K-1} \prod_{i=1}^k \left(1 - \frac{\mu_{N-i}}{L}\right) \left(\|z_l^{N-k} - z_l^*\| + B_{N-k} \|z_l^*\| \right) .$$

Proposition 2.5 reveals multiple stages in the Jacobian estimation. First, one can see that if all iterates used for the back-propagation lie on the support S^* , the Jacobian estimate has a quasi-linear convergence, as shown in the following corollary.

Corollary 2.6. Let $\mu > 0$ be the smallest eigenvalue of $D_{:,S^*}^\top D_{:,S^*}$. Let $K \leq N$ be the back-propagation depth and let $\Delta_N = F_S(z_N, D) - F_S(z^*, D) + \frac{L}{2} \|z_N - z^*\|$. Suppose that $\forall n \in [N - K, N]$; $S_n \subset S^*$. Then, we have

$$\|J_l^* - J_l^N\| \leq \left(1 - \frac{\mu}{L}\right)^K \|J_l^*\| + K \left(1 - \frac{\mu}{L}\right)^{K-1} \|D_l\| \frac{4\Delta_{N-K}}{L^2} .$$

Once the support is reached, ISTA also converges with the same linear rate $(1 - \frac{\mu}{L})$ and thus, the gradient estimate g_N^2 converges almost twice as fast as g_N^1 , as $\mathcal{O}(K(1 - \frac{\mu}{L})^{2K})$, similarly to [Ablin et al., 2020].

Second, Proposition 2.5 shows that $\|J_l^* - J_l^N\|$ may increase when the support is not well-estimated, leading to a deterioration of the gradient estimate. This is due to an accumulation of errors materialized by the sum in the right-hand side of the inequality, as the term $B_N \|z^*\|$ may not vanish to 0 as long as $S_N \not\subset S^*$. Interestingly, once the support is reached at iteration $S < N$, the errors converge linearly towards 0, and we recover the fast estimation of g^* with g^2 . This result suggests that a low number of iterations might better estimate the Jacobian and the overall gradient when far from the support.

Figure 1 confirms the linear convergence of J_l^N once the support is reached. Moreover, the behavior in early iterations might be poor when the number of iteration grows, leading to exploding gradient, as shown in the second case. In this case, using a small number of iterations or truncated back-propagation may be necessary to prevent accumulating errors from the previous iterations.

2.2 Numerical evaluation of the gradient and truncated back-propagation

The principal interest of unrolled algorithms is to use them with a small number of layers – or, equivalently, a small number of iterations. Figure 2 shows that DDL can outperform AM using only 30 iterations of FISTA on a problem of image denoising. Therefore, it is of interest to empirically quantify the impact of back-propagation depth on the quality of the gradient estimate. Let $\langle g_N^1, g^* \rangle$ and $\langle g_N^2, g^* \rangle$ be the angles between the estimates g_N^1 and g_N^2 , and the true gradient g^* , where $\langle x, y \rangle = \frac{\text{Tr}(x^\top y)}{\|x\| \|y\|}$. In Figure 2, we display the relative difference $\frac{\langle g_N^2, g^* \rangle - \langle g_N^1, g^* \rangle}{1 - \langle g_N^1, g^* \rangle}$ depending on iterations N and back-propagation depth K . When the back-propagation goes too deep, the performance of g_N^2 decreases compared to g_N^1 . On the contrary, there is a gain in the first 30 iterations, especially with truncated back-propagation, which makes the estimate more stable, as stated in Proposition 2.5.




Figure 2: Optimal loss values of DDL and AM for patches from a noisy image depending on the number of iterations (left). Average relative difference between $\langle g_N^2, g^* \rangle$ and $\langle g_N^1, g^* \rangle$ on the optimization path for patches from a noisy image (right).

3 Analysis vs Synthesis

In the previous section, theoretical studies highlighted that unrolling might accelerate the gradient computation to minimize G compared to AM for Synthesis. However, this is not sufficient to guarantee that the gradient descent solution will be better in practice. In particular, performances critically depend on the problem structure, and it is of interest to evaluate the differences between Analysis and Synthesis in the framework of unrolled algorithms. Here, we empirically compare both formulations for prior learning without access to ground truth data regarding prior recovery, quality of the reconstruction, and behavior of the objective function. First, we compare the reconstruction performance of networks with a small number of layers, typically around 20, on an inpainting task with real images. Then, we check the ability of DDL and DPDPL to identify a synthetic prior on degraded data. The computations have been performed on a GPU NVIDIA Tesla V100-DGXS 32GB using PyTorch [Paszke et al., 2019].¹

Optimization. For UN and Conv+UN, the optimization is performed with projected gradient descent, and matrix multiplications are replaced with convolutions in Pytorch when necessary. For UNTF, the new point on the manifold is computed with a Cayley transform as proposed by Wen and

¹Code is available at <https://github.com/bmalezieux/plipy>.




Figure 3: Inpainting on a color image of size 200×200 , with 50% lacking pixels. Each RGB channel is processed separately.

Yin [2013] and Li et al. [2017]. We rely on full-batch gradient descent and line search for the sake of precision and robustness of our experiments, as done in [Ablin et al., 2019].


Improvement of precision. As stated before, using a low number of iterations allows for efficient and stable computations, but this can make the sparse code less precise. To compensate for imprecise representations, one can learn the steps sizes of (F)ISTA and Condat-Vu, in order to accelerate the convergence, as proposed by Ablin et al. [2019] for LISTA. To avoid poor results due to large degrees of freedom in unsupervised learning, we propose a method in two steps to refine the initialization of the prior before relaxing the constraints on the step sizes:

1. We learn the prior with fixed step sizes given by convergence conditions. For Synthesis, the step size is $\frac{1}{L}$ where $L = \|AD\|^2$. For Analysis, we take $\sigma = 1$ and $\tau = \frac{1}{\frac{1}{2}\|A\|^2 + \sigma\|\Gamma^\top\|^2}$. Lipschitz constants are computed at each gradient step, outside of the scope of the network graph.
2. Then, once convergence is reached, we jointly learn the step-sizes and the prior. Both are still updated using gradient descent with line search to ensure stable optimization. This improves the convergence of the networks towards $z_N^*(D)$ or $x_N^*(\Gamma)$. Note that steps sizes learning may lead to poor results when the data is too noisy.

3.1 Reconstruction on real data

To illustrate the performance of unrolled algorithms for prior learning, we consider an unsupervised inpainting task. The original image is degraded with additive Gaussian noise to get an SNR of 10dB, and then a fraction of its pixels are randomly removed. We learn priors and recover the image with unrolled algorithms using 20 iterations of FISTA or Condat-Vu. For UN and UNTF, the image is decomposed into patches, while convolutions can be applied directly to the whole image. Figure 3 displays the best results for each method. Synthesis-based methods achieve better image reconstruction compared to their Analysis counter-part. Convolutions improve the performances for Analysis, especially with small kernels – size 4 in that case, compared to 10 for Synthesis – emphasizing the importance of the selected constraints set. These results also demonstrate the ability of Synthesis to recover a good prior in the case of compressed or lacking information. Finally, Synthesis leads to more realistic reconstructions than TV-based methods. Additional results are provided in Figure A.2 in the appendix for a gray level image with different rates of lacking pixels and several patch dimensions.

The ability of gradient descent to find adequate local minima strongly depends on the structure of the problem. To quantify this, we evaluate the variation of PSNR depending on the percentage of lacking pixels for 50 random initializations in the context of convolutional prior learning. Figure 4 shows that Synthesis is more robust to random initialization, and almost all local minima are similar in terms of reconstruction quality. On the contrary, Analysis suffers from a significant number of poor local minima, and the quality of reconstruction highly depends on the initialization. Extra experiments in Figure A.3 show



Synthesis and Analysis minima





Figure 4: PSNR depending on the rate of lacking pixels without noise for 50 random initializations with convolutions (top). Synthesis is more robust than Analysis. Comparison of the loss landscapes in 1D (middle) and 2D (bottom). Analysis (left) is poorly conditioned compared to Synthesis (right) and suffers from the bad local minima.




Figure 5: Prior recovery in the context of denoising depending on the SNR and the dimension of y . DPDPL performs well only with the UNTF constraint when the prior is almost square.

similar behavior for different noise intensity (SNR) for fixed missing pixel ratio. We propose to study the loss landscape for Analysis and Synthesis with the help of visualization techniques presented in Li et al. [2018]. The 3D landscape is displayed in Figure 4 using the Python library K3D-Jupyter². We also compare the shapes of local minima in 1D by computing the normalized values of the loss along a line between two local minima. These representations of the loss landscape clearly show that Synthesis is much smoother than its Analysis counterpart and the different minima have similar performances in this case. Synthesis locally behaves like a convex function while Analysis landscape is much steeper with poor local minima.

3.2 Prior recovery on synthetic data

We now evaluate the prior recovery performance of unrolled algorithms on synthetic data. For Synthesis, a sparse code z is drawn from a Bernoulli Gaussian distribution of mean 0, variance σ^2 , and Bernoulli parameter p . The signal is computed as $D_{ref}z$. For Analysis, we adopt the data generation process proposed in Elad et al. [2007]. The sparsity s of the signal is drawn from a Bernoulli distribution of parameter p , and $L - s$ rows of Γ_{ref}^\top are chosen such that the corresponding sub-matrix is full-rank.

The generated signal is $u - (\Gamma_{ref}^\top)^\dagger \Gamma_{ref} u$ with $u \sim \mathcal{N}(0, \sigma^2)$. The recovered dictionary, or prior, evaluation should reflect its ability to generate the same signals as the ground truth data. We compare the atoms using their correlation and denote as C the cost matrix whose entry i, j compare the atoms i of the first prior and j of the second. We define a sign and permutation invariant metric $S(C) = \max_{\sigma \in \mathfrak{S}_L} \frac{1}{L} \sum_{i=1}^L |C_{\sigma(i), i}|$, where \mathfrak{S}_L is the group of permutations of $[1, L]$. This metric, proposed in [Moreau and Gramfort, 2020] for convolutions, corresponds to the best linear sum assignment on the cost matrix C , and it can be computed using the Hungarian algorithm.

Denoising. The data are generated with a UNTF dictionary of 50 atoms. We compare the maximal score for several values of λ and depth (between 5 and 50 iterations) depending on the Signal to Noise Ratio (SNR) and the dimension of the observation. The SNR is defined as $10 \log_{10}(\frac{\sigma_a^2}{\sigma_b^2})$ where σ_b^2 is the variance of the noise. We assume that the number of atoms is known, and the results are provided in Figure 5. DDL behaves like standard Dictionary Learning, implemented in scikit-learn [Pedregosa et al., 2012]. The experiment confirms that UN does not work. Moreover, UNTF works only when the measurement dimension is close to the dimension of the signal and for a large SNR. We also compare the score depending on λ and the number of iterations in Figure A.4 in the appendix and find that unrolled algorithms perform well with a small number of layers, as anticipated in Section 2.2.

Compressed sensing. The data are generated with an orthogonal dictionary of size 50 in a compressed sensing scenario. The results are presented in Table 1. DDL outperforms DPDPL, especially when the size of the measurements decreases. Analysis seems poorly conditioned compared to

²Package available at <https://github.com/K3D-tools/K3D-jupyter>.

Table 1: Orthogonal dictionary recovery for compressed sensing with SNR = 20dB

Algo	dim. $y = 50$	dim. $y = 40$	dim. $y = 30$	dim. $y = 20$	dim. $y = 10$
DDL	0.98	0.89	0.76	0.61	0.36
DPDPL	0.96	0.78	0.57	0.43	0.37




Figure 6: Inpainting task on a noisy image generated using seven letters {N, e, u, r, I, P, S} with 50% of pixels removed. DDL recovers a dictionary (left) and the image (right).

Synthesis, and the UNTF constraint is not sufficient to recover the orthogonal prior. Compared to other regularization techniques in inverse problems, the main interest of prior learning is to provide meaningful insights into the data. This can be achieved with the method studied in this paper, but the results depend on the structure of the data.

Convolutional Dictionary Learning. We generate a text image from seven letters {N, e, u, r, I, P, S}, and degrade it by removing pixels and adding noise. We compare the quality of reconstruction of TV and DDL in Figure 6, for a dictionary of 30 atoms of size 20×20 . TV cannot adapt to the structure of the data, while DDL with convolutions learns a dictionary of patterns from the measurements and recovers the image simultaneously. The quality of the recovery highly depends on the amount of available information. Figure 7 (left) shows the average recovery score of CDL with initialization from patches of the image. The dictionary can be well-recovered for a large fraction of missing pixels, given that the SNR is not too low. A comparison to random initialization is provided in Figure A.5 in appendix. When there is not enough information, initialization from patches does not work.

Convolutional prior learning with Analysis.

Analysis is not adapted to recover patterns but can be applied to signals generated from recurrent relations such as Partial Differential Equations (PDEs) – such as autoregressive models or textures. For instance, Analysis has been used in Kitić et al. [2015] to regularize inverse problems involving physical signals with PDEs. In fact, a discrete scheme from an ODE is equivalent to a convolution. Let’s consider the function $x : t \rightarrow A \sin(\omega t)$. It satisfies $\ddot{x} + \omega^2 x = 0$, and the associated Euler discretization is $x_{n+1} + (dt^2\omega^2 - 2)x_n + x_{n-1} = 0$, where dt is the step size. In Figure 7, we measure the ability of convolutional prior learning with Analysis to recover a filter close to $[1, dt^2\omega^2 - 2, 1]$, which corresponds to this discretization, by estimating the frequency $f_{CPL} = \frac{\omega}{2\pi}$ given noisy data. The algorithm is successful when the sampling rate is high enough. However, the question of the kernel identifiability arises: while the algorithm denoises the signal, it fails to recover meaningful kernels in 2D without additional constraints.




Figure 7: Average recovery score in Deep CDL depending VS SNR (dB) and rate of lacking pixels (%) (left). $MSE(f_0, f_{CPL})$ VS sampling rate (Hz) and signal frequency (Hz) (right).

4 Discussion

Prior knowledge of the signal plays a key role in inverse problems resolution, especially without access to ground truth data. This work showed that unrolled algorithms offer multiple advantages in dictionary and prior learning from observed data. They can be faster than alternating minimization, and they perform well on various kinds of inverse problems even with a small number of layers or iterations thanks to automatic differentiation. However, their behavior highly depends on the structure of the optimization problem they solve.

Future work. The principal limit of dictionary and prior learning in unsupervised inverse problems is that the amount of available information largely determines the quality of the result. Up to our

knowledge, no study has been carried out on the theoretical properties of Dictionary Learning on degraded observed measurements. It would be of interest to quantify the impact of the dimension reduction on the identifiability of the dictionary. Moreover, unrolled algorithms are able to perform prior learning with a few iterations only, as we empirically demonstrated in this paper. Providing stronger theoretical elements would be necessary to confirm our first findings. Finally, developing and analyzing stochastic optimization algorithms for such methods would allow to scale up prior learning to larger data sets.

References

- P. Ablin, T. Moreau, M. Massias, and A. Gramfort. Learning step sizes for unfolded sparse coding. In *Advances in Neural Information Processing Systems*, pages 13100–13110, 2019.
- P. Ablin, G. Peyré, and T. Moreau. Super-efficiency of automatic differentiation for functions defined as a minimum. In *Proceedings of the 37th International Conference on Machine Learning*, pages 32–41, 2020.
- A. Agarwal, A. Anandkumar, P. Jain, and P. Netrapalli. Learning sparsely used overcomplete dictionaries via alternating minimization. *SIAM Journal on Optimization*, 26(4):2775–2799, 2016.
- M. Aharon, M. Elad, and A. Bruckstein. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54:4311 – 4322, 2006.
- Q. Bertrand, Q. Kloppenstein, M. Blondel, S. Vaïter, A. Gramfort, and J. Salmon. Implicit differentiation of lasso-type models for hyperparameter optimization. In *International Conference on Machine Learning*, pages 810–821. PMLR, 2020.
- A. Brifman, Y. Romano, and M. Elad. Turning a denoiser into a super-resolver using plug and play priors. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 1404–1408. IEEE, 2016.
- A. Chambolle and T. Pock. Learning consistent discretizations of the total variation. <https://hal.archives-ouvertes.fr/hal-02982082/>, 2020.
- S. H. Chan, X. Wang, and O. A. Elgandy. Plug-and-play admm for image restoration: Fixed-point convergence and applications. *IEEE Transactions on Computational Imaging*, 3(1):84–98, 2016.
- X. Chen, J. Liu, Z. Wang, and W. Yin. Theoretical linear convergence of unfolded ista and its practical weights and thresholds. *Advances in Neural Information Processing Systems*, 2018.
- L. Condat. A primal–dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms. *Journal of Optimization Theory and Applications*, 158(2):460–479, 2013.
- J. M. Danskin. *Theory of Max-Min and Its Application to Weapons Allocation Problems*. Springer Berlin Heidelberg, Berlin/Heidelberg, 1967.
- I. Daubechies, M. Defrise, and C. Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constrains. *Communications on Pure and Applied Mathematics*, 57, 2004.
- C.-A. Deledalle, S. Vaïter, J. Fadili, and G. Peyré. Stein unbiased gradient estimator of the risk (sugar) for multiple parameter selection. *SIAM Journal on Imaging Sciences*, 7(4):2448–2487, 2014.
- T. Dupré la Tour, T. Moreau, M. Jas, and A. Gramfort. Multivariate convolutional sparse coding for electromagnetic brain signals. *Advances in Neural Information Processing Systems*, 31:3292–3302, 2018.
- M. Elad. *Sparse and redundant representations: from theory to applications in signal and image processing*. Springer Science & Business Media, 2010.
- M. Elad, P. Milanfar, and R. Rubinstein. Analysis versus synthesis in signal priors. *Inverse Problems*, 23:947, 2007.

- A. Gramfort, M. Kowalski, and M. Hämmäläinen. Mixed-norm estimates for the m/eeg inverse problem using accelerated gradient methods. *Physics in medicine and biology*, 57:1937–61, 2012.
- K. Gregor and Y. LeCun. Learning fast approximations of sparse coding. *International conference on machine learning*, pages 399–406, 2010.
- R. Gribonval, R. Jenatton, and F. Bach. Sparse and spurious: dictionary learning with noise and outliers. *IEEE Transactions on Information Theory*, 61(11):6298–6319, 2015.
- R. Grosse, R. Raina, H. Kwong, and A. Y. Ng. Shift-Invariant Sparse Coding for Audio Classification. *Cortex*, 8:9, 2007.
- B. D. Haeffele and R. Vidal. Global optimality in tensor factorization, deep learning, and beyond. *arXiv preprint arXiv:1506.07540*, 2015.
- M. Jiu and N. Pustelnik. A deep primal-dual proximal network for image restoration. *IEEE Journal of Selected Topics in Signal Processing*, 15(2):190–203, 2021.
- S. Kitić, L. Albera, N. Bertin, and R. Gribonval. Physics-driven inverse problems made tractable with cosparsity regularization. *IEEE Transactions on Signal Processing*, 64(2):335–348, 2015.
- E. Kobler, A. Effland, K. Kunisch, and T. Pock. Total deep variation for linear inverse problems. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7549–7558, 2020.
- B. Lecouat, J. Ponce, and J. Mairal. A flexible framework for designing trainable priors with adaptive smoothing and game encoding. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein. Visualizing the loss landscape of neural nets. In *Advances in neural information processing systems*, pages 6389–6399, 2018.
- Y. Li, S. Ding, Z. Li, X. Li, and B. Tan. Dictionary learning in the analysis sparse representation with optimization on stiefel manifold. In *IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 1270–1274. IEEE, 2017.
- J. Liu and X. Chen. Alista: Analytic weights are as good as learned weights in lista. In *International Conference on Learning Representations (ICLR)*, 2019.
- J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*, 11, 2009.
- S. Mallat. *A Wavelet Tour of Signal Processing*. Academic press, 2008.
- T. Moreau and J. Bruna. Understanding neural sparse coding with matrix factorization. In *International Conference on Learning Representation (ICLR)*, 2017.
- T. Moreau and A. Gramfort. DicoDile: Distributed convolutional dictionary learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- B. A. Olshausen and D. J. Field. Sparse coding with an incomplete basis set: A strategy employed by \protect{V1}. *Vision Research*, 37(23):3311–3325, 1997.
- A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, and G. Louppe. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12, 2012.
- G. Peyré and J. M. Fadili. Learning analysis sparsity priors. In *Sampta’11*, pages 4–pp, 2011.

- A. Ribes and F. Schmitt. Linear inverse problems in imaging. *IEEE Signal Processing Magazine*, 25(4):84–99, 2008.
- J. Rick Chang, C.-L. Li, B. Póczos, B. Vijaya Kumar, and A. C. Sankaranarayanan. One network to solve them all—solving linear inverse problems using deep projection models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5888–5897, 2017.
- Y. Romano, M. Elad, and P. Milanfar. The little engine that could: Regularization by denoising (red). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844, 2017.
- M. Scetbon, M. Elad, and P. Milanfar. Deep k-svd denoising. *arXiv preprint arXiv:1909.13164*, 2019.
- A. Shaban, C.-A. Cheng, N. Hatch, and B. Boots. Truncated back-propagation for bilevel optimization. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1723–1732. PMLR, 2019.
- J.-L. Starck. Sparsity and inverse problems in astrophysics. *Journal of Physics: Conference Series*, 699, 2016.
- J. Sun, Q. Qu, and J. Wright. Complete dictionary recovery over the sphere i: Overview and the geometric picture. *IEEE Transactions on Information Theory*, 63(2):853–884, 2016.
- W. Tang, E. Chouzenoux, J.-C. Pesquet, and H. Krim. Deep transform and metric learning network: Wedding deep dictionary learning and neural networks. *arXiv preprint arXiv:2002.07898*, 2020.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B*, 58:267–288, 1996.
- B. Tolooshams, S. Dey, and D. Ba. Deep residual autoencoders for expectation maximization-inspired dictionary learning. *IEEE Transactions on Neural Networks and Learning Systems*, PP:1–15, 2020.
- B. C. Vũ. A splitting algorithm for dual monotone inclusions involving cocoercive operators. *Advances in Computational Mathematics*, 38(3):667–681, 2013.
- Z. Wen and W. Yin. A feasible method for optimization with orthogonality constraints. *Mathematical Programming*, 142(1-2):397–434, 2013.
- M. Yaghoobi, S. Nam, R. Gribonval, and M. E. Davies. Constrained overcomplete analysis operator learning for cosparse signal modelling. *IEEE Transactions on Signal Processing*, 61(9):2341–2355, 2013.
- J. Zarka, L. Thiry, T. Angles, and S. Mallat. Deep network classification by scattering and homotopy dictionary learning. *arXiv preprint arXiv:1910.03561*, 2019.

A Extra figures and experimental results




Figure A.1: Illustration of LISTA for Dictionary Learning with initialization $z_0 = 0$ for $N = 3$. $W_D^1 = \frac{1}{L}(AD)^\top$, $W_D^2 = (I - \frac{1}{L}(AD)^\top AD)$, where $L = \|AD\|^2$. The result $z_N(D)$ output by the network is an approximation of the solution of the LASSO.




Figure A.2: Shows the maximal PSNR obtained for an inpainting task on a grayscale image of size 128×128 , depending on the dimension of the patches and the percentage of pixels removed. Several values of λ and number of atoms are evaluated for several runs. 20 iterations are unrolled in this example. Synthesis and Analysis with UNTF constraint perform better with large patches, while Analysis with convolutions is better conditioned with small kernels. Synthesis performs better in general, even with a large rate of lacking pixels.




Figure A.3: Shows the PSNR obtained on an inpainting task on a grayscale image of size 128×128 with 50% pixels removed, depending on the SNR (dB) for 50 random initializations with convolutions and 20 layers. The results are similar to those observed in Figure 4: Synthesis is more robust to random initialization, even with a low SNR.





Figure A.4: Show the prior recovery score depending on λ and the number of iterations for an orthogonal dictionary of 50 atoms in a context of denoising with small noise value (SNR = 20 dB). Unrolled algorithms reach good performance level with a small number of iterations, between 20 and 30 depending on the value of λ . If the value of λ is high enough, the convergence is very fast. For small values of λ , unrolled algorithms perform better than standard Dictionary Learning. In this simple case, DPDPL with UNTF works well.



h!

Figure A.5: Shows the average recovery score in Deep CDL for 10 random initializations or 10 initializations from patches, depending on the SNR (dB) and the rate of lacking pixels (%), for a fixed value of λ and 20 layers. CDL performs well when the level of noise is not too high. Both initialization methods are equivalent, except when there is not enough information in the image. In this case, initial patches are almost zero and lead to poor results.




Figure A.6: Shows the loss landscapes for Analysis (left) and Synthesis (right). Synthesis loss landscape is locally convex and very smooth, making the optimization easy. Analysis loss landscape is non-smooth and non-convex, and unequal local minima can be found by gradient descent depending on the initialization.

B Proofs of theoretical results

This section gives the proofs for the various theoretical results in the paper.

B.1 Proof of Proposition 2.1.

Proposition 2.1. *Let $U^* = \operatorname{argmin}_U G(U)$ and $U_N^* = \operatorname{argmin}_U G_N(U)$, where N is the number of unrolled iterations. We denote by $K(U^*)$ a constant depending on U^* , and by $C(N)$ the convergence speed of the algorithm, which approximates the inner problem solution. We have*

$$G_N(U_N^*) - G(U^*) \leq K(U^*)C(N) .$$

We take the example of Synthesis. Let $G(D) \triangleq F_S(z^*(D), D)$ and $G_N(D) \triangleq F_S(z_N(D), D)$ where $z^*(D) = \operatorname{argmin}_{z \in \mathbb{R}^L} F_S(z, D)$ and $z_N(D) = \text{FISTA}(D, N)$. Let $D^* = \operatorname{argmin}_D G(D)$ and $D_N^* = \operatorname{argmin}_D G_N(D)$:

$$G_N(D_N^*) - G(D^*) = G_N(D_N^*) - G_N(D^*) + G_N(D^*) - G(D^*) \quad (12)$$

$$= F_S(z_N(D_N^*), D_N^*) - F_S(z_N(D^*), D^*) \quad (13)$$

$$+ F_S(z_N(D^*), D^*) - F_S(z(D^*), D^*) \quad (14)$$

By definition of D_N^* :

$$F_S(z_N(D_N^*), D_N^*) - F_S(z_N(D^*), D^*) \leq 0 \quad (15)$$

The convergence rate of FISTA for a fixed dictionary D is $\frac{K(D)}{N^2}$. Therefore:

$$F_S(z_N(D^*), D^*) - F_S(z(D^*), D^*) \leq \frac{K(D^*)}{N^2} \quad (16)$$

Hence:

$$G_N(D_N^*) - G(D^*) \leq \frac{K(D^*)}{N^2} \quad (17)$$

The proof is similar for Analysis. The convergence rate of Condat-Vu is $O(\frac{1}{N})$.

B.2 Proof of Proposition 2.2

Proposition 2.2. *Let $D \in \mathbb{R}^{n \times L}$. Then, there exists a constant $L > 0$ such that for every number of iterations N*

$$\|g_N^1 - g^*\| \leq L \|z_N(D) - z^*(D)\| \quad (6)$$

We have

$$F_S(z, D) = \frac{1}{2} \|ADz - y\|_2^2 + \lambda \|z\|_1 \quad (18)$$

$$\nabla_2 F_S(z, D) = A^\top (ADz - y) z^\top \quad (19)$$

$z_0(D) = 0$ and the iterates $(z_N(D))$ converge towards $z^*(D)$. Hence, they are contained in a closed ball around $z^*(D)$. As $\nabla_2 F_S(\cdot, D)$ is continuously differentiable, it is locally Lipschitz on this closed ball and there exists a constant $L(z^*(D))$ depending on $z^*(D)$ such that

$$\|g_N^1 - g^*\| = \|\nabla_2 F_S(z_N(D), D) - \nabla_2 F_S(z^*(D), D)\| \quad (20)$$

$$\leq L(z^*(D)) \|z_N(D) - z^*(D)\| \quad (21)$$

B.3 Proof of Proposition 2.3.

Proposition 2.3. *Let $D \in \mathbb{R}^{n \times L}$. Let S^* be the support of $z^*(D)$, S_N be the support of z_N and $\tilde{S}_N = S_N \cup S^*$. Let $R(J, \tilde{S}) = J^+(\nabla_{z,z}^2 f(z^*, D) \odot \mathbb{1}_{\tilde{S}}) + \nabla_{D,z}^2 f(z^*, D) \odot \mathbb{1}_{\tilde{S}}$. Then there exists a constant $L > 0$ and a sub-sequence of (F)ISTA iterates $z_{\phi(N)}$ such that for all $N \in \mathbb{N}$:*

$$\begin{aligned} \exists g_{\phi(N)}^2 \in \nabla_D f(z_{\phi(N)}, D) + J_{\phi(N)}^+ \left(\nabla_z f(z_{\phi(N)}, D) + \lambda \partial_{\|\cdot\|_1}(z_{\phi(N)}) \right) \text{ s.t. :} \\ \left\| g_{\phi(N)}^2 - g^* \right\| \leq \left\| R(J_{\phi(N)}, \tilde{S}_{\phi(N)}) \right\| \|z_{\phi(N)} - z^*\| + \frac{L}{2} \|z_{\phi(N)} - z^*\|^2 \end{aligned} \quad (8)$$

We have

$$g_N^2(D) \in \nabla_2 f(z_N(D), D) + J_N^+ \left(\nabla_1 f(z_N(D), D) + \lambda \partial_{\|\cdot\|_1}(z_N) \right) \quad (22)$$

We adapt equation (6) in Ablin et al. [2020]

$$g_N^2 = g^* + R(J_N, \tilde{S}_N)(z_N - z^*) + R_N^{D,z} + J_N^+ R_N^{z,z} \quad (23)$$

where

$$R(J, \tilde{S}) = J^+ \left(\nabla_{z,z}^2 f(z^*, D) \odot \mathbb{1}_{\tilde{S}} \right) + \nabla_{D,z}^2 f(z^*, D) \odot \mathbb{1}_{\tilde{S}} \quad (24)$$

$$R_N^{D,z} = \nabla_D f(z_N, D) - \nabla_D f(z^*, D) - \nabla_{D,z}^2 f(z^*, D)(z_N - z^*) \quad (25)$$

$$R_N^{z,z} \in \nabla_z f(z_N, D) + \lambda \partial_{\|\cdot\|_1}(z_N) - \nabla_{z,z}^2 f(z^*, D)(z_N - z^*) \quad (26)$$

As z_N and z^* are on \tilde{S}_N

$$\nabla_{D,z}^2 f(z^*, D)(z_N - z^*) = \left(\nabla_{D,z}^2 f(z^*, D) \odot \mathbb{1}_{\tilde{S}_N} \right) (z_N - z^*) \quad (27)$$

$$J^+ \left(\nabla_{z,z}^2 f(z^*, D)(z_N - z^*) \right) = J^+ \left(\nabla_{z,z}^2 f(z^*, D) \odot \mathbb{1}_{\tilde{S}_N} (z_N - z^*) \right) \quad (28)$$

As stated in Proposition 2.2, $\nabla_D f(\cdot, D)$ is locally Lipschitz, and $R_N^{D,z}$ is the Taylor rest of $\nabla_D f(\cdot, D)$. Therefore, there exists a constant $L_{D,z}$ such that

$$\forall N \in \mathbb{N}, \left\| R_N^{D,z} \right\| \leq \frac{L_{D,z}}{2} \|z_N(D) - z^*(D)\|^2 \quad (29)$$

We know that $0 \in \nabla_z f(z^*, D) + \lambda \partial_{\|\cdot\|_1}(z^*)$. In other words, $\exists u^* \in \lambda \partial_{\|\cdot\|_1}(z^*)$ s.t. $\nabla_z f(z^*, D) + u^* = 0$. Therefore we have:

$$R_N^{z,z} \in \nabla_z f(z_N, D) - \nabla_z f(z^*, D) - \nabla_{z,z}^2 f(z^*, x)(z_N - z^*) + \lambda \partial_{\|\cdot\|_1}(z_N) - u^* \quad (30)$$

Let $L_{z,z}$ be the Lipschitz constant of $\nabla_z f(\cdot, D)$. (F)ISTA outputs a sequence such that there exists a sub-sequence $(z_{\phi(N)})_{N \in \mathbb{N}}$ which has the same support as z^* . For this sub-sequence, $u^* \in \lambda \partial_{\|\cdot\|_1}(z_{\phi(N)})$. Therefore, there exists $R_{\phi(N)}^{z,z}$ such that

1. $R_{\phi(N)}^{z,z} \in \nabla_z f(z_{\phi(N)}, D) + \lambda \partial_{\|\cdot\|_1}(z_{\phi(N)}) - \nabla_{z,z}^2 f(z^*, x)(z_{\phi(N)} - z^*)$
2. $\|R_{\phi(N)}^{z,z}\| \leq \frac{L_{z,z}}{2} \|z_{\phi(N)} - z^*\|^2$

For this sub-sequence, we can adapt Proposition 2 from Ablin et al. [2020]. Let $L = L_{D,z} + L_{z,z}$, we have

$$\exists g_{\phi(N)}^2 \in \nabla_D f(z_{\phi(N)}, D) + J_{\phi(N)}(\nabla_z f(z_{\phi(N)}, D) + \lambda \partial_{\|\cdot\|_1}(z_{\phi(N)})), \text{ s.t. :} \quad (31)$$

$$\|g_{\phi(N)}^2 - g^*\| \leq \|R(J_{\phi(N)}, \widetilde{S_{\phi(N)}})\| \|z_{\phi(N)} - z^*\| + \frac{L}{2} \|z_{\phi(N)} - z^*\|^2 \quad (32)$$

B.4 Proof of Theorem 2.4.

Theorem 2.4. *At iteration $N + 1$ of ISTA, the weak Jacobian of z_{N+1} relatively to D_l , where D_l is the l -th row of D , is given by induction:*

$$\frac{\partial(z_{N+1})}{\partial D_l} = \mathbb{1}_{|z_{N+1}|>0} \odot \left(\frac{\partial(z_N)}{\partial D_l} - \frac{1}{L} \left(D_l z_N^\top + (D_l^\top z_N - y_l) Id_n + D^\top D \frac{\partial(z_N)}{\partial D_l} \right) \right) . \quad (9)$$

$\frac{\partial(z_N)}{\partial D_l}$ will be denoted by J_l^N . It converges towards the weak Jacobian J_l^* of $z^*(D)$, whose values are

$$J_l^* S^* = -(D_{:,S^*}^\top D_{:,S^*})^{-1} (D_l z^{*\top} + (D_l^\top z^* - y_l) Id_n) S^* , \quad (10)$$

on the support S^* of z^* , and 0 elsewhere. Moreover, $R(J^*, S^*) = 0$.

We start by recalling a Lemma from Deledalle et al. [2014].

Lemma B.1. *The soft-thresholding ST_μ defined by $ST_\mu(z) = \text{sgn}(z) \odot (|z| - \mu)_+$ is weakly differentiable with weak derivative $\frac{dST_\mu(z)}{dz} = \mathbb{1}_{|z|>\mu}$.*

Coordinate-wise, ISTA corresponds to the following equality:

$$z_{N+1} = ST_\mu \left(\left(I - \frac{1}{L} D^\top D \right) z_N + \frac{1}{L} D^\top y \right) \quad (33)$$

$$(z_{N+1})_i = ST_\mu \left((z_N)_i - \frac{1}{L} \sum_{p=1}^m \left(\sum_{j=1}^n D_{ji} D_{jp} \right) (z_N)_p + \frac{1}{L} \sum_{j=1}^n D_{ji} y_j \right) \quad (34)$$

The Jacobian is computed coordinate wise with the chain rule:

$$\frac{\partial(z_{N+1})_i}{\partial D_{lk}} = \mathbb{1}_{|(z_{N+1})_i|>0} \cdot \left(\frac{\partial(z_N)_i}{\partial D_{lk}} - \frac{1}{L} \frac{\partial}{\partial D_{lk}} \left(\sum_{p=1}^m \left(\sum_{j=1}^n D_{ji} D_{jp} \right) (z_N)_p \right) + \frac{1}{L} \frac{\partial}{\partial D_{lk}} \sum_{j=1}^n D_{ji} y_j \right) \quad (35)$$

Last term:

$$\frac{\partial}{\partial D_{lk}} \sum_{j=1}^n D_{ji} y_j = \delta_{ik} y_l \quad (36)$$

Second term:

$$\frac{\partial}{\partial D_{lk}} \sum_{p=1}^m \sum_{j=1}^n D_{ji} D_{jp} (z_N)_p = \sum_{p=1}^m \sum_{j=1}^n D_{ji} D_{jp} \frac{\partial(z_N)_p}{\partial D_{lk}} + \sum_{p=1}^m \sum_{j=1}^n \frac{\partial D_{ji} D_{jp}}{\partial D_{lk}} (z_N)_p \quad (37)$$

$$\frac{\partial D_{ji} D_{jp}}{\partial D_{lk}} = \begin{cases} 2D_{lk} & \text{if } j = l \text{ and } i = p = k \\ D_{lp} & \text{if } j = l \text{ and } i = k \text{ and } p \neq k \\ D_{li} & \text{if } j = l \text{ and } i \neq k \text{ and } p = k \\ 0 & \text{else} \end{cases} \quad (38)$$

Therefore:

$$\sum_{p=1}^m \sum_{j=1}^n \frac{\partial D_{ji} D_{jp}}{\partial D_{lk}} (z_N)_p = \sum_{p=1}^m (2D_{lk} \delta_{ip} \delta_{ik} + D_{li} \delta_{pk} \mathbb{1}_{i \neq k} + D_{lp} \delta_{ik} \mathbb{1}_{k \neq p}) (z_N)_p \quad (39)$$

$$= 2D_{lk} (z_N)_k \delta_{ik} + D_{li} (z_N)_k \mathbb{1}_{i \neq k} + \sum_{\substack{p=1 \\ p \neq k}}^m D_{lp} (z_N)_p \delta_{ik} \quad (40)$$

$$= D_{li} (z_N)_k + \delta_{ik} \sum_{p=1}^m D_{lp} (z_N)_p \quad (41)$$

Hence:

$$\begin{aligned} \frac{\partial (z_{N+1})_i}{\partial D_{lk}} &= \mathbb{1}_{|(z_{N+1})_i| > 0} \cdot \left(\frac{\partial (z_N)_i}{\partial D_{lk}} - \frac{1}{L} (D_{li} (z_N)_k + \right. \\ &\quad \left. \delta_{ik} \left(\sum_{p=1}^m D_{lp} (z_N)_p \right) + \sum_{p=1}^m \sum_{j=1}^n \frac{\partial (z_N)_p}{\partial D_{lk}} D_{ji} D_{jp} - \delta_{ik} y_l \right) \end{aligned} \quad (42)$$

This leads to the following vector formulation:

$$\frac{\partial (z_{N+1})}{\partial D_l} = \mathbb{1}_{|z_{N+1}| > 0} \odot \left(\frac{\partial (z_N)}{\partial D_l} - \frac{1}{L} \left(D_l z_N^\top + (D_l^\top z_N - y_l) Id_m + D^\top D \frac{\partial (z_N)}{\partial D_l} \right) \right) \quad (43)$$

On the support of z^* , denoted by S^* , this quantity converges towards the fixed point:

$$J_l^* = -(D_{:,S^*}^\top D_{:,S^*})^{-1} (D_l z^{*\top} + (D_l^\top z^* - y_l) Id_m)_{S^*} \quad (44)$$

Elsewhere, J_l^* is equal to 0. In the case of denoising, $n = m$ and Id_m can be replaced by Id_n . To prove that $R(J^*, S^*) = 0$, we use the expression given by (43)

$$J^* = \mathbb{1}_{S^*} \odot \left(J^* - \frac{1}{L} \left(\nabla_{D,z}^2 f(z^*, D_l)^\top + \nabla_{z,z}^2 f(z^*, D)^\top J^* \right) \right) \quad (45)$$

$$J^* - \mathbb{1}_{S^*} \odot J^* = \frac{1}{L} \mathbb{1}_{S^*} \odot \nabla_{D,z}^2 f(z^*, D_l)^\top + \mathbb{1}_{S^*} \odot \nabla_{z,z}^2 f(z^*, D)^\top J^* \quad (46)$$

$$0 = J^{*+} (\nabla_{z,z}^2 f(z^*, D) \odot \mathbb{1}_{S^*}) + \nabla_{D,z}^2 f(z^*, D) \odot \mathbb{1}_{S^*} \quad (47)$$

$$0 = R(J^*, S^*) \quad (48)$$

B.5 Proof of Corollary B.2.

Corollary B.2. *At iteration N of ISTA, the Jacobian of $D_l \rightarrow z_{N+1}(AD)$ where D_l is row l of D , can be written as a function of $(J_{N+1}^l(AD))_{1 \leq l \leq L}$ where $J_{N+1}^l(AD)$ is the Jacobian computed in Theorem 2.4 applied to the matrix AD :*

$$\frac{\partial (z_{N+1}(AD))}{\partial D_l} = \sum_{i=1}^L A_{i,l} J_{N+1}^i(AD)$$

$z_{N+1}(AD)$ can be written as a function of all rows of AD :

$$z_{N+1}(AD) = z_{N+1}((AD)_1, \dots, (AD)_L) \quad (49)$$

Then the chain rule can be applied as follows:

$$\frac{\partial z_{N+1}(AD)}{\partial D_l} = \sum_{i=1}^L \frac{\partial (AD)_i}{\partial D_l} \frac{\partial z_{N+1}((AD)_1, \dots, (AD)_L)}{\partial x_i} \quad (50)$$

$$= \sum_{i=1}^L A_{i,l} \frac{\partial z_{N+1}((AD)_1, \dots, (AD)_L)}{\partial x_i} \quad (51)$$

$$= \sum_{i=1}^L A_{i,l} J_{N+1}^i(AD) \quad (52)$$

B.6 Proof of Proposition 2.5 and Corollary 2.6

Proposition 2.5. *Let N be the number of iterations and K be the back-propagation depth. We assume that $\forall n \geq N - K$, $S^* \subset S_n$. Let $\bar{E}_N = S_n \setminus S^*$, let L be the largest eigenvalue of $D_{:,S^*}^\top D_{:,S^*}$, and let μ_n be the smallest eigenvalue of $D_{:,S_n}^\top D_{:,S_{n-1}}$. Let $B_n = \left\| P_{\bar{E}_n} - D_{:, \bar{E}_n}^\top D_{:, S^*}^\dagger P_{S^*} \right\|$, where P_S is the projection on \mathbb{R}^S and D^\dagger is the pseudo-inverse of D . We have*

$$\|J_l^N - J_l^*\| \leq \prod_{k=1}^K \left(1 - \frac{\mu_{N-k}}{L}\right) \|J_l^*\| + \frac{2}{L} \|D_l\| \sum_{k=0}^{K-1} \prod_{i=1}^k \left(1 - \frac{\mu_{N-i}}{L}\right) \left(\|z_l^{N-k} - z_l^*\| + B_{N-k} \|z_l^*\| \right).$$

We denote by G the matrix $(I - \frac{1}{L} D^\top D)$. For z_N with support S_N and z^* with support S^* , we have with the induction in Theorem 2.4

$$J_{l,S_N}^N = (GJ_l^{N-1} + u_l^{N-1})_{S_N} \quad (53)$$

$$J_{l,S^*}^* = (GJ_l^* + u_l^*)_{S^*} \quad (54)$$

where $u_l^N = -\frac{1}{L}(D_l z_N^\top + (D_l^\top z_N - y_l) Id_m)$ and the other terms on \bar{S}_N and \bar{S}^* are 0.

We can thus decompose their difference as the sum of two terms, one on the support S^* and one on this complement $\bar{E}_N = S_N \setminus S^*$

$$J_l^* - J_l^N = (J_l^* - J_l^N)_{S^*} + (J_l^* - J_l^N)_{\bar{E}_N}.$$

Recall that we assume $S^* \subset S_N$, and let's study the terms separately on S^* and $\bar{E}_N = S_N \setminus S^*$. These two terms can be decompose again to constitute a double recursion system,

$$(J_l^N - J_l^*)_{S^*} = G_{S^*} (J_l^{N-1} - J_l^*) + (u_l^{N-1} - u_l^*)_{S^*} \quad (55)$$

$$= G_{S^*,S^*} (J_l^{N-1} - J_l^*)_{S^*} + G_{S^*,\bar{E}_{N-1}} (J_l^{N-1} - J_l^*)_{\bar{E}_{N-1}} + (u_l^{N-1} - u_l^*)_{S^*}, \quad (56)$$

$$(J_l^N - J_l^*)_{\bar{E}_N} = (J_l^N)_{\bar{E}_N} = G_{\bar{E}_N} (J_l^{N-1} - J_l^*) + G_{\bar{E}_N,S^*} J_l^* + (u_l^{N-1})_{\bar{E}_N} \quad (57)$$

$$= G_{\bar{E}_N,S^*} (J_l^{N-1} - J_l^*)_{S^*} + G_{\bar{E}_N,\bar{E}_{N-1}} (J_l^{N-1} - J_l^*)_{\bar{E}_{N-1}} + (u_l^{N-1} - u_l^*)_{\bar{E}_N} + \left((u_l^*)_{\bar{E}_N} - D_{:, \bar{E}_N}^\top D_{:, S^*} (D_{:, S^*}^\top D_{:, S^*})^{-1} (u_l^*)_{S^*} \right).$$

We define as $\mathcal{P}_{S_N, \bar{E}_N}$ the operator which projects a vector from $olsiE_N$ on $(S_N, olsiE_N)$ with zeros on S_N . As $S^* \cup \bar{E}_N = S_N$, we get by combining these two expressions,

$$(J_l^N - J_l^*)_{S_N} = G_{S_N, S_{N-1}} (J_l^{N-1} - J_l^*)_{S_{N-1}} + (u_l^{N-1} - u_l^*)_{S_N} \quad (59)$$

$$+ \mathcal{P}_{S_N, \bar{E}_N} \left((u_l^*)_{\bar{E}_N} - D_{:, \bar{E}_N}^\top D_{:, S^*} (D_{:, S^*}^\top D_{:, S^*})^{-1} (u_l^*)_{S^*} \right)$$

Taking the norm yields to the following inequality,

$$\|J_l^N - J_l^*\| \leq \|G_{S_N, S_{N-1}}\| \|J_l^{N-1} - J_l^*\| + \|u_l^{N-1} - u_l^*\| + \left\| (u_l^*)_{\bar{E}_N} - D_{:, \bar{E}_N}^\top D_{:, S^*} (D_{:, S^*}^\top D_{:, S^*})^{-1} (u_l^*)_{S^*} \right\|. \quad (60)$$

Denoting by μ_N the smallest eigenvalue of $D_{:, S_N}^\top D_{:, S_{N-1}}$, then $\|G_{S_N, S_{N-1}}\| = (1 - \frac{\mu_N}{L})$ and we get that

$$\|J_l^N - J_l^*\| \leq \prod_{k=1}^K \left(1 - \frac{\mu_{N-k}}{L}\right) \|J_l^{N-K} - J_l^*\| + \sum_{k=0}^{K-1} \prod_{i=1}^k \left(1 - \frac{\mu_{N-i}}{L}\right) \left(\|u_l^{N-k} - u_l^*\| + \left\| (u_l^*)_{\bar{E}_{N-k}} - D_{:, \bar{E}_{N-k}}^\top D_{:, S^*}^\dagger (u_l^*)_{S^*} \right\| \right). \quad (61)$$

The back-propagation is initialized as $J_i^{N-K} = 0$. Therefore $\|J_i^{N-K} - J_i^*\| = \|J_i^*\|$. Moreover $\|u_i^{N-k} - u_i^*\| \leq \frac{2}{L} \|D_l\| \|z_i^{N-k} - z_i^*\|$. Finally, $\left\| (u_i^*)_{\bar{E}_{N-k}} - D_{:, \bar{E}_{N-k}}^\top D_{:, S^*}^{\dagger\top} (u_i^*)_{S^*} \right\|$ can be rewritten with projection matrices $P_{\bar{E}_{N-k}}$ and P_{S^*} to obtain

$$\left\| (u_i^*)_{\bar{E}_{N-k}} - D_{:, \bar{E}_{N-k}}^\top D_{:, S^*}^{\dagger\top} (u_i^*)_{S^*} \right\| \leq \left\| P_{\bar{E}_{N-k}} u_i^* - D_{:, \bar{E}_{N-k}}^\top D_{:, S^*}^{\dagger\top} P_{S^*} u_i^* \right\| \quad (62)$$

$$\leq \left\| P_{\bar{E}_{N-k}} - D_{:, \bar{E}_{N-k}}^\top D_{:, S^*}^{\dagger\top} P_{S^*} \right\| \|u_i^*\| \quad (63)$$

$$\leq \left\| P_{\bar{E}_{N-k}} - D_{:, \bar{E}_{N-k}}^\top D_{:, S^*}^{\dagger\top} P_{S^*} \right\| \frac{2}{L} \|D_l\| \|z_i^*\| \quad (64)$$

Let $B_{N-k} = \left\| P_{\bar{E}_{N-k}} - D_{:, \bar{E}_{N-k}}^\top D_{:, S^*}^{\dagger\top} P_{S^*} \right\|$. We have

$$\|J_i^N - J_i^*\| \leq \prod_{k=1}^K \left(1 - \frac{\mu_{N-k}}{L}\right) \|J_i^*\| + \frac{2}{L} \|D_l\| \sum_{k=0}^{K-1} \prod_{i=1}^k \left(1 - \frac{\mu_{N-i}}{L}\right) \left(\|z_i^{N-k} - z_i^*\| + B_{N-k} \|z_i^*\| \right). \quad (65)$$

We now suppose that the support is reached at iteration $N - s$, with $s \geq K$. Therefore, $\forall n \in [N - s, N]$ $S_n = S^*$. Let $\Delta_n = F_S(z_n, D) - F_S(z^*, D) + \frac{L}{2} \|z_n - z^*\|^2$. On the support, F_S is a μ -strongly convex function and the convergence rate of (z_N) is

$$\|z^* - z_N\| \leq \left(1 - \frac{\mu}{L}\right)^s \frac{2\Delta_{N-s}}{L} \quad (66)$$

Thus, we obtain

$$\|J_i^N - J_i^*\| \leq \prod_{k=1}^K \left(1 - \frac{\mu_{N-k}}{L}\right) \|J_i^*\| \quad (67)$$

$$+ \frac{2}{L} \|D_l\| \sum_{k=0}^{K-1} \prod_{i=1}^k \left(1 - \frac{\mu_{N-i}}{L}\right) \left(\|z_i^{N-k} - z_i^*\| + B_{N-k} \|u_i^*\| \right) \leq \prod_{k=1}^K \left(1 - \frac{\mu_{N-k}}{L}\right) \|J_i^*\| \quad (68)$$

$$+ \frac{2}{L} \|D_l\| \sum_{k=0}^{s-1} \left(1 - \frac{\mu}{L}\right)^k \left(\|z_i^{N-k} - z_i^*\| \right) + \frac{2}{L} \|D_l\| \left(1 - \frac{\mu}{L}\right)^s \sum_{k=s-1}^{K-1} \prod_{i=s-1}^k \left(1 - \frac{\mu_{N-i}}{L}\right) \left(\|z_i^{N-k} - z_i^*\| + B_{N-k} \|u_i^*\| \right) \leq \prod_{k=1}^K \left(1 - \frac{\mu_{N-k}}{L}\right) \|J_i^*\| \quad (69)$$

$$+ \frac{2}{L} \|D_l\| \sum_{k=0}^{s-1} \left(1 - \frac{\mu}{L}\right)^k \left(1 - \frac{\mu}{L}\right)^{s-1-k} \frac{2\Delta_{N-s}}{L} + \frac{2}{L} \|D_l\| \left(1 - \frac{\mu}{L}\right)^s \sum_{k=s-1}^{K-1} \prod_{i=s-1}^k \left(1 - \frac{\mu_{N-i}}{L}\right) \left(\|z_i^{N-k} - z_i^*\| + B_{N-k} \|u_i^*\| \right) \leq \prod_{k=1}^K \left(1 - \frac{\mu_{N-k}}{L}\right) \|J_i^*\| \quad (70)$$

$$+ \|D_l\| \left(1 - \frac{\mu}{L}\right)^{s-1} s \frac{4\Delta_{N-s}}{L^2} + \frac{2}{L} \|D_l\| \left(1 - \frac{\mu}{L}\right)^s \sum_{k=s-1}^{K-1} \prod_{i=s-1}^k \left(1 - \frac{\mu_{N-i}}{L}\right) \left(\|z_i^{N-k} - z_i^*\| + B_{N-k} \|u_i^*\| \right) \quad (71)$$

Corollary 2.6. Let $\mu > 0$ be the smallest eigenvalue of $D_{:,S^*}^\top D_{:,S^*}$. Let $K \leq N$ be the back-propagation depth and let $\Delta_N = F_S(z_N, D) - F_S(z^*, D) + \frac{L}{2} \|z_N - z^*\|$. Suppose that $\forall n \in [N - K, N]; S_n \subset S^*$. Then, we have

$$\|J_l^* - J_l^N\| \leq (1 - \frac{\mu}{L})^K \|J_l^*\| + K(1 - \frac{\mu}{L})^{K-1} \|D_l\| \frac{4\Delta_{N-K}}{L^2}.$$

The term $\frac{2}{L} \|D_l\| (1 - \frac{\mu}{L})^s \sum_{k=s-1}^{K-1} \prod_{i=s-1}^k (1 - \frac{\mu_{N-i}}{L}) (\|z_l^{N-k} - z_l^*\| + B_{N-k} \|(u_l^*)\|)$ vanishes when the algorithm is initialized on the support. Otherwise, it goes to 0 as $s, K \rightarrow N$ and $N \rightarrow \infty$ because $\forall n > N - s, \mu_n = \mu < 1$.

C Iterative algorithms for sparse coding resolution.

ISTA. Algorithm to solve $\min_z \frac{1}{2} \|y - ADz\|_2^2 + \lambda \|z\|_1$

Algorithm 1 ISTA

y, A, D, λ, N
 $z_0 = 0, n = 0$
 Compute the Lipschitz constant L of $(AD)^\top(AD)$
while $n < N$ **do**
 $u_{n+1} \leftarrow z_n - \frac{1}{L}(AD)^\top((AD)z_n - y)$
 $z_{n+1} \leftarrow ST_{\frac{\lambda}{L}}(u_{n+1})$
 $n \leftarrow n + 1$
end while

FISTA. Algorithm to solve $\min_z \frac{1}{2} \|y - ADz\|_2^2 + \lambda \|z\|_1$

Algorithm 2 FISTA

y, A, D, λ, N
 $z_0 = x_0 = 0, n = 0, t_0 = 1$
 Compute the Lipschitz constant L of $(AD)^\top(AD)$
while $n < N$ **do**
 $u_{n+1} \leftarrow z_n - \frac{1}{L}(AD)^\top((AD)z_n - y)$
 $x_{n+1} \leftarrow ST_{\frac{\lambda}{L}}(u_{n+1})$
 $t_{n+1} \leftarrow \frac{1 + \sqrt{1 + 4t_n^2}}{2}$
 $z_{n+1} \leftarrow x_{n+1} + \frac{t_n - 1}{t_{n+1}}(x_{n+1} - x_n)$
 $n \leftarrow n + 1$
end while

Condat-Vu. Algorithm to solve $\min_x \frac{1}{2} \|y - Ax\|_2^2 + \lambda \|\Gamma^T x\|_1$

Algorithm 3 Condat-Vu

y, A, Γ, λ, N
 (τ, σ) s.t. $\frac{1}{\tau} - \sigma \|\Gamma^\top\|^2 \geq \frac{\|A\|^2}{2}$
 $p_0 = 0, d_0 = 0, n = 0$
while $n < N$ **do**
 $p_{n+1} \leftarrow p_n - \tau A^\top(Ap_n - y) - \tau \Gamma d_n$
 $d_{n+1} \leftarrow \text{prox}_{\sigma \lambda (\|\cdot\|_1^*)}(d_n + \sigma \Gamma^\top(2p_{n+1} - p_n))$
 $n \leftarrow n + 1$
end while
